TUCK SCHOOL OF BUSINESS AT DARTMOUTH
GLASSMEYER/MCNAMEE CENTER FOR DIGITAL STRATEGIES
Case #6-0009

EMC: Creating a Storage-centric World

While able to navigate technical and business roadblocks throughout the prior decade, by 2002, EMC Corporation faced a new set of challenges to its business model. This new set of challenges included the rise of distributed networking, whereby storage would become less centralized, in addition to the precipitous fall in demand due to the slowing global economy. Could EMC rely on its secrets for success from the past to propel its growth well into the future?

EMC History

Formerly college roommates, Richard Egan and Roger Marino (the "E" and "M" in EMC) founded EMC in 1979, as a supplier of add-on memory boards. However, by 1989, the company moved into the information storage market. In 1990, EMC introduced a product line specifically designed to provide storage systems based on an array of small, commodity hard disk drives for the mainframe market.[1] In 1995, the company created the first storage system capable of simultaneously supporting all the major computer operating systems and thus pioneered so-called “open storage.”

This system was the first version of a storage area network (SAN; see Appendix A). The shift to SANs significantly changed the storage industry, as companies moved from just producing hard drive arrays to creating networks and software to manage storage. The innovation’s success was boosted by the Internet boom, which generated both enterprise and public data that required the efficiency and security provided by SANs.

As storage became a more critical part of an IT system, total cost of ownership (TCO) became a clearer driver in differentiating the value of the different network options. The TCO concept engendered evaluation of how much it would cost to operate the storage system. A June 19, 2001, study by Merrill Lynch and McKinsey & Company estimated the

[1] Wood, Leslie. “EMC – A Company Profile,” *Network Storage Forum*, January 2002. http://www.enterprisestorageforum.com/technology/features/article/0,,10564_951881,00.html

This mini-case was prepared by Jonathan Kwoh (T'02—MBA Fellow, Center for Digital Strategies) of the Tuck School of Business at Dartmouth under the supervision of Visiting Assistant Professor Melissa M. Appleyard. It was written as a basis for class discussion and not to illustrate effective or ineffective management practices. The authors gratefully acknowledge the support of the Glassmeyer/McNamee Center for Digital Strategies, which funded the development of this case. CDS Case #02011. Version: January, 2002.

following costs for different storage architectures: $0.84/Megabyte (MB) for Direct Attached Storage (DAS), $0.38/MB for Storage Area Network (SAN), and $0.35/MB for Network Attached Storage (NAS).[2] (See Appendix A for a comparison of these options.) The results stood in contrast to surveys that indicated that IT executives were reluctant to move from DAS to SAN due to the associated costs. Don Swatik, Vice President of EMC's Global Alliances, observed that EMC used TCO analysis when conveying the benefits of EMC's systems to customers.[3]

EMC CEO Michael Ruettgers claimed in a May 2000 speech that the storage elasticity of demand was 4 to 1, meaning that a 25% drop in cost would increase customer demand by 100%—in effect doubling demand. EMC rode the storage economics to create incredible revenue and profit growth.

EMC's Change in Technology Focus over Time

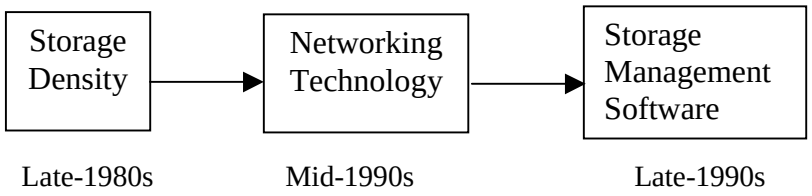

Storage Convergence

"Infinite storage" complemented by "infinite bandwidth" was the foundation of a storage-centric future envisioned by EMC. The "infinite storage" thesis was based on the rapid increase in storage capacity. Magnetic storage density had growth at 60% a year through the 1990s. Exploring additional alternatives, EMC invested in such future technologies as holographic, molecular, and optical storage. Jim Rothnie, the CTO of EMC, observed that in 2000, the cost of storage was 30 cents per MB and was anticipated to drop to a penny per MB in 2005.[4]

The "infinite bandwidth" claim was driven by the incredible growth in fiber capacity and the potential of ubiquitous broadband last-mile access. According to Rothnie, the economies of scale drove towards consolidation of data with the only barrier being the cost and bandwidth availability of the network. Thus, broadband introduction would likely create rapid consolidation of data into large storage warehouses. Swatik compared the relationship of storage and bandwidth to the microprocessor and operating system, claiming that

[2] *Computer Technology News*. Speech by Jim Rothnie, July 2001. www.wwpi.com/lead_stories/01_07_30/Rothnie.html
[3] Author interview with Don Swatik, EMC Vice President of Global Alliances, January 23, 2002.
[4] *Computer Technology News*, op. cit.

improvements in one element would drive development of the other element to make use of the additional capability.

This trend of communications and storage consolidation affected the hardware that customers demanded. In 2000, Reuttgers projected that by 2003-2005, storage would be accessed from the network rather than included inside every computer.[5] Such a future would play to EMC's strong suit, as a more consolidated storage system with distributed access would require the powerful storage management software that EMC was a leader in. By 2000, information storage software constituted $1.44 billion of the company's $8.9 billion annual revenue.[6] EMC's product offerings in 2002, highlighted the company's focus on this vision. Their Automated Information Storage concept represented the kernel for an automated storage management software platform that would coordinate access to stored data. The complementary E-Infostructure Developers program was focused on ensuring the system had the interoperability to work across all vendors' products. By early 2002, setbacks in the telecommunications sector were expected to slow the proliferation of information, requiring EMC to reassess demand projections for its products and services.

Competitors

By 2002, EMC's competitors included major server manufacturers like IBM and Sun as well as direct storage companies like Hitachi. On the software side, pure-play storage management software companies like Veritas posed a challenge. Software accounted for approximately 20% of EMC's revenues, and EMC expected software to grow to 30% of sales.

Challengers

Developments in the storage industry by 2002, indicated that interactions between storage economics and bandwidth capacity were already at work. The following examples demonstrate the trends that had emerged.

Storage Service Providers

In 1998, StorageNetworks pioneered the Storage Service Provider (SSP) concept. By 2002, over fifty SSPs were in operation. The idea behind SSPs was that a single storage array for multiple customers would exploit economies of scale in terms of maintenance as well as enhanced reliability. Implicitly, these SSPs were promoting the very gains described by EMC's strategists.

The SSP business model had not come to dominate the storage market by 2001. SSPs appealed to small companies like dot-coms that had massive data needs but didn't have the resources to handle the data. Thus, the failure of such Internet companies hurt the growth of

[5] *The Exchange*. "Storage Moves Front and Center," Perspective Column, May 2000.
[6] EMC company web site, http://www.emc.com/about/corp_profile/index.jsp?openfolder=all, January 27, 2002.

the SSP business. Because of their ability to capture gains from economies of scale internally and their concerns regarding security and dependability, large enterprises did not subscribe to SSPs.[7]

SSPs tended to partner with "collocation facilities" (see Appendix A), using them as junctions for customer access to the storage arrays as well as physical location of the arrays. There was a potential threat that the "colos" would subsume storage as they moved from managing the physical facility to managing the whole service package.

Content Delivery Networks

Content Delivery Networks (CDNs) bucked the trend toward data consolidation. The idea behind CDNs was that download performance could be improved by locally storing data at ISPs. By 2002, two of the big players in CDN were Akamai Technologies and Digital Island. The need for CDNs came about because of the insufficient bandwidth in the Internet backbone. If the "infinite bandwidth" espoused by EMC CTO Jim Rothnie were to arrive, then the need for CDNs would be expected to decline.

The CDN companies have exhibited little concern regarding the possibility of free bandwidth. Akamai's Avi Freedman, vice president and chief network architect, noted in 2001 that "Sure, there is plenty of fiber in the backbone, including dark fiber that no one is using. But that misses the point. The Internet is a screwed-up place ever since it stated going commercial… Backbone networks don't connect to each other well because they don't like each other. [As a result,] the problem of universal good connectivity is going to be with us for many years."[8]

Freedman's quote alluded to the peering arrangements that allowed backbone providers to exchange data across their networks. The arrangements worked poorly in part because the backbone providers touted network performance as a differentiator, which discouraged them from working toward perfectly smooth exchanges.

Overall Trends

Growth in the storage sector grew 120% per year from 1998 to 2000. In 2001, Computer Technology Review projected growth in the 50% range.[9] This projected slowdown in the sector's growth rate would therefore be expected to dampen EMC's own growth.

On the other hand, other developments could potentially increase demand for high-end storage solutions. Even though the dot-com bubble deflated, a 2001 study conducted by the School of Information Management and Systems (SIMS) at the University of California,

[7] Schmelling, Sarah. "Storage spurs the latest gold rush," *Upside Magazine*, January 2001. http://www.upside.com/texis/mvm/story?id=3a2d52da2f0
[8] Draenos, Stan. "Web storage will keep data moving," *Upside Magazine*, January 2001. http://www.upside.com/texis/mvm/story?id=3a3129645
[9] *Computer Technology Review* (2001). "Total Cost of Ownership Will Drive Storage Bottom Line: Look for these 10 Happenings in 2002," November.

Berkeley,[10] predicted that humankind would generate more information between 2001 and 2003 than was created in the previous 300,000 years combined. In 2001, six exabytes (EB= 1 billion MB) was generated worldwide. Expansion of data volumes would require increased sophistication in data management. In addition to adeptly managing ever-increasing data volumes, customers required features like data recovery from their data management systems. The terrorist attacks on September 11, 2001, brought disaster recovery to the forefront of company planning. Swatik noted that "Information is their [customers'] core asset. Information is becoming the heart and soul of every business," and emphasized that the capabilities exhibited by EMC solutions have grown in importance.[11]

Who Will Manage Storage?

In this storage-centric world, EMC may make the tools to manage the storage, but another company may run the networked storage service and yet other companies, the backbone owners, would be ultimately responsible for guaranteeing that data travel through the networks correctly.

Several other possible players may stand between EMC and dominance of the storage market. Economies of scale that accompany the centralization of storage encouraged the formation of large collocation facilities, which wield their own market power. By 2002, collocation facilities already were expanding their capabilities to provide networked storehouses.

Other likely players were the incumbent local exchange carriers (ILECs, for example the former Bell operating companies) who could possibly add storage as a service and have the heft to create economies of scale. Generally, telephone networks were not suited for storage; however, the efforts to map Fibre Channel to IP and Ethernet protocols may allow usage of the existing networks. If Ethernet-compatible standards take off, Yipes and other IP / Ethernet network providers could play a big role in storage networking.

Overall by early 2002, the question remained as to whether EMC's growth could be maintained. With the slowdown in both the storage and telecommunications sectors throughout 2001, EMC reflected on the time horizon required to fulfill its vision of a "storage-centric" world.

[10] Wood, op. cit.
[11] Fisher, Andrew. "Understanding Storage," *Financial Times*, November 23, 2001. http://specials.ft.com/understandingstorage/FT33ELQBYTC.html

Appendix A: Storage Technology Primer and EMC's Offerings

Storage Architectures

Prior to 1995, storage was attached directly to a server (described as Directly Attached Storage (DAS) configuration) and the business of storage was focused on shrinking hard drive size. In 1995, EMC pioneered the creation of networks specifically designed to support transfer of storage data between servers. This configuration, known as Storage Area Network (SAN), led to the creation of a network that was optimized for connecting multiple servers and providing long distance storage transfer capabilities. This allowed for better protection of data as well as better storage utilization as data could be stored across multiple storage devices as opposed a single unit. If a server were on a SAN, the files managed on the server would be backed up and distributed across several hard drives, allowing a more efficient use of hard drive space as data from other servers would be stored there as well. When accessing the data, the user would not know where the data had been actually stored.

EMC's focused on providing SAN-compatible servers and SAN management software, but the storage network not only required servers but also SAN-compatible network switches to make the SAN run efficiently. By 2002, Brocade had 40% of the SAN-compatible network switch market and McData (spunoff from EMC) was a trailing competitor.

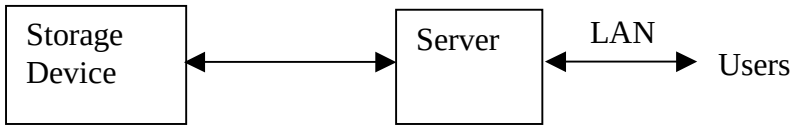

Directly Attached Storage

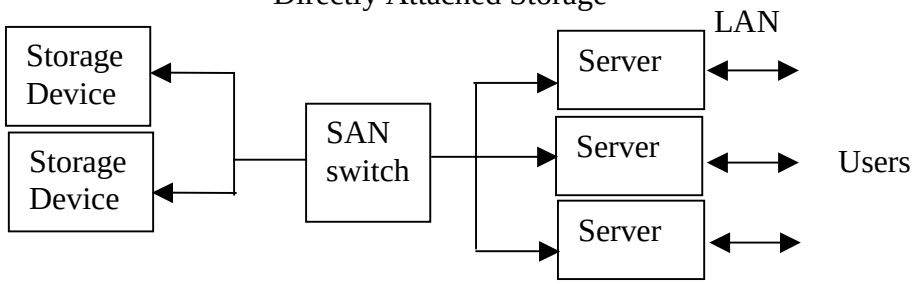

Storage Area Networks

The main alternative to SAN was Networked Attached Storage (NAS). Where SAN was optimized to transfer data between many storage devices, the NAS architecture was designed to allow sharing of data between many end-user terminals attached to the network. A NAS device would be assigned an IP address and thus absorbs some of the file management functionality of a traditional server. By 2002, the NAS market was much smaller than the SAN market, with the NAS products being both less expensive and less sophisticated. In December of 2000, EMC introduced the CLARiiON product line to challenge Network

Applicance's 40% share of the NAS market. EMC VP Don Swatik believed that the distinction between NAS and SAN would be eliminated as EMC's SAN products increasingly incorporated NAS-like capabilities like file-sharing and easy attachment. He noted that EMC's terminology for its product offerings migrated to the new label of "network information storage" in order to emphasize that EMC's product line had moved beyond the SAN/NAS categorization.

Standards

Like everything network-related, standards were a critical technology consideration. The leading transmission standard for SAN was Fibre Channel architecture. Fibre Channel enabled high-speed transmission of data files over distances of up to six miles. The drawback with Fibre Channel was that the standard was not directly compatible with other traditional networking protocols like IP. Thus, access to a Fibre Channel SAN was typically done via a "collocation facility," where an interface between the two standards could easily be handled. EMC was one of the leaders in crafting the Fibre Channel standards.

In the SAN segment, there were other important standards in development. iSCSI "serialized" the SCSI interface and allowed the use of IP for transmission of storage data files. For example, Nishan Systems developed a Storage over IP (SoIP) solution that eliminated incompatibility issues.[12]

These developments indicated that the capabilities of SAN and NAS were converging to create an IP-friendly SAN capability with the NAS-like data-sharing capabilities. In early 2002, EMC's Swatik dismissed the "hype of IP derivatives [standards]," noting that Fibre Channel was currently the only dependable SAN standard.[13] EMC dealt with the IP issue by making its equipment capable of converting data from Fibre Channel to IP for very long distance transmission.

[12] Schmelling, op.cit.
[13] Author interview with Don Swatik, EMC Vice President of Global Alliances, January 23, 2002.