CHAPTER 1

INDUSTRY LANDSCAPE: STORAGE COSTS AND CONSUMPTION

Just as the effects of the recent economic downturn have been universally felt across all sectors and industries, likewise do the principle concepts discussed in this chapter—the commoditization of hardware and storage utilization efficiencies—apply to all IT environments, regardless of the size or the nature of the business application. This chapter sets the stage for understanding the storage network as a value-add to the firm insofar as it is capable of alleviating the management and financial burdens associated with direct-attached storage (DAS).

Networked storage offers significant business advantages over DAS, and the impact of these benefits can be quantified and measured. To understand the nature of the business benefits of networked storage, a brief, general discussion of overall IT spending and the specifics of storage spending is required and provides a basis for the remainder of the analysis performed in later chapters.

This chapter covers the following topics:

- Storage management
- Implementing a storage vision
- The commoditization of hardware
- The impact of industry trends and legislation on storage consumption
- Storage utilization, storage yield, and the Cost of Poor Quality (COPQ)

Storage Management Matters

In May, 2003, author Nicholas Carr garnered much attention with a *Harvard Business Review* article on the strategic worth of information technology. The article's provocative title, "IT Doesn't Matter," bespoke Carr's argument that the commoditization of information technology solutions has essentially depleted the strategic advantage of information technology as a whole. In "IT Doesn't Matter," Carr states succinctly, "What makes a resource truly strategic—what gives it the capacity to be the basis for a sustained competitive advantage—is not ubiquity, but scarcity."[1] Carr points to innovations, such as electricity and rail transportation, which offered competitive advantages to early adopters, but whose value diminished over time as the use of these technologies became common place.

In 2004, Carr expanded his position in his book, *Does IT Matter? Information Technology and the Corrosion of Competitive Advantage*, in which he urges

readers to decrease IT spending, to avoid being an early adopter whenever possible, and to focus on "vulnerabilities" instead of "opportunities" where critical services are at risk.[2]

Carr could not be more accurate. It is also important to understand, however, that investment in storage networks allows firms to decrease storage spending and focus on service vulnerabilities. In addition, Fibre Channel SANs are well past the early adopt phase. Investment in storage networking technologies (not just Fibre Channel, but IP-based storage solutions as well) can help companies become more efficient and therefore more competitive.

NOTE Everett Rogers originally outlined the concept of the early adopter in his work *The Diffusion of Innovations*. Detailed discussion of Rogers' work and how it applies to product adoption life cycles follows in Chapter 4, "How It Should Be Done: Implementation Strategies and Best Practices."

Understanding competitive forces is a fundamental premise of business leadership. Harvard Business School professor and author Michael Porter is a renowned expert on strategy and competition. He has written extensively on the nature of competition between rival firms and nations. Porter's groundbreaking essay, "How Competitive Forces Shape Strategy," was first published in 1979; twenty-five years later, Porter's "Five Forces," as they have come to be known, still aptly describe the interplay between rival firms' strategic endeavors.

As Porter outlined, the five main forces shaping competition between firms in similar industries are the following:

- Buyer bargaining power
- Supplier bargaining power
- The threat of substitute products
- Rivalry
- Barriers to entry

In his essay, Porter lists "economies of scale" and "cost disadvantages independent of size" as two of the major sources of "barriers to entry."[3] Although

"learning curves," "experience curves," and "economies of scale" are concepts typically applied to manufacturing environments, these concepts also have distinct applications in IT, relative to the management of IT assets, and storage assets in particular.

Without a doubt, one of the most significant vulnerabilities facing companies today is the state of enterprise storage, now in overwhelming disarray following the deployment at breakneck speed of over two million DAS external disk units worldwide between the years of 1999 and 2003. The total number of DAS versus networked storage units sold between 1999 and 2003 is shown in Table 1-1.

Table 1-1 *Worldwide External Non-OEM Factory Revenue ($M) and Shipments, 1999-2003 (Source: IDC, 2004)*[4]

Worldwide External Non-OEM Factory Revenue ($M) and Shipments, 1999-2003					
	2003	**2002**	**2001**	**2000**	**1999**
	$M	**$M**	**$M**	**$M**	**$M**
DAS	$5504	$5932	$9357	$14,452	$13,773
DAS Units	270,379	298,264	425,255	509,667	503,608
Networked*	$8087	$7165	$7838	$7299	$4368
Networked Units	128,599	141,148	140,902	138,455	74,215

*Denotes SAN and NAS storage.

Jon William Toigo outlined the storage management problem facing IT managers in his book, *The Holy Grail of Storage Management*, published in 2000. Toigo stated clearly and early on that corporate IT departments would face serious challenges in the coming years with managing data storage. The need for online or near-online data and the lack of a rational strategy for dealing with storage growth indicated that in a short amount of time, companies would have their hands full of storage problems.[5] Few in corporate IT today are in a position to disagree with Toigo.

Storage networks allow firms to drive down operational costs and increase economies of scale to remain competitive. At the same time, storage networks allow firms to address critical business vulnerabilities. Although storage networks alone do not magically solve all storage-related problems, a networked storage infrastructure does help increase operational and utilization efficiencies, which ultimately lowers the overall storage total cost of ownership (TCO).

NOTE Storage networks do not intrinsically solve the problems related to data and information management, but in later chapters I demonstrate how economies of scale with regard to storage management (and the cost advantages of increased storage utilization) have a significant impact on the firm's bottom line.

NOTE The ubiquity of information technology resources in corporate datacenters underscores the drop in prices for IT products and the diminished magnitude of the capital outlays required to build an enterprise-level IT infrastructure. This ubiquity is the tangible evidence—the hangover, if you will—from the party that heralded the advent of the New Economy.

Implementing a Storage Vision

Now, more than ever, companies are adopting storage networks as fundamental building blocks of a storage vision that addresses the capacity, utilization, and management issues related to data storage. This broader strategy or vision is designed to:

- Reduce the overhead associated with providing storage solutions
- Maximize critical business continuance capabilities
- Increase the performance and flexibility of the overall data storage infrastructure

A storage vision begins with the migration to storage networks, and proceeds with the decommissioning of DAS. A storage vision also requires the classification of environments into tiers and the creation of a service-level framework to measure the efficacy and performance of storage-related deliverables. A storage vision culminates in the ability to provide storage services in a utility-like fashion. The net effect of a storage vision is an overall lower TCO for storage.

The need for low-cost, highly-available storage solutions, coupled with the high demand for long-distance replication functionality (spurred by legislation and security concerns), has helped to increase sales of Fibre Channel (FC) storage networking and optical transport products. This increase in product sales occurred even as disk revenues fell dramatically in 2001 and 2002. The management burden of DAS and the difficulties of managing heterogeneous storage on an FC storage network have led to an increased interest in IP storage networks. It is the belief of some vendors that a strong Fibre Channel infrastructure facilitates the adoption of Internet Small Computer Systems Interface (iSCSI), Fibre Channel over IP (FCIP), and IP over Fibre Channel (IPFC).

The soft economic climate of the last three years has fostered the realization that not every application requires five-star accommodations. Application environments are now consolidated to conserve resources. Likewise, business processes are now modified to provide service-level management (SLM) frameworks that match an application's needs to the most appropriate and cost-efficient storage solution. Service-level management is an increasingly important concept in storage management, and, as shown in the case studies in Part II, "Case Studies," SLM forms the framework around which solid storage visions are currently built.

Five years ago, the typical IT department was asked to provide the most expensive server and disk solutions for every conceivable array of applications. At that time, it was customary for IT departments to provide storage capacity based on poorly scoped application requirements. Then, it was acceptable for IT to serve strictly as a cost center. Those days are over. Now, the focus is on cutting costs at a time when legislation and competition actually create new requirements and drive increased costs. In addition, data storage is growing at such a rate that cutting costs without a storage management strategy is almost impossible.

To understand the importance of a storage vision, it is necessary to look at broader trends in the market. An analysis of the overall storage and IT spending rates for the last several years is illustrative of the current storage management headache facing today's IT decision maker.

Irrational Exuberance

It is no secret that corporate spending on information technology hardware, software, and services has slowed dramatically in recent time. If the drop-off in IT spending was dramatic, the run-up previous to the decline was equally spectacular.

No doubt, times have changed and just as electronic commerce and web technologies have matured, business leaders now understand the importance of value case analysis, and are returning to Net Present Value (NPV) and return on investment (ROI) as methods for validating new IT investments.

The "irrational exuberance" in the securities markets of the late 1990s, noticed as early as 1996 by Federal Reserve Chairman Alan Greenspan, presented significant hurdles to planners, analysts, and those in charge of charting the path of the U.S. economy.[6] This exuberance was fueled in part by Y2K and in part by the multi-million dollar IT budgets burning a hole in the pockets of both Fortune 500 companies and start-ups alike. These firms together shared the collective aim of gaining both a long-term boost in productivity and a competitive edge in the marketplace. The churn-and-burn mentality of the start-ups and dot-coms led to massive capital purchases, inflating the revenues of almost every high-tech company in the value chain.

Table 1-2 clearly shows that one of the primary areas to benefit from exuberant IT spending during this time frame was external disk storage, as highlighted by the increases in vendor revenues between 1999 and 2000.

Table 1-2 *Worldwide External Disk Storage Systems Non-OEM Factory Revenue ($M) and Units, 1999-2003 (Source: IDC, 2004)*[7]

Worldwide External Disk Storage Systems Non-OEM Factory Revenue ($M) and Units, 1999-2003					
	2003	2002	2001	2000	1999
	$M	$M	$M	$M	$M
External Disk Storage Systems	$13,591	$13,097	$17,195	$21,751	$18,141
Units	398,978	439,412	566,157	648,121	577,823

These figures are sufficiently eye-opening in that they highlight the marked increase and then sudden decline in overall revenues. Aside from highlighting a precipitous drop in margins, the unit numbers in Table 1-2, coupled with the percentage of DAS sold worldwide during the same timeframe (as shown in Table 1-1), indicate that there is a mountain of DAS currently deployed.

NOTE As shown in Table 1-1, 87 percent of supplier revenues in 1999 and 78 percent of supplier revenues in 2000 were from sales of DAS solutions.

As is well documented by now, the "damn-the-ROI" mentality prevailed in IT spending until a series of events accelerated the well-known recent economic downturn.

Macro Sources of Economic Downturn

With capital spending trending downwards, many firms began to report disappointing revenues in late 2000 and early 2001. Of those reporting declines, arguably one of the most significant was Cisco Systems.

On February 7, 2001, Cisco Systems missed its quarterly earnings estimates for the first time in almost three years. Cisco Systems, technology bellwether, and long-time advocate of the virtual close-a process that allowed earnings snapshots to be retrieved at any time to provide guidance to its leadership-came up short of analysts' per share expectations for the second quarter of fiscal year 2001. The subsequent write-down of $2.2 billion worth of Cisco inventory sent shockwaves through its supply chain and had a deleterious effect across the industry.[8] Companies in many sectors questioned their capability to forecast sales and profitability, shareholders suffered, and visibility into U.S. economic recovery became even murkier.

On September 11, 2001, terrorists attacked the World Trade Center and the Pentagon, killing almost 3000 people. The New York Stock Exchange, NASDAQ, and AMEX exchanges were closed for four days. An already shaky U.S. economy found itself against the ropes, and the United States prepared for a multi-front war. Subsequent foreign intervention (in Afghanistan and eventually Iraq) dampened hopes that an economic upswing was imminent, and six cuts in the federal fund rate (one each in the remaining months of 2001 after the September 11 attacks, and one each in 2002 and 2003), shown in Figure 1-1, indicated that the Federal Open Market Committee (the Federal Reserve) saw little sign of economic revival, equally thwarting hopes of a recovery.

Figure 1-1 *Federal Fund Rate Cuts Since January 1, 2000*

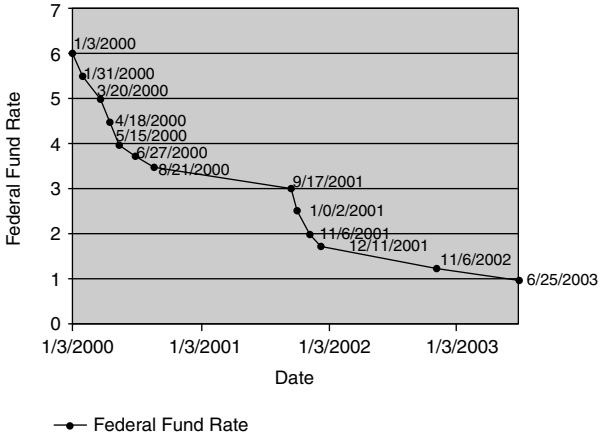

NOTE On June 30, 2004, the Federal Open Market Committee raised the federal fund rate by one-quarter of one point—its first rate hike in four years.

Data from the Bureau of Economic Analysis highlights the effect of decreased spending for electronic hardware on the U.S. Gross Domestic Product (GDP) during the period in question (shown in Figure 1-2).[9]

Figure 1-2 *Electronic and Electric Equipment Manufacturing Contribution to U.S. GDP 1992-2001*

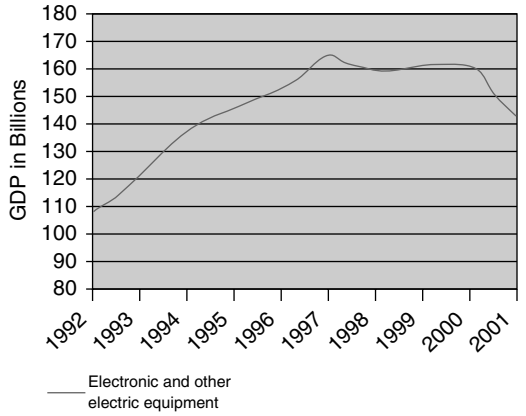

Analysis of the effects of economic growth related to the "New Economy" began in October, 2001, when the McKinsey Global Institute released its study titled "U.S. Productivity Growth 1995-2000." This study indicated that although IT spending increased between 1995 and 2000, IT was just one of several factors (including innovation, cyclical demand, and competition) contributing to U.S. productivity growth during this time frame.[10]

Some point to this "productivity paradox," as it has come to be called, as highlighting the failure of corporate spending on IT products and services to lead to a tangible increase in sustained output of U.S. companies. Although this point is debatable, what is clear is the subsequent decrease in profits for major storage and server vendors, indicating the commoditization of both the disk and the host.

NOTE Whether or not the economic downturn was officially a recession, there seems to be little doubt at this point. In July, 2003, the National Bureau of Economic Research (NBER) issued a report stating that the last U.S. recession ended in November, 2001. The NBER's Business Cycle Dating Committee, which tracks the timelines of U.S. business cycles, pinned the length of this most recent recession, which began in March, 2001, at eight months, three less than the post-World War II average of 11 months.[11]

Commoditization of Hardware

So, what were the long-term effects of the "New Economy?" This is the subject of heated debate; however, economists J. Bradford DeLong and Lawrence Summers provide us with an insight into one aspect of the "New Economy:" competition. In an address at Kansas City's Federal Reserve Bank Symposium in August, 2001, DeLong and Summers argued that the long-term effects of the technological advances of the "New Economy" would not be the creation of "scale-related cost advantages," but the creation instead of a more level playing field, making competition itself "more effective."[12]

Obviously DeLong and Summers refer primarily to the supply side of the economic equation. Accordingly, they state, "Competitive edges based on past reputations, or brand loyalty, or advertising footprints will fade away. As they do so, profit margins will fall: Competition will become swifter, stronger, more pervasive, and more nearly perfect. Consumers will gain and shareholders will lose."[13]

The commoditization of disk and server hardware is therefore a visible symptom of stronger and more perfect competition, and certainly the consumer in most circumstances benefits from increased purchasing power. The question, however, remains: Do the consumers gain a true advantage? Not if the commoditized assets are poorly utilized, which, when dealing with storage, is more often than not, the unfortunate case.

The Disk as Commodity

As noted, IT spending on the whole declined dramatically between the years 2000 and 2003, and the effect on the disk storage industry has been punishing. IT spending began to suffer in some cases as early as 1999, but the disk storage industry shows itself to be a lagging indicator of decreased corporate spending. The delayed decline of storage revenues were due in part to poor utilization efficiencies, which buffeted disk spending by forcing companies to purchase more storage.

The increased revenues for disk storage systems between the years of 1999 and 2000 (as shown in Table 1-2) were primarily due to three factors:

- Spending on Y2K-related infrastructure
- Continued demand for web and electronic commerce applications
- Increases in the number of complex enterprise resource planning and supply-chain management installations

As a major manufacturer of disk storage systems and a provider of disk-related software and services, the annual revenues for the Hopkinton, Massachusetts-based firm, EMC Corporation, provide an excellent snapshot of disk spending for the two years on either side of Y2K (shown in Figure 1-3).

Figure 1-3 *EMC Annual Revenues from 1998–2002*[14]

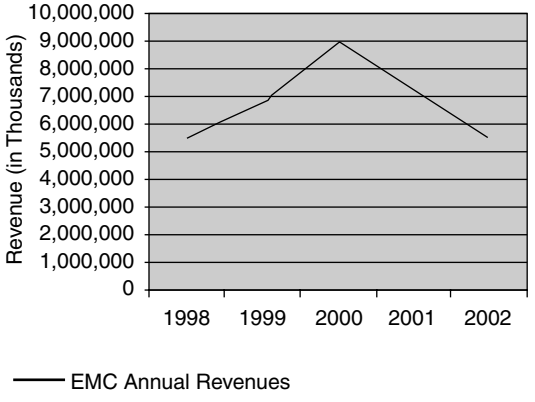

What is more germane to the premise of this discussion, however, is the breakdown of revenue by line of business at EMC Corporation from 2000 to 2003 (see Figure 1-4). These figures show at a glance the growing shift in focus from disk sales to revenue generation through software and services. This shift indicates further commoditization of disk storage.

Figure 1-4 *Percentage of EMC Revenue by Line of Business*[15]

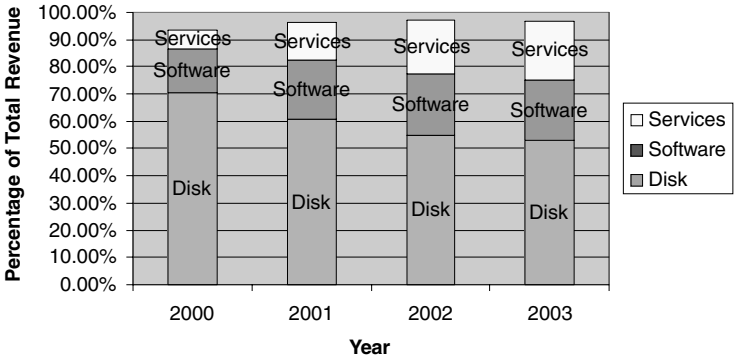

NOTE Note that the software figures exclude revenues from Legato and Documentum because the acquisition of these two firms occurred midway through the 2003 financial year (in October and December, 2003, respectively).

It is important for the purposes of this discussion not only to note the contraction in revenues and units sold for disk storage systems, as previously shown, but also to note the steady increase in sales of Fibre Channel infrastructure components as outlined in Table 1-3.

Table 1-3 *Worldwide Fibre Channel Switches and HBAs Factory Revenue ($M), 2000–2003 (Source IDC, 2004).*[16]

Worldwide Fibre Channel Switches and HBAs Factory Revenue ($M), 2000-2003				
	2003	2002	2001	2000
WW FC Switches and HBAs	$1673	$1448	$1346	$1181

Three consecutive years of growth in the Fibre Channel switch market point to a shift from DAS to SAN infrastructure and, as intelligence moves onto the storage network, the commoditization of disk storage continues.

The Server as Commodity

Server sales decreased dramatically between 2000 and 2002. The number of server units sold during this same time frame indicates the increased commoditization of the host.

Table 1-4 shows the decrease in worldwide server revenues between 2000 and 2002, echoing the trend evidenced previously in the storage numbers. These figures also show the relative plateau in the number of units shipped during this same period. The 2003 spike in revenues and units sold highlights decreased margins and the move to lower-priced, rack-mountable server platforms, further illustrating the trend toward commoditization.

Table 1-4 *Worldwide Server Factory Revenue ($M) and Units Shipped from 1999–2003 (Sources: IDC 2004 Release)*[17]

Worldwide Server Factory Revenue ($M) and Units Shipped					
	2003	2002	2001	2000	1999
	$M	$M	$M	$M	$M
WW Total	$46,131	$44,649	$50,496	$61,675	$57,708
Units	5,281,231	4,442,690	4,276,119	4,369,840	3,761,141

NOTE The increased popularity of rack-mountable servers stems from the capability to provide enterprise-class service for a fraction of the cost and the footprint. More rack-mountable server deployments, however, translate into increased datacenter costs in terms of management efficiencies.

Blade servers have the potential to alleviate the pain points associated with implementing significant numbers of rack-mountable servers. Rack-mountable servers tend to overwhelm the datacenter with network, power, cooling, and storage demands, whereas blade servers can utilize shared datacenter resources to cut back on the capacity consumed per deployment.

Additional discussion of blade servers and virtualization of disk and CPU resources follows in Chapter 5, "Maximizing Storage Investments."

It is reasonable to assume that the adoption of the Linux operating system has also contributed to the commoditization of the server.

An application infrastructure based on the Linux operating system has a significantly lower TCO than one based on a proprietary operating system. The increased customer adoption rate of Linux will continue to apply pressure on companies whose revenues are derived from sales of proprietary operating systems and enterprise-class servers.

The Impact of Competition on IT

As previously seen, a number of industry trends and current events have helped to shape IT spending over the last three years. The worldwide storage market (which includes sales of disk, software, and storage-related services) now constitutes more than $40 billion of business and, despite the commoditization of

disk storage, showed signs of recovery in 2003. Table 1-5 highlights recent increases in revenues for storage software and services.

Table 1-5 *Worldwide Storage Software and Services Market ($M), 1999–2003 (Source: IDC, 2004)*[18]

Worldwide Storage Software and Services Market ($M), 1999-2003					
	2003	2002	2001	2000	1999
Storage Software	$6621	$5730	$6157	$6113	$4640
Storage Services	$23,360	$21,171	$20,552	$19,501	$17,250

While disk sales notched a nearly four percent increase in 2003 (see Table 1-2), sales of storage software and services were more robust. Sales of storage software in 2003 increased nearly 16 percent over 2002, whereas sales of storage services increased more than 10 percent over the previous year.

As worldwide server revenues indicate, after two consecutive years of double-digit declines, server sales also experienced a slight uptick of 3 percent in 2003 (see Table 1-4).

Recent increases in IT spending are now driven by companies seeking both a competitive edge in the marketplace and compliance with new laws and regulations stemming from current events.

Despite the lack of a clear correlation between IT spending and profitability, increased competition in every sector forces companies to continue to seek ways to use IT to create a competitive advantage. The advantage of IT, as envisioned during the dot-com era, no longer exists. IT is now correctly seen as the framework around which solid business strategies are built.

These strategies are executed primarily through one or more of these four tactics:

- Electronic commerce
- Enterprise resource planning (ERP)
- Supply chain management (SCM)
- Customer relationship management (CRM)

Each of these solutions requires significant investment in storage infrastructure, server resources, and disaster-recovery capabilities.

Electronic commerce has profoundly changed the face of the retail industry. Many traditional brick-and-mortar businesses are now "brick-and-click" businesses that require around-the-clock availability. These environments demand high performance servers and terabytes of storage to house logs for millions of website hits and queries.

In August of 2003, The Economics and Statistics Administration of the U.S. Census Bureau announced an uptick in e-commerce retail spending of 27.8 percent over the previous year. These figures indicate a 214 percent increase in e-commerce retail sales since the fourth quarter of 1999.[19] Similarly, figures for the first quarter of 2004 showed an increase of 28.1 percent over the first quarter of 2003.[20]

Although consumer retail e-commerce is growing, the overwhelming majority of e-commerce transactions are still between businesses. According to data from the U.S. Census Bureau, 93.3 percent of e-commerce sales in 2001 and 93.9 percent in 2000 were business-to-business (B-to-B) transactions. This percentage, however, reflects only 7.3 percent of all traditional B-to-B and business-to-customer (B-to-C) shipments and revenues (or only $1,066 billion out of $14,572 billion) in the United States. [21] These statistics indicate there is still significant room for growth in e-commerce markets.

Similar to e-commerce environments, ERP applications (which often connect to complex e-commerce infrastructures) demand high performance and high availability, as well as replication capabilities. In addition to increased requirements for enterprise class storage, these environments also need massive amounts of storage for the frequent migration and redeployment of application code-trees.

SCM and CRM environments—although extremely different from each other in terms of business functions—have similar high availability requirements and strict performance specifications. As electronic commerce continues to show steady growth, demand planning and call-center capabilities, along with other SCM and CRM functions, become more important. Manufacturing and sales support functions, therefore, now have availability requirements comparable to those of ERP and e-commerce systems.

In addition, the discussion of the business application of IT is incomplete without some mention of email, the foundation of business communication in the 21st century. Although most emails are relatively small in size, the inclusion of attachments in emails greatly increases the burden on backup and storage

environments. Regulatory trends and recent legislation dictate the lengthening of email retention policies, which in turn increases the TCO for storage.

Finally, a surge in demand for business process management (BPM) software and ERP software add-ons designed to simplify the process of compliance with complex regulations indicates businesses are coming to terms with compliance with new legislation.

The Impact of Legislation on IT

Several new regulations and acts of legislation will likely increase the corporate data growth rate and will most assuredly change the way companies manage storage resources.

As the following examples indicate, the standard operating procedures of business are changing, and the impact of these changes on financial services and healthcare business systems will be significant, particularly on downstream functions of the storage value chain (offsite data storage, storage-related professional services, and so on).

Regulation Fair Disclosure

When the bull market of the late 1990s was capped off with the NASDAQ crash and the deflation of the Internet bubble, a sobering and humbling string of corporate scandals surfaced just in time to keep the bad news flowing. With the passage of Regulation Fair Disclosure in 2000, the SEC instituted, at least on paper, the first in a long series of efforts designed to limit the ability of the firm and its management to run amok.

In this particular case, "Reg. FD," as it came to be known, outlined a process for limiting publicly traded companies' exposure to the likelihood of insider training. Although Reg. FD forced companies to make the same quality of data available to both analysts and the public simultaneously, authorities did not seek to actively prosecute violators until mid-2001 when cease-and-desist actions were levied against several companies for both intentional and unintentional violations of the regulation. Of course, these actions were obscured a short time later by the activities surrounding the MCI-WorldCom and Enron scandals.

Sarbanes-Oxley

To provide stringent guidelines for corporate governance and in direct response to the debacles at MCI-WorldCom and Enron, the United States Congress passed the Sarbanes-Oxley Act in the summer of 2002. In addition to requiring senior corporate officers to certify financial reports (section 302), blocking personal loans to executive personnel (section 402), and forcing the documentation of internal processes and controls (section 404), the Sarbanes-Oxley Act has potentially far-reaching ramifications to the way companies manage data.

One section of the law—section 409—has the potential to cause significant disruption in current data management policies. In particular, section 409 requires enabling real-time disclosure of pertinent financial data. The impact of this legislation on businesses is such that requirements for storage capacity are likely to increase. Interest in content-addressed storage (CAS) has already increased primarily because of its capability to provide easy access to archived data based on key words and content-specific retention requirements. Compliance with Sarbanes-Oxley will increase sales of networked storage and CAS devices in the near-term.

Health Insurance Portability and Accountability Act of 1996

In addition to Sarbanes-Oxley, the Health Insurance Portability and Accountability Act of 1996 (HIPAA), which serves to make available to every patient in the United States his or her own medical records ("Protected Health Information"), creates a standard interface for the transfer of medical data to ensure privacy and security. HIPAA also establishes measures of accountability in the healthcare industry. Not only does HIPAA complicate backup and retention procedures, however, but it also increases storage consumption rates. As the compliance dates approach, and even the smallest healthcare offices are required to demonstrate some disaster contingency capabilities, storage sales will increase.

NOTE It is difficult to imagine the amount of data comprising the "Protected Health Information" of every prescription, dental, and medical record of every U.S. citizen. HIPAA requires healthcare providers to keep multiple copies of and lengthen the retention periods for every billing and medical record for every person receiving healthcare services in the United States. That's a lot of data!

Numerous other updated regulations, as well as the aggressive enforcement of laws already on the books, will cause even more headaches for storage consumers while serving to buffer disk manufacturers from a more precipitous drop in revenues.

Title 21

The Title 21 Code of Federal Regulations (21 CFR Part 11), which was issued by the U.S. Food and Drug Administration and announced in August, 2002, promises to streamline the current process for the manufacturing of pharmaceutical products. At the same time, however, Title 21 mandates extension of periods for record-retention. In all likelihood, Title 21 will cause wrinkles in data management programs in most pharmaceutical companies and potentially increase sales of disk devices over time.

Securities and Exchange Commission

Also in August, 2002, the Securities Exchange Commission, together with the Board of Governors of the U.S. Federal Reserve and the Office of the Comptroller of the Currency, issued a request for comments on a draft of a white paper outlining "sound practices" designed to strengthen the infrastructure of the U.S. financial markets.

As a review of lessons learned from the September 11 attacks, the document outlines the lack of controls and processes required to increase business continuance capabilities. The document highlights the need for "rapid recovery

and timely resumption" of "critical operations" in the event of catastrophic loss of local or regional disruption, particularly with regard to "core clearing and settlement" organizations whose outages present a "systemic risk" to the stability of the market as a whole.[22]

The "sound practices" outlined in the document specify the need for identification of critical services and the testing of recovery systems in as timely and as cost-effectively a manner as possible. Most significant in the document is the recommendation that companies providing core financial services implement backup strategies so that time-to-recover and distance functions extend well beyond the capabilities provided by current solutions. The draft of the white paper suggests distances of 200–300 miles for "out-of-region" backup facilities and recovery time objectives of typically same-day at a minimum, if not within a few hours.[23]

In April, 2003, the interagency group published the white paper and a summary of comments received. Although it is obvious that there is rigorous debate about what constitutes achievable recovery objectives and realistic distance requirements, based on cost-benefit analyses and the technical capabilities of solutions currently available on the market, one thing is clear: The writing is on the wall. The SEC suggests that organizations that perform "core clearing and settlement functions" continue to work toward having these "sound practices" implemented as soon as the end of 2004, and companies that "play significant roles" (those companies who settle or clear five percent or more of the market) should have similar guidelines implemented within three years of the release of the paper.[24]

In December, 2002, the New York Stock Exchange (NYSE), the National Association of Securities Dealers (NASD), and the SEC, in a joint legal action, levied a total of $8,250,000 in fines against five broker-dealers (Salomon Smith Barney Inc., U.S. Bancorp Piper Jaffray Inc., Goldman, Sachs and Company, Morgan Stanley & Co. Inc, and Deutsche Bank Securities Inc.) for failure to keep communications, in this case electronic mail, in an "accessible place" for the two years stipulated by SEC Rule 17a-4 of the 1934 Securities Exchange Act.[25]

Without a doubt, governmental and corporate IT centers find that living with disaster recovery policies and becoming or remaining compliant with current and updated legislation requires a cohesive storage vision to avoid runaway costs and a management nightmare comprised of expensive and poorly utilized storage.

Current and future legislation is also in a position to further decrease storage utilization, thereby increasing the storage TCO for many consumers. The following section explains the impact of poor utilization.

Utilization and Yield

A fundamental piece of the storage TCO equation is utilization and its direct correlation to what can be referred to as the storage yield. If one assumes that the average company used at best 50 percent of their storage assets between 1999 and 2002 (which is itself a conservative number), then, based on the worldwide revenues shown in Table 1-2, we can estimate that over $35 billion dollars in storage assets went unutilized during that time.

NOTE In this section, I borrow two terms from different fields—the COPQ (from Total Quality Management and Six-Sigma) and yield (from manufacturing and agriculture)—and I apply those terms to the discussion of storage utilization.

Storage utilization is the most important storage management issue today: Poor utilization wastes millions of dollars a year in unused storage assets.

Understanding utilization is crucial for the introduction of ROI, Net Present Value (NPV), and TCO in Chapter 3, "Building a Value Case Using Financial Metrics." This material is required to build the financial models with which the business case for storage networks can be justified.

A close analysis of storage yield and the COPQ demonstrates how increased utilization helps lower the overall storage TCO.

The Cost of Poor Quality and the Storage Problem

The Cost of Poor Quality, in terms of quality and yield management, typically refers to the costs associated with poor or undesirable performance of a product over the course of its economic usefulness.[26]

A high COPQ implies higher manufacturing, operations, and labor costs, and consequently, lower revenues. Couching the value of an IT solution in terms of quality management, the COPQ can be said to be the dollar value of how a product, service, or solution performs relative to its expectations. In terms of financial analysis, this figure equates to a negative ROI.

Just as the buildup of IT capacity and subsequent downturn was the outcome of macroeconomic events, the move to storage networks is part of many corporations' efforts to raise their storage yield over time and lower the COPQ (and the TCO) for their storage infrastructure.

Storage Yield

In manufacturing operations, the term yield refers to the ratio of good output to gross output.[27] In storage operations as in manufacturing, the yield is never be 100 percent as there is always be some waste. The goal of a storage vision is to increase not only storage yields, which can be measured in dollars or percent of labor, but also to increase operational yields (or "good output") as much as possible. Ultimately, a storage vision built on a storage utility model helps increase a company's storage yield, the amount of storage capacity allocated and then used efficiently to create and sustain business value.

A tiered storage infrastructure is required to fully increase storage yield and gain true economies of scale. In Table 1-5, each tier has a different capability model and different direct and indirect costs associated with it. The goal is for the COPQ to be as insignificant as possible (shown here as a percentage of $1,000,000 in revenue), and ideally for the accompanying tiers to be appropriately matched to the level of business impact or business revenue of the associated applications. A typical tiered storage infrastructure might look something like this:

- Tier One—Mirrored, redundant storage devices with local and remote replication
- Tier Two—RAID-protected, non-redundant storage devices with multiple paths

- Tier Three—Non-protected, non-redundant, near-line storage devices (for example, SATA drives used as a tape replacement)

Table 1-6 *Cost of Poor Quality as a Percentage of $1,000,000 of Revenue for 1000 GB*

Storage Type	Cost per MB	GB	Total Cost	Allocated	Utilized
Tier 1	$0.05	1000	$51,200	80%	75%
Tier 2	$0.03	1000	$30,720	80%	60%
Tier 3	$0.01	1000	$10,240	80%	90%
Storage Type	**Allocated Yield**	**Utilized Yield**	**Realized Yield**	**Cost of Poor Quality**	**COPQ % of Revenue**
Tier 1	$40,960	$30,720	60%	$20,480	2.05%
Tier 2	$24,576	$14,746	48%	$15,974	1.60%
Tier 3	$8192	$7373	72%	$2867	0.29%

As seen, a low storage yield has a corresponding high COPQ and indicates an overall higher total cost of storage ownership. A more complete discussion of tiered storage solutions (and Information Lifecycle Management) is presented in Chapter 5.

NOTE The difference between allocated and utilized storage is discussed in the section titled "Utilization."

Obstacles Inherent in DAS

As the predominant storage architecture to date in terms of terabytes deployed, DAS has served the storage needs for millions of environments around the globe. Small Computer Systems Interface (SCSI), DAS is a standard, reliable method of presenting disk to hosts. DAS also presents many challenges to the end user including failover and distance limitations, as well as the increased expense associated with poor utilization.

Failover Limitations

Although some DAS environments are Fibre Channel, large storage environments in open systems datacenters have historically been direct-attached SCSI. SCSI is a mainstream technology that has worked well and has been widely available since the early 1980s. SCSI provided the necessary throughput and was robust enough to get the job done. One disadvantage, however, has always been the inability of the UNIX operating system and most databases to tolerate disruptions in SCSI signals, thus limiting the capability to failover from one path to another without impact to the host. In addition, logical unit number (LUN) assignments are typically loaded into the UNIX kernel when the system is booted up, requiring allocation or de-allocation of storage from the host to be planned during an outage window. If the storage unit in question is shared between different clients with mismatched service-level agreements and different maintenance windows, then negotiating an outage window quickly becomes a hopelessly Sisyphean task.

Distance Limitations

Another significant factor hampering the flexibility of SCSI DAS is that SCSI is limited in its capability to transfer data over significant distances. High Voltage Differential (HVD) SCSI can carry data only up to 25 meters without the aid of SCSI extenders. This limitation presents difficulties for applications requiring long-distance transfer, whether for the purposes of disaster recovery planning, application latency, or just for the more physical logistics of datacenter planning.

Expense

Aside from the technical limitations of DAS, the primary drawback of DAS is, without a doubt, its expense. Ultimately, the storage frames themselves constitute a single point of failure, and to build redundancy into direct-attached systems, it is often necessary to mirror the entire frame, thereby doubling the capital costs of implementation and increasing the management overhead (and datacenter space) required to support the environment.

The expense of DAS also stems from poor utilization rates. A closer look at the two primary types of storage utilization further illustrates the nature of the cost savings inherent in networked storage solutions.

Utilization

When considering the impact of managing storage in general and the financial disadvantages of DAS in particular, the primary variable to monitor is storage utilization. Poor utilization leads to a decreased storage yield and a high COPQ, whereby the storage capital asset purchased to provide a service, fails to perform at an optimal level.

Storage utilization has been a marketing hot-button since Fibre Channel SANs began to gain momentum, and as such, utilization is now laden with many different meanings, all of which are often (and unfortunately) used interchangeably. To prevent further confusion, I prefer to use Jon William Toigo's terminology of efficiencies. Toigo clearly delineates between "allocation efficiency" (broadly referred to as "utilization") and "utilization efficiency," which typically reflects the effects of storage usage policies.[28]

Allocation Efficiency

Due to the physical constraints of the solution, DAS environments are intrinsically susceptible to low "allocation efficiency" rates that cost firms money in terms of unallocated or wasted storage. Let us look at one example of the financial impact of poor allocation efficiency.

Imagine a disk storage system (containing 96 73-GB disk drives) with six four-port SCSI (or Fibre Channel) adapters capable of supporting up to 24 single-path host connections. This system is capable of providing approximately 7008 GB of raw storage, or 3504 GB mirrored. Under most circumstances, hosts have at least two paths to disk, so this particular environment can support a maximum of twelve hosts. In a typical scenario, shown in Figure 1-5, this frame hosts the storage for a small server farm of six clustered hosts (12 nodes).

Figure 1-5 *Sample DAS Configuration*

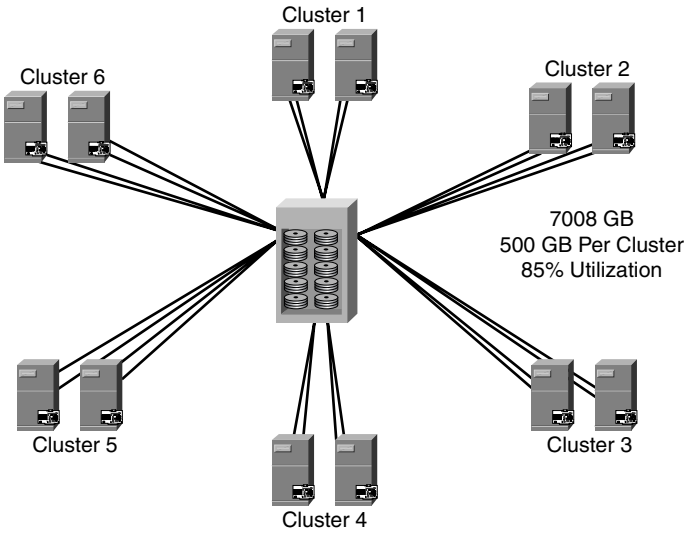

If each cluster hosts six similar applications using 500 GB each (an allocation efficiency rate of 85 percent), almost 500 GB remains unallocated due to the frame's port limitations. With a purchase price of $0.10 per MB, or $100 per GB, there is a loss of $105,120.00 associated with the unutilized disk on that frame.

NOTE Keep in mind that as this frame is formatted for mirroring, the value of the unallocated storage is the cost of the total non-mirrored storage. In other words, the 500 GB of unallocated storage is still 1 TB raw, which must be valued at its purchase price. Also note that use the $0.10 per MB for quick math. The average purchase price of the disk might be significantly lower.

This loss can be considered the COPQ and reflects the costs of additional storage required to provide the expected capacity. An allocation efficiency rate of 85 percent for a DAS environment, however, is significantly higher than the normal average. In June, 2001, a joint study published by McKinsey & Company

and Merrill Lynch's Technology Group, titled "The Storage Report—Customer Perspectives & Industry Evolution," estimated the average utilization rate for DAS environments to be 50 percent.[29]

Fred Moore of Horison Information Strategies has an even more dismal view of allocation efficiency. According to Moore, surveys of clients across various industries indicate allocation efficiencies of 30–40 percent for UNIX and Linux environments and even less for Windows environments, which Moore says frequently see allocation efficiency rates as low as 20 percent.[30]

Using the same environment shown in Figure 1-5 as an example, if the allocation efficiency is only 50 percent, then the loss widens significantly to $350,400, or half the purchase price of the frame. Figure 1-6 shows the costs associated with poor utilization in this environment.

Figure 1-6 *Utilization Rate and Associated Costs—Cash Basis*

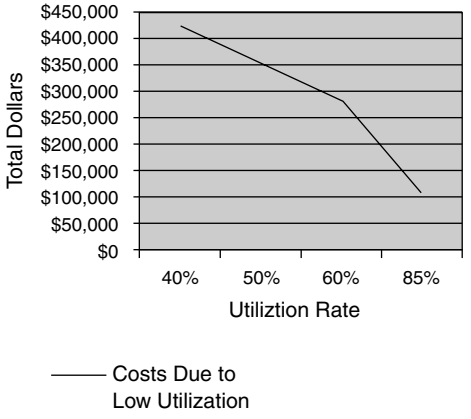

Most firms depreciate the cost of storage over the course of its useful life (assuming the storage is purchased and not leased), so the actual COPQ might vary according to depreciation schedules.

Given the rapid progress of technological advancement, in most cases, depreciation is carried out over three years. If the straight-line method of depreciation is used over a period of three years, the asset value or purchase price of the frame is divided by three with the assumption that one-third of its usefulness

is consumed each year. The impact of the loss, or the COPQ, is then spread across the span of the economic usefulness of the asset. In other words, one third of the COPQ affects the firm's bottom line each year.

Low utilization does not increase or decrease the estimated life of the hardware, nor does this loss change the asset's value in accounting terms. Low utilization does, however, decrease the storage yield of the asset and increases the COPQ, which, in turn, increases the overall TCO. Regardless of the method of depreciation used, poor utilization detracts from the firm's bottom line.

Whether or not the storage units themselves are depreciated, the net effects of poor allocation efficiency are similar: Low allocation efficiency increases the rate of frequency of additional storage purchases. A real life parallel is buying a full tank of gas and being able to use only half of the purchased fuel. As long as you need to drive the car, you will need to purchase more fuel. If more fuel is not consumed, you will be forced to stop at the gas station more often.

Similarly, as long as the firm operates, it needs to purchase storage. The idea that a firm can delay purchasing storage indefinitely by constantly increasing the utilization rate is, to put it bluntly, misinformed. The long-term key to financial success in terms of storage management is optimizing storage usage to minimize the frequency and magnitude of storage purchases. A high allocation efficiency rate helps decrease the size and number of storage purchases, as does a high utilization efficiency rate.

NOTE Capacity on-demand programs are alternative procurement strategies aimed at alleviating the frequency and number of storage purchases. Although these "pay-as-you-go" methods are quite successful at easing the purchase and planning process, they do little to address the rate of consumption or poor utilization found in many environments.

Utilization Efficiency

There might be environments in which the allocation efficiency is at a desirable rate, but the allocated storage is misused, unusable, abandoned, or even hoarded. This is what Toigo refers to as poor utilization efficiency, whereby the

storage itself might be highly allocated, but poorly utilized. In fact, in many open systems environments in which the storage capacity is efficiently allocated, utilization efficiency might be extremely poor, with many applications needlessly consuming data that is rarely, if ever, used.

To resolve these types of issues, a targeted program or project aimed at reclaiming allocated—but lost or poorly used storage—is needed. A project of this magnitude requires a significant time investment and an energetic executive sponsor who is capable of ensuring the proper alignment of goals and initiatives. A storage reclamation project also requires extensive use of a combination of off-the-shelf storage resource management (SRM) software and home-grown scripts dedicated to tracking storage consumption.

Despite the many obstacles that are known factors in implementing DAS, the majority of disk units sold in the last five years are still connected to hosts in a direct-attached fashion. Most companies—even early adopters of storage networking technologies—are still in the implementation phase of building SANs, and therefore have at least a partial mix of SAN and DAS technologies in the datacenter.

Although it is difficult to determine the exact percentage of DAS and SAN storage currently installed world-wide, estimates based on the sales of disk and Fibre Channel gear indicate that nearly three quarters of all disk storage units installed still utilize the direct-attached architecture. As shown in Table 1-1, DAS storage units made up nearly 70 percent of all storage sales in 2003 (with NAS and SAN storage together comprising approximately 30 percent). As these figures indicate, there is still a long way to go before the majority of storage environments currently deployed are networked storage solutions.

In addition to the recently installed DAS, a mountain of DAS that was purchased during the market upswing and it still carries a sizable net book value. As shown in Table 1-1, nearly one million DAS units were shipped between 2001 and 2003, indicating significant depreciation expense for customers when considering the corresponding low utilization rate (and the high COPQ) for DAS.

Conclusion

Networked storage offers strategic benefits over DAS by providing significant cost advantages. Networked storage creates economies of scale and increases operational efficiencies, which reduce the TCO. Networked storage also provides the requisite technology for resolving many storage-related issues, such as the need for long distance data replication.

The current process of managing heterogeneous storage on a storage network is a complex one that at this point in time diminishes operational economies of scale. As storage management software matures, however, operational efficiencies related to managing heterogeneous storage will increase and the TCO for heterogeneous storage will decrease.

Utilizing a tiered storage infrastructure allows more granular management of costs associated with storage management and will in the future (as storage management software matures) lower the overall TCO for storage.

If the charge is then for the CIO to provide the following conditions, then a prime area of concentration for IT departments should be to enable a storage vision that addresses the firm's storage vulnerabilities:

- Scalable, cost-efficient storage solutions that increase the availability of mission-critical business information
- Sound recoverability to business operations in the event of a disaster
- Flexible environments that increase productivity through increased uptime
- Increased business value via cost avoidance and the decreased frequency of hardware procurements through increased yield

A storage vision begins with the consolidation and decommissioning of DAS (and its migration to a networked storage infrastructure) and ends with a framework of cost-effective, tiered storage solutions that are tailored to support applications with storage as a utility-like service.

Summary

This chapter discussed how the commoditization of storage hardware and the poor utilization rates endemic to DAS environments lead to poor storage yield, a correspondingly high COPQ, and an increased TCO for storage. Storage networks provide strategic advantages by lowering IT costs and by eliminating the waste associated with poor storage utilization. Storage networks also address vulnerabilities by increasing data availability.

In the following chapter, I outline how the tangible and intangible benefits of networked storage, primarily increased utilization and increased availability, can lead to a reduction in the TCO for storage.

End Notes

[1] Reprinted by permission of *Harvard Business Review*. From "IT Doesn't Matter" by Nicholas Carr, May 2003. Copyright 2003 by the Harvard Business School Publishing Corporation; all rights reserved.

[2] From "Does IT Matter? Information Technology and the Corrosion of Competitive Advantage" by Nicholas Carr, 2004. Copyright 2004 by the Harvard Business School Publishing Corporation; all rights reserved.

[3] From "How Competitive Forces Shape Strategy" by Michael E. Porter. *Harvard Business Review* March/April, 1979. Copyright 1979 by Harvard Business School Publishing; all rights reserved.

[4] IDC 2004. Framingham, MA.

—"Worldwide Disk Storage Systems, 2003-2008 Market Forecast and Analysis: Conservatism Persists, but Opportunities Abound," IDC #31663, forthcoming.

—"Worldwide Disk Storage Systems Forecast and Analysis," 2003-2007, IDC #30247, October 2003.

—"Worldwide Disk Storage Systems Forecast and Analysis," 2002-2006, IDC #28261, December 2002.

[5] Toigo, Jon William. "The Holy Grail of Data Storage Management," Upper Saddle River, New Jersey: Prentice-Hall PTR, 2000.

[6] Greenspan, Alan. "The Challenge of Central Banking in a Democratic Society," Annual Dinner and Francis Boyer Lecture of The American Enterprise Institute for Public Policy Research, Washington, D.C. December 5, 1996. Accessed July 6, 2004, http://www.federalreserve.gov/BOARDDOCS/SPEECHES/19961205.htm.

[7] IDC, 2004.

—"Worldwide Disk Storage Systems 2003-2008 Market Forecast and Analysis: Conservatism Persists, but Opportunities Abound," IDC #31663, forthcoming.

—"Worldwide Disk Storage Systems Forecast and Analysis," 2003-2007, IDC #30247, October 2003.

—"Worldwide Disk Storage Systems Forecast and Analysis," 2002-2006, IDC #28261, December 2002.

[8] Berinato, Scott. "What Went Wong at Cisco?" *CIO Magazine*, August, 2001.

[9] The Bureau of Economic Analysis at the U.S. Department of Commerce tracks the U.S. Gross Domestic Product by analyzing a series of major accounts of U.S. goods. The most up-to-date analysis can be found on the BEA website at http://www.bea.doc.gov/.

[10] McKinsey Global Institute, "US Productivity Growth 1995-2000: Understanding the Contribution of Information Technology Relative to Other Factors." Washington, D.C.: McKinsey & Company, October 2001.

[11] National Bureau of Economic Research, Business Cycle Dating Committee. "The NBER's Recession Dating Procedure," October 21, 2003. Accessed, July 17, 2003, http://www.nber.org/cycles/recessions.html.

[12] DeLong, Summers, "The 'New Economy': Background, Historical Perspective, Questions, and Speculations." *Federal Reserve Bank of Kansas City, Economic Review*, Fourth Quarter, 2001.

[13] ibid.

[14] EMC Corporation. Form 10-K, Annual report for the fiscal year ended: December 31, 2002. Retrieved June 3, 2004 from EDGAR.

[15] EMC Corporation. Form 10-K, Annual report for the fiscal year ended: December 31, 2003. Retrieved June 28, 2004 from EDGAR. Form 10-K, Annual report for the fiscal year ended: December 31, 2002. Retrieved June 3, 2004 from EDGAR.

[16] IDC. Framingham, MA.

—"Worldwide Fibre Channel Host Bus Adapter 2004-2008 Forecast and 2003 Vendor Shares," IDC #30925, March 2004.

—"Worldwide Fibre Channel Host Bus Adapter Forecast and Analysis, 2003-2007," IDC #29414, May 2003.

—"Worldwide Fibre Channel Switch 2004-2008 Forecast and 2003 Vendor Shares," IDC #31362, June 2004.

—"Worldwide Fibre Channel Hub and Switch Forecast and Analysis, 2003-2007," IDC #29882, August 2003.

[17] IDC. Framingham, MA. 2004. Source: IDC Worldwide Quarterly Server Tracker, Q1 2004 release.

[18] IDC. Framingham, MA. 2004.

—"Worldwide Disk Storage Systems 2003-2008 Market Forecast and Analysis: Conservatism Persists, but Opportunities Abound," IDC #31663, forthcoming.

—"Worldwide Storage Software Forecast and Analysis, 2003-2007," IDC #29983, August 2003.

—"Worldwide Storage Software Forecast and Analysis, 2002-2006," IDC #27477, June 2002.

—"Worldwide and U.S. Storage Services 2004-2008 Forecast: The Opportunity Shifts," IDC #31042, March 2004.

—"Worldwide and U.S. Storage Services Forecast and Analysis, 2003-2007," IDC #28992, March 2003.

[19] U.S. Census Bureau. "Retail 2Q, 2003 E-Commerce Report," http://www.census.gov/mrts/www/current.html. August 22, 2003. Retrieved September 2003.

[20] U.S. Census Bureau. "Retail 1Q, 2004 E-Commerce Report," http://www.census.gov/mrts/www/current.html. May 21, 2004. Retrieved August 2004.

[21] U.S. Census Bureau. "United States Department of Commerce E-Stats," March 19 2003, http://www.census.gov/estats. Retrieved September 2003.

[22] Securities Exchange Commission. "Interagency Concept Release: Draft Interagency White Paper on Sound Practices to Strengthen the Resilience of the U. S. Financial System; Rel. No. 34-46432." Retrieved August 2003. SEC. August 2002. http://www.sec.gov/rules/concept/34-46432.htm.

[23] Ibid.

[24] Securities Exchange Commission. "Interagency Paper on Sound Practices to Strengthen the Resilience of the U.S. Financial System; Release No. 34-47638," http://www.sec.gov/news/studies/34-47638.htm. Retrieved August 2003. SEC. August 2002.

[25] Securities Exchange Commission. "SEC News Digest, Enforcement Proceedings," http://www.sec.gov/news/digest/12-03.txt Retrieved August 2003. SEC. December 3 2002.

[26] There are reams of quality management and Six-Sigma resources available on the web but a good basic definition of COPQ can be found here: http://software.isixsigma.com/dictionary/Cost_of_Poor_Quality-63.htm

[27] Christian Terwiesch, Roger E. Bohn, "The Economics of Yield Driven Processes," Working Paper, Wharton Operations and Information Management Department, The Wharton School, 1998.

[28] Toigo Partners International, LLC. "Storage Management 2003: "The 'Two Towers' of Storage Pain Part One: Provisioning," Toigo Partners International, LLC. 2003: pp2–3.

[29] Merrill Lynch, McKinsey & Company. "The Storage Report – Customer Perspectives & Industry Evolution," Merrill Lynch, McKinsey & Company, June 19, 2001: p 40.

[30] Moore, Fred. "Horison Information Strategies: Storage Facts, Figures, Estimates, and Rules of Thumb," Copyright Fred Moore, 2003.