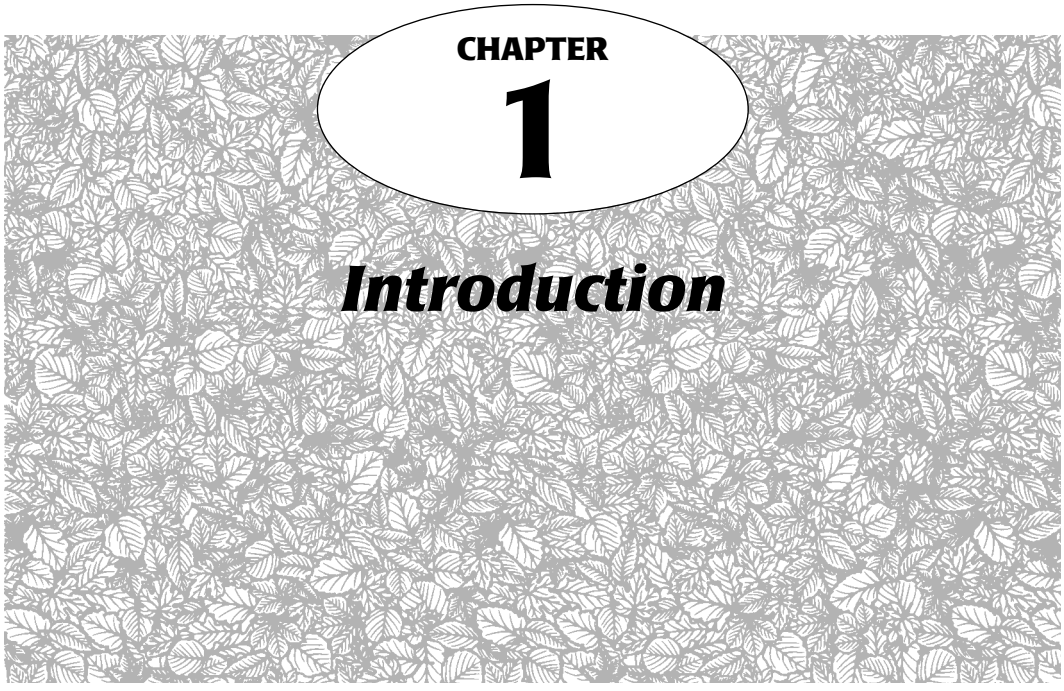

CHAPTER

1

Introduction

Given the high costs associated with enterprise storage in most companies, it might be easy to suspect that some dark and sinister hand is at work. By some analyst accounts, corporate IT departments are spending upward of 60 cents of every dollar budgeted annually for IT hardware to purchase storage platforms. This comes at a time when the per megabyte hardware cost for storage is declining at a rate of about 50 percent per year and the capacity of disk drives is growing (and has been since the late 1990s) at about 120 percent per annum. Given such trends, one might reasonably ask the question, "If storage is getting cheaper, why is it consuming more of my budget today than ever before?"

CONSPIRACY THEORY

If you had a conspiratorial mindset, you might suspect that there was a secret organization at work, a kind of storage vendor cartel laboring in the background to garner a greater and greater percentage of IT spending for the wares of its members. For the sake of argument, let's call them the En-

gineers for the Accelerated Total Depletion of Information Storage Components (EATDISC).[1]

With a bit of imagination, you can envision the group. Think organized criminal underworld, a la SPECTRE in the 1960s James Bond movies or the Mafiosi capo regimes in contemporary gangster films or cable TV's *The Sopranos*.

You might imagine this group meeting in some isolated alpine retreat, possibly situated in a facility carved out of a jagged rock face somewhere in the Colorado Rocky Mountains, where it would be near the development shops of many storage companies. Picture an assemblage of stone-faced men (and maybe one or two women) wearing $1,500 suits, all seated with their cigars and brandies around a great mahogany table. They listen attentively to one presentation after another, each one describing the progress of various nefarious programs and initiatives intended to bring about the global domination of IT spending by data storage technologies.

Imagine that EATDISC has cultivated "friends" in the technology trade press, where they purchase the bulk of the advertising, and in the industry analyst community, where they buy "friendly" coverage for a small sum of a few hundred thousand dollars per year—small, because it is an infinitesimal fraction of the $25 billion in revenues that disk-based storage was (conservatively) expected to generate in 2002, and smaller still as disk-based storage purchasing climbs to more than $30 billion by 2004.[2] For an investment of what amounts to pocket change, EATDISC virtually ensures that no negative press will surface about its members' products and that the two primary sources of information for end-users regarding their technology options are tightly controlled.

Imagine further that these powerful figures control vast cadres of value-added resellers (VARs) and integrators who recommend EATDISC products as "solutions" to their customers based not on how effectively they will meet actual application requirements, but what profit margin the recommended solution will yield to the reseller/integrator. Their rationale is simple: Why sell a customer one or two components that will deliver a rich, redundant, and manageable solution to his or her immediate problem when you could sell a complex, multicomponent storage area network (SAN) that the customer could "grow into" over time, and that, by coincidence, delivers immediate profits to the reseller and vendor?

Savvy customers would see through such a scam, you think? Not if the customer depends on a "trusted solution provider"—a reseller or integrator—to chart his or her strategic path.[3] Not if the customer has been convinced by analysts' reports and trade press articles that SANs are strategic. Not if the IT manager or chief technology officer or chief infor-

mation officer finds his or her decision making undercut by vendor or reseller salespersons who do "end runs" around the company technologists and cultivate sales directly with nontechnical chief financial officers or chief executive officers.

Imagine still further that the organizations responsible for developing "open storage standards" find their efforts stymied by their very openness, which permits representatives of EATDISC member companies to sit on, and in some cases even to chair, standards development committees. The vendor representatives could obfuscate progress on any initiative that might cost their firms "value discriminators," resisting common standards that might that contribute directly or indirectly to the erosion of their company's respective market share leadership. The irony of this situation is that, at the same time as the vendor representative obstructs progress on standards development, the representative's company can legitimately claim to be "actively engaged in the open standards development process," intent as it is upon providing less proprietary solutions for its customers.

Last but not least, imagine that EATDISC was as expert at manipulating the legal system as it was at manipulating data bits. What if large member companies leveraged loopholes in the patent and trademarks registration process to file "blanket patents" that covered all technological development in a given area—even development based on ideas not yet thought of—then used its legal authority to stifle innovation that EATDISC did not sanction?

In short, what if such a cartel worked earnestly behind the scenes to ensure that customer hearts and minds and dollars were continuously cultivated to support the acquisition of proprietary and half-baked technologies, premised upon the flimsiest and most untenable of business value propositions, that resisted common management and required forklift upgrades every one or two years?

Taken collectively, the above scenario would certainly provide an explanation for why something that is increasingly a commodity—like disk-based storage technology—is costing organizations more of their IT budget than ever before. Without a doubt, many industry events over the past few years would appear to validate some or all of the conspiracy theory advanced above.

But, of course, the existence of EATDISC is purely paranoiac fantasy. In reality, such a cartel would require levels of discipline and cooperation that the storage industry has never been able to manifest. The possibility of a real-world EATDISK runs afoul of long-standing fears among storage technology vendors that prevent such a cooperative cartel from ever appearing.

THE STORAGE INDUSTRY: FEAR OF STANDARDS AND COMMODITIZATION

In point of fact, data storage remains one of the few areas of information technology characterized by a paucity of standards—whether open or *de facto*. In this, data storage stands in stark contrast to other technology sectors, where well-defined standards have either:

1. Been articulated by open standards organizations like American National Standards Institute (ANSI), Internet Engineering Task Force (IETF) and adopted across literally all vendor product lines, or
2. Imposed by a single vendor with a commanding market share leadership position (e.g., *de facto* standards).

An example of an open standards-based technology is the network interface card (NIC). In the NIC market, all manufacturers must adhere to a set of formal open standards if they want to sell their products. As a result of open standards, end-users can pick and choose products from the more than 100 products and still achieve their basic goal: connection to an Ethernet network operating the TCP/IP protocol suite, the dominant network protocol in the world today.

Because of open standards for Ethernet and TCP/IP, any NIC card will do the basic job of providing network attachment. Customers can, with confidence, buy a NIC from any vendor, install it in their desktop computer or server, and be sure that their system is ready to be connected to a network.

Of course, to differentiate their products from those of their competitors, some vendors "add value" to their NICs by adding features such as "Wake on LAN" or special TCP offload engine (TOE) chips, but the basic standards-based functionality is available whether the consumer buys the enhanced NIC or its less expensive "bare bones" cousin.

In the final analysis, consumers are the beneficiaries of open standards because such standards level the playing field between vendors, enable the "apples-to-apples" comparisons of different products, and provide customers with assurances of baseline compatibility that, in turn, frees them to consider options for purchase on the basis of feature, function, and price. The presence of open standards in the network space is owed to the work of a generation of network technology start-ups who sought to unseat a previous generation of proprietary network technology vendors, and also of a unified consumer voice that voted with their checkbooks for open standards-based products.

Standards breed commoditization, of course. And commoditization of technology exerts a downward pressure on product prices. That is why a gigabit Ethernet switch that cost $3,000 per port in late 2000 dropped precipitously to about $450 per port in late 2001 and today hovers at around $300. As more standards-compliant switches came to market that provided essentially the same business value, most customers preferred products that were less expensive to comparable products that were more expensive. The result of this purchasing pattern has been price erosion, another feature of open standards-based technology and one that nearly always works to the benefit of the consumer.

Another example of a technology segment "blessed" (from the consumer's standpoint, at least) by standards is the PC market. As of this writing, about 80 percent of personal computer operating systems (OS) bear the logo of Microsoft. The Windows OS is a *de facto* standard in desktop computing. Of course, Windows was never formally ratified as an open standard by an official standards body, but, by sheer presence in the market, the Microsoft OS dominates all contenders.

The result of a *de facto* standard is similar to that of an open standard in at least one respect: It fosters compatibility. The dominance of Microsoft, in turn, drives most independent software vendors (ISVs) to develop application software to run on the platform. For example, this tax season, an end-user can go to a local computer software store and select from among several personal tax preparation software packages from several ISVs, any of which will run on his or her Windows-based PC.

Whether or not you prefer the Microsoft platform is beside the point. The fact that nearly every desktop application software package will operate under the Windows operating system means that the application software products compete on a level playing field imposed by a *de facto* standard. Vendors must add features and functions to their products, or offer licenses under an attractive pricing scheme, if they want to compete. Again, the beneficiary of competition in a standards-based realm is the consumer.

Storage lacks both open and *de facto* standards. And the industry resists virtually any effort to impose meaningful standards (whether by a standards-making body or by one of its own members) that might compromise proprietary advantage for the vendor. Insights offered by many storage industry insiders underscore this position.

For example, a representative of one of the industry's largest disk array vendors tells the tale of how his management's biggest fear is that there will be a repeat, in the storage array business, of what happened in the disk drive business in the mid-1990s. In the late 1980s and early 1990s,

disk drives featured proprietary technologies ranging from vendor-specific interfaces and low-level drive formats to specialized actuator arm and read/write head designs. This diversity ensured that Brand X products worked very differently from Brand Y and that consumers would be locked in to a particular vendor's technology. As a consequence, drive prices remained very high and margins were lucrative for disk drive manufacturers.

With the arrival of the Small Computer Systems Interface (SCSI) standard, however, this picture changed. SCSI provided a standards-based mechanism for communicating with storage devices and had the effect of standardizing drive electronics and interfaces. SCSI was not the only standard for disk drive interfaces—other standards emerged at about the same time, including IDE/ATA. However, the overall impact of standards was to reduce disk drive products to the level of commodity or stock items.

Prices eroded dramatically as low-cost, standards-compliant drives flooded the market from mass producers. According to the source, the greatest fear at his company—a large integrated disk array manufacturer—is that it could happen again. Originally, the taxonomy offered by researchers at the University of California at Berkeley to describe various Redundant Arrays of Independent Disks (so-called RAID levels) was suspected of introducing a commoditizing effect in storage arrays. However, as documented in *The Holy Grail of Data Storage Management*, vendors treated the UC-Berkeley definitions less as a standard taxonomy than as a useful point of reference for describing their varied implementation of disk redundancy schemes. Some even coined RAID levels to describe their products that weren't in the UC Berkeley taxonomy at all. However, this isn't the end of the story.

The same fear of standards-based commoditization has echoed and re-echoed in the public statements of leading storage industry spokespersons over the past couple of years. At a storage networking conference in spring 2001, the then-CEO for a market share leading company in storage arrays noted that the reason his company was transitioning from a hardware-focused to software-focused model was to leverage the only differentiator that remained to distinguish his products from those of his competitors: namely, intellectual property captured in the form of software on his company's arrays. Without software, said the CEO, there was no difference between his products and those of his nearest competitor: "We both just sell a box of Seagate hard drives." That statement underscored the fear of lost differentiation in the market as storage hardware became less proprietary and more commoditized.

As standards drive commoditization, vendors know that they must find new ways to realize proprietary technology-based margins from otherwise commoditized hardware technology. Sometimes, this takes the form of resistance to standards efforts altogether.

According to a spokesperson for a Fibre Channel "SAN"[4] switchmaker in 1999, "A real storage area network—one based on truly open standards and capable of supporting heterogeneous storage arrays (i.e., arrays from different product manufacturers)—will probably never appear in the market. If a real SAN did appear, Joe's JBODs (an abbreviation for Just a Bunch of Disks) would perform exactly the same as high-end, 'name brand,' proprietary arrays, which cost ten times as much. The 'name brand' vendor's customers would quickly realize that they had been paying way too much for storage."[5]

"Under those circumstances," the fellow continued, "do you really think that the leading vendor would allow his arrays to be included in an open SAN?"

The question crystallizes much of what has Balkanized early SAN standards development efforts, as well as bringing into sharp focus the concerns that vendors have about open standards and their commoditizing effect on products.

Another Industry Concern: The Rise of a Dominant Player

Arguably, the only thing that strikes more fear into the hearts of storage technology companies than open standards-based product commoditization is the fear that one of their peers will achieve such market dominance that it will be able to impose its own technology as a *de facto* standard on the entire industry—à la Microsoft in the desktop world.

Evidence of this concern can be inferred from the brief history of the Storage Networking Industry Association (SNIA). SNIA is a quasi-standards group comprised of vendors of storage technology products and services—the brainchild of several well-intentioned (and possibly idealistic) storage industry analysts and insiders who perceived a need, in the absence of standards for storage networks, to provide a forum for "co-opetition" (a blend of cooperation and competition) among storage vendors.

Among other things, SNIA was set up to foster informal agreements between vendors on their implementations of networked storage technology so that the entire storage networking initiative would not fall prey to petty vendor infighting. The interoperability difficulties in the Fibre Channel SANs of the late 1990s were proof positive to the founders that some sort of SNIA-like body was needed.

The fact is, however, that SNIA never received the full backing of the industry until a major storage vendor, EMC Corporation, announced its own competing alliance of companies, the Fibre Alliance, to advance the vendor's own view of how storage networks should work. Rather than allow a single vendor achieve the status of a broker of *de facto* standards, EMC's competitors turned their attention to SNIA and joined up with the fledgling organization in droves.[6]

This phenomenon repeated itself in 2002 when a number of vendors, troubled by the unilateral move of one of their peers (EMC, again) to advance its own model for universal storage management, decided to counter the initiative by supporting a common management approach under development at SNIA: Common Information Model/Web-based Enterprise Management (CIM/WBEM). Until that time, vendors had paid only lip service to CIM, making no real provisions for CIM-based management in their own platforms or products. However, the fear of an emerging *de facto* standard, in the form of EMC's AutoIS initiative and WideSky "middleware," compelled vendors to "take the pledge" at the Storage Networking World conference in Palm Desert, California, in Spring 2002 and to promise CIM-manageable platforms by year's end![7]

Whether or not the industry keeps its pledge to embrace a CIM-based management standard remains to be seen, of course. Cynics suggest that once the specter of single vendor dominance in universal storage management recedes, CIM developers at SNIA will see considerable backsliding among the newly converted. No one, after all, really wants a universal storage management standard.

In the past, those who suggested that consumers might benefit from an open cross-platform standards-based storage management methodology were met by cold stares from storage manufacturers—reminiscent of Anthony Hopkin's character, Hannibal Lecter, in the film version of the Thomas Harris novel, *Silence of the Lambs*. At the risk of belaboring the metaphor, in one scene from the popular film, the hero, a female FBI agent named Clarice Starling, passes a questionnaire to Lecter that is intended to gather information about the serial murderer's psychology for the FBI's behavioral science database. Hopkins, as Lecter, levels a cold, unwavering stare at Starling and asks in a calm and calculated tone whether she believes she can really "dissect" him with such a blunt instrument (i.e., the survey).

This is exactly the response that one could have expected from a storage manufacturer prior to the spring 2002 CIM announcement: It was heresy for anyone to suggest that there was enough common ground among all of the 17-odd thousand storage platform products in the mar-

ket that would enable their management by a common management methodology such as CIM.[8] The very idea was as heretical to storage vendors as it once was to Compaq and other PC and server manufacturers when the Desktop Management Task Force (DMTF) first tried to apply the CIM strategy to PC management in the early 1990s.

THE ROAD TO UNIFICATION?

Of course, despite the deep-seated fears of commoditization and of the emergence of a single market share leader, there have been some cooperative development efforts—occasions when storage vendors seemed to rise a bit above the quagmire of storage industry politics. In some instances, cooperation between competitors was the result of self-interest and a response to clearly articulated consumer demands—particularly, when such demands emanate from *Fortune* 100 customers, who, analysts claim, collectively account for more than 60 percent of annual storage industry revenues.

An example of this dynamic was the development of Fibre Channel over IP (FCIP), a protocol for tunneling Fibre Channel–based storage commands and data between distant "SAN islands" using a TCP/IP network. In late 2000, several large enterprise customers—common to both Brocade Communications Systems, a Fibre Channel fabric switch vendor, and Cisco Systems, an IP networking technology goliath—demanded that their vendors work together to deliver such a tunneling protocol to interconnect small Fibre Channel SANs at geographically dispersed locations. An announcement of a joint initiative between the companies met surprised looks from attendees at the press conference where it was made. Only weeks before, spokespersons for Cisco, who advocated the use of IP networks, rather than Fibre Channel, to provide the plumbing for storage area networks, had criticized the Fibre Channel Industry Association's announcement of a forthcoming two-speed (or 2 Gigabit-per-second) version of its protocol, calling it "yesterday's bandwidth tomorrow." Cisco posited that 10-speed gigabit Ethernet, at 10 gigabits per second, would appear in the market shortly, mitigating the value of the accelerated FC protocol.

However, on an October morning in 2000, cooperation was the theme, and prior competitive utterances by the companies were described as water under the bridge. Representatives of Cisco and Brocade appeared together at the press event and described the synergies between the two companies and their desire to work together to meet the needs of their common customers. FCIP was developed from the initiative, re-

ferred to the IETF for consideration as a standard, and work began on creating a blade for a Cisco switch to support FC tunneling.

In June 2002, however, the relationship between the two companies was severed by Cisco with all the trappings of an ugly Hollywood divorce. The IP networking equipment manufacturer claimed that the goal of developing a tunneling protocol that could be supported on all Fibre Channel switches remained valid, but that the implementation of the protocol had been corrupted by Brocade, which had helped engineer the solution so that it would only work with its own switches. Still, an IETF open standards protocol was born of the initiative.

Another example of industry cooperation was BlueFin, a joint undertaking of several erstwhile storage market competitors intended to advance the cause of Fibre Channel SAN interoperability. Bluefin's creators (who comprised most major storage platform manufacturers in 2002, as well as several prominent storage management software vendors) held what could be described as secret meetings for about six months and, at a cost of nearly $500,000, produced an open standards-based mechanism that enabled the products of different vendors to "discover" each other in the same Fibre Channel SAN (fabric). BlueFin was hailed as a huge breakthrough in heterogeneous Fibre Channel SAN development, and one that laid the groundwork for subsequent management services that are not supported in the Fibre Channel protocol itself.[9]

The jury is still out on the true meaning and importance of BlueFin, which has since been turned over to SNIA for further development as "SMI-S," but it is reasonably certain that the vendors involved in its development were compelled to cooperate both by the slower-than-expected adoption rate of FC SANs and also by often-heard complaints of their largest customers, many of whom had fielded heterogeneous storage platforms over the years and saw the lack of support of Fibre Channel SANs for heterogeneous configurations as an impediment to implementation of the technology. Now, as then, few companies have been willing to do a "forklift upgrade" of their existing storage investment in order to deploy a homogeneous Fibre Channel SAN. Most savvy consumers vocalize deep concerns over the ultimate impact of homogeneous storage architecture: dependency on a single vendor.

These and other examples testify to the fact that, to storage vendors, the only thing more frightening than commoditization, and single vendor dominance, is consumer revolt. In the current economic reality, one in which cash-strapped consumers are putting off storage acquisitions and looking for ways to leverage storage networking and storage management technologies to curb IT costs, the pressure is on the vendor commu-

nity to facilitate their needs. The cost of failure could be the decision by the consumer to give its business to an eager-to-please competitor—and competitors are in no short supply.

YESTERDAY'S RESELLERS, TOMORROW'S COMPETITORS

In fact, new storage start-ups are appearing at a more rapid pace than ever before. In many cases, these start-ups are not strangers to storage, but, to the contrary, are staffed and managed by former insiders or resellers of an established vendor.

As former insiders, the newcomers know the foibles of their previous employers and of their products. As solution architects and engineers, many resellers have the skills to develop comparable product offerings to those of leading vendors that offer full "backward compatibility" with the products of their former suppliers, while offering improved functionality at a lower price point.

Take, for example, the case of a Newark, Rhode Island-based reseller of NAS solutions. In 2002, the company moved from the role of reseller to the role of competitor in response to the draconian marketing practices initiated by equipment suppliers, who were themselves attempting to address the general economic slowdown. Today, the reseller's product is every bit as capable as the leading products in the NAS market, but at a fraction of the cost.

As happened over and over again after February 2001, the storage technology supplier, a publicly-traded firm, sought to improve its quarterly earning reports by redrawing its sales territories and reassigning its most lucrative customer accounts (heretofore managed by resellers) to its own internal direct sales force. The move, which cost the Rhode Island reseller considerable income, fostered a desire to create a competitive product that was "plug and pin" compatible with the supplier's product, but redressed certain "margin-enhancing" features of the supplier's product to cut its expense.

For example, the supplier was selling as part of its NAS solution disk drives priced at three times the cost of the drives on the open market. The supplier ensured that only its drives could be used with its network-attached storage "head" (the thin server component of a NAS) by placing a proprietary block format on the drives. Consumers were denied access to the drive formatting utility that could be used to format less expensive drives, readily obtained on the open market, for use with the vendor's platform.

By removing the special formatting step and using commodity disks, the reseller was able to reduce the cost for his "new" product, while at the same time ensuring that customers who already had an investment in his former supplier's products could preserve that investment and use their older drives and trays with his new NAS head. Such a scenario is repeating itself over and over in the industry today—to the chagrin of brand-name vendors.

While the business motives that drive some vendors to seize control of the most lucrative accounts from their resellers are understandable, the practice is somewhat short-sighted. Virtually every industry observer who has monitored the emergence of networked storage has noted the close link between NAS and SAN adoption rates and the sales education efforts of channel partners (resellers/integrators.) Of the 11,000 or so SANs deployed in the market at the time of this book, the preponderance have been sold to customers by reseller/integrators who have been better situated to educate consumers and to dedicate the necessary time to account cultivation than the minimalist direct sales forces that most equipment manufacturers would permit. Vendors ignore this relationship between channels and NAS/SAN adoption at their own peril. What's more, it is unwise to anger channel partners who possess the engineering talent and the knowledge of their products and access to the vendor's own parts suppliers, and who therefore have the ability to manufacture competitive products at reduced cost to the consumer.

A further gap in the wisdom of the vendors who "disrespected" their channel partners was evidenced by subsequent failure of direct sales in many cases. Direct sales forces are an expense that many vendors have found themselves unable to bear as the economy worsened. It was not unusual, therefore, to see renewed cultivation of channel partnerships in 2002, following massive layoffs of direct sales staff that the vendors could no longer afford to keep on the payroll.

In the final analysis, economic downturns have the twin effect of 1) bringing to the surface increased price sensitivity among consumers (and, in many cases, an increased willingness to try comparable products that cost less money) and 2) opening the door for plug and pin compatible products that provide viable alternatives for the established vendor's products. Some of the brand names of today, including EMC and parts of Hitachi Data Systems, were born of just such economic downturns in the 1980s that increased the vulnerability of dominant vendor IBM to competitive products.

The bottom line is that economic challenge breeds competition and competition, in turn, fuels consumer perceptions of commoditization by

providing numerous "generic" alternatives for expensive "brand-name" technologies.

ABOUT THIS BOOK

All of the above simply sets the stage for the chapters that follow. The storage technology business provides a backdrop that explains some of the confusion that keeps the emergence of a true storage infrastructure, one defined by common management and automated, policy-driven intelligence, a "holy grail."

In the following pages, you will learn of the technological trends that are forcing storage into networked topologies and the challenges and opportunities that are created by this shift. In particular, you will look at the state of disk-drive technology futures, the inevitable outcome of limits to growth in disk-based storage and the drivers that this provides for a networked storage paradigm.

You will further examine the claims of networked storage advocates and evaluate their applicability in the face of current technology capabilities and limitations. Much of the hyperbole around networked storage is patently false: pure "marketecture" that consumers need to see for what it is in order to avoid ill-advised or ill-fated buying decisions. (Marketecture is the author's coined phrase describing an otherwise valid technological concept or solution that has been placed in the hands of a vendor's marketing department.)

You will focus on practical examples, drawn from the real world, to demonstrate what of the burgeoning network storage technology works, what doesn't, and what's needed. And you will look at the application-storage nexus in detail to describe best practices for developing effective storage strategy. As of this writing, the network storage industry appears to have created confusion in many minds about the proper focus for data storage, its design, and its management. Given the primordial state of current SAN technology, it is obvious that storage management strategies that focus on managing storage device capacity utilization and component operations are inadequate to the task of storage management. The application must drive the management of storage since it provides the metric that is ultimately used to discern how well the storage infrastructure itself is managed.

Finally, you will look at fruitful new areas of storage technology development, emphasizing object-oriented schemes for data classification and lifecycle management as the *sine qua non* of storage management.

Managing data in a networked storage setting is ultimately the only effective way to manage storage itself.

Bottom line: if you are looking for a primer about data storage component and platform technology, the author invites you to consult *The Holy Grail of Data Storage Management*, which is also available from this publisher. If you understand the components and you are ready to start becoming strategic about storage while solving immediate and evolving problems, those that contribute the lion's share of the cost to your constrained IT budgets, you are in the right place.

ENDNOTES

1. The idea of EATDISC first appeared in an October 2001 *Enterprise Systems* magazine column, written by the author.
2. These market revenue estimates are mean numbers derived from numerous industry analyst market estimates—each of which, by itself, strikes the author as suspect for a wide range of reasons.
3. Based on the author's research, reported in various trade press publications, U.S. government agencies and departments largely rely on integrators and resellers to define their strategic goals. While less prevalent in the private sector, the outsourcing of "strategic planning" to "trusted solution providers" appears to be a growing trend in the commercial sector, as well as organizations cutting their IT staff and management in response to the negative economic conditions prevailing at the time of this writing.
4. Using the terms Fibre Channel and SAN in connection is commonplace in the industry, though oxymoronical in technical terms. Fibre Channel is a serial storage interconnect that can be deployed in a switched fabric topology. It is not, however, a network protocol and cannot be used to create a network—only a fabric of point-to-point connections switched at high speed. Readers should be aware of this very important difference, which will be discussed in greater detail later in this book.
5. The name of the quoted party is being withheld at his request, though the author cited the speaker in a trade press article published in 1998. Since making the statement, the speaker's company has engaged in a strategic business relationship with the very hardware array manufacturer that was cited in the statement. The speaker asked to have his name removed from the quote (despite the fact that it is one of the more intelligent things anyone ever said about impediments to SANs in the industry) because, in the present business nexus, it was a "career limiting quotation."
6. EMC also joined SNIA, but has been accused from time to time of being an "undependable advocate" of SNIA initiatives—especially when they were at

cross-purposes to EMC's own interests. The same criticism can be leveled at virtually every one of SNIA's members at one time or another.

7. Taking the pledge is a description derived from various self-help groups popular around the time of the SNW announcement of CIM support. Developed by the author in a humorous column for Enterprise Systems magazine written shortly after the press event, it suggests that vendors approached the podium one after another, admitted that they had only supported proprietary management schemes up to that point, and pledged to support an open standards-based approach, CIM, going forward. As in most self-help programs, some "backsliding"—that is, deviation from the pledged path or course of action—is to be expected.
8. In fact, CIM advocates, such as Mark Carlson, a network storage engineer for Sun Microsystems who also oversees CIM development at SNIA, would argue that CIM is not a one-size-fits-all management method. Each device can have a unique Managed Object Format (MOF) that describes the customized features and functions unique to that device. Still, the industry resists the notion that any universal management method, including CIM, could ever capture the "secret sauce" diversity in 17,000-plus competing storage products. Conceding that it could might create the impression in consumer minds that no important differentiators existed between products bearing different brand names.
9. As per the previous note, Fibre Channel was not designed to serve as the plumbing of a SAN. It lacked, as documented in a Road Map whitepaper written by this author on behalf of the Fibre Channel Industry Association in 2001, "fabric services" for in-band device discovery, management, security, etc. that a real network protocol would have provided. BlueFin, if nothing else, adds an important fabric service that should have been provided in any real SAN plumbing solution from the outset.