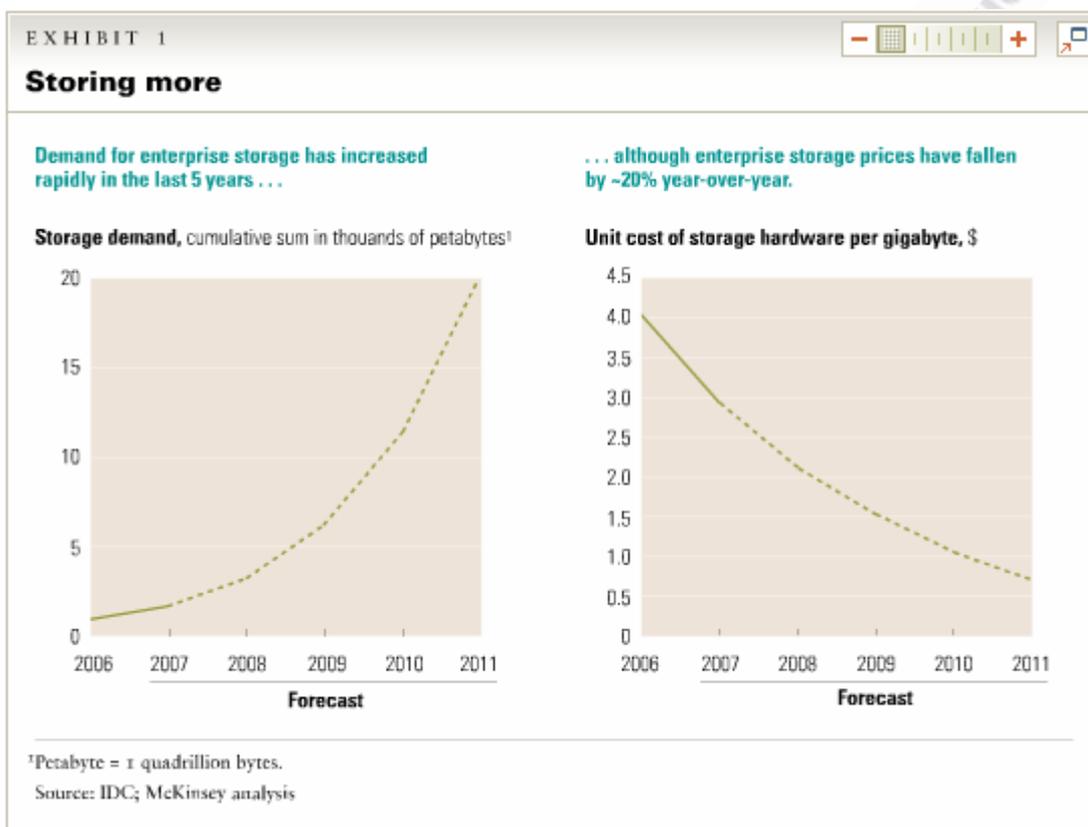


Meeting the demand for data storage

James M. Kaplan, Rishi Roy, and Rajesh Srinivasaraghavan

As the information storage needs of many large enterprises grow and become more complex, IT executives must have better policies to guide their efforts.

For many companies, data storage has become one of the fastest-growing parts of the IT budget, thanks to enterprise-wide transactional systems, massive data warehouses, and explosive growth in e-mail traffic. While the storage market broadly has been growing at 7 percent a year, in some large enterprises, however, networked data storage—disk drives, tape systems, specialized network gear, and the people and software to manage these—has grown by 20 percent or more annually, even where IT budgets have barely budged. The demand for storage has grown by more than 50 percent annually in recent years, even faster than the rapidly decreasing unit cost of storage (Exhibit 1). If storage costs continue their rapid rise, they could make it harder for companies to store and exploit new forms of data—such as tick-level financial data, digital images in life sciences, or video in media companies.



Large enterprises must manage storage more efficiently if they are to exploit opportunities created by new forms of information—such as more detailed financial data, digital images in life sciences, or video in media companies. Unfortunately, it's a daunting challenge for IT managers to combat the rising volume, cost, and complexity of storage. They must grapple with a wide range of decisions about storage architecture, design, operations, and performance. And business users—the internal customers of IT—commonly ask for more than they need because they aren't accustomed to making decisions about the trade-offs between storage cost and quality. (In many cases, storage decisions result from policies designed by legal and regulatory teams, which focus primarily on risk mitigation.) Some IT managers, in turn, err on the side of caution to safeguard against the personal costs they would bear if data were lost. As a result, companies store noncritical data on expensive configurations or keep too many copies of information in too many places.

IT must implement practices and policies that clearly define the storage options available to business users and that help them set priorities based on storage economics and business

needs.¹ The starting point is often a mind-set shift: IT managers need to engage with business leaders in a more systematic way on the trade-offs of investments in risk, performance, and cost. This approach includes providing business customers with a set of standard, productized storage configurations that define clear service levels and unit costs. By working closely with the business, IT leaders can help their internal customers make better decisions to balance the competing demands of storage, cost, resiliency, and performance. What's more, they can take important steps to streamline the process for making new storage requests, planning future demand, and procuring hardware and software. One large enterprise reduced the cost of meeting its demand for storage by 40 percent; it was then able to use the savings to support new capabilities and applications that the business would soon require.²

Struggling to manage storage

New applications, more complex business analytics, and the need to meet regulatory requirements are major contributors to the demand for additional storage capacity. Pharmaceutical companies, for example, are considering digital-storage options for all of their imaging data (such as MRIs)—a change that could speed getting new compounds to market but doubles the storage needs of these companies. Similarly, a hedge fund's storage costs could increase exponentially as it captures and stores more granular "tick level"³ trading data to meet financial regulations and develop new trading strategies.

In many enterprises, the storage environment has grown so rapidly that it has outrun the IT department's ability to manage it effectively. As a result, many organizations are experiencing a common set of problems.

Too many copies of data

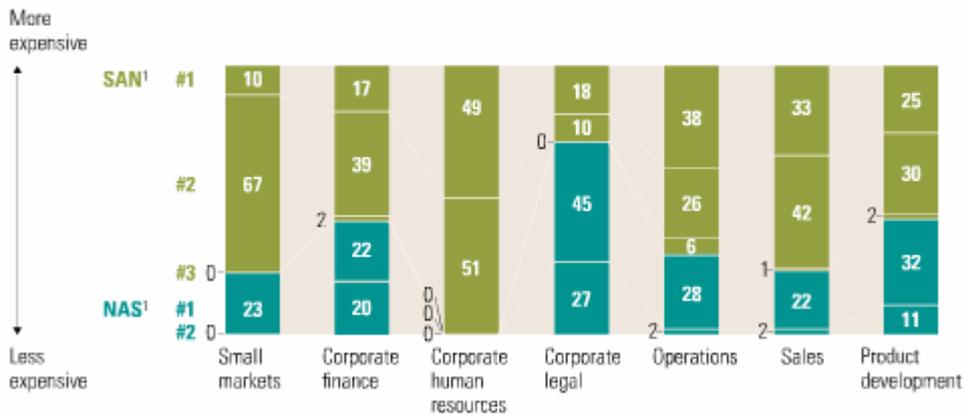
In many cases, companies find it quicker and easier to back up information by making multiple copies. But across a large organization, this approach increases storage volumes dramatically as companies keep too many copies of business data. An electronics manufacturer, for example, kept as many as 12 copies of its core production applications, each with its own full copy of data, thereby tripling the storage needs for these applications. The problem is further complicated by rampant replication. Network storage derives its resiliency, in part, from the replicas of underlying business data it creates, but multiple replicas increase the demand for storage. An 11-terabyte database can require 170 terabytes of raw storage for its replicas, depending on its configuration. One pharmaceutical company stored ten copies of clinical-trial information indefinitely when four would have met its legal obligations. Moreover, very few companies archive and purge outdated information aggressively, either because they have an aversion to risk or they believe it's not worth the expense, further adding to the storage burden.

Placing data on inappropriate storage configurations

For every storage decision, managers have to consider the type of network, the size of drives, and the required mirroring and replication.⁴ Each choice has implications for storage cost, but these decisions are often made haphazardly and can place data on higher-performance storage alternatives than some applications require (Exhibit 2). For example, one electronics manufacturer stored 70 percent of its test and development data on expensive storage options, even though much of this data typically didn't have stringent requirements for uptime or performance.

Inconsistent storage decisions

Average storage by division for disguised example of Fortune 50 enterprise, %



¹SAN = storage-area network; NAS = network-attached storage.

Source: IDC; McKinsey analysis

Buying excess capacity

Well-managed companies use 80 percent or more of their available storage, but in others that figure hovers around 40 to 50 percent. One large IT organization used only 50 percent of its storage capabilities. Some of its individual storage systems were at just 10 to 20 percent of capacity, and one of its businesses utilized only 33 percent of the entire amount of storage it had requested. One could argue that storage needs always catch up with capacity, so having a bit of excess isn't so bad. But as storage consumes a growing share of IT budgets, the costs of maintaining excess capacity are becoming material.

Why effective management is difficult

Poor storage decisions stem from a variety of underlying issues, most notably time pressure, information gaps, and breakdowns in the relationship between business users and IT. Frequently, these challenges are exacerbated by a capability gap: the storage manager is promoted for technology skills rather than for business and financial acumen. This focus on technology limits the manager's ability to understand the total costs of storage and to communicate those costs and the related trade-offs in business terms to the rest of the company.

Time pressure

It's easy for IT managers to make mistakes when they don't have time to explore the long-term implications of their storage decisions. Managers under cost pressures are sometimes swayed by a particular solution's quick turnaround time. This situation leads to an ad hoc deployment of networked storage in response to individual business requests. Even where all stakeholders have agreed on a single vendor, multiple configurations abound.

Information gaps

Many organizations underestimate costs by failing to include in the analysis hidden expenses such as network hardware, short-term backup, and long-term archiving. It's easy to miscalculate because storage costs are relatively complex and depend on pricing options, configuration decisions, and labor productivity across multiple organizations, including the vendor, IT, and the business. Furthermore, many IT organizations lack the basic data required to manage a complex storage environment. We've seen large enterprises that didn't measure storage use or growth, track storage requests by business, or map storage volumes to applications. The absence of this type of information makes it all but impossible to optimize

storage—for example, by releasing unused capacity, removing unnecessary duplicates and replicas, and tightly aligning capacity and growth in demand.

Breakdowns in the relationship between business users and IT

The process for requesting and allocating storage is often inefficient or too complicated. Application developers working with business units sometimes avoid the process altogether by overestimating the storage needed to fulfill a request, resulting in low utilization rates. Even when everyone is fully engaged in the process, it can be difficult to understand business requirements—as in the pharmaceutical example mentioned earlier, where a storage team created too many replicas of clinical-trial data because it never received accurate requirements. Similarly, when business units request storage options, central IT organizations sometimes misunderstand and fail to deliver exactly what their internal customers need. A business unit CIO might ask for a low-cost storage option to hold temporary, analytical, or nonproduction data, but often the storage group doesn't offer a serious low-cost option. One storage team developed a service catalog intended to provide the business unit CIO with several choices, but the least expensive practical option was only 20 percent cheaper than the most expensive path.

Managing storage better

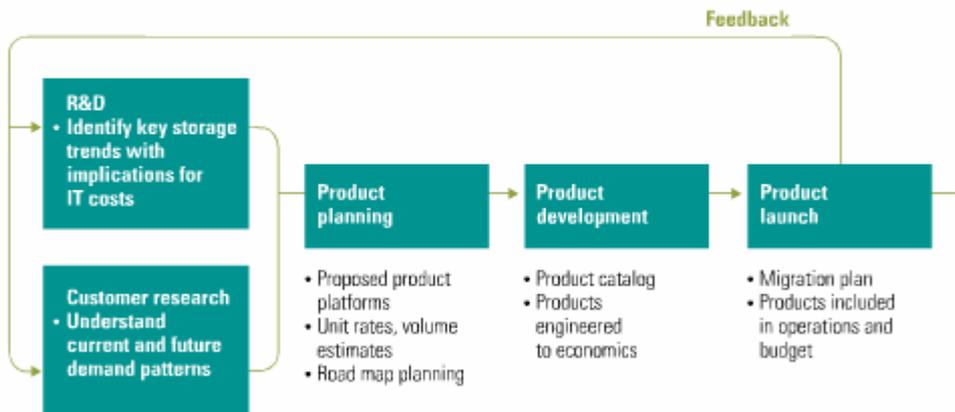
So how can IT departments improve the way they manage data and deliver solutions to meet the rising demands of the business without allowing costs to spiral out of control? Offering a standard set of options and clearly communicating their value to the business can help storage managers accomplish this goal.

Present a standard set of options

Adopting a standard set of options and working with the business to refine and align those options to match business needs can help large enterprises manage storage better (Exhibit 3). Storage organizations should articulate these offerings as end-to-end storage solutions, at guaranteed service levels, rather than technologies (Exhibit 4). This gives the storage manager the freedom to provide the business with the storage in the most efficient way. For example, where companies have developed too much storage capacity or saved too many replicas, the storage organization can work to rationalize those storage volumes, freeing space for new growth. Carefully aligning storage volumes to requirements can result in 30 to 40 percent of volumes being moved to lower cost storage options. Increasingly, leveraging storage virtualization—software that makes it easier to manage storage across individual frames—is a key element in making better use of available storage.

A product-development approach to storage

Illustrative example

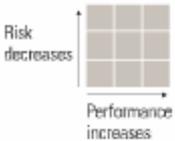


Source: IDC; McKinsey analysis

A utilização deste artigo é exclusivo para

Service-level catalog

Disguised example of global medical-products company

Storage tier ¹	Risk/performance profile 	Unit cost, \$ per gigabyte of business storage	Sample metrics			
			Performance metrics	Risk		
			I/O transfer rate, megabytes per second	Seek time (minutes)	Monthly full backups	Reliability ²
Platinum SAN <ul style="list-style-type: none"> • SAN array with FC drives • Mirrored, remote replication; business continuity volume 		44	200	3.5	4	Five 9s
Silver SAN <ul style="list-style-type: none"> • Low-end SAN array with FC drives • RAID 6 		14	200	3.5	4	Four 9s
NAS <ul style="list-style-type: none"> • NAS array with SATA drives • RAID 5 		9	150	8.5	4	Four 9s
Near-line archive <ul style="list-style-type: none"> • Low-end NAS array with SATA drives • RAID 6 		5	150	8.5	1	Three 9s

¹SAN = storage-area network; NAS = network-attached storage; FC = Fibre Channel; I/O = Input/Output; RAID = Redundant Arrays of Independent Disks; SATA = Serial Advanced Technology Attachment.

²Five 9s means reliable 99.999% of time, or ~5 minutes downtime per year; four 9s means 99.99%, or ~52 minutes downtime; three 9s means 99.9%, or ~9 hours downtime.

Source: IDC; McKinsey analysis

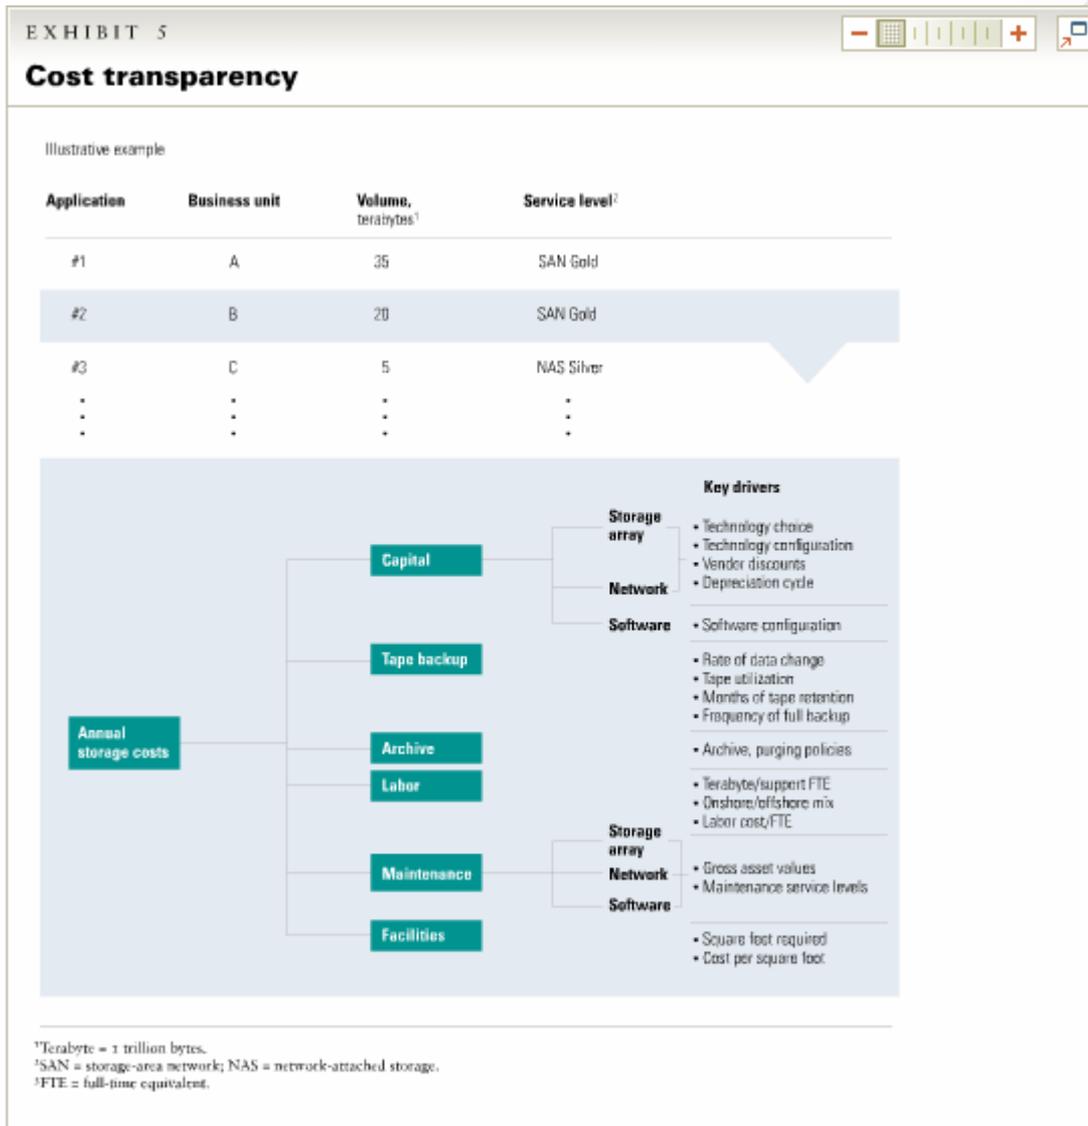
As needs and technology evolve, storage managers might have to generate new options that then become part of the standard menu. Sometimes these new options will be relatively inexpensive. For example, storage solutions based on near-line technologies (that is, backing up information with devices close to the machines rather than online at remote locations) can provide very low-cost storage for data that does not have to be accessed frequently. In other cases, these new options may be more expensive—and better able to provide superior performance for critical data. Solutions that employ solid-state disk technology, for example, might be necessary for applications such as financial trading, which may require immediate and frequent access.

Communicate in business terms

As offerings are developed, storage managers will have to improve the way they communicate the economics of these options—specifically, they'll need to translate technical terms into language that executives are familiar with. Business users struggle to understand storage terminology and often correlate it with factors they can control, such as the amount of information they're storing, the reliability with which they need it stored, their ability to access it, and the length of time it must be retained. IT managers should present options in business terms, such as the amount of business data stored, rather than raw disk space consumed. They should provide line-by-line transparency of storage costs (Exhibit 5) and offer a wider range of options—for example, the highest-priced solution should probably be three to four times the cost of the least expensive. Here, too, storage vendors may have a role to play in

helping to explain the economics of storage to internal business customers. Articulating options in language executives can understand will help build confidence in the storage organization.

It's important for storage managers to explain service levels clearly and commit to meeting them. Without this information, many customers will select the least risky and highest-cost option. To encourage developers and business users to make effective trade-offs, the storage team must be able to define how much longer it would take to recover data stored on one service compared with another. They should back this information up with historical metrics on recoverability (mean time to recovery) and availability (mean time between failures) for the different offerings. Articulating the types of applications (for example, enterprise transactional systems, data warehouses, intranet sites) and corresponding storage service levels will also help drive storage volumes to the right service.



Support and standards

Organizational change of this magnitude depends on support from senior leadership, who must be committed to these efforts and communicate their importance to the company through ongoing dialogue and performance measures. The organization must implement incentive structures based on metrics supplied by internal customers, such as service levels, as well as targets the IT organization sets for itself, such as unit cost of operation. Senior executives should also work with IT managers (including storage managers) to help them better understand the business of the organization they serve. Some IT organizations, for example, have encouraged high-performing IT managers to develop business skills through part-time training.

In addition, it's critical to create a standard approach to coordinate the planning process for IT with the business's annual and quarterly planning cycle. Managers can then use this plan to direct other key IT processes, including procurement, storage rationalization, and the commissioning and decommissioning of stored data. Since users often overestimate their needs, new demand-management systems should include performance incentives for accuracy in forecasts and for compliance with estimates.

Disponível em: <<http://www.mckinseyquarterly.com>>. Acesso em 25/6/2008.

A utilização deste artigo é exclusivo para fins educacionais.