July 2003

Tivoli® software

IBM TotalStorage®

IBM storage solutions:
Reducing cost and complexity

Contents

Executive summary

In many environments, today's storage costs, as a percentage of IT capital expenditures, have grown to the point that storage now represents the largest portion of capital spending. These complex storage environments—with multiple device types, implementations of advanced functions like replication, and device and storage management software interfaces—also make it more difficult for administrators to effectively manage storage resources—driving management costs up as well.

The goal for IT managers concerned with this trend toward higher capital and management costs in their storage environments is to find ways to reduce the complexity that their administrators must deal with and to reduce overall storage expenditures. (See Figure 1).

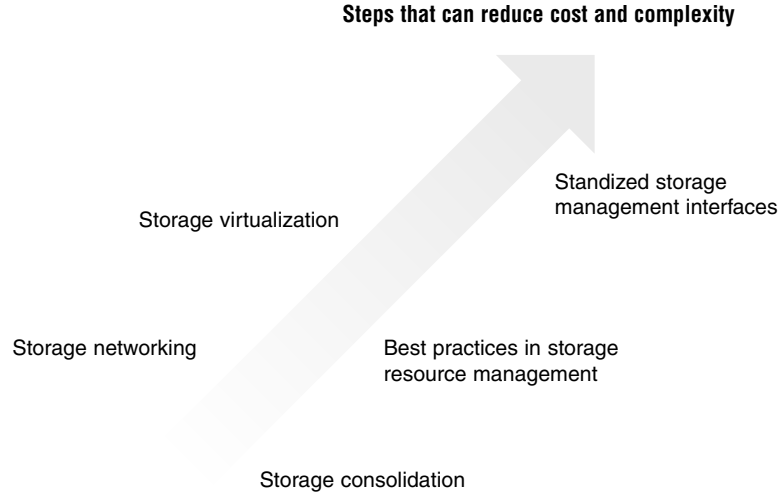

Figure 1. IBM storage solutions are designed to help IT managers take steps to improve the management of existing storage resources while reducing the complexity of their storage infrastructure.

To help IT managers reduce the cost and complexity in their storage environments, IBM is delivering solutions that:

- *Implement best practices for improving the management of existing storage resources.*
- *Introduce virtualization into the storage environment to reduce the complexity that administrators must deal with.*
- *Exploit common, open interfaces to help reduce cost now while maintaining flexibility in the future.*

Highlights

- *IBM Tivoli Storage Resource Manager*
- *IBM TotalStorage Virtualization Family*
  - *IBM TotalStorage SAN Volume Controller*
  - *IBM TotalStorage SAN File System*
  - *IBM TotalStorage SAN Integration Server*

Organized under both the IBM Tivoli® software and IBM TotalStorage® brands, IBM is working to deliver comprehensive storage software and hardware solutions that help IT managers evolve into an on demand operating environment.[1] This paper focuses on the following current and upcoming members of IBM's comprehensive storage portfolio (see Highlights to the left) and a selected list of their integrated capabilities for reducing cost and complexity.

Complexity breeds cost

IT managers seldom list "I increased complexity" among their most significant accomplishments. No, in fact, most IT managers strive to reduce the complexity that their administrators must deal with regularly. The reason is that this complexity is the breeding ground for errors that can cause application outages and for undesirable management overhead that can drive up cost.

Complexity can creep into a storage environment for many reasons:

- *There may be a multivendor procurement strategy—the goal being to encourage competition and drive down capital costs. The unfortunate side effect, however, is a storage environment with devices having different configuration interfaces and different replication service application programming interfaces (APIs) and functions. This combination can lead to an increase in management cost and a decrease in advanced functions and application availability.*
- *Computing environments from different departments, or even different companies—in the case of a corporate merger or outsourcing—may have to be consolidated. Having evolved separately, it is likely that these environments use different host operating systems, storage hardware, storage networking equipment and storage management software. The resulting management overhead can make it difficult to realize the cost reductions that drove the consolidation in the first place.*
- *Heterogeneous storage infrastructure components or storage management software may be introduced as the result of tactical decisions driven by short- or medium-term cost pressures. Strategic decisions by IT managers to switch storage vendors may also result in short-term complexity while the transition is in process.*

Highlights

Whatever the reason, complex environments with multiple device types, multiple implementations of advanced functions like replication, and multiple device and storage management software interfaces make it more difficult for administrators to manage storage resources effectively. It also makes it more difficult for businesses to leverage the capabilities of their storage resources for improving application availability. Overall costs go up and effective use of storage assets goes down.

Reduce cost and complexity with IBM storage solutions

IBM is working to deliver solutions that help IT managers dramatically reduce the complexity that their administrators must deal with while driving cost out of the storage environment. Organized under both the IBM TotalStorage and IBM Tivoli software brands, these integrated solutions are modular in design, each component with well-defined functions and open interfaces, making it possible for IT managers to construct, upgrade or enhance their on demand operating environments one component at a time.

Improve the management of the resources you have

Effective storage resource management includes four evolutionary steps—Identification, Evaluation, Control and Prediction.

Storage costs, as a percentage of IT capital expenditures, have been growing at a steady pace for several years. In many IT environments, storage now represents the largest portion of capital spending. One technique that IT managers have used for driving cost out of their storage environments is to manage this expensive resource using the information and automation available from storage resource management software tools. Suggested best practices for effective implementation of storage resource management tools include four evolutionary steps – Identification, Evaluation, Control and Prediction. *IBM Tivoli Storage Resource Manager* software is designed to deliver the functionality that IT managers can use in implementing these best practices.

Identification

As the first step in effective storage resource management, you can use Tivoli Storage Resource Manager to perform an automated identification of the storage resources in an infrastructure, deliver an analysis on how effectively those resources are being used and highlight areas that are at risk of causing application outages (see Figure 2).

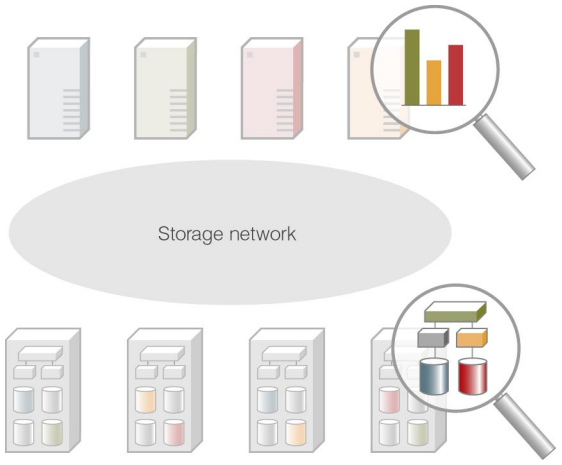

Figure 2. IBM Tivoli Storage Resource Manager automatically identifies and evaluates storage device, file system and database resources in an environment.

Tivoli Storage Resource Manager includes a series of preconfigured storage-asset inventory reports that first enable IT managers to verify that the expected storage assets are available and online. Next, IT managers can use capacity-utilization reports to quickly profile the usage of their storage assets. The results are sometimes unexpected and can point to areas of significant opportunities for improving asset effectiveness. In a recent study,[2] IT managers indicated that their average storage utilization in open-system environments was 48%—well below the utilization that they required. Often times, this low utilization is the result of the lack of effective tools for allocating and manipulating physical storage, which can be improved by introducing virtualization—a topic discussed later in this paper. However, with this much free space in the environment, it is easy to overlook the fact that all resources do not conform to the average. Some assets don't come close to the average, perhaps hovering around 10 - 20% utilization

Highlights

and, of greater concern, others can be dangerously high, topping 90 - 95% utilization. Tivoli Storage Resource Manager identifies these "at-risk" file systems and disk volumes as being potentially close to causing an application outage due to out-of-space error conditions.

Another category of at-risk assets are disk subsystems that are surfacing errors to the hosts whose data they store. With this storage network-wide view of subsystem errors, IT managers can quickly identify application data that could be proactively moved to avoid application outages.

Evaluation

As the second step in implementing best practices for managing storage resources, evaluation represents a deeper level of analysis about the file allocation and usage patterns that are driving storage-resource utilization. File-level evaluation uncovers several categories of files that, if deleted or archived, can potentially represent significant reductions in the amount of data that must be stored, backed up and managed:

Many categories of files—including orphan, obsolete and misused file types—can be managed effectively to reduce the amount of data that must be stored and backed up.

- *With a dynamic work force, IT environments are subject to constant change as users of the environment come and go—sometimes leaving behind data that was unique to an individual. Orphan files that are owned by a user who is no longer active in the environment can represent a drain on storage resources.*
- *With normal business processes, there are files that get referenced every day—some that are part of weekly activities and others that are part of monthly, quarterly or annual reporting cycles. Obsolete files that reside on primary storage without having been accessed for periods longer than that are often candidates for archiving or deletion—freeing up storage capacity and reducing the amount of data that has to be managed.*
- *Many businesses today have policies against storing personal data or certain kinds of frequently misused data on corporate storage assets. Some common file types included in this group are .mp3 music files, .bmp graphic images and .avi video clips. Enforcing these existing policies on an environment-wide basis can again result in a significant reduction in the amount of data that must be stored and, in some cases, can potentially limit the liability incurred when employees use corporate assets to store questionable material.*

Tivoli Storage Resource Manager can help IT managers evaluate these and many other categories of files, potentially resulting in a better understanding of the dynamics that are driving storage utilization.

Control

The next step in managing existing storage resources is to take the knowledge gained during identification and evaluation and begin to develop policies for automated control. These policies are customizable with actions that can include centralized alerting, distributed responsibility and fully automated response (see Figure 3).

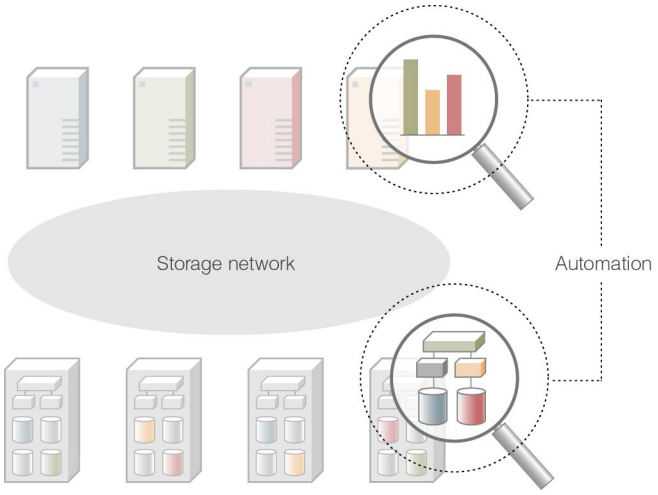

Figure 3. Tivoli Storage Resource Manager gives IT managers automated control over storage resources.

In many cases, there are events where the immediate attention of an administrator may be useful in avoiding larger problems—such as finding a new disk defect, a file system that has grown over 95% utilization or a key application directory that has crossed some size threshold. The Tivoli Storage Resource Manager alerting infrastructure is fully integrated with the *IBM Tivoli Enterprise Console* for centralized, enterprise-wide event management. Alerting can also be integrated with other systems management infrastructures through Simple Network Management Protocol (SNMP) or can be directed to individual administrators through e-mail.

Highlights

Establishing policies for automated control of storage resources helps IT managers maintain an efficient and available storage environment.

Establishing disk-usage limits, or quotas, for individuals and work groups is an effective technique for controlling storage growth. Like other events, quota-violation alerts can be managed centrally. However, it is often a more effective practice to distribute quota-violation alerts directly to the owners of the files. This encourages individuals to take responsibility for proper use of corporate storage assets. In addition to quotas, which look at total usage, another effective practice can be for IT managers to put constraints on certain categories of files, like obsolete files and many of the misused file types discussed earlier. Alerts for this type of violations can also be sent directly to individual users, further encouraging them to take personal responsibility for effectively managing storage assets.

Policy violations can also trigger automated responses. For example, IT managers may choose to automatically archive obsolete or orphan data to tape, freeing the online storage assets. Tivoli Storage Resource Manager is fully integrated with the *IBM Tivoli Storage Manager* archive function to deliver this level of automation. Tivoli Storage Resource Manager also includes the capability to trigger a custom script giving IT managers the ability to integrate triggered actions with other products as well.

With Tivoli Storage Resource Manager, IT managers can also automate the provisioning of new storage capacity in response to a file system exceeding a utilization threshold. For example, if a file system that houses a key business application exceeds 95% utilization, Tivoli Storage Resource Manager can identify unused storage in the existing storage network, assign that storage to the at-risk host and extend the file system onto the new storage in a completely automated operation.

In environments where storage is offered as a service—perhaps when a central IT organization charges individual business lines for IT services or when storage services are outsourced—service-level reporting and chargeback can also be effective tools in controlling storage utilization growth. Tivoli Storage Resource Manager gives IT managers the ability to produce reports for much of the information and many of the actions discussed thus far—setting a common base for service-level reporting. Additionally, *Tivoli Storage Resource Manager for Chargeback* assists IT managers in associating the cost of storage with individual users, departments, locations or other defined groups within the enterprise.

Prediction

A fourth step that can be taken to manage storage resources effectively is to use historical information to predict future growth and future at-risk conditions. Data collected during identification and evaluation can be retained in the Tivoli Storage Resource Manager repository and used to provide trending analysis. These trends can be used to track positive progress against storage resource management goals and they can be used to assist IT managers in understanding where—and when—new capacity may be needed. In environments that have not taken the important step of consolidating storage resources on a storage network, capacity trend analysis can also be used to identify applications that could benefit from this consolidation—setting in place the physical infrastructure required to move underutilized storage resources to areas where new capacity is forecasted (see Figure 4).

Figure 4. Tivoli Storage Resource Manager helps IT managers understand where—and when—new capacity may be needed.

These four evolutionary steps in storage resource management—Identification, Evaluation, Control and Prediction—can help IT managers drive cost out of their existing storage infrastructure. Enhancing the infrastructure with virtualization can drive further cost out of the environment and reduce the complexity that administrators must deal with.

Highlights

A consolidated storage network provides the necessary foundation to derive value from storage virtualization.

Introduce virtualization

Virtualization is a term used to describe the basic concept of separating the logical representation of a resource from its physical implementation. One common example of virtualization is a virtual private network (VPN). When a remote mobile computer user connects to a central computing infrastructure through a VPN, the logical representation is that the computer is directly and securely connected to the corporate network. However, the physical implementation of that connection may be through a local Internet service provider (ISP), across the public Internet and through a corporate firewall. The virtualization in the VPN connection masks this complexity from the remote user. It also enables the remote user to move from dial-up to high-speed connection, from home to hotel, or from country to country and still present a consistent "directly connected" logical view to the applications on his or her mobile computer. By separating the logical representation from the physical implementation, physical network infrastructure changes may be made almost at will without any change in the operation of host applications.

The concept of virtualization can also be introduced into the storage environment to help reduce the level of complexity that administrators must deal with while insulating host applications from storage infrastructure changes that, traditionally, resulted in downtime. As mentioned before, using storage networks to consolidate storage resources can be an important step in reducing the cost of a storage environment. Storage networks also serve as the base upon which the value of virtualization can be introduced into your storage environment.

Highlights

Block virtualization provides a consistent view of disk storage resources to host systems and applications.

Block virtualization

The first level of storage virtualization deals with removing the complexity of managing disk devices. This is accomplished by inserting a block virtualization layer between the host systems and the physical storage devices (see Figure 5).

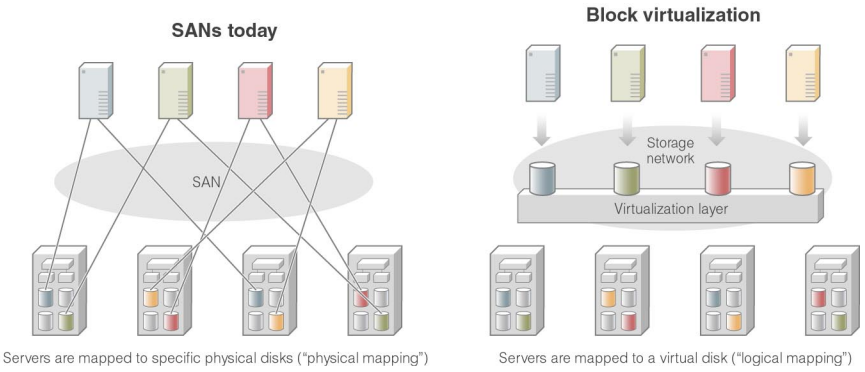

Figure 5. Block virtualization reduces complexity by insulating host systems and applications from changes in the physical storage infrastructure.

The block virtualization layer provides a consistent view of disk storage resources to host systems and applications. This virtualized view is then logically mapped to the physical storage, which means that—like the VPN—the physical infrastructure may be changed almost at will, typically without host systems or applications needing to know that changes are being made. For example, an administrator can move data from one physical storage device to another—even if the two devices are different models or from different vendors—without host systems needing to be aware of the change and without the need for planned application outages. Storage hardware from different vendors can be maintained in the storage environment—by design, as the result of consolidating existing data centers or as a short-term transition from one strategic partner to another—while still presenting a common and consistent virtualized view of storage to the host systems.

Another characteristic of block virtualization is virtualized replication services. Today, many high-end storage controllers include some form of local or remote replication capability, or both. IT managers routinely use replication capability to assist in data protection. Local replication in particular is also used to speed application development and testing—a practice whose usage peaked during recent "mass change" events like Year 2000 (Y2K) and the euro conversions. Analysts have noted that use of replication services is not nearly as pervasive as it could, or should be. It has been suggested that one of the driving reasons for this is the inherent complexity and cost involved in replication. This comes from two realities. First, different vendors implement replication services in different ways—different levels of functionality and different interfaces—making it complex to implement and reducing IT manager flexibility in the future. Second, traditional replication services are expensive. Most implementations are on high-end disk controllers and require that both the primary and the replicated copy of data be on the same type of high-end controller. Block virtualization has the potential to solve both of these challenges.

First, the block virtualization layer can provide a single, common method for performing both local and remote replication. IT managers can confidently integrate applications and processes with this common set of interfaces without the looming threat of the interface changing when physical storage devices are changed. With this new flexibility, administrators are free to swap or mix-and-match physical storage without impacting the operation of business processes. Second, because the replication service is virtualized, IT managers have the flexibility to choose high-end storage for the primary copy of data and low-cost storage for the replicated copy—driving cost out of the implementation of replication services.

IBM offers block virtualization in two packages, both part of the *IBM TotalStorage SAN Virtualization Family.*

Highlights

The first package, the *IBM TotalStorage SAN Volume Controller*, is designed for IT managers who have already taken the first step of consolidating storage resources on a storage network. The SAN Volume Controller is an appliance that IT managers can use to insert block virtualization into existing storage networks, insulating host systems from physical storage changes, and driving cost out of both managing storage resources and implementing replication services.

The second package, the *IBM TotalStorage SAN Integration Server*, is designed for IT managers who have not taken the initial step of consolidating storage resources on a storage network. The SAN Integration Server includes the foundational components of a virtualized storage environment. The SAN Integration Server integrates, at manufacturing time, the SAN Volume Controller with IBM TotalStorage SAN switch, IBM TotalStorage FAStT storage, IBM Tivoli Storage Area Network Manager, rack, and cabling to deliver a complete SAN environment.

File virtualization

The second level of storage virtualization deals with removing the complexity of managing files. This is accomplished by implementing a common file-management utility in the storage network (see Figure 6).

File virtualization introduces a utility view of file resources that host systems "plug-in" to, giving them visibility to all file resources and the ability to share data.

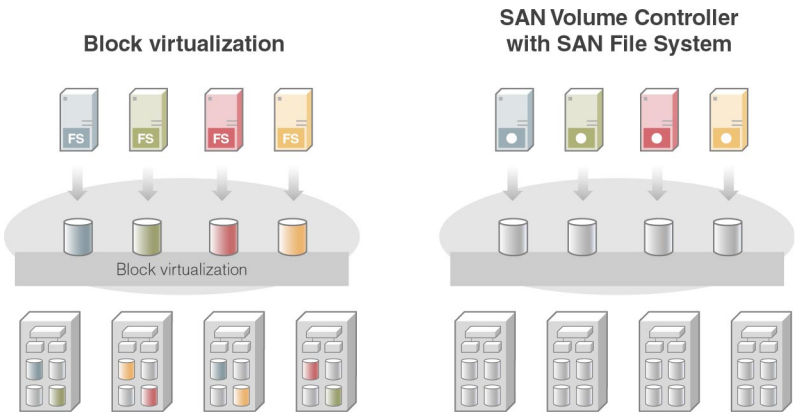

Figure 6. File virtualization has the potential to help reduce cost by eliminating the need for duplicate files and enabling policy-based data placement on different quality-of-service storage.

Highlights

The common file utility also delivers policy-based placement of files so that administrators can ensure proper utilization of different quality-of-service storage assets.

The *IBM TotalStorage SAN File System*, planned for availability in late 2003, has been designed to enable host systems to "plug-in" to a common SAN-wide file structure. With the SAN File System, files and file systems are no longer owned by individual computers, they are viewed as a centralized IT utility. This utility view of file resources has several positive benefits for reducing cost and complexity.

First, data needed by different applications no longer has to be duplicated. For example, many IT environments run enterprise resource planning (ERP), human resources and various custom applications. Although all the systems that these applications run on may be connected to a consolidated storage network, the individual file systems are not shared. In order to make this data available to other systems—for consolidated business reports, for example—processes would need to be built for extracting data from the various platforms, shipping it across a TCP/IP network and managing retention of the duplicate data. By introducing a common file structure, SAN-connected systems are given "visibility" to the same files—enabling the sharing of that data. Data extracts no longer need to be copied around the network. A reporting system can directly read the data created by the various business applications running on other platforms. This provides the opportunity for significantly reducing the cost of creating, storing, backing up and managing duplicate data.

Second, administrators no longer need to deal with the complexity of mapping individual host file systems to specific logical volumes for appropriate data placement. In many IT environments, business drivers dictate the specific placement of files to achieve a desired quality of service. Some examples include database files that need the availability delivered by a Redundant Array of Independent Disks (RAID) and the recoverability delivered by remote replication; database log files that need to be physically separated from the primary database; shared work directories that need the availability of a RAID but not remote replication; and individual home directories that only require low-cost storage. In the past, data placement has been directly tied to the physical mapping of the host file system. Data placement across a consolidated group of tens or hundreds of servers has meant that administrators had to individually visit every server and carefully map host file systems to the appropriate type of volumes.

Highlights

With the common file utility supported by the SAN File System, this level of complexity can be removed. IT managers will be able to use a centralized set of policies to govern where data is placed. When applications or individuals on any of the consolidated host systems create a file, the SAN File System uses these policies to direct the new file to the appropriate storage. By implementing this SAN-wide file utility with its central point of control, the complexity of dealing with individual file systems and their physical mapping to storage volumes has the potential to be dramatically reduced.

Standardize management interfaces

Another area of focus that IT managers can leverage to reduce cost and complexity in their storage environments is standardizing management interfaces – specifically the programmatic interfaces that storage infrastructure components and storage management software use to communicate with each other. IT managers are well familiar with the value of a common management interface. Consider the SNMP in traditional networking environments. Can anyone imagine implementing the scale of traditional networks that we enjoy today if there were no standard way for network management tools to discover, classify, monitor and configure the network infrastructure components? What if, instead of a common standard, individual vendors had developed their own proprietary mechanisms for performing these functions? IT managers would be forced to work around the complexity introduced by multiple interfaces and would not enjoy the flexibility to construct their traditional network environment from heterogeneous components.

Exploiting common, open interfaces like SMIS can help IT managers reduce cost now while maintaining flexibility in the future.

Fortunately, a common management interface for storage networks is evolving very rapidly. The interface, Storage Management Interface Specification (SMIS) is the result of a collaborative effort involving many of the leading storage vendors and is overseen by the Storage Networking Industry Association (SNIA).[4] SMIS reduces the complexity of dealing with multiple, vendor-proprietary management APIs and gives IT managers flexibility in making future storage infrastructure and storage management software decisions.

IBM is part of a growing list of storage vendors who are delivering products that implement the SMIS interface. The three primary products discussed in this paper—IBM Tivoli Storage Resource Manager, IBM TotalStorage SAN Volume Controller and IBM TotalStorage SAN File System (planned for availability in late 2003)—along with the IBM TotalStorage Enterprise Storage Server® (ESS) have been delivered (or are planned at initial availability to be delivered) with SMIS interface capability. When choosing new components to evolve their storage environment, IT managers should consider components that are actively delivering SMIS functionality.

Conclusion

Reducing complexity in the storage environment is a significant challenge that can directly contribute to both capital and management cost reductions. To help IT managers achieve these reductions, IBM is working to deliver open storage solutions that implement best practices for improving the management of existing storage resources, introduce virtualization into the storage environment and exploit common, open interfaces to help IT managers reduce cost now while maintaining flexibility in the future. These integrated solutions are offered in modular components, so IT managers can construct, upgrade or enhance their on demand operating environments one component at a time.

For more information about IBM e-business on demand™ and storage, go to **ibm.com**/ondemand, **ibm.com**/totalstorage and **ibm.com**/tivoli

IBM storage solutions:
Reducing cost and complexity

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

IBM, the IBM logo, e-business on demand, Enterprise Storage Server, Tivoli and TotalStorage are trademarks of International Business Machines Corporation in the United States, other countries, or both.

Other company, product or service names may be trademarks or service marks of others.

[1] For a more complete discussion of the value of an on demand operating environment, see the IBM white paper *IBM Storage Solutions: Evolving to an on demand operating environment.*

[2] For more information, see the IBM white paper *IBM Storage Solutions: A total cost of ownership study.*

[3] For more information about SMIS, see the IBM white paper *IBM Storage Solutions: Delivering the value of an open environment.*

[4] For more information about the Storage Networking Industry Association, visit www.snia.org.

IBM Systems Group
5600 Cottle Road
San Jose, CA 95136
U.S.A.

07/03

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This information could include technical inaccuracies or typographical errors. IBM may make improvements and/or changes in the product(s) and/or program(s) at any time without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead. It is the user's responsibility to evaluate and verify the operation of any non-IBM product, program or service.

THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR INFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted according to the terms and conditions of the agreements (e.g., IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. IBM is not responsible for the performance or interoperability of any non-IBM products discussed herein.