Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com

WHITE PAPER

IBM TotalStorage Software: Building Storage Solutions in Alignment with Current and Future Business Requirements

Sponsored by: IBM

Richard L. Villars

October 2004

Executive Summary

In today's ever more competitive business environment, the key challenge for CIOs and other IT executives remains aligning IT investments with general business needs. The three key elements in successful efforts to align IT and business needs are:

- Consolidating/simplifying IT infrastructure and operations to more effectively deliver needed applications and services to business units
- Implementing robust business continuity for both mission-critical and general applications as well as the underlying information that businesses generate
- Providing simpler but more intelligent systems for managing the creation, protection, retention, and ongoing use of business information

Today, many companies are only just beginning to recognize the challenges that meeting these key elements pose for existing IT architectures and IT organizations. In their continued effort to meet the sometimes conflicting cost, service-level, business continuity, and new application demands of business executives, IT managers need to adopt a new strategy for deploying server and storage solutions. These new solutions must provide better performance, higher reliability, greater capacity, and lower cost of operations while making it easier to reuse existing assets and provision new systems as needed.

This IDC white paper examines emerging technologies with a focus on storage technologies that allow companies to develop an architecture that can grow and adapt to meet current and future business requirements. These include the deployment of a tiered storage infrastructure, virtualization, and integrated server/storage management. It also examines several existing tiered storage and storage management products within IBM's TotalStorage software product family.

This white paper also discusses how IT managers can use these products (SAN Volume Controller, SAN File System, TotalStorage Productivity Center, and the Tivoli Storage Manager) to develop an effective and future-proof storage architecture that better serves internal business units, improves application reliability, and streamlines existing management policies and procedures.

Three Key Goals Drive IT Investment Strategies

It is a well-known adage in the IT industry that CIOs and other IT executives must focus on aligning IT investments with general business needs. The key, of course, is figuring out what those needs are today and anticipating how they will evolve over time. In our discussions with executives about their efforts to align IT with business goals, three goals consistently emerge:

- **Consolidated, simplified infrastructure and operations:** IT organizations need to enhance their ability to deploy, scale up, and upgrade general IT and business application services to meet service level agreements made with business units.

- **Robust business continuity:** IT organizations must ensure available and reliable information systems for both mission-critical and general applications (e.g., email) as well as improve time to recovery in case of failures.

- **Improved information life-cycle management (ILM):** IT organizations must provide solutions that allow business managers to quickly organize, access, and mine all corporate information, throughout its life cycle, to meet regulatory compliance mandates and improve general business analytics.

These three goals are the primary driving forces behind current IT investment strategies. IT organizations are consolidating and simplifying their operations, strengthening their business continuity plans, and sharpening their ILM capabilities.

Consolidation and the Demand for Better Service

From the perspective of IT executives, consolidation is consistently one of the most widely cited goals driving their overall IT strategies and investment decisions.

For the past six years, IDC repeatedly asked IT managers at medium-sized and large U.S. companies if they were currently undertaking any IT consolidation efforts. As Figure 1 shows, the percentage of companies engaged in consolidation efforts has consistently increased from 1999 to 2004. Today, one can easily say that IT consolidation is the norm for most companies. However, between the late 1990s and today, the underlying goals driving consolidation are changing.

FIGURE 1

IT Consolidation Efforts, 1999, 2002, and 2004

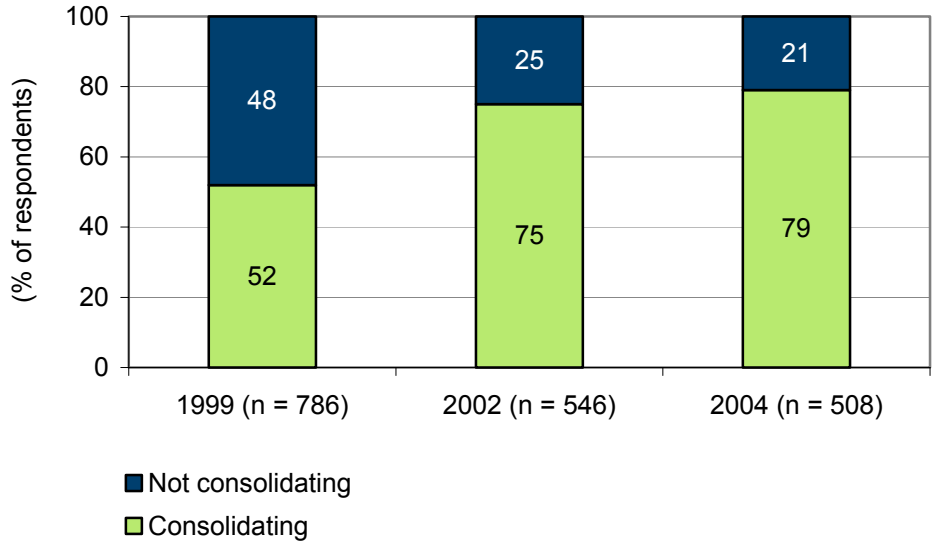

Notes:

Consolidation initiatives increase with company size.

Central IT is the main driver.

Business units have significant influence at large companies.

Source: IDC, 2004

Enterprises have a wide range of goals for their consolidation efforts, as shown in Figure 2. Most people assume that consolidation is driven solely by desires to reduce purchasing and ongoing maintenance costs for equipment and software in the face of ever-tightening IT budgets. Although operational savings remains an important goal, IDC has learned from many CIOs that cost control can no longer be the dominant issue.

FIGURE 2

Importance of Leading Drivers of IT Consolidation Effort

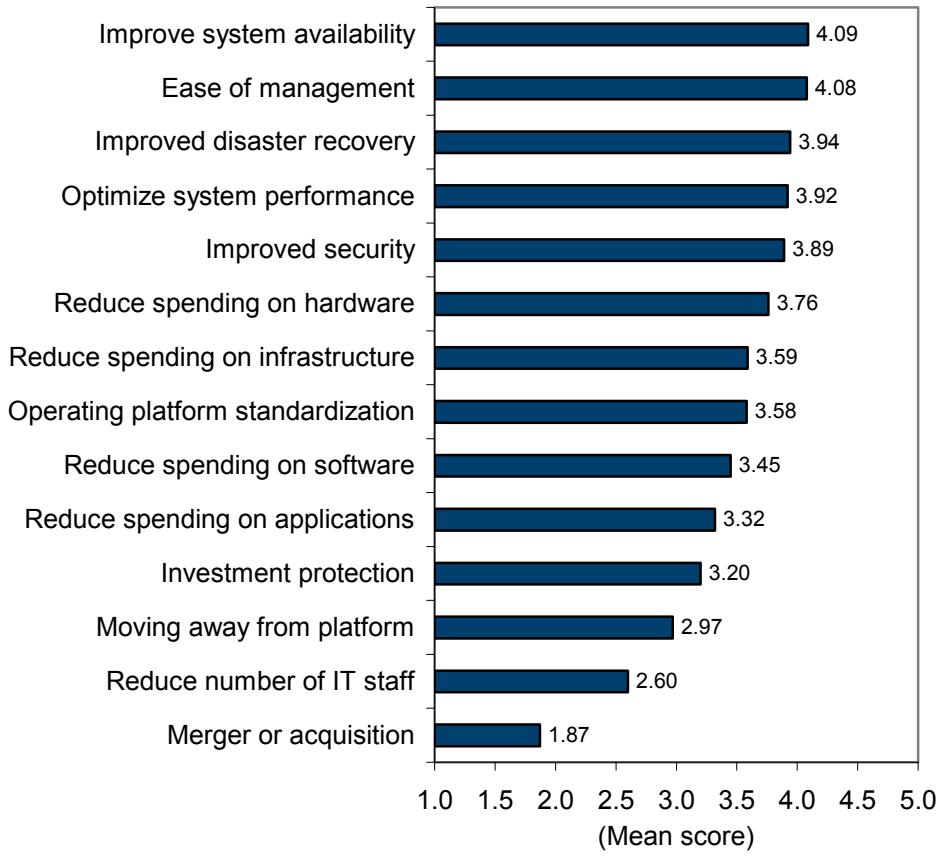

n = 397

Note: Importance was rated on a scale of 1 to 5, with 5 being critical.

Source: IDC, 2004

When business executives request IT support, the call for services is not couched in terms of servers supported, terabytes stored, or even applications offered; rather, business executives focus their requests on numbers of concurrent users or transactions, improvements in application availability, and time to new feature deployment (see Figure 3).

FIGURE 3

Goals of IT Consolidation

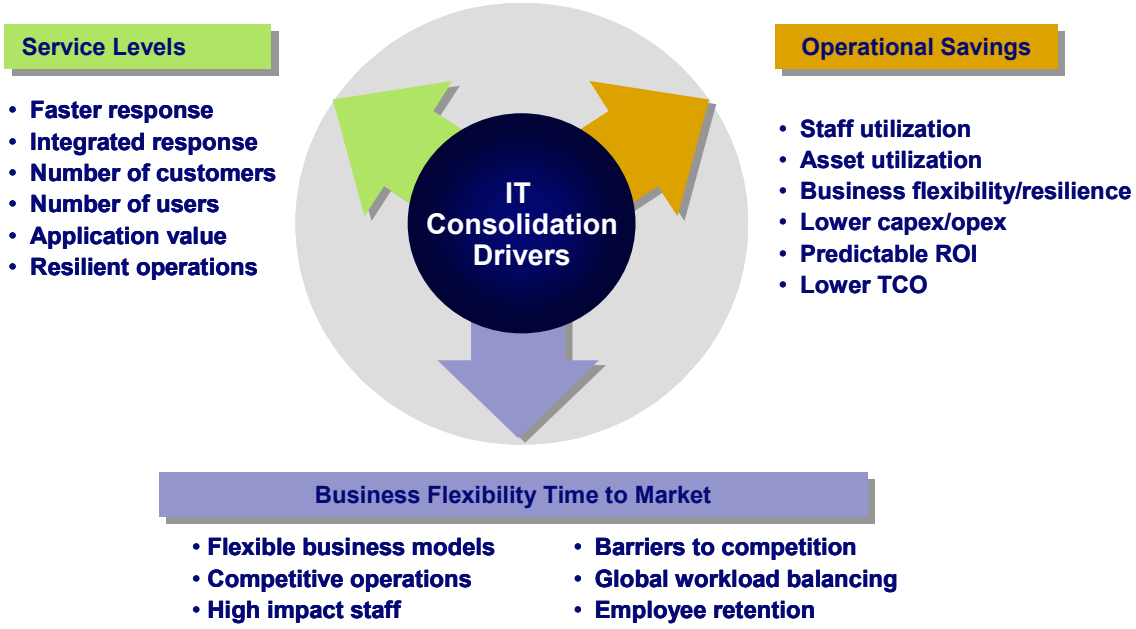

Source: IDC, 2004

As IT executives evaluate strategies for meeting these service-level requests and time-to-market requirements, they often find that the way current IT solutions are purchased, deployed, and managed precludes success.

Inconsistency is the key barrier to improved service by IT organizations. The historic deployment of IT solutions in stovepipes, whether based on server type, application set, or business units, led to wasted capacity, duplicated functions, and inconsistency in policies for security and data protection. Even with unlimited resources, IT executives can't meet growing service-level requirements if they must continue to operate such uncoordinated environments.

Infrastructure Simplification Is the Key to Consistency

Consolidation is one strategy that IT organizations are adopting to bring consistency to the enterprise. However, IT executives also know that consolidation rarely if ever means consolidating to a single type of server platform, operating system, or storage system. The complexity of an enterprise in terms of its scale and scope, its need to reap maximum benefit from past investments, and the reality that IT products are optimized for specific environments mean that consolidation is never "one size fits all."

The key to successful consolidation is simplifying the infrastructure that connects all IT elements and manages the interaction of those elements.

IDC spoke with a vice president of IT at a major health insurance company, and he emphasized several of these points. His company currently deploys a number of server and storage solutions in support of different applications. Historically, each storage system was directly connected to each application server, leading to inconsistencies in utilization and storage provisioning. As part of consolidation efforts, the company deployed a storage area network (SAN) for simplified connection of all storage and server assets. This SAN also included systems that provide the enterprise with the ability to *virtualize* volumes for different applications. The IT department now has a single administrative system for provisioning new storage to applications while maximizing the utilization of existing storage assets. It also leverages this function to reduce downtime associated with server upgrades or replacements.

Areas where IT executives need to focus their attention as they seek to simplify infrastructure and implement consistency include:

- Physical connectivity of servers and storage systems (i.e., disk and tape) to ensure that data can be moved quickly and efficiently around the enterprise
- Logical virtualization of volumes (disk systems) and file systems to make server and storage provisioning more efficient
- Common data replication and backup between disparate systems to provide centralized backup and recovery
- Automation and orchestration of server and storage provisioning and administration tasks to reduce operating expenses

Robust Business Continuity: More than Disaster Recovery

Business continuity is no longer just about planning for the next system failure, network outage, or local catastrophe. Enterprises must develop business continuity plans that allow them to operate more efficiently and react quickly to changing business conditions.

In ongoing conversations with IT managers, IDC finds that business continuity issues include reducing downtime, minimizing backup windows, and speeding application recovery. Less frequently mentioned but of growing concern for IT managers is finding time for planned downtime because many businesses operate continuously or nearly continuously.

Planned downtime is needed because servers need upgraded processors, operating systems require patches, and application software must be enhanced and modified to keep it in alignment with business processes. Evenings and weekends have traditionally been times when systems could be taken down without consequence, but today's customer relationship management (CRM) and supply chain management applications can extend the concept of business hours to 24 x 7.

> IDC spoke with an IT manager at a large telecommunications company who reported that managing server upgrades and the migration of storage to new arrays were her most critical concerns. Downtime at any time of the day has a significant impact on revenue and customer satisfaction. As a result, the IT staff needed a solution that would let them add or replace either servers or storage without having to reconfigure all interconnected systems. They leveraged a combination of volume virtualization and data replication to deploy new storage and migrated volumes to the new system without disrupting ongoing server operations. This solution also allowed them to rapidly, and non-disruptively, add new server capacity when critical applications processing requirements grew. This system enables the IT organization to significantly reduce planned downtime, accelerate adoption of more cost-effective storage systems, and react more quickly to changing business conditions.

Business volatility is a more recent and increasingly important business continuity concern. Enterprises continually struggle to gauge marketplace behavior and the demands employees, customers, and business partners will place on IT systems. Miscalculations can be very costly in terms of wasted effort and wasted resources. Overly optimistic business forecasts encourage IT planners to increase capacity. When business lags, IT investments are underutilized and the total cost of ownership (TCO) rises proportionally. Conversely, pessimistic business forecasts lead to smaller investments in IT systems. When business booms, IT becomes a limiting factor, opportunity costs are incurred due to lost business, and TCO rises as the need to satisfy demand outstrips the normal focus on efficiency.

Balancing Recovery, Integrity, and Cost

Best practices in business continuity planning dictate that the degree of resiliency (i.e., recovery time) should be provided in proportion to the cost or consequences of downtime; the value of continued, timely access to information; and the value of the data itself.

To meet customer and employee needs more effectively, IT organizations must collect and store more extensive information (e.g., medical records, product design histories, purchasing records, customer correspondence, images, and videos). They must also assess the value of integrating rich content (i.e., audio, image, and video files) into existing sales, marketing, and customer service solutions as well as internal applications that improve business processes and efficiency.

All this information, both traditional transaction data and newer fixed-content reference data, needs to be available when employees, partners, and customers need it. Companies are storing more information for longer periods and need to guarantee the integrity of that information for both business governance and regulatory compliance reasons.

> IDC spoke with an IT executive from the European subsidiary of a large consumer and business image creation and product company. The company was launching an image archiving service for

consumer customers and had the stated goals of never losing an image and always providing rapid access (seconds) to those images. It quickly became apparent to the IT team that this service would require a system that was solely disk based (no tape backup), allowed for the rapid and continuous addition of storage capacity, and provided a scalable file system that could support millions of users and hundreds of terabytes of files. The company is implementing a tiered storage solution that virtualizes storage assets for rapid data migration, data replication, and capacity expandability as well as a virtualized file system for seamless migration of infrequently accessed files to lower-cost storage assets.

Managing an increasingly diverse array of applications with different availability, performance, capacity, and cost requirements poses a significant challenge for IT managers as they seek to develop policies for data protection and business continuity. Areas where IT executives should focus their attention as they seek to develop enterprise wide business continuity standards include:

- The deployment of a tiered storage architecture across multiple remote sites to better match storage system performance and capacity to different data sets and to improve system resiliency
- Logical virtualization at volume level (disk systems) and file systems to improve storage system utilization and reduce the cost of maintenance and provisioning
- Policy-based placement, movement, and deletion of data over its life cycle to meet business data retention needs and to align infrastructure costs with information value and access requirements
- Common data replication and backup procedures among disparate systems to enhance system resilience and gain consistency in the storage architecture
- Automation of data replication, retention, and migration tasks to both improve backup and restore capabilities and reduce operating costs

Information Life-Cycle Management: Better Alignment of IT and Business Goals

At the confluence of storage consolidation through infrastructure simplification and improved enterprise-wide business continuity planning is the emerging issue of information life-cycle management. Consolidation of storage infrastructure makes it possible for IT administrators to manage the storage and replication of all enterprise data based upon preset policies. Enterprise-wide business continuity planning provides a foundation for classifying data based upon specific business requirements, the need for availability, and time to recovery. ILM involves leveraging both of these capabilities to help IT managers identify and manage data based upon its value to business executives and end customers.

To date, much of the discussion related to ILM has focused on helping companies meet emerging compliance regulations relating to information retention. Compliance has certainly been a justification for many IT executives to begin evaluating ILM

solutions. However, the long-term value of ILM will be as a strategy for enhancing enterprises' use of all their gathered information.

IT and business executives often have completely different perspectives when assigning value to and making decisions about the retention of data. ILM is rapidly emerging as the framework within which these historically different viewpoints can be reconciled to a common understanding of overall corporate needs. ILM will also allow both IT and business executives to more systematically assign values to different types of information, thereby making it easier to make sound investment decisions about the underlying systems that process and store all of this information.

Finding the Right Storage Solutions

In their continued effort to meet the sometimes conflicting cost, service-level, and business continuity demands of business executives, IT managers require a number of storage solutions that address specific requirements, be they better performance, higher reliability, greater capacity, or lower cost. In response, the industry has seen a proliferation of storage technologies that address specific customer requirements. This section looks at a number of these in more detail and highlights key issues IT managers need to consider when evaluating such solutions.

Tiered Storage

IT managers can now choose from an array of different storage systems, including:

- High-end storage arrays that provide high performance and reliability in a single, centrally managed system
- Midrange storage arrays that allow IT managers to add storage capacity for applications in smaller increments
- Low-end storage arrays that enable the sharing of storage capacity across multiple servers in small businesses and at remote sites

In addition, the past year has seen the emergence of new classes of storage systems that leverage lower-cost and larger-capacity disk technologies (e.g., Serial ATA) to significantly boost array capacities while reducing capacity costs. Today, incorporating capacity-oriented storage systems into existing environments is a high priority for many companies.

Accidentally or purposely, companies that employ these new storage solutions are constructing a tiered storage infrastructure. IT managers need solutions that allow them to intelligently provision, reallocate, and protect storage assets to meet the needs of many different applications.

IT managers need a supplier that:

- Understands how tiered storage can be configured and scaled to meet changing capacity and performance requirements

- Delivers a broad portfolio of storage system infrastructure (e.g., interconnect technology, management software) that can be tuned to meet the different characteristics of storage assets within a tiered storage architecture

- Provides a common set of tools capable of managing different storage assets that populate the tiers

Virtualization

LUN, volume, and file management across an ever-expanding range of applications, servers, and storage arrays remains one of unsatisfactory areas of storage management.

Any modifications (e.g., adding new servers and storage systems, migrating data) require IT staff to perform a complex series of manual reconfigurations on both servers and storage arrays. Reconfiguration consumes valuable IT staff resources, is fraught with the potential for error and unplanned downtime, and contributes to increasingly long lead times in application deployment and modification. One especially galling consequence of these limits, in the past, was the need for IT managers to significantly overprovision storage in anticipation of long-term application growth. Over-provisioning was often in anticipation of needs 12 to 18 months in the future.

Disk Virtualization and Volume Management

LUN and volume management is the area where the box-by-box limitations of homogenous systems and the even greater inconsistency across heterogeneous systems are major contributors to the consistent underutilization of disk capacity. By placing volume virtualization in a networked controller, IT managers introduce a consistent logical layer of abstraction between LUNs and volumes on all storage systems.

This approach to block virtualization promises a number of benefits:

- Centralized and consistent management of all volumes within the datacenter and, ultimately, across the enterprise

- Reduced downtime attributable to planned outages (e.g., time spent on storage expansion and migration and on server upgrades)

- Improved storage resource utilization, making it more practical to add storage on demand

- Common replication of data across different classes of storage devices for both data migration and data protection purposes

Common, Scalable File Systems

Shifting LUN and volume management to the logical level is an important step in enabling the creation of flexible and scalable pools of storage for enterprise applications, but it addresses only a portion of the problem. IT managers must still

deal with a diverse array of servers with individual file systems and databases that are linked to storage assets, even if they are all physically connected via a SAN.

The next step for enterprises is to deploy a solution that combines the file virtualization of existing network attached storage (NAS) with the scalability and performance capabilities of SAN environments. Such a solution provides file-serving front ends based on shared or distributed file systems, significantly reducing the cost of managing file access to large volumes of information. A centralized global file system should permit policy-based placement and movement, over time, of files across multiple tiers of storage to automatically align data value with infrastructure cost.

Delivering Virtualization

The advantages of the virtualization of logical volume and file management infrastructure hold great promise for IT managers struggling to get more from current and future server and storage consolidation efforts. This is especially true for IT managers that are also seeking to leverage virtualization technologies on servers to enable more rapid, on-demand allocation of processing resources.

Delivering on the promise of storage consolidation and virtualization, however, requires the providers of such solutions to address several important customer concerns. IT managers need:

- Flexibility in deployment options (e.g., array based, network based, and server based) without sacrificing consistency and commonality of functions
- Scalability that is simple, quick, and cost-effective, with minimal disruption to ongoing operations
- Volume and file virtualization that is rapid and painless and accompanied by data and application migration services across heterogeneous server and storage environments

Data Replication and Replication Management

Over the next two years, another key to turning the idea of tiered storage into a usable reality for IT managers will be making it easier to automatically replicate and/or migrate data across the different storage systems within the tiered storage environment. With today's storage systems, providing common replication across many different storage systems is often a daunting if not impossible task.

Today, data replication solutions, such as mirroring or point-in-time copy (snapshot) for offline backup are often limited to a single array family. Remote copy to another array, though possible if the appropriate software is loaded on each array, is commonly limited to like, and often very expensive, systems. Storage system diversity translates into growing administrative overhead and reduced interoperability.

The deployment of common data replication functions across different storage systems, in conjunction with volume virtualization, leads to major improvements in data movement, protection, and recovery. For example, IT managers can take

advantage of installed systems for reuse as targets in data-replication actions. Secondary storage systems can be put to use for local rapid recovery and long-term storage of fixed data without sacrificing system integrity. The IT organization can establish consistent data replication and recovery policies across all storage systems.

IT managers need solutions that address the need for common replication at three separate levels:

- Support for a broad portfolio of replication capabilities across a range of storage systems
- Provision of these replication functions on a number of different platforms ranging from single arrays to networked storage controllers, so that enterprises have greater flexibility to deploy the solutions that best meet their needs
- Delivery of replication solutions that are integrated with existing backup/storage management software systems to enable enterprises to organize and monitor data protection processes

Integrated Storage Management

Delivering effective management for tiered storage requires more than common control of the physical devices. IT managers need to be able to create classes of storage based on performance, capacity, cost, and level of data protection across multiple physical systems. These classes of storage can then be assigned as application resources based upon unique workload requirements.

At this logical level, the administrative concerns for IT managers shift more to areas such as process automation, capacity planning, policy management, and application-level performance monitoring. Effective tiered storage management solutions, therefore, must deliver capabilities that enable workflow systems for identifying and automating repetitive administrative tasks.

Equally important, these solutions must deliver on these promises without:

- Requiring wholesale replacement of existing management tools and processes
- Reducing the reliability of existing data protection processes
- Introducing additional management complexity that further taxes staff resources

The remainder of this white paper looks in more detail at one company, IBM, that IDC believes is in a strong position to deliver solutions for tiered storage, virtualization, common data replication, and integrated storage management. It assesses the company's TotalStorage Open Software Family of products and discusses the challenges that IBM must address to meet its goal of fulfilling enterprises' evolving storage needs.

IBM's TotalStorage Software Solutions

IBM is committed to the delivery of storage hardware and software solutions that address its customers' needs for more effective provision of storage services, enhanced business continuity, and management of the information life cycle. The company provides solutions with its IBM TotalStorage family of hardware and software products as well as with a number of management solutions from Tivoli.

TotalStorage hardware products provide companies with a range of storage options, from the highly scalable and reliable DS8000, DS6000, and ESS arrays for high-end datacenters to the modular DS4000 series (formerly FAStT) for open systems and secondary storage. IBM also offers low-cost entry-level systems, the DS400 and DS300 arrays, for smaller businesses. The company offers a broad range of enterprise tape and tape automation solutions for the high-end and midrange customer. In addition, IBM supports SAN fabric switches and directors to connect their storage systems to servers ranging from the largest mainframe to low-end Intel servers. IBM supports switches and directors from Brocade, Cisco, CNT, and McDATA.

IBM's TotalStorage Open Software Family delivers a wide range of products to help companies build, operate, and manage tiered storage environments built on IBM hardware as well as systems from other storage vendors. In the past two years, IBM introduced several solutions that focus on disk (volume) and file virtualization, heterogeneous data replication, and integrated storage management. These network-based storage services solutions include TotalStorage SAN Volume Controller, TotalStorage SAN File System, and TotalStorage Productivity Center. These software virtualization capabilities complement the virtualization included in IBM's DS8000 enterprise disk system that can logically partition this system into two virtual storage systems.

In addition, IBM offers a solution for managing the organization, movement, and protection of information throughout its life cycle, the Tivoli Storage Manager family. Tivoli Storage Manager products address storage management issues including recovery management, hierarchical storage management, and archive management.

The remainder of this white paper looks at SAN Volume Controller, SAN File System, TotalStorage Productivity Center, and Tivoli Storage Manager in more detail.

SAN Volume Controller

The first targets for IBM's storage solution are enterprises that want to consolidate diverse existing storage assets into an easily shared and consistently managed pool, thereby providing more flexible and scalable service levels for different applications. These solutions must easily scale without sacrificing storage performance or reliability.

IBM's TotalStorage SAN Volume Controller (SVC) is a complete software and hardware solution of two eServer xSeries Pentium IV processor–based servers (i.e., storage engines), each with a 4GB read/write cache. The two storage engines are tightly coupled into a node, and these nodes can be combined into clusters that allow

for scaling of virtualization as storage environments grow. With the release of the fourth generation of SVC software, Version 1.2.1, IBM increases the size of a cluster to eight nodes, enabling customers to manage larger environments and double throughput. This latest release also increases the number of virtual disks that SVC can manage to 4,096.

SAN Volume Controller also includes traditional disk array functions such as data replication (i.e., snapshot and remote mirroring) to help enhance business continuity. These functions operate above the disk arrays, within the storage network, to enable customers to apply the functions across disk arrays obtained from multiple vendors. Now customers can replicate from high-performance, highly available, high-cost disk storage to lower-performance, lower-cost disk storage, such as devices using Serial ATA drives. The data replication in SAN Volume Controller includes support for consistency groups that coordinate copies of related LUNs (i.e., for applications such as databases, which use multiple LUNs) for FlashCopy and for Metro Mirror, which is a peer-to-peer remote copy capability.

Server connections supported on SAN Volume Controller include AIX, Windows 2000/2003, Red Hat and SuSE SLES Linux, Sun Solaris, and HP-UX (including AIX, HACMP, and Windows MSCS clustering environments). SAN switches supported include those from Brocade, Cisco, CNT, and McDATA. Supported storage arrays in the first release include IBM's DS 4000 (FAStT) and ESS products. SAN Volume Controller also supports arrays from other storage vendors including HDS (9200), HP (MA8000, EMA12000, EMA16000, and EVA 3000/5000), and EMC (Symmetrix 8000 series; CLARiiON CX200, CX400, CX600, and FC4700) offerings. SAN Volume Controller also works with IBM's Tivoli Storage Manager and Storage Resource Manager so companies can manage SVC-controlled storage assets within their overall system management framework.

IBM is also delivering the SAN Volume Controller storage software embedded within two Caching Services Modules for the Cisco MDS 9000 family of switches and directors. With the addition of Cisco's network storage controller option, IBM can now offer customers a choice in configuration options without sacrificing consistent volume management and data replication services. Customers that adopt the Cisco option can also leverage server connectivity using both Fibre Channel (FC) and iSCSI, which permits servers outside the datacenter to be a part of a single, global, shared storage system.

SAN File System

In November 2003, IBM began delivering its TotalStorage SAN File System. Much as SAN Volume Controller merges disparate storage devices into a single block-level pool of storage capacity, SAN File System provides a global namespace with secure file sharing across heterogeneous types of application servers. SAN File System provides a common pool of storage used by the application servers and provides policy-based storage and data management.

The SAN File System client that runs on application servers uses the Unix virtual file system (VFS) interfaces or the Windows installable file system (IFS) interfaces to connect into each server's file system. A SAN File System client accesses network-

connected SAN File System metadata servers to obtain the exact location of the required file within the storage pools. Supported application servers include the following operating systems: AIX 5.1 (32-bit), AIX 5.2 and 5.3 (32 and 64-bit versions), Windows 2000 Server, Windows 2000 Advanced Server, Windows Server 2003 (Standard and Enterprise editions), Red Hat Enterprise Linux 3.0 (AS, ES, WS), Solaris 9 (64-bit), and SuSE Linux Enterprise Server 8.0 (32-bit).

The second component is the SAN File System metadata server cluster that maintains a global namespace with the physical and logical location of all files and volumes on supported storage systems. It also includes storage pools based on business needs, policy-based file provisioning, and storage quota management. FlashCopy provides file-based, space efficient backups, and automated volume drain provides non-disruptive storage volume migration.

IBM is issuing its third generation of SFS software (version 2.2) that adds hierarchical storage management (HSM) support with policy-based migration of files between storage pools. It also incorporates policy-based deletion of files, adds user and group security checking between Windows and Unix environments, and introduces support for hosts and storage attached with iSCSI (in addition to Fibre Channel). The combination of FC and iSCSI permits a single global file system that can be shared by applications within a datacenter as well as those connected by IP networks beyond the datacenter.

Because reliability and scalability are critical for this centralized metadata store, SAN File System supports a cluster of two to eight Metadata servers. The Metadata servers are tightly linked and continuously monitor each other's health via an internal watchdog timer. If one of the Metadata servers fails, the workload is automatically picked up by the other Metadata servers in the cluster.

TotalStorage Productivity Center

As the provider of a broad spectrum of storage hardware and software products, IBM also needs to offer companies storage management solutions that better match the diverse needs of its customers.

In March 2004, IBM began delivering its TotalStorage Productivity Center to improve storage capacity utilization, availability, and administrator productivity through tight integration of diverse storage management applications focused specifically on storage infrastructure management. This solution provides a set of storage management products for managing IBM's own storage hardware and software products as well as SMI-S–enabled products from other storage vendors. By integrating these products, IBM is offering a single viewpoint on a customer's entire storage infrastructure. The IBM TotalStorage Productivity Center comprises a user interface and the following components:

- **IBM TotalStorage Productivity Center for Data.** Automated identification of the storage resources and analysis of how effectively those resources are being used, including evaluation at the volume and file level

- **IBM TotalStorage Productivity Center for Fabric.** Automated management of SAN-attached devices (switches and HBAs) including multivendor switch zone provisioning and visualization of the SAN topology

- **IBM TotalStorage Productivity Center for Disk.** Centralized point of control for disk configuration and performance monitoring

- **IBM TotalStorage Productivity Center for Replication.** Simplified and automated storage replication environments including the ability to ensure that data on multiple related volumes across storage subsystems remains consistent

With the release of version 2.1, IBM is also providing tighter integration with Tivoli Storage Manager to enable the automated activation of backup or archive of key files when utilization or performance thresholds are violated. IBM has also begun to further integrate these products by using a "suite install" approach, making it easier for storage administrators to quickly deploy this suite and gain value from it. IBM offers flexible licensing options so customers can chose the functions that they want to deploy at any given time without sacrificing the ability to make future enhancements.

TotalStorage Productivity Center with Advanced Provisioning

A significant solution for storage provisioning and automation is TotalStorage Productivity Center with Advanced Provisioning, which integrates IBM TotalStorage Productivity Center with Tivoli Provisioning Manager to deliver more integrated server and storage provisioning and automation workflows. By integrating Tivoli Provisioning Manager with the TotalStorage Productivity Center, IBM's strategic intent is to provide integrated storage and server provisioning. Workflow, the regular steps required to complete a process, is a common ingredient in business systems. For key business processes, information systems route data to those individuals who share responsibility for a process, track progress, and expedite completion of all tasks in a timely and efficient manner.

First announced in the spring 2004, IBM has added further workflow support for storage management. Workflow for storage management means routing information from database administrators to system and storage administrators for typical storage processes such as expanding storage capacity for a growing database.

For storage management tasks such as relocating data in a multitier storage architecture, fine granularity is needed. TotalStorage Productivity Center with Advanced Provisioning is a fine-grained provisioning tool that provides workflows for provisioning storage at the volume level. Policies about data replication, performance requirements, access authorization, and data retention may all apply to a file containing enterprise financial operating data. These same policies will differ for a file containing reference data about product pricing, for example. To optimize the use of a tiered storage architecture without overwhelming storage administrators, the automation now being incorporated into IBM's solutions will be critical.

Tivoli Storage Manager

The final component of IBM's tiered storage, business continuity, and information life-cycle management strategies is the IBM Tivoli Storage Manager (ITSM). The Tivoli Storage Manager family of offerings is already widely deployed in enterprises. It provides centralized, automated data protection that can help reduce the risks associated with data loss while also helping to reduce administrative complexity, manage costs, and address compliance with regulatory data-retention requirements.

TSM provides an integrated solution that includes backup, data archiving, disaster preparation, disaster recovery, HSM, and space management. It provides these capabilities across the full range of storage systems (primary disks arrays, secondary disk arrays, and tape) within tiered storage environments. Backup policies can be automatically set and defined at virtually any level of the hierarchy: single file, directory, volume, client computer, or group of clients based on business needs.

In the latest release of TSM (Version 5.3), IBM announced a number of additional enhancements to this already mature product. These included:

- Single sign-on for TSM administrators to enhance security
- An enhanced user interface that simplifies administration and makes it easier to monitor TSM servers, configure devices, and report activities across multiple TSM environments
- Greater flexibility and granularity in scheduling of backup jobs
- More robust disk-to-disk backup and data reclamation processes
- Enhanced support for NDMP to enable consistent support for backup of network attached storage (NAS)

With these enhancements, TSM provides a centralized control point for managing the movement of files from one storage device to another to help conserve valuable disk space for live data as well as archiving data in an alternate location. Through the use of a relational database, ITSM tracks data at a file, byte, or block level and backs up only data that is new or changed. Similarly, when performing a restore, only the desired version of the file is restored. This avoids full plus incremental/differential backups, which can lead to tremendous time and cost savings.

Tivoli Storage Manager for Data Retention

One critical subarea of ILM for many IT managers today is data retention to meet regulatory compliance mandates. TSM's existing policy-based data management capabilities can already be used to help organizations meet many regulatory requirements of government and industry agencies. Many new regulations, however, require additional safeguards on data retention.

Tivoli Storage Manager for Data Retention facilitates compliance with stringent regulatory requirements in a flexible and function-rich manner. It helps manage and simplify the retrieval of the increasing amounts of data that organizations must retain

to meet strict records retention regulations. TSM for Data Retention makes the deletion of data before its scheduled expiration extremely difficult. Short of physical destruction of the storage media or server or deliberate corruption of data or deletion of the Tivoli Storage Manager database, Tivoli Storage Manager for Data Retention will not allow data on the storage managed by the Tivoli Storage Manager-Extended Edition server to be deleted before its scheduled expiration date. The system also allows IT managers to leverage content management and archive applications that apply business policy management for ultimate expiration of archived data at the appropriate time.

Several IBM content management and archive products — IBM DB2 Content Manager, IBM DB2 Records Manager, IBM DB2 CommonStore, and the TotalStorage DR 550 — use the TSM API as the repository for managing content and come bundled with TSM.

IDC Analysis

As companies seek to more effectively leverage existing storage investments and take advantage of new tiered storage infrastructure, they need storage virtualization, replication, and management solutions that simplify IT infrastructures, enhance business continuity, and enable information lifecycle management.

Opportunities

With this latest series of product enhancements, IBM is boosting the scalability and flexibility of its storage solutions while also expanding support for heterogeneous systems and improving overall manageability.

IBM's ongoing consolidation and repackaging of its storage management software products is an important move for the company. As IBM's customers seek to more effectively leverage existing storage investments and take advantage of new technologies within a tiered storage infrastructure, they will need storage management solutions that simplify demonstrative tasks and provide a more coherent view of the overall environment.

IBM's decision to employ TotalStorage Productivity Center as the common data collection and repository system for management information as well as tighter integration with TSM lays a solid foundation for further enhancements in administrator usability. IBM must continue to enhance the links between TotalStorage Productivity Center and existing archive, content, and data management systems.

Challenges

To effectively meet IT managers' continually evolving needs including their increasing focus on ILM, however, IBM must do more to address several specific issues.

The most important issue for IBM remains educating IT managers on its new systems (SAN Volume Controller, SAN File System, and TotalStorage Productivity Center) for organizing and managing storage assets. IBM and its storage business partners must continue to identify general business and industry-specific applications that allow potential customers to evaluate the functions promised by the IBM TotalStorage Open

Software family. IBM needs to continue focusing on data migration, consolidation, and business continuity projects in which customers can reap immediate benefits without exposing all of their storage assets.

IBM must also ensure that its own applications provide more than simple "integration on the glass." Deeper integration includes employing consistent processes and workflows for completing tasks that cross devices and applications; leveraging common databases for configuration and performance information; and employing standard tools for job scheduling, reporting, and security. Continued integration and data sharing between TotalStorage Productivity Center with Advanced Provisioning and Tivoli Storage Manager is the area where IBM needs to pay particular attention.

The most critical long-term challenge for IBM is to ensure that customers can effectively leverage the existing tie between its storage product lines with the growing array of electronic content management systems, including its own content management product lines. These will be the foundation for companies' evolving ILM strategies. Tight integration to the underlying storage environment will be a key requirement for companies that want to effectively scale their ILM systems.

Conclusion

Today, many companies are only just beginning to recognize the challenges that changing business conditions pose for existing IT architectures and IT organizations. In their continued effort to meet the sometimes conflicting demands of business executives, IT managers need to adopt a new strategy for deploying server and storage solutions. These new solutions based upon the concepts of tiered storage, virtualization, and ILM can provide better performance, higher reliability, greater capacity, and lower cost of operations while making it easier to reuse existing assets and provision new systems as needed.

Deployed effectively, IBM's TotalStorage hardware and software products provide a solid foundation for building a future-proof storage architecture that better serves internal business units, improves application reliability, and streamlines existing management policies and procedures.

Copyright Notice