

# Storage Management

Best Practices

**Tivoli**

## **Storage Management Best Practices**

### ***Introduction***

The exploding growth of corporate data combined with the falling price of storage has created both an opportunity and a challenge for information technology (IT) managers. More affordable storage technology allows IT managers to buy more storage devices for rapidly increasing volumes of corporate data, but managing these expanding storage networks becomes a complex, resource-intensive task. In many organizations, storage management is executed without strategy, reducing cost-effectiveness and efficiency.

Applying the discipline of storage management, combined with the appropriate technology and a well-crafted set of storage management best practices, can provide significant business value by helping enterprises increase revenues and decrease costs. This paper discusses common approaches to storage management as well as best practices that can be used to enhance the business value of both storage technology and stored data. It also discusses the key functionality to consider when selecting a storage management product.

### ***The Most Valuable Asset***

The expanding range of IT devices, platforms, and applications in use across the enterprise complicates the storage management picture. Data can reside in geographically diverse locations and on technologically disparate devices; managing these resources across the corporate information grid is no easy task.

But the complexities of managing heterogeneous technology networks are not the IT manager's

most pressing challenge—for, despite the user community's growing understanding of distributed computing, more than anything users require a responsive, available application environment. Applications have the most direct effect on the corporate bottom line, enabling employees to efficiently communicate, handle tasks, and work to expand revenues and decrease costs.

Businesses depend heavily on their employees' ability to access data and process it through an application. Poor storage management practices can act as a stone dropped into the enterprise's financial pool, creating a potentially disastrous ripple effect throughout the business processes. Unavailable data spreads harmful effects to corporate applications, rendering them ineffective; this can slow or even halt business operations, impairing the enterprise's financial success in one continuous stream.

Consider the example of an online auction service that joins buyers and sellers electronically. Through a Web application, the service collects data from thousands of auction participants; if the application becomes inaccessible, the company's business transactions stop, cutting off cash flow. Without protecting the data that feeds the applications, the online auction service is out of business.

### ***Developing a Strategic Storage Management Approach***

IT managers must develop a strategic approach and several best practices that protect their most important asset: the data supporting applications. Many storage management tools are available to protect data through routine backups and centrally

manage the information grid, but such tools alone are insufficient for organizations preparing to move competitively into the 21st century. The strategic approach, a distributed storage management solution, ensures not only that the data is backed up but also that the entire information grid can be recreated in business-need priority in case of disaster. By implementing a strategic storage management approach, IT managers can do the following:

- Minimize the resources consumed for storage management operations by transferring and storing the least amount of information necessary to protect the information grid
- Extend the life of their current network infrastructure and processing power
- Produce greater returns with lower media costs on the company's investment in secondary tape resources

These benefits can be achieved in a centrally managed environment by selecting the right storage management solution and implementing the best practices described in the following section.

### **Storage Management Practices**

Every enterprise relies on a set of storage management practices. Some common practices might not be the most effective or efficient methodology when viewed as a corporate strategy. In contrast, best practices help corporations improve not only the effectiveness of a storage management strategy but also its efficiency. For example, improving the efficiency of a task might enable faster execution, but improving the effectiveness of the process might

involve eliminating unnecessary steps and automating the remaining steps. Delivery of maximum business value requires improving both the effectiveness and the efficiency of the storage management strategy.

### **Selecting the Most Efficient and Cost-effective Storage Management Solution**

Many vendors offer distributed backup products, but only Tivoli® Storage Manager was developed from the ground up as a comprehensive storage management solution. More than just a distributed backup product, Tivoli Storage Manager is a distributed data backup, recovery, and storage management solution. Unlike other storage products, Tivoli Storage Manager is built on a common, highly portable code base that is consistent across all Tivoli Storage Manager server platforms. This common code base enables Tivoli Storage Manager to manage thousands of desktop clients or a single large database server.

### **Tivoli Storage Manager Strengths**

Tivoli Storage Manager was designed as an enterprise-wide storage management application that focuses its resources on recovery. Four key architectural features unique to Tivoli Storage Manager make it the most efficient and intelligent distributed storage management solution on the market: its database, a progressive backup methodology, collocation, and reclamation.

### **Database**

The specially designed Tivoli Storage Manager database retains information about all client system and user files, business policies, disaster recovery

## Comparing Storage Management Techniques

### **Common Practices**

- Installing a tape drive on each server, an expensive option that does not provide high reliability
- Manually monitoring or neglecting to monitor backup jobs, which might create incomplete backup tape sets
- Managing tapes for each server, which requires significant manual effort to remove the previous backup and insert tapes for the next backup
- Manually moving numerous daily backup tapes away from the server for disaster recovery, creating not only a staffing issue but also a resource constraint during restoration
- Delaying restores to obtain the latest backup, then waiting for tapes to be scanned until the needed file is found
- Putting off the process of disaster recovery planning, which can create lengthy recovery delays in case of disaster and potentially jeopardize the company's ability to do business

### **Best Practices**

- Using a central database of files under storage management, which enables central monitoring and control and reduces the amount of data stored
- Backing up only changed data, which improves backup time, reduces network demands, and saves on storage resources
- Relocating backup tapes for offsite storage to support disaster recovery efforts (the storage manager should provide a tape "pull" list and a directory of the tape contents)
- Allowing both end users and LAN administrators to perform restores; empowering end users to restore only their own files saves time and effort for IT staff
- Automatically restarting restore operations in case of network failure during a restore operation
- Implementing a storage management policy as needed, across heterogeneous platforms
- Archiving data using a centrally managed archival/retrieval system, so that end users and applications benefit from robust records-retention services without the overhead of direct management of these activities

The above comparison shows how implementing best practices can provide significant value to your business.

information, and the scheduling of client and administrative tasks. This database retains information called metadata, which means data that describes data. The flexibility of the Tivoli Storage Manager database enables customers to define storage management policies around business needs for individual clients or groups of clients. Client data attributes such as storage destination, number of versions, and retention period can be assigned at the individual file level and stored in the database.

The Tivoli Storage Manager database also ensures reliable storage management processes.

To maintain data integrity, the database uses a recovery log to roll back any changes made if a storage transaction is interrupted before it completes. This is known as a two-phase commit. Also, the Tivoli Storage Manager database and recovery log both can be mirrored for availability, providing automatic volume switching after a media failure. In the unlikely event of a Tivoli Storage Manager database recovery, operators can restore the database to the exact point of a failure by rolling the recovery log forward after restoring from the latest database backup.

### ***Progressive Backup Methodology***

The Tivoli Storage Manager architecture uses an intelligent backup methodology that provides efficiencies during both the backup and restore of client data. During the initial client backup, Tivoli Storage Manager backs up all eligible files, creating a full backup. Subsequently, files are backed up again only if they are new or have changed since the last backup. Tivoli Storage Manager maintains a pointer in its database to the latest version of each file for each client, eliminating the need for another full backup to consolidate the files into a single image.

Other backup products require an initial full backup, followed by regular incremental or differential backups (usually once a day), and then additional periodic full backups (usually once a week). This less efficient backup method results in redundant weekly full backups of files that haven't changed, wasting both network and media resources. The multistep restore process of such products requires restoration of the last full backup, along with more recent incremental or differential backups, to recover the latest version of a file or an entire system.

### ***Collocation and Reclamation***

The collocation and reclamation features are powered by the intelligence of the Tivoli Storage Manager database. **Collocation** consolidates all data belonging to a particular client onto the same tape or group of tapes. Thus, although a client might perform progressive backups over a long period, the number of tapes used by a given client remains small. **Reclamation** is the process by which the Tivoli Storage Manager server monitors the utilization of tapes in its inventory.

As backup versions expire, the server tracks the amount of free space on tape volumes. When the amount of free space reaches a customer-defined threshold, the Tivoli Storage Manager server automatically mounts that tape and moves the remaining data to another eligible tape, consolidating and defragmenting data stored on the tape. The combination of collocation and reclamation enhances the efficiency and speed of client data recovery, because the client data can be restored from the smallest number of tapes possible.

Along with the unique architectural features described above, other exclusive functions of the Tivoli Storage Manager base product further differentiate it from its competitors.

### ***Media Management***

Tivoli Storage Manager provides sophisticated media management capabilities that enable IT managers to do the following:

- Track multiple versions of files (including the most recent version)
- Respond to online file queries and recovery requests
- Automatically move files to the most cost-effective storage media
- Expire backup files that are no longer needed
- Recycle partially filled volumes

Tivoli Storage Manager provides these capabilities for all backup volumes, including on-site volumes inside tape libraries, volumes that have been checked out of tape libraries, and on-site and offsite copies of the backups.

Tivoli Storage Manager provides a powerful media management facility that can be used to create multiple copies of all client data stored on the Tivoli Storage Manager server. Enterprises can use this facility to back up primary client data to two copy pools: one stored in an offsite location, and the other kept on site for possible recovery from media failures. If a file in a primary pool is damaged or resides on a damaged volume, Tivoli Storage Manager automatically accesses the file from an on-site copy if it is available or indicates which volume needs to be returned from an offsite copy.

Tivoli Storage Manager also provides a unique capability for reclaiming expired space on offsite volumes without requiring the offsite volumes to be brought back on site. Tivoli Storage Manager tracks the utilization of offsite volumes just as it does for on-site volumes. When the free space of offsite volumes reaches a determined reclamation threshold, Tivoli Storage Manager uses the on-site volumes to consolidate the valid files onto new volumes, then directs the new volumes to be taken off site. When the new tapes arrive off site, Tivoli Storage Manager requests the return of the original off-site volumes, which can be reused as scratch volumes.

#### **Enterprise Administration**

Tivoli Storage Manager Enterprise Administration enables an administrator to manage multiple Tivoli Storage Manager servers from any platform in the enterprise using a Web-based interface. Tivoli Storage Manager servers can be deployed near their data for high-performance backups to locally attached devices with minimal network usage, which means

the product does not sacrifice network performance for central control. Unlike other distributed backup products, Tivoli Storage Manager Enterprise Administration does not create a master/slave single-point-of-failure dependency between the managing server and the managed servers. Managed servers remain independent and can function normally even if the “managing” server is unavailable.

Tivoli Storage Manager Enterprise Administration provides centralized management, facilitates consistent backup policies, and provides enhanced protection of critical data by electronically vaulting data between servers. Administrators can centrally define common policies and configuration information at one Tivoli Storage Manager server and then propagate them to other Tivoli Storage Manager servers. This multitiered inheritance structure speeds the dissemination of policy changes and ensures accuracy and consistency of policies within groups. Additional automated functions relieve administrators from repetitive, time-consuming tasks by enabling them to run server scripts from the Enterprise Administration Console.

#### **Instant Archive**

Tivoli Storage Manager can execute a restore from almost any point in time the customer chooses. An extension of the point-in-time restore capability, instant archive allows the customer to create a virtual full backup (or client archive) from data already stored at the Tivoli Storage Manager server. Simply put, the Tivoli Storage Manager server executes a restore operation specified by the customer. The difference is that the data is not sent

back to the client, but is sent instead to another tape or CD on the Tivoli Storage Manager server; customers can create a full backup tape without actually making a full backup. Customers who use this function are often preparing for a disaster recovery situation, creating baseline copies of a system for long-term archival, or getting ready to restore a remote or mobile computer. Instant archive allows all these capabilities by using data already stored on the Tivoli Storage Manager server without having to move that data across the network.

#### ***Rapid Recovery***

Rapid recovery is an integrated attribute of the instant archive media described above. Besides a full copy of information from a given point in time, these pieces of media also contain all inventory information necessary so they can perform a restore without interacting with the Tivoli Storage Manager database. The Tivoli Storage Manager client also can read the instant archive media directly, which means that in case of a disaster, or in remote areas without sufficient bandwidth to transmit a full restore over a network, the instant archive media can be removed from the Tivoli Storage Manager server and mounted directly to the client machine. This enables a rapid full restore without any interaction with the Tivoli Storage Manager server or the network.

#### ***Mobile Backup: Adaptive Differencing Technology***

Tivoli Storage Manager provides a new, patented technology called Adaptive Differencing that dynamically transfers client data at a byte, block, or file level based on its size. By definition, implementing a

mobile backup means an administrator must manage backup and restore services for machines that are rarely seen; this makes it crucial to require minimal interaction with the machine. With the mobile backup capability, administrators can enable and disable client functions, such as Adaptive Differencing and data encryption, remotely from the Tivoli Storage Manager administrative interface. Using this same interface, administrators can also remotely monitor the success or failure of backup operations.

#### ***SAN Tape Resource Sharing***

The Tivoli Storage Manager SAN tape resource sharing capability delivers immediate benefits by reducing the traffic on the IP network and allowing shared utilization of resources over a storage area network (SAN). SANs remove the overhead commonly found with slow, overworked communication networks and facilitate quicker access time. Tape library and drive resources are more efficiently used because they can be shared by multiple Tivoli Storage Manager servers across the SAN.

#### ***LAN-free Data Transfer***

LAN-free data transfer provides an alternative path for moving data between the Tivoli Storage Manager client and server. LAN-free data transfer exploits the SAN path by enabling the Tivoli Storage Manager client to back up and restore data directly to and from SAN-attached storage, which is shared between the Tivoli Storage Manager server and client and managed by the server. The existing local area network (LAN) connection is used to exchange control information, such as policy information and metadata about the

objects being backed up, but the data transfer uses the SAN to write directly to the storage media.

### **Tivoli Storage Manager Complementary Products and Features**

Tivoli Storage Manager is the *only* choice for complete integrated storage management solutions in mixed-platform environments. Tivoli Storage Manager is more powerful, more modular, and more flexible than ever. The following is a brief description of optional Tivoli Storage Manager products and features.

#### ***Disaster Recovery Manager***

The Tivoli Storage Manager Disaster Recovery Manager simplifies the disaster-recovery planning process by generating an automated recovery plan file based on a predetermined recovery strategy. Companies can use this auditable recovery plan to certify the recoverability of their distributed environment. The recovery plan contains information and procedures necessary to restore the key Tivoli Storage Manager server and the clients managed by that server.

Sample information contained in the recovery plan for restoring the Tivoli Storage Manager server includes:

- Installation-specific recovery instructions
- A list of the volumes needed, including their offsite location
- Devices required to read the tapes
- Space requirements
- Copies of necessary configuration files
- Macros for automating the recovery process

Sample information contained in the recovery plan to assist in the recovery of clients includes:

- Client machine requirements and location information
- Business priorities associated with the client machines
- Descriptions and locations of client-bootable media

#### ***Tivoli Space Manager***

Tivoli Space Manager, a separate product, requires Tivoli Storage Manager as a prerequisite. It uses hierarchical storage management to automatically and transparently migrate rarely accessed files to Tivoli Storage Manager storage; the most frequently used files remain in the local file systems. By migrating rarely accessed files to server storage, Tivoli Space Manager frees administrators and users from manual file-system pruning by ensuring that sufficient free storage is always available at their workstation or file server, deferring the need to buy additional disk storage.

#### ***Tivoli SANergy™***

Tivoli SANergy™ provides the power to share SAN-based storage arrays, file systems, and files across multiple systems simultaneously. Tivoli SANergy provides all the file sharing capabilities of a LAN-based file server without the performance-limiting overhead of LAN protocols or file server bottlenecks. Tivoli SANergy reduces, or entirely eliminates, the need for multiple file systems on the SAN storage. It allows multiple servers to access the same storage, file systems, and files simultaneously so that backups and other large movements of data can



occur entirely over the SAN instead of the LAN. With Tivoli SANergy and a SAN, customers can both simplify their storage management and remove a high percentage of data traffic from their LAN. SANs and Tivoli SANergy are the long-term solution to the short-term problems that network attached storage tries to solve.

#### ***Tivoli Decision Support for Storage Management Analysis***

Tivoli Decision Support for Storage Management Analysis provides event, performance, and general health reporting and analysis on Tivoli Storage Manager deployments, whether small or large. It helps Tivoli Storage Manager administrators make informed decisions about storage management implementation.

Tivoli Decision Support, a separate Tivoli product, enables users to make business decisions based on historical data. Decision Support Guides, a set of best-practice guides, can be used to analyze and display data about applications and products, presenting information in a variety of text and graphical formats that enable a user to drill down for details about a particular aspect of an environment. The product maintains historical data and makes trend analysis possible.

Tivoli Decision Support for Storage Management Analysis includes the following features:

- Provides health, performance, and capacity guides for analyzing Tivoli Storage Manager deployments
- Allows users to make business decisions based on historical trending data
- Uses Decision Support Guides to analyze and display data about applications and products

- Presents information in a variety of text and graphical formats, enabling a user to drill down to get details about a particular aspect of an environment
- Requires the Tivoli Decision Support product, but not Tivoli Framework (the supported relational database management systems are DB2®, Microsoft® SQL, and Oracle)

#### ***Tivoli Storage Manager: Information Management for a Connected World***

Tivoli Storage Manager offers sophisticated functionality to reduce the total cost of ownership of distributed storage management in the following key areas:

##### ***Tape Library Slots Usage***

When Tivoli Storage Manager consolidates files onto fewer volumes, more slots become available for other tapes. Customers can keep more data and more scratch volumes in the library. Using scratch tapes allows a given growth cushion to be met with fewer tape volumes.

##### ***Operator Time***

Because Tivoli Storage Manager can support a backup system using fewer tape volumes, less operator time is spent checking volumes into and out of a tape library. Also, when Tivoli Storage Manager detects a media-read error in a backup copy of data, it automatically requests the data from the on-site or offsite backup copy of that data; the operator does not have to look up the location of the backup copy. Because Tivoli Storage Manager keeps track of

all backup volumes, both on-site and offsite, operators do not have to spend time manually keeping track of volume names and locations.

#### **Administration**

Tivoli Storage Manager also saves administrator time in keeping track of backup files as they move from volume to volume during tape (or optical) reclamation. Also, the Tivoli Storage Manager Enterprise Administration centralized control feature reduces the overall IT cost and workload, enabling the customer to add Tivoli Storage Manager servers without adding administrators.

#### **Media**

Using its database, progressive backup methodology, collocation, and reclamation, Tivoli Storage Manager can support a given amount of backed up and archived data using fewer volumes than other backup products, reducing the cost of media.

#### **Media Rotation/Migration**

Tivoli Storage Manager uses its storage hierarchy migration automation to migrate data from one media type (such as 8 mm) to another media type (such as DLT). Using this capability, Tivoli Storage Manager also can move data from old volumes to new volumes of the same media type, freeing the administrator from tracking the volumes or files.

#### **Managing Offsite Tape Volumes**

Tivoli Storage Manager tracks files on offsite tape volumes that expire because of age or version number. The product initiates reclamation of tape space automatically without retrieving offsite

volumes, which protects the offsite copies and reduces the volume and cost of offsite storage.

#### **Copying Backups for On-site and Offsite Storage**

Tivoli Storage Manager automates scheduling for the copying of backed up and archived data for both on-site and offsite storage. Administrators and operators are freed from managing this data on a volume-by-volume basis, which saves time and reduces the chance of error.

#### **Keeping Disaster Recovery Plans Current**

Enterprises that back up some data every day also must update the disaster recovery plan to reflect the daily tape volume serial numbers. Tivoli Storage Manager tracks all this information, consolidating it with other information stored in the disaster recovery plan and reducing the cost of manually updating the plan. Tivoli Storage Manager can even schedule offsite shipment of the daily plan.

#### **Conclusion**

As data storage management grows more sophisticated, it also becomes more complicated. The key to properly managing a complex storage environment is to approach it strategically. This paper discusses a number of ways to begin this approach, which can make storage management easier, more efficient, and more effective. Further, this paper discusses several benefits that come from adopting a strategic storage management solution. Once implemented, this solution gives IT managers the best chance for increasing the effectiveness and value of their business.



**Manage. Anything. Anywhere.™**

© IBM Corp. 2001. All Rights Reserved.

IBM, Tivoli, DB2, SANergy, and Manage. Anything. Anywhere. are trademarks or registered trademarks of International Business Machines Corporation or Tivoli Systems Inc. in the United States, other countries, or both. Microsoft is a registered trademark of Microsoft Corporation in the United States, other countries, or both. Other company, product, and service names may be the trademarks or service marks of others.

0201

[www.tivoli.com](http://www.tivoli.com)