

For IT administrators: Tivoli storage management  
January 2002

**Tivoli** software

The IBM logo, consisting of the letters 'IBM' in a bold, sans-serif font, with each letter formed by eight horizontal stripes of varying lengths, creating a striped effect. The logo is white on a black background.

**Achieving cost savings  
through a true storage  
management architecture**

Contents
<b>2 Introduction</b>
<b>3 Enterprise storage management architecture</b>
<b>4 Network backup with flat-file catalog</b>
<b>5 Network storage management with database and recovery log</b>
<b>6 Unique features enabled by Tivoli Storage Manager architecture</b>
<b>7 Policy management</b>
<b>8 Progressive/incremental backup</b>
<b>10 Storage hierarchy</b>
<b>12 Tape reclamation</b>
<b>13 Collocation</b>
<b>14 Conclusion</b>
<b>15 IBM software integrated solutions</b>
<b>15 To learn more</b>
<b>15 Tivoli software from IBM</b>

## Introduction

Many popular data storage applications are misleadingly labeled “enterprise storage management” products when in fact their capabilities are limited primarily to automating backup and recovery routines. It’s an important distinction, because different approaches to enterprise storage management can drive significant differences in the total cost of ownership, particularly for midrange to large computing environments. This paper explains how an intelligent storage management software product such as Tivoli® Storage Manager can help deliver substantial savings and productivity improvements over less advanced or comprehensive offerings.

How substantial is “substantial”? A recent case study\* demonstrated that one customer could achieve up to a fivefold reduction in tape hardware requirements by exploiting the Tivoli Storage Manager architecture. The customer’s actual and projected tape hardware savings—measured in terabytes stored—are indicated in Figure 1. (“Tape hardware” costs include tape drives, tape library infrastructure, media costs and drive maintenance.) The study showed that similar reductions could be achieved with respect to the network bandwidth required for transferring backup data.

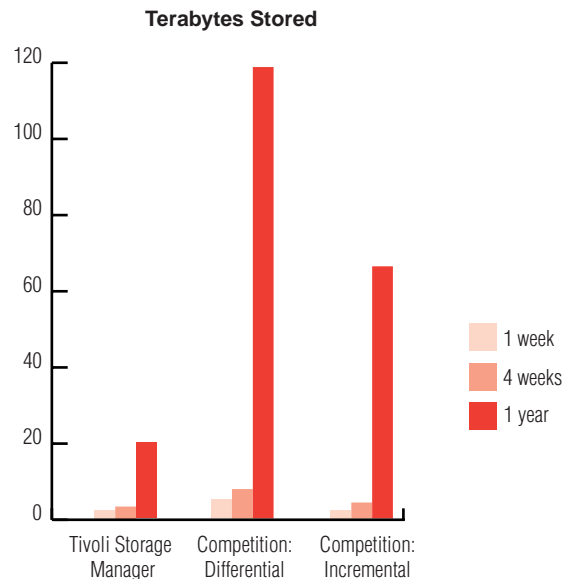


Figure 1

\*Information provided in 2001 by a Tivoli customer, a major financial services corporation based in the United States.

---

## Highlights

---

These savings were modeled on a combination of features made possible by the Tivoli Storage Manager product architecture, as detailed in this technical brief. (Upon request, projected benefits can be modeled for the specifics of your enterprise environment. Please contact your Tivoli account representative for details. In addition, the results of a wider customer survey can be found in the white paper “Accelerating Return on Investment for Storage,” written by Enterprise Management Associates and available on the Tivoli Web site at [tivoli.com/resource\\_center/maximize/data\\_protect/roi\\_storage.html](http://tivoli.com/resource_center/maximize/data_protect/roi_storage.html).

It is important to note that the true cost of a storage management solution cannot be measured adequately by point-in-time benchmark comparisons of product performance. Each product’s total cost of ownership also must take into account its performance over a period of years as storage demands continue to grow.

Of course, reductions in tape hardware and network bandwidth requirements are just the beginning. The Tivoli Storage Manager intelligent storage management architecture delivers other key advantages as well, including potential savings in administrator time, speed of data recovery, and reliability and integrity of backups. As storage management needs grow more complex, the benefits of the product’s flexible, scalable architecture become even more important in helping minimize operating costs while improving the ability of IT departments to respond quickly to business demands.

***Selecting the right architecture is crucial to achieve a comprehensive solution.***

### **Enterprise storage management architecture**

Two alternative designs are available for an enterprise storage management architecture in today’s marketplace. Both approaches are client/server architectures that offer centralized backup capabilities over a network—but only one architecture qualifies as a comprehensive storage management solution, with capabilities that go well beyond those of a backup manager for enterprise computing resources.

To understand how a true storage management architecture can offer significant benefits such as hardware cost savings, improved recovery performance and reliability, and enhanced administrator and user productivity, we need to examine the differences between the two approaches—including their unique features and their potential advantages and disadvantages.

### Approach #1—network backup with flat-file catalog

Most backup applications offered today employ an architecture that grew out of traditional backup techniques developed in the 1980s. This architecture uses a network to collect the data to be backed up from various enterprise computing resources. A simple flat-file index is employed to track the location of backup objects in the enterprise's storage devices. Some applications require full backups; others can perform incremental or differential backups. However, the flat-file catalog is common to all such backup management products (see Figure 2).




Figure 2

This approach offers three primary advantages:

- *It is easy to understand and is based on well-known data storage concepts, such as “grandfather-father-son” file identification techniques.*
- *It is cost-effective in environments where the storage management requirements are simple, such as small departmental networks or single-system environments.*
- *It is supported by most database backup applications. (However, this is no longer always the case as new technologies are introduced, including block-level incremental backup and server-free backup.)*

The simplicity of the flat-file catalog approach also creates several disadvantages for system administrators:

- *File-level policy management is not possible, because backup objects are grouped in “full volume” or “incremental” sets of files.*
- *Target devices for backups are limited to tapes.*

---

---

**Highlights**

---

---

- *Data integrity may be compromised (for example, through the creation of “orphan files” due to unexpected stoppages during the backup process).*
- *Recovery times are lengthened unnecessarily when multiple copies of the same files are restored over the network.*
- *If a failure occurs during restore processing (for example, due to a bad tape or broken network connection), the restore must be restarted at the beginning of the full-plus-incremental sequence.*
- *The storage repository managed by the index is limited to handling backup objects. It cannot be used for other purposes, such as handling archive objects or hierarchical storage management (HSM) objects or providing a storage management service for other applications, such as document content management.*

***Tivoli Storage Manager architecture was designed to excel in complex environments.***

**Approach #2—network storage management with database and recovery log**

Tivoli Storage Manager employs a unique architecture—network storage management with relational database and recovery log—specifically designed for managing the data storage needs of complex computing environments. Developed in the 1990s, the Tivoli Storage Manager architecture provides a single, sophisticated storage management system that can be exploited by virtually all storage management applications. It supports not only backup and recovery applications but also disaster recovery, HSM, archiving, document management and the management of data objects generated by custom or internally




Figure 3

developed applications (see Figure 3). The Tivoli Storage Manager approach offers several potential advantages over the flat-file catalog architecture:

---



---

## Highlights

---

- *High data integrity (and therefore, reliable storage), utilizing two-phase commit transaction processing*
- *Policy-based management capabilities, including file-level*
- *The flexibility to handle multiple types of data objects in a consolidated storage management service (for example, backup, archive, HSM, subfile backups, user-defined objects)*
- *The ability to exploit “disk pools” and “tape pools” in a true hierarchy and virtualization of storage resources*
- *The ability to use “progressive/incremental” backup techniques for file system data*
- *The ability to fill storage tapes to maximum capacity and to reclaim “expired” space on storage devices*
- *The ability to restart a failed restore at the point of failure, utilizing checkpoint/restart*
- *Exceptionally fast recovery times, driven by non-redundant file restore processes*

One possible disadvantage of the Tivoli Storage Manager architecture is that the configuration process can appear more complex than that required for more basic enterprise backup applications. For example, to gain the maximum benefit from the Tivoli Storage Manager architecture, it's important that storage management policies be planned in conjunction with service level agreements. As a result, the Tivoli Storage Manager architecture may offer greater functionality than is needed for environments with simple or limited backup requirements.

***These five capabilities can work together to help lower costs.***

### **Unique features enabled by Tivoli Storage Manager architecture**

The Tivoli Storage Manager architecture enables five important capabilities that can work in combination to help lower the cost of ownership for storage management resources in complex computing environments:

- *Policy management*
- *Progressive/incremental backup*
- *Storage hierarchy*
- *Tape reclamation*
- *Collocation*

---

---

### Highlights

---

---

***The Tivoli solution can enhance policy management.***

Along with helping improve backup and recovery performance, these five capabilities also can help enhance the performance of other storage management applications, such as archiving, HSM and disaster recovery.

***Policy management***

A data storage environment consists of three types of resources: computers, data and rules. Computers contain the data that must be backed up or stored; the rules specify how the stored data is to be treated. For example, in the case of backup, rules determine how many versions of different data objects should be kept, where they should be stored, for how long and so on.

The Tivoli Storage Manager architecture helps enhance the ability to define policies for managing both data objects and rules. Under the Tivoli Storage Manager architecture, backup policies can be defined at any level of the hierarchy: single-file, directory, file-system, volume, client computer or group of clients. Depending on the needs of the business, policies can be set as simple defaults, or they can be highly tuned and customized to help keep costs low while still delivering on the stringent requirements of specific service level agreements.

Figure 4 shows the high-level policy management structure enabled by the architecture of Tivoli Storage Manager:

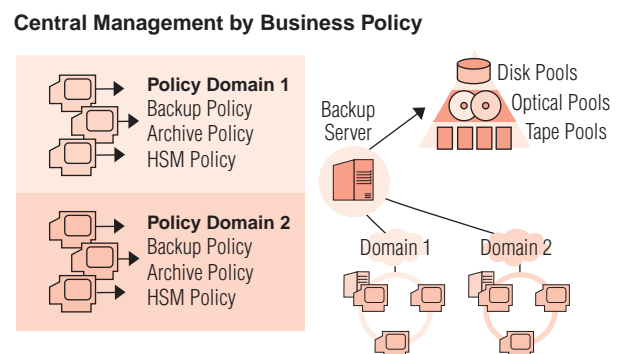


Figure 4

---

**Highlights**

---

A policy-based approach is the key to effective storage management. After the storage administrator has invested time in identifying and planning for business needs, the software tool should provide policy-setting capabilities that let the administrator exactly match the requirements. After policies are defined, the data can be managed automatically. Furthermore, if business needs change, then a simple central update can be made to the Tivoli Storage Manager policies, helping facilitate a fast change in how storage routines are managed.

Tivoli Storage Manager provides a set of policy components that are defined in a hierarchy. This hierarchy lets administrators provide logical groupings of data at virtually any level of granularity. The hierarchy also makes it possible to provide separately managed services for different data types, such as backup, archive and HSM.

A policy-based approach to storage management can help deliver additional value in several specific ways:

- *File-level management capabilities allow administrators to finely tune data storage needs against business requirements.*
- *The cost of storage infrastructure can be minimized through the use of accurate policy definitions, helping avoid the waste of resources common to a “one size fits all” approach.*
- *Stored data can be placed on the storage devices that exactly meet business needs (for example, critical small files can be targeted for storage on online disk devices, thereby helping enable fast recovery).*
- *Automation can greatly reduce the cost of administration.*
- *Change controls are simplified, which can help improve business availability.*

***Backup capabilities can enhance cost savings.***

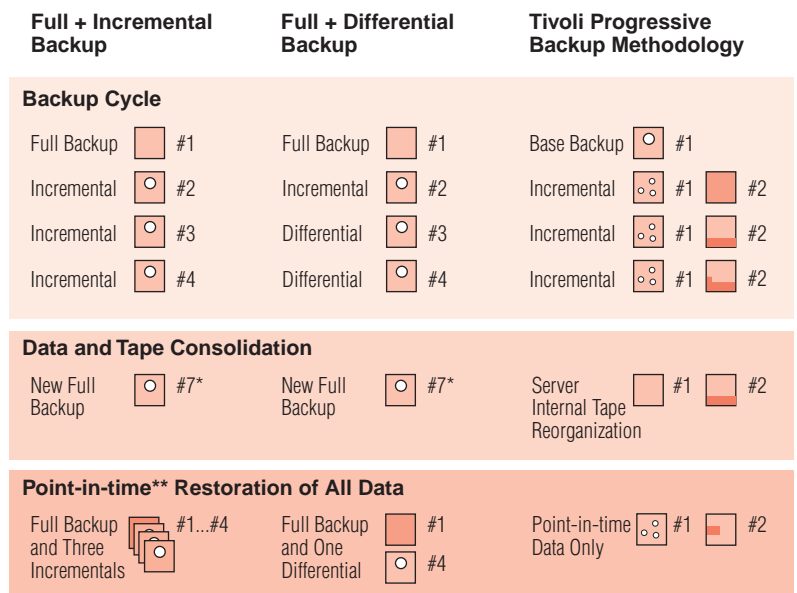
***Progressive/incremental backup***

An important feature enabled by the Tivoli Storage Manager architecture is progressive/incremental backup. When applied to all file-system data, this feature can create the potential for significant savings in tape hardware and network bandwidth. These savings can be achieved because significantly less backup data needs to be transferred over the network and stored than with traditional backup and recovery applications.



The progressive/incremental backup feature is made possible by the ability of Tivoli Storage Manager to do file-level tracking in the database and recovery log. With the progressive technique, only incremental backups are required after the first full backup is completed. This drives the potential for other benefits as well. For example, the Tivoli Storage Manager tape reclamation feature can reorganize data on tape to maximize the utilization of storage space. An optional tape collocation feature keeps data grouped in logical sets, helping enable fast restores. In addition, when a restore is required, it is not necessary to transfer the full backup plus the differential data (a combination that often contains multiple copies of the same files); instead, the restore process transfers only the actual files required for a full restore. In effect, Tivoli Storage Manager makes it possible to assemble a full backup for almost any point in time the administrator wants to specify, thereby helping improve overall restore performance.

Figure 5 illustrates the difference between the traditional cycles of backup and recovery and the progressive technique enabled by Tivoli Storage Manager.



\* Assuming five incrementals between full backups  
\*\* After the third incremental

Figure 5

---

## Highlights

---

The product's unique file-level progressive backup technique delivers value by helping prevent unnecessary backups of unchanged data. This capability allows storage resources to be used more efficiently. Administrators can reduce or consolidate recovery tape sets because they are no longer storing redundant data. They can also achieve faster recovery times because they are not restoring multiple versions of the same file.

This technique also avoids the need for periodic full backups, because only changed files need to be backed up. This change can drive significant savings in backup time, number of tapes used, network traffic and size of backup servers required. Because less data must be transferred over the network, the Tivoli Storage Manager progressive backup technique can also help improve the availability of other applications that rely on the network.

***Tivoli Storage Manager can utilize available space across a variety of media.***

### ***Storage hierarchy***

The ability to manage different types of data (backup, archive and HSM) in a single hierarchy of storage pools is one of the key features that distinguishes a true storage management architecture from that of a backup and recovery application. Typically, backup applications support only tape volumes connected to a backup server; as a result, the data cannot be moved between different devices or managed automatically by predetermined policies.

Tivoli Storage Manager provides a much more flexible architecture, one that "virtualizes" storage space across a variety of media. The storage space on the Tivoli Storage Manager server can be split into a hierarchy of storage pools. Each storage pool can be created from a combination of disk, optical, tape or robotic storage devices locally connected to the server or accessible through a storage area network (SAN).

In the Tivoli Storage Manager storage pool hierarchy, many different policies can be defined to store data in the right place. When a client backs up, archives or migrates data to the server, the data is stored in a primary storage pool. This pool can be defined as any type of device (for example, disk or tape).

In addition to using the primary storage pool, Tivoli Storage Manager lets administrators configure storage pools to provide an optimal combination of performance throughput and data permanence. In many cases, keeping client data on tape media is an essential requirement, especially for backup—but taking the backups directly to tape might not always yield the best system performance. For example, it might be preferable to back up to disk in cases where more clients are requesting backups simultaneously than the available tape drives can support. Some backup applications address this requirement by multiplexing client data across many tapes. This approach can improve backup performance, but it also has a tremendous impact on the time required for restore, because the data must be recovered from a fragmented set of data objects.

The Tivoli Storage Manager approach helps avoid such problems by using a hierarchy of storage pools. If the first disk pool reaches capacity before the backup completes, Tivoli Storage Manager can automatically move or “migrate” files to the next storage pool in the hierarchy without affecting the client’s backup operation (see Figure 6).




Figure 6

In summary, by enabling a storage hierarchy, Tivoli Storage Manager can drive business value in several ways:

- *Storage resources are used more efficiently, helping reduce the need to buy new hardware.*
- *Automation of key tasks within the storage pool hierarchy helps reduce the management burden for administrators.*

---

## Highlights

---

- *Restore time can be minimized, because Tivoli Storage Manager can place data accurately on storage devices that are faster than the tape drives used by traditional backup applications.*
- *Because Tivoli Storage Manager can “virtualize” storage space across a variety of media, it can reduce the potential for “out of space” failures, helping avoid a problem commonly experienced with traditional backup applications.*

**The Tivoli solution has a unique tape reclamation feature.**

### **Tape reclamation**

In optimizing their data storage requirements, administrators face the key challenge of using tape media efficiently. Often a particular tape volume will contain files that expire on different dates. As a result, when these files reach their expiry date, “virtual” empty spaces begin to appear on the tape volume; this fragmentation wastes space on the tapes and slows the restore process because of the time required to skip over empty spaces. Because tapes are sequential media (that is, they can be written only from beginning to end), it is not possible to rewrite new data into the spaces occupied by expired files.

Tivoli Storage Manager addresses this challenge with an innovative tape reclamation feature used to free up entire tape (or optical) volumes in sequential storage pools. As individual files get marked for expiration, the amount of space that can be reclaimed on a volume increases over time. After this available space reaches a specified threshold, Tivoli Storage Manager automatically initiates a process to reclaim the volume. Remaining active files on the tape volume are rewritten to other tape volumes; then the original volume is returned to scratch. Figure 7 illustrates this process for a single tape volume.




Figure 7

---

### Highlights

---

***Optional features can help IT staff deliver cost-effective service level agreements.***

#### ***Collocation***

Tivoli Storage Manager provides optional features, such as collocation, that are designed to help administrators meet their service level agreements in a cost-effective way. The collocation function is designed to optimize the performance of the storage pool for multiple clients. It gives administrators a way to store all the files belonging to a specific client on a minimal number of sequential access volumes (usually tapes).

The collocation option generally would be used in situations where the client requires a fully optimized recovery time. Collocation also makes it possible to avoid conflicts in the restore process, such as when a single tape volume would have to be mounted to restore data for two different clients. When the collocation feature is used, each client's restore can be completed simultaneously and independently.

Figure 8 illustrates the basic operation of the collocation option as backup data is migrated from one storage pool to the next.

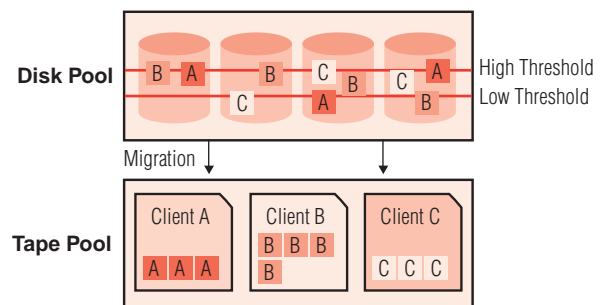


Figure 8

### Conclusion

Storage management is not the simple process it once was, when simple backup tools could be used to efficiently manage the data stored on stand-alone machines. Enormous growth in the amount of data to be stored has given rise to complex new environments, such as storage networks and an ever-increasing range of storage devices.

The unique architecture of Tivoli Storage Manager is designed to meet these new challenges. The database and recovery log technology at the heart of Tivoli Storage Manager helps enable several features that can help drive new storage efficiencies and lower your total cost of ownership. Figure 9 summarizes those features.

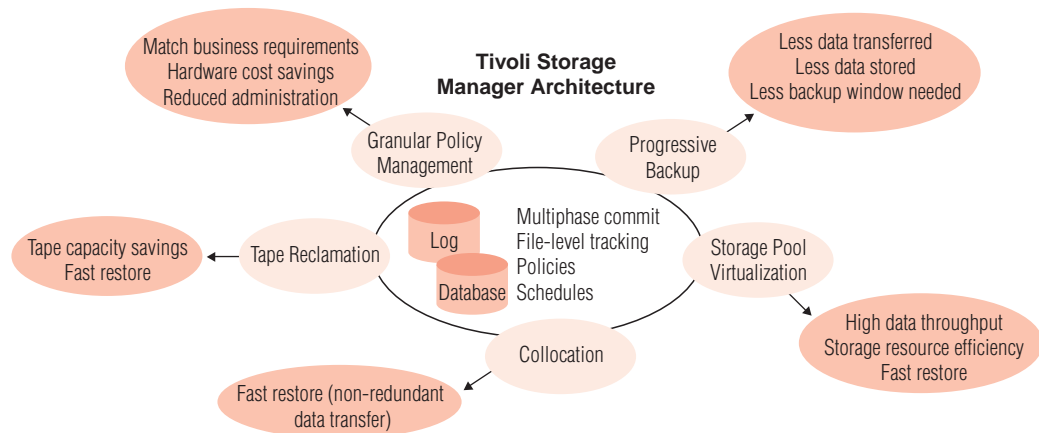


Figure 9

Unlike many enterprise storage management products that are, in fact, little more than enhanced backup and recovery applications, Tivoli Storage Manager is specifically designed to efficiently manage storage resources in an enterprise computing environment. The product's flexible architecture provides an integrated, intelligent storage management solution for backup and recovery, disaster recovery, archiving and other vital storage tasks. The architecture has been shown to lead to significant cost savings and improved restore performance, particularly as storage management needs grow more complex. Therefore, Tivoli Storage Manager should be thoroughly evaluated by any IT organization concerned with minimizing its total cost of ownership for storage management solutions.

### **IBM software integrated solutions**

The Tivoli storage management solution supports a wealth of other offerings from IBM software. IBM software solutions give you the power to achieve your priority business and IT goals.


- *DB2® software helps you leverage information with solutions for data enablement, data management and data distribution.*
- *Lotus® software helps your staff be productive with solutions for authoring, managing, communicating and sharing knowledge.*
- *Tivoli software helps you manage the technology that runs your e-business infrastructure.*
- *WebSphere® software helps you extend your existing business-critical processes to the Web.*

### **To learn more**

For information on the Tivoli storage management solution and integrated solutions from IBM, contact your IBM sales representative or visit [info.tivoli.com/storageforsuccess](http://info.tivoli.com/storageforsuccess)

### **Tivoli software from IBM**

An integral part of the comprehensive IBM e-business infrastructure solution, Tivoli technology management software helps traditional enterprises, emerging e-businesses and Internet businesses worldwide maximize their existing and future technology investments. Backed by world-class IBM services, support and research, Tivoli software provides a seamlessly integrated and flexible e-business infrastructure management solution that connects employees, partners and customers.



© Copyright IBM Corporation 2002

IBM Corporation  
Software Group  
Route 100  
Somers, NY 10589  
U.S.A.

Printed in the United States of America  
01-02  
All Rights Reserved

IBM, Tivoli, the e-business logo, the IBM logo, the Tivoli logo, DB2 and WebSphere are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

Lotus is a registered trademark of Lotus Development Corporation and/or IBM Corporation.

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

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program or service is not intended to state or imply that only that IBM product, program or service may be used. Any functionally equivalent product, program or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program or service.

The Tivoli home page on the Internet can be found at **tivoli.com**

The IBM home page on the Internet can be found at **ibm.com**

 Printed in the United States on recycled paper containing 10% recovered post-consumer fiber.