

Archiving for enhanced cost control, operational efficiency and compliance

*Facilitating automated data migration from IBM Tivoli
Storage Manager to IBM Information Archive*



Contents

- 2 Executive summary
- 2 Keeping pace with growing data storage needs
- 3 Offloading data from Tivoli Storage Manager
- 4 IBM Information Archive: The next-generation archive system
- 5 The synergy between Tivoli Storage Manager and IBM Information Archive
- 7 Key benefits of the data migration utility
- 8 For more information
- 8 About Tivoli software from IBM

Executive summary

As explosive data growth poses more and more storage-related challenges, it becomes increasingly important for today's organizations to find efficient and effective ways to meet those challenges.

Organizations that have archivable data stored in backups and that rely on IBM® Tivoli® Storage Manager as their storage management platform can benefit by offloading some of that data to IBM Information Archive, a flexible, scalable and simple archiving solution. To aid in this process, IBM offers the Data Migration Utility for Tivoli Storage Manager to

IBM Information Archive. This utility facilitates the identification and migration of data from the backup and archive environment in Tivoli Storage Manager to the archive environment provided by IBM Information Archive. This type of migration can result in improved backup and recovery performance, improved compliance, and reduced backup storage capacity requirements and costs.

This paper describes the storage challenges posed by data growth today, particularly with regard to data backup and archive requirements, and describes the advantages of migrating selected data to the IBM Information Archive solution. It also explains how using the data migration utility can facilitate the migration process—through the identification of data to be moved, and the automated migration of that data.

Keeping pace with growing data storage needs

Data growth can create tremendous IT complications for organizations. Compliance costs can soar, for example, when IT departments have to continually meet and manage ever-expanding data retention requirements associated with government regulations. Operational and infrastructure costs can also rise when organizations are forced to invest in more storage capacity and to dedicate more IT resources to storage activities and administration, including backup and archiving. At the same time, the demands that data storage places on the infrastructure are likely to reduce IT efficiency and productivity: Applications may perform poorly, backup and restore processes can take an inordinate amount of time, and information is likely to be increasingly difficult to find.

In response to these challenges, organizations are recognizing that not all data is of equal value, and, as a result, not all data should be stored in the same way. They're identifying ways to categorize data appropriately and store it accordingly, keeping some of it readily available while relegating to archive the data that must be retained despite less frequent access. For example, organizations will want to keep active production data immediately available and accessible, which includes making regular backups that can be used to recover information in the event of unexpected data loss or corruption. At the same time, they are better served by storing older data that is rarely used (but that might be required on demand, such as for litigation purposes, to demonstrate compliance with data retention regulations or for historical data analysis) in active archives that are designed specifically for that purpose and that enable immediate access. Archiving data in this manner moves it out of the active storage environment, where the presence of too much data may negatively affect the performance of applications, the productivity of employees, and the efficiency of operations. The question then becomes, what is the most effective and cost-efficient way to archive the data?

Offloading data from Tivoli Storage Manager

For some organizations, it's common practice to keep archivable data in the form of backups within the Tivoli Storage Manager environment. But the backup function of Tivoli Storage Manager works best as a means of keeping copies of active data

that can be accessed quickly in the event of a loss. In the backup process, data is frequently overwritten as it changes, to keep the most current information available for recovery. This is ideal for backup purposes, but less suitable for archiving, which is intended to preserve a data record over a long period of time.

One alternative to using backups for long-term retention is to use the archive function within Tivoli Storage Manager. But today, there is an even more effective alternative available. IBM Information Archive is designed specifically to meet the needs of organizations that want to archive data cost-effectively while maintaining easy access. It supports compliance and auditability, and it provides the capability to search for data by metadata or by content, so that the data can be found easily long after it has been archived.

With its extensive backup and recovery features—including collocation for faster recovery, offsite copying for disaster protection, and integrated capabilities for protecting virtualized environments—Tivoli Storage Manager is an outstanding platform for data backup. However, IBM Information Archive offers advanced features for longer-term data retention and access, including customizable levels of protection and retention to meet specific requirements, such as legal hold orders from courts or regulators. It also allows organizations to easily search and index all archived information. This makes it an ideal choice for long-term storage.

IBM Information Archive: The next-generation archive system

Archiving is an effective process for managing inactive or infrequently accessed data that still has value, while providing the ability to preserve, search and retrieve the information during the period for which it needs to be retained. IBM Information Archive helps organizations meet these requirements, by providing a data storage platform that:

- Offers a simple-to-use, highly reliable and available appliance to enable long-term data retention.
- Scales capacity up to 402 TB (usable capacity) and also includes seamless attachment to tape for scaling to petabytes.
- Accommodates both structured and unstructured content.
- Provides data protection policies that facilitate compliance with stringent data retention requirements, including those found in the most heavily regulated environments.
- Allows the use of standard data transmission interfaces such as Network File System (NFS) and Common Internet File System (CIFS)—enabling application use or retrieval of previously stored information—and provides access via the IBM System Storage® Archive Manager (SSAM) interface.
- Supports user-defined metadata and enables efficient search of data content and/or metadata.
- Supports efficient policy-driven retention and tiered storage management.

IBM Information Archive can accommodate up to three virtual archives (or “collections”) per appliance. Each collection has a customizable protection level to meet requirements for basic, intermediate or maximum protection:

- Basic protection enables the most flexibility for managing data retention needs, with the ability to delete documents before they expire, and to increase and decrease retention periods as needed.
- Intermediate protection places control with a storage administrator, who can increase and decrease retention periods and increase protection to a higher level as needed, but restricts deletion of documents until they expire.
- Maximum protection imposes the strictest safeguards on data to meet stringent regulatory and governance requirements. Documents cannot be deleted until the end of their retention period, and retention periods can only be increased, not decreased. After enabling this level of protection, it cannot be modified to a lower level. Collections with this level of protection can never be deleted; files will be removed as their retention period expires.



Figure 1. IBM Information Archive enables customizable protection levels to meet specific data protection requirements.

The synergy between Tivoli Storage Manager and IBM Information Archive

The Data Migration Utility for Tivoli Storage Manager to IBM Information Archive is designed to increase the synergy between Tivoli Storage Manager and IBM Information Archive and to facilitate the migration of data. It can copy all files or a subset of files from the Tivoli Storage Manager archive repository to IBM Information Archive. It can also help identify files

in the Tivoli Storage Manager backup environment that would be better managed as archives and copy those files to IBM Information Archive. It supports Microsoft Windows data only (CIFS), i.e., files that have been backed up or archived via the Windows Tivoli Storage Manager backup/archive client.

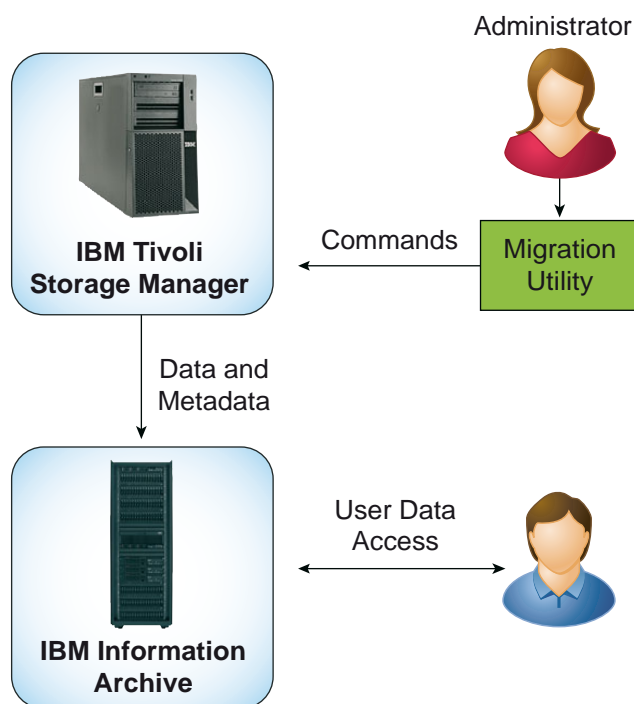


Figure 2: The Data Migration Utility for Tivoli Storage Manager to IBM Information Archive is designed to facilitate the migration of data from IBM Tivoli Storage Manager to IBM Information Archive.

User-defined rules

The data migration utility allows storage administrators to define rules for selecting backup and archive data from Tivoli Storage Manager servers. For example, the utility can scan through files using a rule such as “all files with TSM node name server1.” Rule criteria and data searches are based on Tivoli Storage Manager attributes such as backup or archive data, client (node) name, management class, archive description, file owner, file path and filename. Rules can be easily added, deleted or edited, and they can be suspended or resumed as needed.

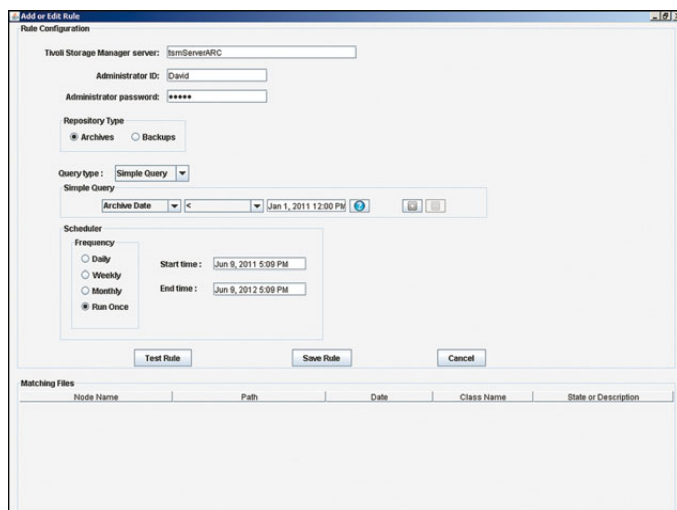


Figure 3: The data migration utility allows storage administrators to easily define rules for selecting backup and archive data from Tivoli Storage Manager servers.

Data migration

After identifying candidate files based on user-defined rules, the migration utility finds those files, extracts them from the backup or archive repository, and copies them to IBM Information Archive. When the utility copies these files, it preserves important Tivoli Storage Manager information, such as the original Tivoli Storage Manager node name and source server, backup/archive date, and management class.

Option to delete copies

The utility does not automatically delete files from the Tivoli Storage Manager repository once they have been copied to IBM Information Archive. However, once the data has been migrated, the utility produces a list of files that have been copied so the storage administrator can delete them from the backup or archive repository. This option to delete allows organizations to reclaim valuable storage space and creates the potential for improving backup performance. Importantly, it gives organizations the flexibility to control the delete process in a way that meets their business needs. This process can also help reduce capacity-based licensing costs for organizations using the Tivoli Storage Manager Suite for Unified Recovery, or who have converted their Tivoli Storage Manager deployment to capacity-based licensing.

Rule ID	Start Time	End Time	Frequency	Repository	Server No.	Administrat.	Query Criteria	Status	Copy Progress	Commit Progre
4	2011-06-07 23:10:00.0	2011-07-07 23:08:00.0	Daily	backups	tsm_local	pin	(Management Class LIKE STANDARD)	✓		
	2011-06-07 23:12:00.0	2011-07-07 23:08:00.0	ExecuteOnce	backups	tsm_local	pin	(Node Name LIKE %PINZHOU%) AND (File Name LIKE %TCT%)	Expired		
14	2011-06-08 17:01:00.0	2011-07-07 21:40:00.0	Daily	archives	tsm_local	pin	(Node Name LIKE %ERIC%)	⚠	Done with 3 Files Retrieved and 1 Res failed	Done with 3 Files Committed and 0 Res failed
15	2011-06-08 17:03:00.0	2011-07-07 23:08:00.0	Daily	backups	tsm_local	pin	(Node Name LIKE %PINZHOU%) AND (Path Name LIKE %LONGDIR%)	✓	Done with 8 Files Restored and 0 Res failed	Done with 8 Files Committed and 0 Res failed

Status

Query Record 14: Committing files in Commit FileList (Node Name LIKE %ERIC%) with queryId 14

Query Record 14: Committing files in Commit FileList Successful for tsm query (Node Name LIKE %ERIC%) with queryId 14

Query Record 14: TSM Crawling done with WARNING for tsm query (Node Name LIKE %ERIC%) with queryId 14

Query Record 15: Scheduling the job for query ID 15

Query Record 15: Starting query Job for query ID 15 with tsm query (Node Name LIKE %PINZHOU%) AND (Path Name LIKE %LONGDIR%) with TimeStamp Wed Dec 31 16:00:00 PST 1995

Query Record 15: TSM crawler start time for query (Node Name LIKE %PINZHOU%) AND (Path Name LIKE %LONGDIR%) with query ID 15

Thu Jun 09 17:03:00 PDT 2011

Query Record 15 Generated file List -> C:\TSMCrawler1\0fileList\tsm_local_PINZHOU\ALMADEN.IBM.COM_e_win_15_2011-06-08_17-02-54

Figure 4: Data Migration Utility Operational Interface

Key benefits of the data migration utility

There are a number of benefits for organizations that use the data migration utility to take data that's currently being stored in Tivoli Storage Manager and move it to IBM Information Archive. They range from making backup and recovery easier and more effective, to improving compliance efforts, to reclaiming primary and backup storage capacity requirements and reducing costs even as storage demands continue to grow.

Improved backup and recovery

Using the data migration utility to move data from Tivoli Storage Manager to IBM Information Archive reduces the Tivoli Storage Manager database size, which can improve

performance and reduce backup times. Some organizations put capacity limits on their Tivoli Storage Manager databases, and reducing the database size can help reduce the number of servers managed. On the recovery side of the equation, backup systems can get clogged with older files that are being kept solely for regulatory compliance or internal governance reasons. Moving these files to IBM Information Archive improves recovery efficiency by reducing the amount of content in the backup system.

Better compliance

Moving files containing private customer data or certain types of corporate financial data into the secure and protected IBM Information Archive environment improves an organization's ability to comply with regulatory requirements governing the privacy and integrity of data. IBM Information Archive also makes it easy to find data after it's archived through rapid search and retrieval that's based either on the metadata stored in the archive for the file or on file content.

Reclaiming backup storage capacity while reducing costs

The data migration utility can run periodically to move data from Tivoli Storage Manager backup and archive repositories to IBM Information Archive, freeing up storage capacity within the Tivoli Storage Manager server for more backups. The cost advantage comes with the ability to backup more data without paying more in Tivoli Storage Manager licensing costs. For organizations that also use Tivoli Storage Manager Suite for Unified Recovery, this approach adds to the cost advantages of that solution, which simplifies licensing by applying a single, unified approach to operating and managing multiple data protection and recovery tools.


For more information

The Data Migration Utility for Tivoli Storage Manager to IBM Information Archive can help you archive data efficiently and cost-effectively. If you are using Tivoli Storage Manager today, there's more reason than ever to explore IBM Information Archive and this new migration utility. To learn more about the utility and how to use it to migrate data from Tivoli Storage Manager to IBM Information Archive, contact your IBM representative or visit: ibm.com/software/data/smart-archive

About Tivoli software from IBM

Tivoli software from IBM helps organizations efficiently and effectively manage IT resources, tasks and processes to meet every-shifting business requirements and deliver flexible and responsive IT service management, while helping to reduce cost. The Tivoli portfolio spans software for security, compliance, storage, performance, availability, configuration, operations and IT lifecycle management, and is backed by world-class IBM services, support and research. For more information on Tivoli software from IBM, visit: ibm.com/tivoli

All client examples cited represent how some clients have used IBM products and the results they may have achieved. Performance data for IBM and non-IBM products and services contained in this document was derived under specific operating and environmental conditions. The actual results obtained by any party implementing such products or services will depend on a large number of factors specific to such party's operating environment and may vary significantly. IBM makes no representation that these results can be expected or obtained in any implementation of any such products or services.



© Copyright IBM Corporation 2011

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

Produced in the United States of America
July 2011
All Rights Reserved

IBM, the IBM logo, ibm.com, and Tivoli are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarks are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at ibm.com/legal/copytrade.shtml

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

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. IBM products are warranted according to the terms and conditions of the agreements (e.g. IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided.

The customer is responsible for ensuring compliance with legal requirements. It is the customer's sole responsibility to obtain advice of competent legal counsel as to the identification and interpretation of any relevant laws and regulatory requirements that may affect the customer's business and any actions the customer may need to take to comply with such laws. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the customer is in compliance with any law or regulation.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program or service is not intended to imply that only IBM's product, program or service may be used. Any functionally equivalent product, program or service may be used instead. Offerings are subject to change, extension or withdrawal without notice.



Please Recycle