

White Paper
November 2009

Tivoli software



Using IBM data reduction solutions to manage more data with less infrastructure

*By Richard Vining, Product Marketing Manager,
Tivoli Storage Software*

Contents

2	<i>Executive summary</i>
3	<i>The growing pressures on storage administrators</i>
4	<i>Surviving the tidal wave of data: Options for effective data reduction</i>
4	<i>Avoiding data duplication</i>
9	<i>Categorizing data for migration and deletion</i>
11	<i>Automating data migration, archival and deletion</i>
16	<i>Deduplicating and compressing what's left</i>
22	<i>Why IBM?</i>
23	<i>Next steps</i>
24	<i>For more information</i>

Executive summary

Organizations are storing and using more data than ever before. Data volume is growing exponentially, and government regulations and competitive pressures are increasing—forcing organizations to retain more data for longer periods of time. But budgets are flat or being cut. And as everyone becomes more dependent on digital information, the costs of losing any of it are increasingly painful.

If your data systems are unavailable, you can lose revenue opportunities; in today's on-demand world, customers will simply go elsewhere if you can't help them immediately. Business disruptions can cause customers and partners to lose confidence in you. Serious outages and breaches are reported in the media all the time, causing embarrassment in the market.

Failure to deliver products or services can result in contract penalties, and failure to deliver reports to government agencies can result in fines. In extreme cases, such as under the Sarbanes-Oxley Act, your CEO or CFO could even go to jail. The bottom line is that organizations need to do a better job of protecting data assets, and as these assets grow and budgets shrink, you need to do more with less. You need smarter solutions.

This paper addresses the challenge of managing and surviving the tidal wave of data and describes various options available from IBM to deal with the challenge by effectively reducing the amount of data that must be managed.

Highlights

Storage administrators are facing a tidal wave of data, along with shrinking backup windows and increased risk of data loss.

The growing pressures on storage administrators

Storage administrators are on the front lines of the “tidal wave of data” battle. As data grows, they are challenged to complete backup operations within established backup windows—even as business requirements are shrinking these allotted times. As a result, some data may not be backed up at all, leaving it dangerously unprotected.

More data in the backup system means it takes longer to recover when something goes wrong. This adds to downtime, which increases the risk of losing revenue opportunities. And of course, if data isn’t properly backed up, it can’t be recovered.

One solution to all this data growth has been to simply buy more storage. The cost of the storage itself has been decreasing over time, so that isn’t so painful. But the costs of housing, powering, cooling and managing all these devices are exploding.

And of course, your business is always changing. Storage administrators need to adapt to any number of changes in their environments, from bringing new applications and data sources online, to assimilating data and systems from mergers and acquisitions, to complying with new corporate and government data management mandates.

IBM can help you build a dynamic storage management infrastructure that will enable you to cope with all of these challenges. We have solutions to help reduce your data storage footprint, and to help you:

- *Reduce your capital and operational costs by reducing your storage requirements.*
- *Improve your application availability and service levels by reducing downtime.*
- *Mitigate the competitive risks associated with losing data in a rapidly changing environment.*

Highlights

IBM takes a holistic approach to data reduction, providing solutions for efficient backups, tiered storage, and data deduplication.

Surviving the tidal wave of data: Options for effective data reduction

IBM takes a holistic approach to data reduction, unlike competitors that offer point solutions to problems that they may in fact be causing. For example, a huge contributor to data growth is the repeated duplication of large amounts of data every time you perform a full backup.

In IBM's holistic approach, one option is to avoid data growth from unnecessary data duplication, by only backing up data that has changed since the last backup. Another option is to determine what different types of data you have and categorize it so that you can manage it most effectively, by moving less frequently-accessed data to lower-cost tiers of storage, and by deleting data that you no longer need or want. This will shorten your backup cycles and improve application performance. A third option is to put automated processes in place, based on policies that meet business requirements and/or service-level agreements, to migrate, archive and delete data. The last option is to compress and deduplicate the data you end up putting into your data protection and retention systems.

Avoiding data duplication

The idea is to treat the cause of the problem, not the symptom. If you have a 10 percent per week data change rate, 90 percent of your data didn't change this week. If you perform a full backup on that data this weekend, you're duplicating almost everything you backed up last weekend. Not only does that take a lot of storage capacity, but it also takes a long time—and these problems only get worse as you create more new data. (It's no wonder that data deduplication products are so popular; they were designed to eliminate all of this duplicate data. And when they claim to reduce your backup storage footprint by 90 percent or more, this is exactly the data that they're talking about.)

Highlights

The IBM Tivoli Storage Manager portfolio enables incremental-only backups that prevent unnecessary data duplication.

But what if you never had to perform a full backup again after the initial one? If you backed up only the new and changed data—always—you wouldn't be creating all that duplicate data that needs an expensive deduplication solution to undo. Shorter backup windows, less storage required, and reduced storage acquisition costs would all be benefits of eliminating that weekly full backup. So would faster restore times, since deduplicated data wouldn't need to be rehydrated in order to be useful.

IBM has smarter solutions that can help prevent the need to perform full backups. The products in the IBM Tivoli® Storage Manager portfolio of recovery management solutions all provide incremental-only backups.

- *IBM Tivoli Storage Manager backs up the files that have changed since the last backup; for larger files, such as huge databases, it can perform subfile backups, copying only the sections of the file that changed.*
- *IBM Tivoli Storage Manager FastBack takes it to the next level, by backing up only the individual blocks of data that change as they are written to disk. And because it performs backups without impacting applications, it can perform them more frequently, which means less data is at risk of loss.*
- *IBM Tivoli Storage Manager FastBack for Workstations continuously protects the data on desktop and laptop computers, again copying only the files that are new or changed.*

Backup Methodologies

Full + incremental

- *Backup – This requires a full backup and then incremental backups over time—usually a full backup each weekend with incremental backups for the following six days. Only data that has changed from the day before is transferred to tape. Then at the end of the week another full backup must be run.*
- *Restore – The full backup must be restored, then each day’s incremental data applied to it. This means that if you have a full backup and three incremental backups of the same file, it will be restored four times. It is a waste of time and money, and introduces risk.*

Full + differential

- *Backup – This requires a full backup and then differential backups over time—usually a full backup each weekend with differential backups for the following six days. This means that all data that has changed since the last full backup will be backed up. If you assume a 10 percent daily change rate, then you will backup 100 percent (full) on the first day, 10 percent on the second, 20 percent on the third, 30 percent on the fourth, 40 percent on the fifth, 50 percent on the sixth, and 60 percent on the seventh. That means that you are backing up 260 percent of your data every week! You’ll need 10 times your production capacity for just a month of backups.*
- *Restore – You would restore the full backup and then the last differential up to the date you were restoring to. This is faster and more reliable than the Full + incremental model, but at the cost of much more storage capacity.*

Highlights

Progressive incremental

- *Backup – This requires a full backup the first time you back up, and then only incremental backups. There are no extra transfers of data, which saves network bandwidth and transfer time, makes backup and restore faster, and can save thousands of dollars in disk and tape costs.*
- *Restore – You select the point in time that you want to restore from, and then restore the necessary files just once. This is much faster than with the other two methods.*

Tivoli Storage Manager's unique progressive incremental backup methodology can save terabytes of backup capacity every month.

Progressive incremental backup with Tivoli Storage Manager

An internal enterprise-class relational database enables Tivoli Storage Manager to perform progressive incremental backups because it tracks each individual file and knows exactly how your computer looked on each day. When a restore is required, only the version of the file needed is restored. Unlike other file-based backup solutions that require you to run periodic (usually weekly) full backups to ensure reasonable recovery times, Tivoli Storage Manager's unique progressive incremental backup methodology never requires you to run another full backup after the first one is done to set the base. The result, as shown in Figure 1, can be a savings of many terabytes of backup capacity every month.

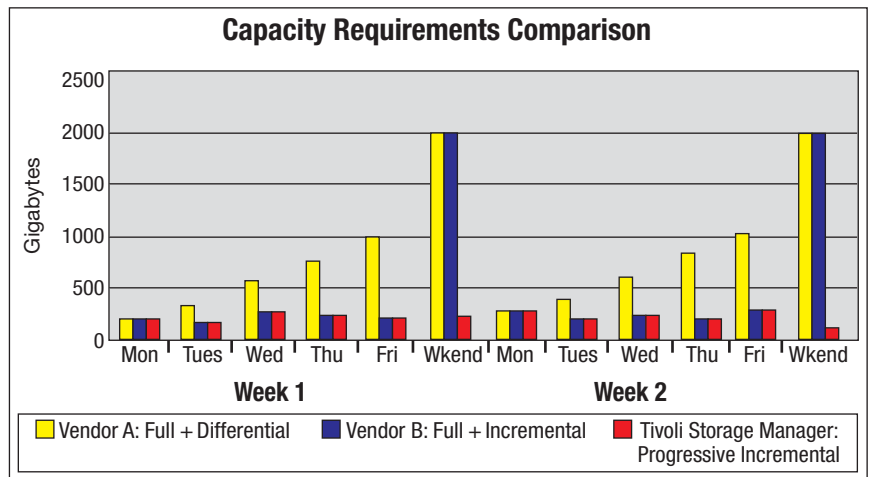


Figure 1: Tivoli Storage Manager's progressive incremental backup methodology can save many terabytes of backup capacity.

Backup Capacity Needed for One Month:	
Vendor A	26 TB
Vendor B	14 TB
Tivoli Storage Manager	7 TB

The analysis shown in Figure 1 starts with 2 TB of data and adds or changes 200 GB per day. The assumption is that a full backup has already been performed to set the base.

- *Full + differential, in yellow, shows that once per week, a full backup is performed, and then on each day between the full backups, all the new data is copied since the last full backup was performed. In this scenario, 26 TB of capacity would be needed to store one month of backups.*

Highlights

- *Full + incremental, in blue, shows that once per week, a full backup is performed, and then on each day between the full backups, only the new data since the last backup is copied. In this scenario, 14 TB of capacity would be needed to store one month of backups.*
- *Tivoli Storage Manager's progressive incremental approach, in red, never requires subsequent full backups. As a result, only 7 TB of capacity is needed to store one month of backups in this scenario.*

Categorizing data for migration and deletion

Another option for reducing the data storage footprint is to assess the different types of data and where they are in the data life cycle. If your organization is like most, you have all your unstructured data in flat file systems, which are probably full of data that you rarely, if ever, need to access. This may include data you are no longer required by law or policy to keep, but that you haven't deleted—such as old e-mails and memos—that could prove costly if discovered in legal proceedings.

Moving data to less expensive tiers of storage and deleting unnecessary data from the environment can reduce the need for primary storage capacity and can make data easier to manage and protect.

The goal is to identify what data can be moved to less expensive tiers of storage, and what data can be deleted entirely from the environment. This will reduce the need to buy more primary storage capacity and make it easier to manage and protect what you have. Backup and restore performance will improve, and it will be easier to prove that you are meeting data retention and expiration policies.

IBM offers IBM Tivoli Storage Productivity Center for Data for this purpose (see Figure 2). This solution reports on where your data is, sorted by access or saved dates, who owns it, the application that created it, and numerous other filters. From the intelligence you gain from these reports, you can

set meaningful policies in your data management software to automatically take the appropriate action on data that shouldn't be clogging up your primary systems. Tivoli Storage Productivity Center for Data can also help identify and eliminate duplicate data, orphan data, temporary data and nonbusiness data.

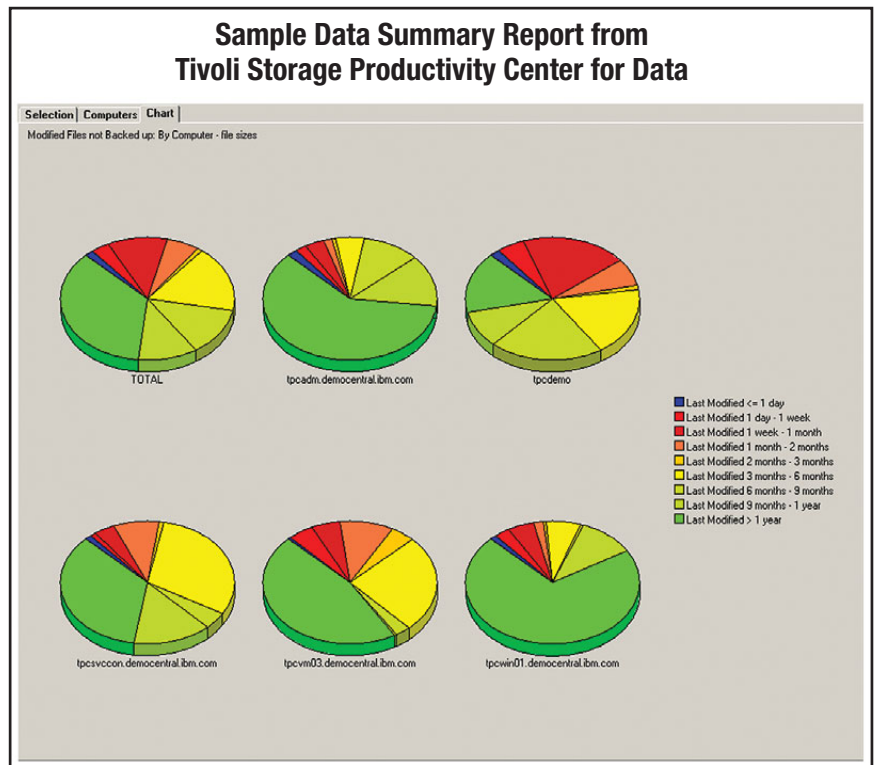


Figure 2: Tivoli Storage Productivity Center for Data can help you better understand your data, enabling you to categorize and manage it more effectively.

Highlights

Tivoli Storage Manager solutions automatically move data to other media based upon policies you set, freeing up valuable disk space for active files.

Automating data migration, archival and deletion

The Tivoli Storage Manager family includes two solutions for automating the migration of data between multiple tiers of storage: IBM Tivoli Storage Manager for Space Management, which is for IBM AIX®, HP-UX, Sun Solaris and Linux® servers, and IBM Tivoli Storage Manager HSM for Windows®, which is for Microsoft® Windows servers. These solutions work transparently in the background, automatically selecting and moving files from primary to secondary tiers of storage based on the policy criteria that you set, such as file size or length of time since a file has been opened. They leave a pointer, or stub file, where the data was originally stored so that users and applications don't need to worry about where the data was moved; the solution transparently reroutes the request for any moved files. No more out of disk space messages will appear. Tivoli Storage Manager solutions automatically move data to the proper media based upon policies you set, freeing up valuable disk space for active files and providing automated access to the migrated files when needed.

Tivoli Storage Manager data migration solutions help customers get control of, and efficiently manage, data growth and its associated storage costs by providing automated space management. These solutions provide the following key features:

- *Storage pool “virtualization” helps maximize utilization of Tivoli Storage Manager storage resources.*
- *Restore management is optimized based on the location of the data in the hierarchy.*
- *Migration is transparent to the users and to applications.*
- *Migrations are scheduled to minimize network traffic during peak hours.*
- *Automatic migrations occur outside the backup window.*

Tivoli Storage Manager data migration solutions not only help you clean up your primary storage systems to help them run more efficiently, they can also be used to easily move data to new storage technologies as they are deployed (see Figure 3). Migrating files to Tivoli Storage Manager also helps expedite restores, because there is no need to restore migrated files in the event of a disaster.




Figure 3: Tivoli Storage Manager space management solutions can help reduce storage and administration costs.

Using Tivoli Storage Manager to enable more efficient use of storage can help you:

- *Improve the response times of file servers by off-loading inactive data.*
- *Move low-activity or inactive files to a hierarchy of lower-cost storage.*
- *More fully exploit existing storage assets.*
- *Reduce backup times and resource usage by focusing on active files only.*
- *Eliminate manual file system cleanup activities.*

Information archive

Archiving is another important data reduction technique for certain types of data. One example of this would be financial reporting data (such as weekly, monthly, quarterly, annual data), that needs to be retained for future trending, requirements or auditing, but does not need to consume valuable disk space where live data should reside. Historical medical records and customer statements also often fit into this category.

Archiving is for long-term record retention. It differs from backup in that it keeps files for a specific amount of time (where backup keeps a certain number of versions of a file) while removing the data from the primary production storage systems completely.

Key features of IBM archiving solutions include:

- *Long-term storage on cost-effective media.*
- *Point-in-time copies that provide revision history and enable auditability.*
- *Data deduplication to remove redundant copies of data.*
- *Retention period and “retention hold” policy enforcement.*
- *Fast expiration processing.*

Using IBM archiving solutions for records retention can help you:

- *Speed file-server recovery times by moving archived files and file archive copies to a hierarchy of lower-cost storage.*
- *Reduce backup times and resource usage by focusing on active files only.*
- *Locate historical information easily using archived files that are indexed with descriptive metadata.*

Highlights

IBM offers a choice of solutions for archiving, depending on customer preferences and requirements for long-term record retention.

IBM offers a choice of solutions for archiving, depending on customer preferences and the applications involved (see Figure 4). Tivoli Storage Manager 6 includes an archiving capability directly integrated into its client backup software. It is policy based, allowing the administrator to set retention times. If the requirement for how long a file must be retained changes, all the administrator has to do is update the policy, and the solution will retroactively update the already archived files; there is no need to restore and rearchive, as some competitive offerings require. Tivoli Storage Manager also offers the option of integrating data from many different applications into your archive repository, and the archive repository can be a virtualized pool of heterogeneous storage systems.

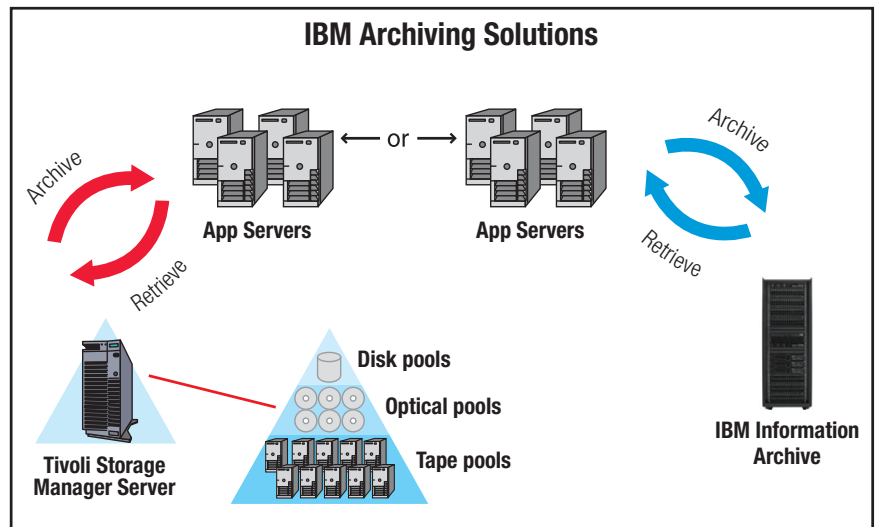


Figure 4: IBM archiving solutions include Tivoli Storage Manager and IBM Information Archive.

Highlights

IBM Optim Data Growth Solution is a unique database archiving solution that transparently migrates unneeded records from database tables to secondary storage.

IBM Information Archive, which contains a specialized version of Tivoli Storage Manager called IBM System Storage™ Archive Manager, is a stand-alone archive appliance that ingests data directly from more than 40 applications including messaging, healthcare and medical imaging, design and engineering, document management, and others.

Database archiving with IBM Optim and Tivoli Storage Manager

IBM Optim™ Data Growth Solution is a unique database archiving solution that transparently migrates unneeded records from database tables to secondary storage. Like Tivoli Storage Manager's space management and archive solutions, Optim provides database and storage administrators with a range of cost and performance benefits.

There are also benefits to using Tivoli Storage Manager in conjunction with Optim, which works seamlessly with Tivoli Storage Manager's application program interface (API) to move archived database records directly into Tivoli Storage Manager's storage hierarchy.

Optim can also be used with other file-based backup/restore products; however, this involves a two-step process to first archive the data and then back it up. When used with Tivoli Storage Manager, Optim automatically archives database records and then uses the API to store/archive data in a Tivoli Storage Manager storage pool hierarchy. With any other file-based backup/restore product, Optim uses standard file operations to store/archive data in a disk-based file system, and then the backup product can backup the file to supported backup media.

Highlights

Data deduplication is designed to reduce data storage capacity requirements by eliminating duplicate data from multiple source systems.

Using Optim and Tivoli Storage Manager together allows you to:

- *Archive data directly to disk or tape or have Optim use Tivoli Storage Manager to automatically migrate it to tape.*
- *Back up Optim archive data incrementally to a Tivoli Storage Manager storage pool that can be managed by Tivoli Storage Manager for local availability, disaster recovery or remote vaulting.*

Deduplicating and compressing what's left

If you're using a backup solution that requires periodic full backups, data deduplication will certainly help you reduce the amount of data you're managing. Otherwise, the reduction ratio you can expect from deduplication will depend on your specific backup environment. Data deduplication is designed to reduce data storage capacity requirements. Some of the duplicate data exists across multiple source systems, but most of the duplication is caused by storing repetitive full backup sets.

The basic idea behind data deduplication is to store just one copy of any data object, and place pointers to the single copy wherever duplicates are eliminated (see Figure 5). Some solutions do this at a file level, so that the files have to be exactly the same to be deduplicated. This is often called single-instance storage (SIS). Other solutions deduplicate data at a fixed or variable block length. IBM's solutions use a blended approach based on the size of the data—file-based for smaller files, and variable block for larger files. Most deduplication solutions run a checksum algorithm against the selected data to create a hash signature, then check to see if that signature has ever been seen before. If it has, the data is discarded and a pointer to the already

stored data is put in its place. A small number of high-end solutions perform a complete byte-level differential comparison of the data to remove all potential for “data collisions,” where two distinct data blocks may share the same hash signature.

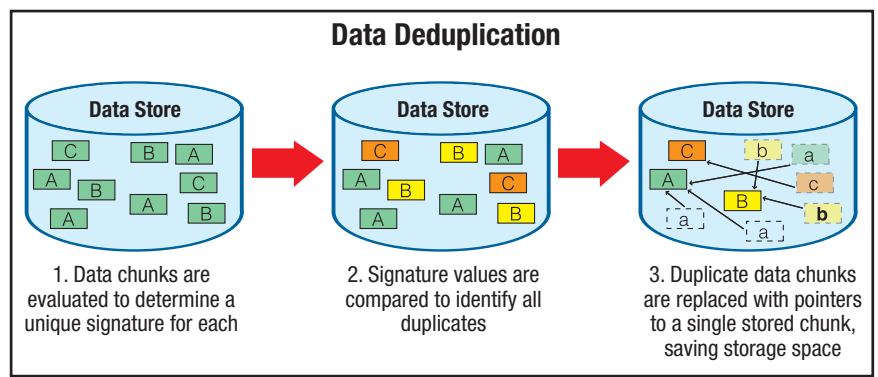


Figure 5: Data deduplication reduces the amount of data that must be managed by creating just one copy of any data object and replacing duplicates with pointers to this single copy.

Data deduplication can and does occur at many points in the data creation and management life cycle. In general, these points of deduplication can be broken into source-side, where the data is created, and target-side, where it is stored and managed. Backup applications, for example, can perform source-side deduplication by not transferring data that has previously been backed up over the LAN or WAN, saving on bandwidth.

On the target side, the most popular use of deduplication is in virtual tape libraries, or VTLs. These disk-based systems emulate tape libraries and drives, but apply deduplication to store equivalent amounts of data on disk very cost-effectively while providing performance advantages over tape. Performing

Highlights

Deduplication capabilities in Tivoli Storage Manager 6 help to reduce recovery times by making it possible to store more backup data on disk rather than tape.

deduplication on tape-based systems is considered to be a bad idea, given the portable nature of tapes and the need to recycle them over time; it would be very difficult to guarantee that you maintain the original data for all of the pointers that are out there.

Today, IBM offers two compelling data deduplication solutions. The Extended Edition of Tivoli Storage Manager 6 includes deduplication capabilities to eliminate duplication of data that has been backed up from multiple production systems. The other solution is the IBM System Storage ProtecTIER® family of deduplication systems for reducing data coming from multiple sources, including Tivoli Storage Manager servers, backups from other backup systems, or archive software solutions.

Deduplicating and compressing with Tivoli Storage Manager

Deduplication capabilities in Tivoli Storage Manager 6 help to reduce recovery times by making it possible to store more backup data on disk rather than tape. It works with data from sources including normal backups, data imported via the Tivoli Storage Manager API, and archive and HSM data. Tivoli Storage Manager deduplicates disk-based data pools as a post-process, so there is no impact on backup performance. After running, it automatically reclaims the storage that has been freed up.

Tivoli Storage Manager already eliminates the most common cause of duplicate data—full backups—so the reduction ratios that can be expected from its deduplication capabilities are fairly modest, averaging about 40 percent. But when combined with its progressive incremental backup approach and built-in data compression, Tivoli Storage Manager's effective data reduction rate is highly competitive with any other solution on the market, as has been detailed in a report from Enterprise Strategy Group.*

Highlights

Tape reclamation is another way that Tivoli Storage Manager increases ROI through intelligent tape utilization. Over time, files on tape will expire, leaving chunks of “dead” space that cannot be reused because the media can only be appended to. Instead of having to wait for all of the data on a tape to expire, a policy can be set up in Tivoli Storage Manager to move, based on a percentage of dead space, the remaining good data onto another tape—thus consolidating the good data and freeing up the other tapes for reuse. Tivoli Storage Manager is also unique in that it will automatically do reclamation of off-site volumes. It does this by creating a new set of tapes to be taken off-site, and then having the old set of off-site tapes brought back for reuse.

Using Tivoli Storage Manager’s tape reclamation capability allows you to:

- *Reduce costs through better utilization of tapes.*
- *Constantly monitor tape utilization.*
- *Establish user-defined reclamation thresholds.*
- *Schedule reclamation to occur at specified times or when free space reaches a threshold.*
- *Return original tapes to the scratch pool.*

IBM System Storage ProtecTIER performs a byte-level differential to ensure that data is a duplicate, ensuring enterprise-class data integrity.

Deduplicating with IBM System Storage ProtecTIER TS7650

IBM System Storage ProtecTIER is a technology leader in performance, scalability, data integrity and reliability. It is the fastest solution on the market in real customer environments. A single ProtecTIER system can easily scale in both performance (1000 MB/sec) and capacity (1 PB). ProtecTIER is one of the few solutions that doesn’t rely on a hash algorithm and instead performs a byte-level differential to ensure data is a duplicate, providing enterprise-class data integrity. And ProtecTIER features all IBM best-of-breed components, unlike the inexpensive OEM’d parts found in competitive products.

ProtecTIER has been proven in very large production environments and is supported worldwide by IBM's services operations. IBM System Storage ProtecTIER TS7650 deduplication solutions range from small (7 TB) to medium (18 TB) to large-scale (36 TB) appliances. And the TS7650G gateway offerings allow you to add the storage of your choice, up to 1 PB. Active-active cluster configurations also provide high-availability capabilities.

Choosing between Tivoli Storage Manager and ProtecTIER for data deduplication

Both Tivoli Storage Manager and ProtecTIER help reduce storage capacity requirements, operational costs, energy usage, and total cost of ownership, and both enable faster data recoveries. The choice of which to use for data deduplication depends on a number of criteria.

IBM recommends using Tivoli Storage Manager data deduplication capabilities under the following conditions:

- *You have a single Tivoli Storage Manager server.*
- *You want to improve Tivoli Storage Manager recovery times by storing more backup data on disk.*
- *You don't have a large amount of duplicate data across the systems protected by multiple Tivoli Storage Manager servers.*
- *You want deduplication operations completely integrated within Tivoli Storage Manager so that you can have the benefits of deduplication without having to acquire separate hardware or software.*
- *You want end-to-end data life-cycle management with minimal data store requirements.*

Use ProtecTIER under the following conditions:

- You are deduplicating across multiple Tivoli Storage Manager (or other backup) servers.
- You need the highest performance up to 1000 MB/sec or more.
- You have a large amount of data and need scalable capacity and performance.
- You need inline deduplication to avoid the operational impact of post processing.
- You don't have Tivoli Storage Manager and are performing weekly full backups.

You can also use Tivoli Storage Manager and ProtecTIER together to enable deduplication at multiple points within a complex data infrastructure (see Figure 6).

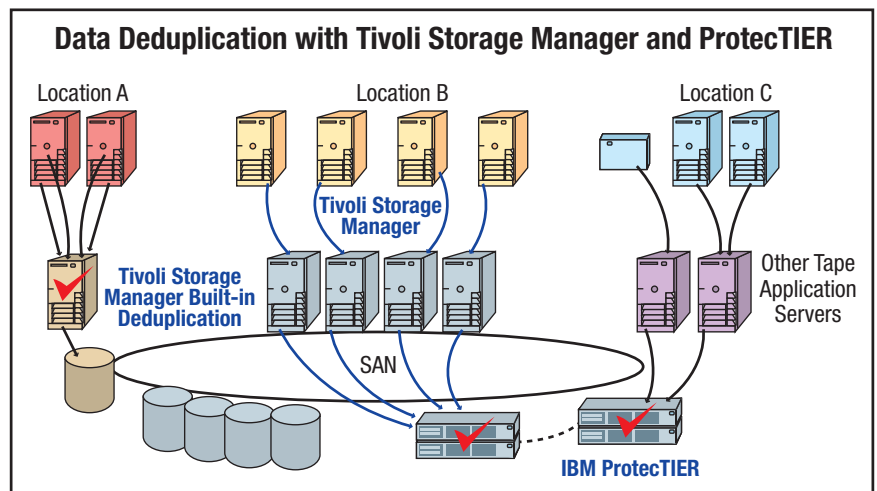


Figure 6: Tivoli Storage Manager and ProtecTIER can be deployed together to enable deduplication at multiple points in the data infrastructure.

Highlights

IBM's broad portfolio of data reduction solutions help solve complex customer issues, while high-quality global support services help protect your investments.

Why IBM?

IBM is the only vendor with a comprehensive set of data reduction technologies. Our broad portfolio of data reduction solutions gives us the freedom to solve customer issues with the most effective technology, while our high-quality global support services help ensure that your investment in data reduction will meet your needs over the long term. IBM is continuing to invest in research and development to further create and deliver the advanced features our customers are requesting.

As shown in Figure 7, IBM offers data reduction through incremental-only backup across a unified recovery management portfolio consisting of Tivoli Storage Manager, Tivoli Storage Manager FastBack, and Tivoli Storage Manager FastBack for Workstations. You may never have to perform a full backup again, anywhere in your extended enterprise. Tivoli Storage Manager also offers best-in-class tape management and data compression capabilities to further reduce your backup and archive storage requirements.

IBM also offers target-side deduplication built into Tivoli Storage Manager 6 and the ProtecTIER TS7650 family of appliances and gateways. Finally, through industry-leading solutions such as IBM Information Archive, IBM provides integrated archive solutions for storing and managing multiple types of information from multiple sources, helping organizations meet the widest range of retention demands.

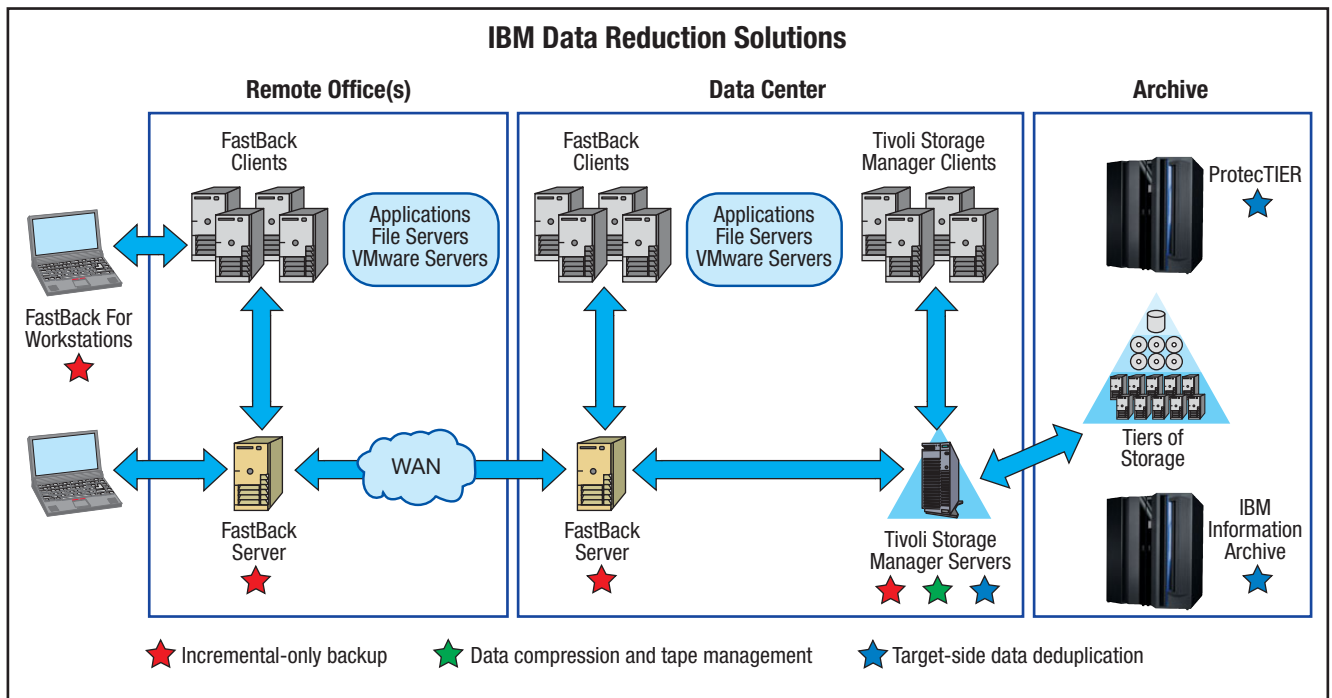


Figure 7: IBM offers data reduction across a unified recovery management portfolio.

Next steps

IBM Global Technology Services and IBM Business Partners stand ready to help you assess your current situation and recommend next steps. We can help you determine which data reduction technologies will have the most cost-effective impact on your operations. You can also ask for a comprehensive ROI analysis using IBM's Business Value Analyst (BVA) tool.



For more information

IBM is uniquely positioned to provide a wide range of storage management solutions, even for mixed-vendor environments. We can help you implement a dynamic storage infrastructure that scales to meet your changing business requirements. To learn more about how IBM Tivoli storage management solutions can help you address data management concerns in your environment, contact your IBM sales representative or IBM Business Partner, or visit ibm.com/software/tivoli/solutions/storage. Additional information can be found at ibm.com/software/tivoli/solutions/reduction

© Copyright IBM Corporation 2009

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

Produced in the United States of America
November 2009
All Rights Reserved

IBM, the IBM logo, ibm.com, ProtecTIER, System Storage 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 trademarked terms 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

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

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

References in this publication to IBM products and services do not imply that IBM intends to make them available in all countries in which IBM operates.

No part of this document may be reproduced or transmitted in any form without written permission from IBM Corporation.

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.

THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. 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.

*Enterprise Strategy Group, Lab Validation Report: IBM TSM 6.1 Unified Data Protection and Recovery Management, April 2009.
www.servicemanagementcenter.com/main/pages/IBMRBMS/OMSA/ShowCollateral.aspx?oid=55497



Recyclable, please recycle

TIW14046-USEN-00