Designed to enable comprehensive data protection in a hierarchy of storage to help optimize cost

IBM Tivoli Storage Manager and IBM Tivoli Storage Manager Extended Edition

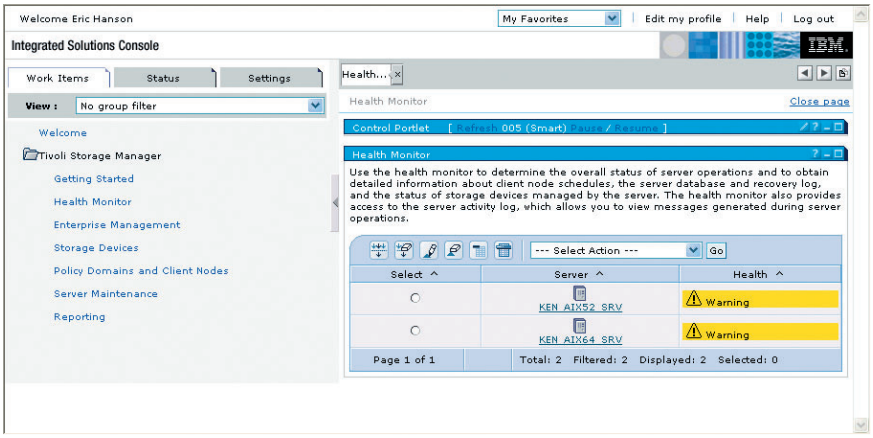

Highlights

- ***Helps manage multiple types of inactive data in a hierarchical repository***
- ***Designed to empower rich function through use of advanced architecture***
- ***Helps lower storage cost through intelligent hierarchy***
- ***Designed to provide centralized, comprehensive management***
- ***Helps reduce network bandwidth through intelligent data movement***
- ***Designed to minimize manual backup, archive and recovery tasks through policy-based automation***
- ***Offers improved backup capabilities through IBM Tivoli® Storage Manager Extended Edition***

Today, an On Demand Business faces an expanding set of data protection challenges and regulatory data retention requirements. The 24x7 availability of mission-critical data and applications is no longer a goal—it's a business necessity. The IBM Tivoli Storage Manager family of offerings is designed to provide centralized, automated data protection that can help reduce the risks associated with data loss while also helping to reduce complexity, manage costs and address compliance with regulatory data retention requirements.[1]

IBM Tivoli Storage Manager enables you to protect your organization's data from failures and other errors by storing backup, archive, space management and bare-metal restore data, as well as compliance and disaster recovery data in a hierarchy of offline storage. Because it is highly scalable, IBM Tivoli Storage Manager can help protect computers running a variety of different operating systems, on hardware ranging from notebooks to mainframe computers and connected together through

ON DEMAND BUSINESS™

the Internet, wide area networks (WANs), local area networks (LANs) or storage area networks (SAN). It uses Web-based management, intelligent data move-and-store techniques and comprehensive policy-based automation that are working together to help increase data protection and potentially decrease time and administration costs.

Managing inactive data in a hierarchy of storage

Every environment has inactive data that it must manage, and in most IT environments, inactive data makes up the bulk of gigabytes stored. In an on demand operating environment, inactive data is managed and stored in a hierarchy of lower-cost storage. This helps to lower costs as automated policies store data on the type of media that best meets that data's longevity, access speed and cost needs.

Using IBM TotalStorage® Open Software Family solutions built on IBM Tivoli Storage Manager software and a wide range of optional, integrated IBM Tivoli Storage Manager modules, you can centrally manage inactive data, help protect your organization from hardware failures and other errors, and better match the value of your data to the most cost-effective storage management practices.

IBM Tivoli Storage Manager Extended Edition

IBM Tivoli Storage Manager Extended Edition is designed to expand IBM Tivoli Storage Manager capabilities to include support for libraries larger than three drives or 40 slots, library sharing, Network Data Management Protocol (NDMP) backup of network attached storage (NAS) devices and automated disaster recovery planning.

Helps manage multiple types of inactive data in a hierarchical repository

The IBM Tivoli Storage Manager advanced architecture, its ability to store data in a hierarchy of varying cost storage and its open application programming interface (API), give it the unique ability to seamlessly manage many types of data and integrate with many functions. IBM Tivoli Storage Manager can store and manage a range of data types, including:

- ***Backup.*** *IBM Tivoli Storage Manager and IBM Tivoli Storage Manager Extended Edition are designed to provide the ability to back up numerous versions of files to the hierarchy of storage. Should an online storage device fail, a data error occur or someone accidentally deletes a file, the chosen version of the data can be easily restored. Various types of backups are offered to meet the user's needs. They include file level, volume level backups, subfile backups, and backup sets.*

- ***Archive.*** *IBM Tivoli Storage Manager and IBM Tivoli Storage Manager Extended Edition are designed to provide the ability to archive files into the hierarchy of storage. The files are kept for a specified amount of time and then expired from the hierarchy of storage.*

- ***Database and application data.*** *IBM Tivoli Storage Manager for Mail, IBM Tivoli Storage Manager for Databases, IBM Tivoli Storage Manager for Enterprise Resource Planning, IBM Tivoli Storage Manager for Application Servers and IBM Tivoli Storage Manager for Hardware are designed to perform online backups of databases and applications. These backups are sent to the hierarchy of storage.*

- ***Space management.*** *IBM Tivoli Storage Manager for Space Management is designed to automatically identify and move low-activity and inactive files to the hierarchy of storage. Should an application try to access these migrated files, the individual files are automatically recalled back to the specified user's machine.*

- ***Data retention data.*** *IBM Tivoli Storage Manager for Data Retention helps meet regulatory requirements by preventing stored data in the hierarchy of storage from being changed or deleted prior to the intended expiration dates.*

- ***Disaster recovery planning.*** *IBM Tivoli Storage Manager Extended Edition is designed to provide built-in and automated disaster recovery planning tools.*

- ***Bare-machine recovery.*** *Bare-machine recovery brings back the system to the state of the last backup, recovering all the operating system changes and customizations. It streamlines and automates the operating system recovery process and eliminates the need for highly skilled professionals to manually reinstall hardware, network configurations and patches. This speeds up the recovery time and is integrated with the backup procedure. The bare machine backups are stored directly in the IBM Tivoli Storage Manager server's hierarchy of storage.*

Designed to empower rich function through use of advanced architecture

IBM Tivoli Storage Manager is powered by an advanced architecture that includes a built-in relational database and transaction log to track the meta-data pertaining to protected data and a hierarchy of storage that holds the actual data. This architecture helps enable IBM Tivoli Storage Manager to provide many advanced features as part of its base functionality and act as a repository for many different types of data.

Helps lower storage cost through intelligent hierarchy

IBM Tivoli Storage Manager implements a hierarchy of lower-cost storage which supports hundreds of devices and that can include low-cost disk systems, such as Serial Advanced Technology Attachment (SATA), automated tape systems and optical libraries. As an added benefit, the data can be compressed and encrypted on its way into the hierarchy. As soon as the data is in the hierarchy, it can be seamlessly moved from one type of storage to another.

Flexible storage pools

IBM Tivoli Storage Manager is designed to allow you to decide where your data should be backed up, archived or space-managed to, and where it should reside. You can also take advantage of its unique capability to automatically migrate data from one storage pool to another based on policies set up to reflect business needs.

Direct to disk

IBM Tivoli Storage Manager is designed to maximize your backup window and minimize your restore time by giving client machines the option to back up, archive and space-manage data initially to a disk storage pool on the IBM Tivoli Storage Manager server. Going directly to disk means the number of client machines moving data at the same time is not gated by how many tape drives you have. Should a disk storage pool threaten to fill up before a backup is complete, the data can automatically start to move to another pool, helping to maximize IT storage infrastructure efficiencies. Multiple disk volumes appear to IBM Tivoli Storage Manager as one large volume, thus removing disk-volume size limitations. Once data is moved to the disk storage pools, policies can help automate migration onto tape or optical storage pools.

Before migrating out to tape or optical storage pools, each data is grouped together so it is written in an organized fashion, thus allowing faster restores.

Collocation

Collocation takes the grouping of individual client-machine data a step further by allowing you to specify that a particular client's data should reside on its own tape or set of tapes. By collocating a group of client's, a client's or a file system's data, you expedite restores, as fewer tapes need to be mounted, less tape positioning needs to occur and contention for tapes during simultaneous data restores does not occur.

Reclamation

Expiring files that reside on tape create dead space, which cannot be rewritten, wasting valuable tape space and causing restores to have to be positioned over that space. IBM Tivoli Storage Manager has an easy solution to this—reclamation. Reclamation is designed to be automatically controlled based on polices you set up. As soon as a tape has a certain percentage of dead space on it, the remaining unexpired files are moved to other tape volumes. The resulting empty tapes are then reused for new data—enabling you to make more efficient use of your data and helping to improve your restore times.

Offsite copies

Disasters that wipe out your entire IT center happen. Tapes break. To protect against these scenarios, IBM Tivoli Storage Manager is designed to allow you to create, using schedules or on demand, copies of your backed-up data. You can then take these copies offsite to protect your backups. IBM Tivoli Storage Manager continues to track the offsite volumes' content, so should a tape need to be brought back onsite because the original tape failed, or because the data on it has expired, IBM Tivoli Storage Manager is designed to automatically notify the administrator. Alternatively you can send the backups automatically over the network to another IBM Tivoli Storage Manager server, thus eliminating the need to manually take tapes offsite.

Migration to new technology

The hierarchy of storage also helps manage the process of automatically moving data to newer media types when it has outlived the usefulness of the current media. This often happens when long-term archive data outlives a particular tape technology. The data does not have to be re-backed up, re-migrated or re-archived to achieve this.

Provides centralized, comprehensive management

IBM Tivoli Storage Manager is designed to simplify administrative tasks and help save the administrator time through a range of robust features.

Central administration

Central administration helps provide server management of a large Tivoli deployment from any IBM Tivoli Storage Manager client platform using robust administrator capabilities.

Web-based user interface

The administration center is designed to provide a leap forward in ease of use through its updates and focus on the new task-oriented user interface, simplification of administration and management, efficient use of resources and its inclusion into the IBM Integrated Solutions Console framework (ISC).

The Administration Center includes a single interface to sign onto and manage multiple IBM Tivoli Storage Manager servers; Wizards to simplify and reduce the time needed to complete more complex task; and Health Monitor, where an administrator might go first thing every day to easily see the status of the server selected, successful and unsuccessful scheduled runs,

query the activity log using search strings and time stamps and see the status of storage devices attached or accessed through the IBM Tivoli Storage Manager server.

Operational reporting

Sending reports about IBM Tivoli Storage Manager operations directly to administrators helps save them time. Operational reporting draws from IBM Tivoli Storage Manager database information to create three types of reports. You can customize the reports to meet your needs or use them as they are. The first report is sent every day to your administrator. It reports information such as the amount of files backed up, missed backups or the number of tapes available. The second report is automatically sent to any computers where the backup failed. The third report is generated any time a particular condition (like low on backup tapes, or a tape drive is offline) is met. Reports can be viewed as e-mails, desktop alerts or pages, through the IBM Tivoli Storage Manager Microsoft® Management console or exported to a Web site.

Enhanced automation functions

These are designed to help eliminate or reduce staff-intensive manual procedures by using server scripts to automate server management operations and policy constructs. This helps to automatically extend the server database and recovery log when administrator-defined thresholds are reached.

Enterprise management

Companies that have multiple IBM Tivoli Storage Manager servers reduce their workload with enterprise management tools. IBM Tivoli Storage Manager configuration and policy information can be defined one time at an IBM Tivoli Storage Manager configuration server and then be propagated to any number of managed IBM Tivoli Storage Manager servers. Enterprise command routing functions are designed to let an administrator use a single command to perform queries and other tasks on groups of networked IBM Tivoli Storage Manager servers, as well as on a single server. Results of the routed commands are returned to the origin of the request. Enterprise logging functions allow IBM Tivoli Storage Manager servers to forward their events and those of their clients to a server designated as the event server, thereby consolidating events from groups of IBM Tivoli Storage Manager servers and clients.

Server-to-server communication

Server-to-server communication helps enable objects to be sent to or received from another server directly. This is designed to allow one server's disaster-recovery copies to be automatically sent over the network to another IBM Tivoli Storage Manager server.

Intelligent data movement

IBM Tivoli Storage Manager includes a number of intelligent data-movement capabilities to save tapes, processing cycles and network bandwidth—and to help decrease backup and restore times.

Progressive backup

IBM Tivoli Storage Manager uses, by default, an intelligent, progressive backup strategy. Only files that have changed or that are new are backed up, helping to reduce or eliminate unnecessary data transfers that rob your network and processors of vital power and productivity. With progressive backup, you do not need to perform wasteful periodic full backups and you can get faster restores because IBM Tivoli Storage Manager needs to restore only the version of the file requested. Progressive backup backs up less data, saving you network bandwidth, tapes and management overhead.

Subfile backup
IBM Tivoli Storage Manager sub file backup technology takes progressive backup one step further by backing up just the changed portion of a file. With sub file backup, the backup-archive client is designed to dynamically determine the most efficient approach for creating backup copies and which bytes, blocks or files changed, delivering improved backup performance over dial-up or WAN connections. This application is designed for mobile computer users and other users with a need to minimize the amount of data transmitted over a network.

Backup sets and rapid recovery
Rapid recovery is designed to provide the ability to create a backup set that consolidates a client's files onto a set of media that is portable and can be directly readable by the client's system for fast, LAN-free (non-networked) restore operations. The portable backup set, which is synthesized from existing backups, is tracked by the IBM Tivoli Storage Manager server. Portable backup sets can be copied to numerous types of media. IBM Tivoli Storage Manager backup-archive clients can directly restore data from the backup set media using standard operating system device drivers. IBM Tivoli Storage Manager server is *not* required to restore the data and no data is moved over the network—making this an excellent tool for remote recoveries or rapid disaster recovery.

Restartable backup and restore
IBM Tivoli Storage Manager gives you the ability to restart from where you left off if backups or restores are interrupted. This saves you critical time by not having to start a backup or restore from the beginning.

Non-disruptive online image backup and restore
For instances in which speed is more important than individual file restores, image backups complement progressive incremental backups to help provide full file-system backup and restores.

Multiple session journal-based backup
Journal-based backup, which is available on Windows, can improve the overall backup time of incremental backups by eliminating the scan time of a file system for file selection. The journal-based backup function can be used to provide concurrent backup of different file systems. The existing journal-based backup capability is enhanced to allow multiple sessions to back up the different file systems by running multiple instances of the backup client or by using the Resource Utilization client option.

Multi-process reclamation and migration
The time needed to perform IBM Tivoli Storage Manager reclamation and migration for sequential access storage pools is reduced when multiple tape drives are configured for these operations. The number of multiple concurrent reclamation and migration processes can be configured for the same sequential-access storage pool to utilize multiple drives, which is in addition to the concurrency that occurs when there are multiple sequential pools. The new concurrency will occur on the volumes within a pool. Disk storage pools and optical storage pools that are formatted for sequential access can also utilize the multiple drives.

API
The IBM Tivoli Storage Manager API is designed to help enable applications to access an IBM Tivoli Storage Manager server for backup, archive or specialized services and utilize our storage hierarchy and the hundred of devices supported. Our no-charge, open API offers critical services to data-intensive applications.

Minimizes manual tasks through policy-based automation

Policies and schedules reduce an administrator's workload and help ensure that the data is being protected in the correct manner.

Policies

IBM Tivoli Storage Manager has an extremely granular policy engine that is designed to automate its capabilities to match your business needs. This policy engine is configurable down to the individual file level. You can set policies as simple defaults or as highly tuned and customized options to help reduce costs and still meet specific service-level agreements. The end result is a *set-it-and-forget-it* experience that makes IBM Tivoli Storage Manager highly suited for *lights-out* implementations. Should you need to make changes to the policies, IBM Tivoli Storage Manager can apply these new policies to the data it's already managing—so you don't have to back up, archive or migrate the data again. After it has been configured, IBM Tivoli Storage Manager can manage data protection for thousands of computers with minimal administrator assistance.

Client schedules

Users do not need to remember to back up their machines. Automated backup schedules do this for them. IBM Tivoli Storage Manager is designed to allow you to set up as few or as many schedules as needed to meet your users' backup needs.

Administrator schedules and scripts

Routine administrative tasks can be scheduled to run during convenient server activity times. These activities can also be scripted to allow the processing of server commands sequentially or in parallel to other tasks.

Scheduling of both client and administrator operations can be done with an external scheduler or with the built-in IBM Tivoli Storage Manager scheduler. The ITSM scheduler provides calendar-based scheduling for increased flexibility of client and administrative schedules. The scheduler can be exploited via administrative commands or a Web interface.

IBM Tivoli Storage Manager Extended Edition capabilities

IBM Tivoli Storage Manager Extended Edition expands IBM Tivoli Storage Manager capabilities through the following functions:

Automatic disaster preparation and recovery scripts

IBM Tivoli Storage Manager Extended Edition can automatically create a recovery plan—an up-to-date text file with detailed recovery steps and automated computer scripts. This plan helps minimize recovery time from a disaster that affects your data center. The disaster-recovery functionality can also track where offsite copies of data are stored: Where in the vault are they? Are they in transport to or from the vault? Are they waiting to be picked up from the library? This level of detail helps expedite disaster recovery because you can rapidly locate the physical tapes that your offsite copies are stored on.

Data protection for NAS appliances through NDMP

IBM Tivoli Storage Manager Extended Edition uses NDMP to help enable high-performance, scalable backups and restores on NAS devices. NDMP data movements are designed to minimize network traffic by transferring data outbound from both the IBM Tivoli Storage Manager client and server. Both full and differential file system image backups and both full and file-level restores are supported.

Small and large tape library support

IBM Tivoli Storage Manager Extended Edition works with more than 500 different offline storage devices from more than a dozen vendors with no limitations on the number of drives or slots.

Tape library sharing between multiple IBM Tivoli Storage Manager servers

With IBM Tivoli Storage Manager Extended Edition, multiple IBM Tivoli Storage Manager servers can share the same data tape library for improved data centralization and asset utilization.

Benefits of IBM Tivoli Storage Manager and IBM Tivoli Storage Manager Extended Edition

Improved application availability

- *Speeds recovery of files, file systems or bare machines*
- *Includes file recovery in a comprehensive data disaster recovery plan*

Enhanced storage resource utilization

- *Designed to optimize backup granularity*
- *Stores backup, archive and space-managed copies in a hierarchy including low-cost storage—the specific level chosen to optimize cost against the recovery-time objective*
- *Reduces backup times and resource usage by focusing on active files*

Enhanced storage personnel productivity

- *Creates a single point of control, administration and security for file recovery*
- *Indexes file archive copies with descriptive metadata to aid administrators in locating historical information*

IBM Tivoli Storage Manager and IBM Tivoli Storage Manager Extended Edition

Supported servers

- *IBM AIX®*
- *HP-UX*
- *Microsoft Windows®*
- *Sun Solaris operating environment*
- *IBM OS/400®*
- *IBM OS/390® or IBM z/OS®*
- *Linux® on IBM @server® pSeries®: SUSE Enterprise Server*
- *Linux on IBM @server xSeries®: Red Hat Enterprise Linux AS*
- *Linux on IBM @server zSeries®: SUSE LINUX Enterprise Server*

Supported clients

- *IBM AIX*
- *HP/UX*
- *Linux*
- *Macintosh*
- *Novell NetWare*
- *z/OS USS and Linux for zSeries*
- *SGI IRIX UNIX®*
- *Sun Solaris operating environment*
- *Tru64 UNIX*
- *Microsoft Windows*

IBM Tivoli Storage Manager and IBM Tivoli Storage Manager Extended Edition support many leading vendor products. The list is dynamic and extensive. Please visit our Web site to access the most current list and other details at:

ibm.com/software/tivoli/products/storage-mgr/platforms.html

This page was left intentionally blank.

This page was left intentionally blank.

For more information

To learn more about IBM Tivoli Storage Manager and IBM Tivoli Storage Manager Extended Edition, visit:

ibm.com/software/tivoli/products/storage-mgr-extended

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

Produced in the United States of America
December 2004

This document could include technical inaccuracies or typographical errors. IBM may make changes, improvements or alterations to the products, programs and services described in this document, including termination of such products, programs and services, at any time and 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 contained in this document is current as of the initial date of publication only and is subject to change without notice. IBM shall have no responsibility to update such information.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. IBM makes no representations or warranties, express or implied, regarding non-IBM products and services.

Performance data for IBM products and services is based on measurements and projections using standard IBM benchmarks in a controlled environment. 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. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the information provided herein.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "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, if at all, 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.

References in this document to IBM products, programs or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM program or product in this document is not intended to state or imply that only that program may be used. Any functionally equivalent program or product that does not infringe IBM's intellectual property rights may be used instead. It is the user's responsibility to evaluate and verify the operation of any non-IBM product, program or service.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:
IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

Note to U.S. Government Users—Documentation related to restricted rights—Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract and IBM Corp.

[1] IBM's customer is responsible for ensuring its own 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.

TSD00614-USEN-02