LAB VALIDATION REPORT

IBM TSM 6.1

Unified Data Protection and Recovery Management

By Tony Palmer
With Brian Garrett and Lauren Whitehouse

April, 2009

Table of Contents

ESG Lab Reports

The goal of ESG Lab reports is to educate IT professionals about emerging technologies and products in the storage, data management and information security industries. ESG Lab reports are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objective is to go over some of the more valuable feature/functions of products, show how they can be used to solve real customer problems and identify any areas needing improvement. ESG Lab's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments. This ESG Lab report was sponsored by IBM.

Introduction

IT organizations are challenged with maintaining application and data availability while satisfying compliance requirements and service level agreements. Relentless information growth is straining IT resources—necessitating greater investments—yet capital and operational expenses are under scrutiny. Given this conflict, IT organizations are seeking efficiency and optimization to regain control. IBM Tivoli Storage Manager (TSM) software is a client/server software solution providing backup/recovery, archive, hierarchical storage management (HSM), and disaster recovery (DR). This ESG Lab Validation focuses on key improvements in the TSM platform that drive greater efficiency: incorporation of the DB2 database for enhanced scalability, availability and performance; embedded data deduplication; enhanced monitoring and reporting; as well as integration with NetApp's Snapshot Difference API and SnapMirror to tape technologies to back up IBM N series and NetApp FAS volumes.

Background

Today, end-users are eyeing new approaches and technologies in hopes of easing the stress on their backup and recovery environments. Employing disk in the backup process has proven to be a powerful weapon for reducing pressure on backup windows caused by ever-increasing data growth. While typically used as a temporary cache for backups on the way to less-costly tape media, the combination of lower disk costs and the desire to comply with recovery SLAs has prompted many organizations to retain backups longer on disk—with some eliminating tape altogether. The practicality and feasibility of leveraging disk in backup and recovery processes changes dramatically when capacity optimization technology, such as data deduplication, is applied.

FIGURE 1. DEDUPLICATION SELECTION CRITERIA

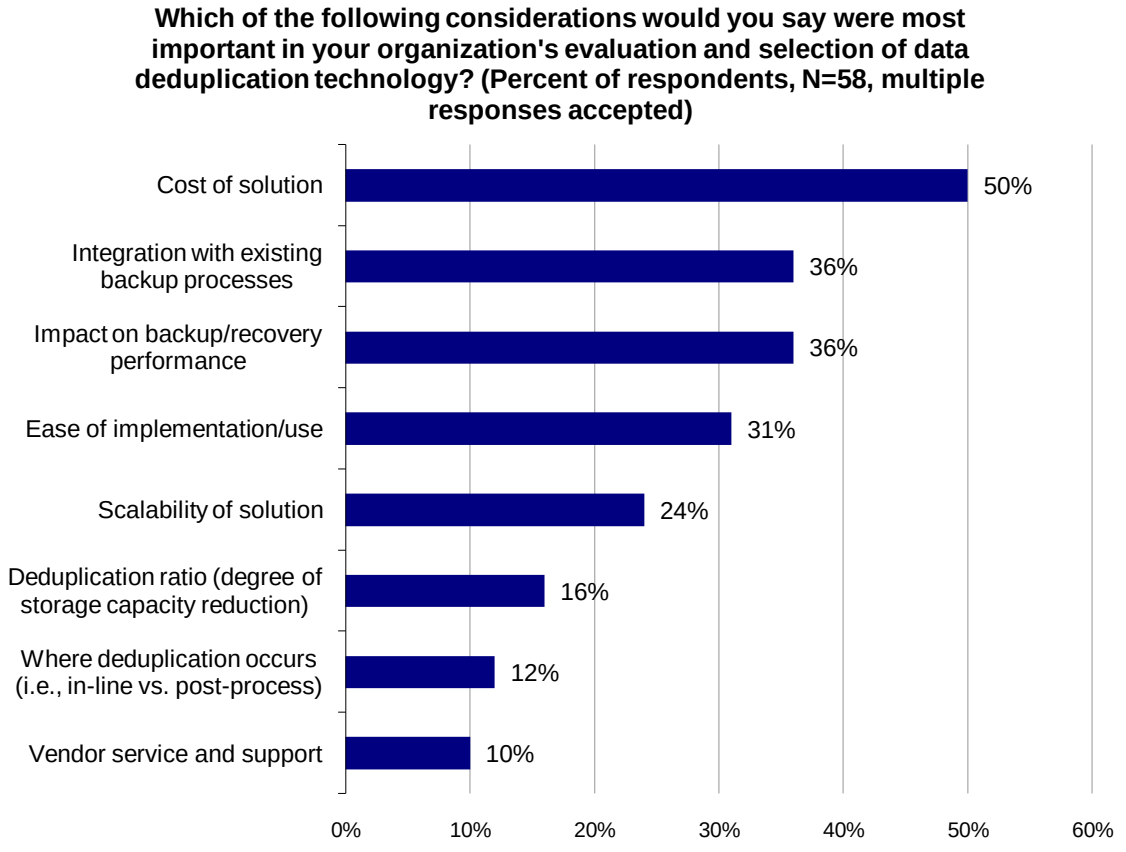

As shown in Figure 1, ESG research found that the top three criteria for selecting deduplication solutions include cost, integration with existing backup processes, and impact on backup/recovery performance.[1] Backup/recovery software that embeds data reduction technology can address all three of these factors handily. IBM TSM 6.1 now has native deduplication capabilities built into its Extended Edition (EE) as a no-cost option. After data is written to the primary disk pool, a deduplication operation can be scheduled to eliminate redundancy at the sub-file level. Data deduplication, as its name implies, identifies and eliminates redundant data.

TSM 6.1 also includes features that optimize TSM scalability and manageability to meet increasingly demanding service levels resulting from relentless data growth. The move from a proprietary back-end database to IBM DB2 improves scalability, availability, and performance without adding complexity; the DB2 database is automatically maintained and managed by TSM. IBM upgraded the monitoring and reporting capabilities to near real-time and completely redesigned the dashboard that provides visibility into the system. TSM and TSM EE include these enhanced monitoring and reporting capabilities at no cost. IBM has made the monitoring and reporting tools compatible with TSM 5.x for existing customers as well.

Figure 2 represents a simplified view of the TSM 6.1 architecture. Each TSM server instance is supported by a DB2 database and a recovery log. A primary disk storage pool often functions as an initial landing area for daily backup jobs. Data is then migrated to a copy storage pool (for DR) and deduplication is applied based on policies set by the administrator. The TSM database tracks the actual location of all backup objects at all times. TSM 6.1 plugs into the DB2 database to provide centralized monitoring and reporting through the GUI.

FIGURE 2. TSM ARCHITECTURE

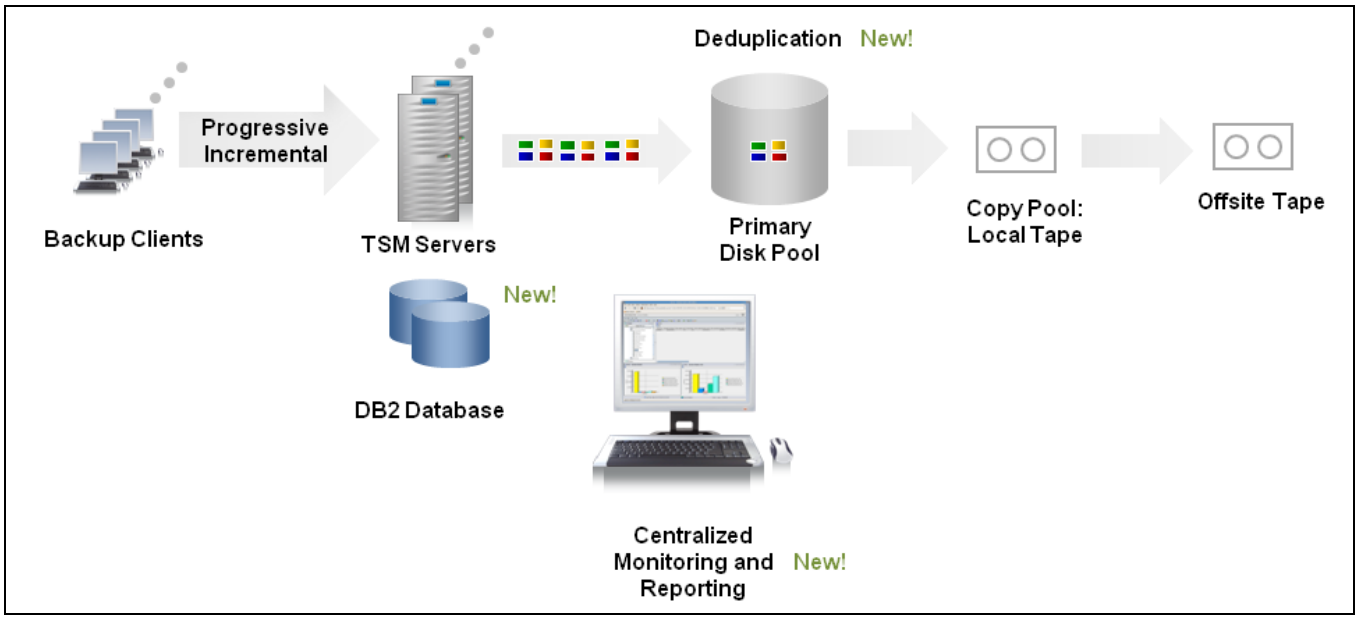

Three enhancements to TSM in version 6.1 streamline backup and recovery in consolidated and virtualized environments. NetApp SnapMirror to tape technology enables customers to use NDMP (Network Data Management Protocol) support to make block level backups of IBM N series and NetApp FAS volumes. NetApp Snapshot Difference API enables TSM to reduce the time required for progressive incremental file backups by eliminating client filesystem scans. Enhanced integration with VMware Consolidated Backup (VCB) automates Virtual Machine backup and recovery processes within the GUI at the image level.

This ESG Lab report examines several of the enhancements to TSM capabilities in version 6.1 designed to drive greater efficiency in backup and recovery processes in greater detail. Deduplication, the DB2 database, monitoring and reporting, and NAS and VMware enhancements were tested by ESG Lab.

[1] Source: ESG Research Report, Data *Protection Market Trends*, January 2008.

ESG Lab Validation

ESG Lab performed hands-on evaluation and testing of IBM TSM 6.1 at IBM's labs in Tucson, Arizona. Validation was performed using a combination of hands-on testing, audits of IBM test environments, and detailed discussions with IBM TSM experts. ESG Lab began with an examination of the new data deduplication functionality integrated into TSM 6.1.

Data Deduplication

For 15 years, TSM has been providing customers with data reduction capability with its progressive incremental backup technology. Traditional backup methodologies use periodic full backups combined with more frequent incremental or differential backups to conserve storage capacity. This approach makes restores more complex and time consuming as users must first restore the most recent full backup, then restore subsequent incremental backups to recover to the current point in time. TSM's progressive incremental backup technology still makes an initial full backup, but all subsequent backups are incremental. When combined with TSM's advanced data management techniques, this enables fast full restores without the complexity of managing multiple backup sets. Users get the speed and recoverability benefits of a daily full backup in addition to the reduced network traffic and data storage requirements of incremental backups.

FIGURE 3. DATA REDUCTION WITH PROGRESSIVE INCREMENTAL BACKUPS AND DEDUPLICATION

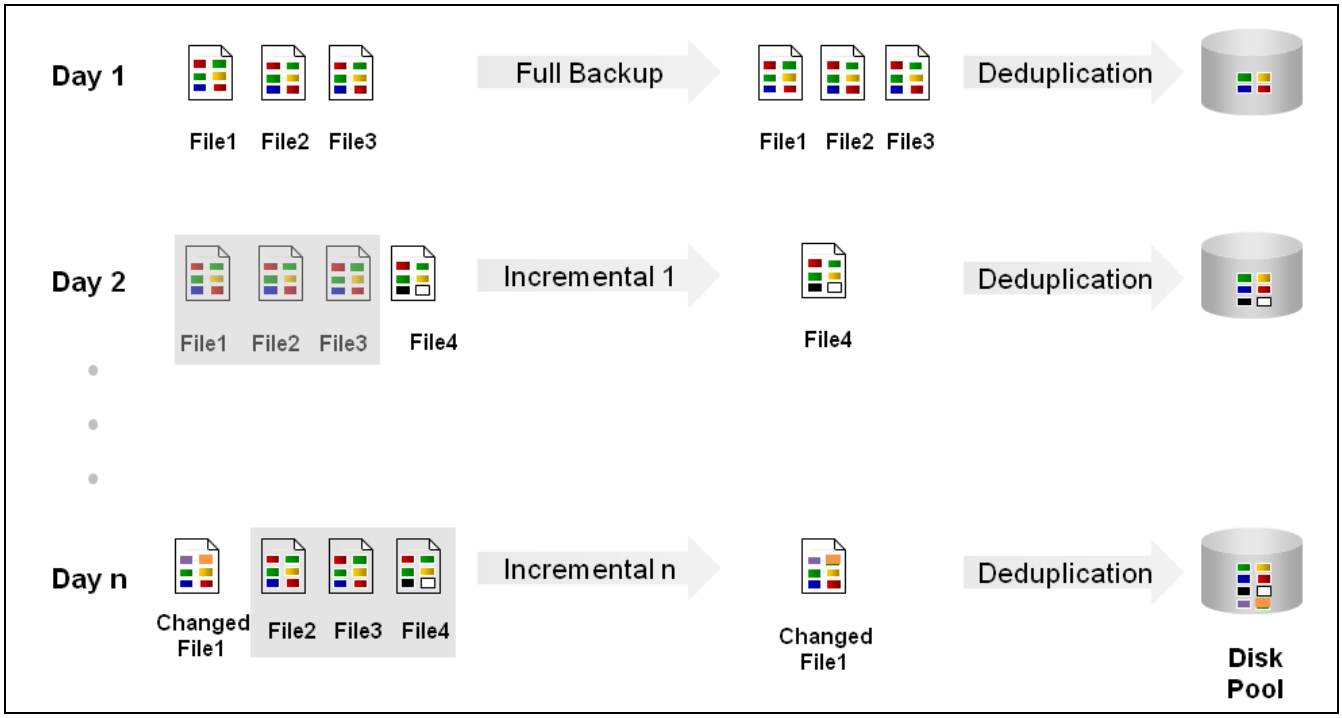

Data deduplication takes this one step further. Figure 3 shows a simplified timeline of backups using both progressive incremental and deduplication technology. Day 1 represents the first full backup, in which all files are sent to TSM. TSM then applies deduplication to the data set, marking redundant blocks of data for removal. A separate process (reclamation) deletes the redundant blocks. Day 2 represents the first incremental backup, in which one new file was added. Just that one file is sent to TSM, which then compares the blocks in the file to the pool of already stored blocks and marks redundant blocks for removal. Day n represents a future incremental after a file has been changed. The changed file is sent to TSM, which again compares the blocks in the file to already stored data in the deduplicated disk pool. The unique blocks are retained and the redundant blocks are

marked for deletion during the next reclamation process. TSM manages the organization and placement of data so that the backup administrator can select any day and perform a full restore to that point in time without having to perform multiple passes of full and incremental or differential backups.

ESG Lab Testing

The TSM management interface was used to examine how deduplication was configured in TSM. Deduplication was enabled on a storage pool using a single command and individual clients were pointed to the deduplicated storage pools using Policy Domains. Figure 4 shows a list of client systems configured for different backup policies.

FIGURE 4. ENABLING DEDUPLICATION

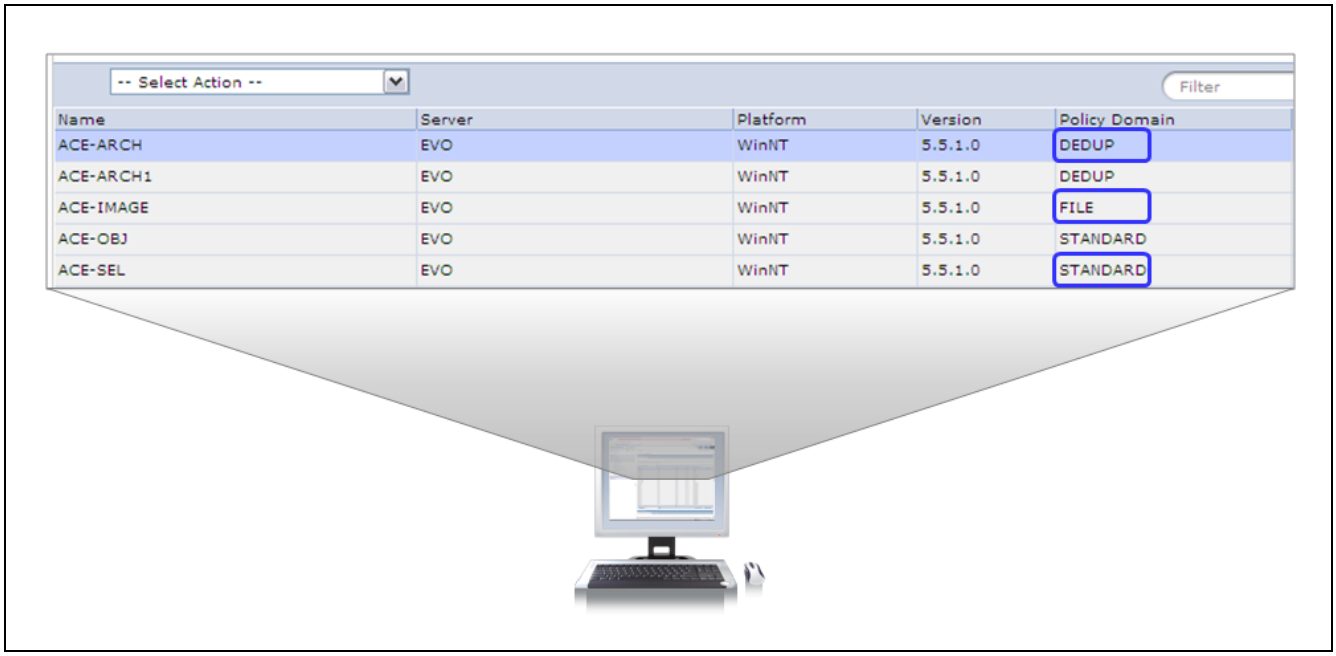

ESG Lab audited IBM's in-house testing to validate the benefit of deduplication in combination with progressive incremental technology.[2] During the ten day period examined, IBM performed nearly 500 backups daily against an 11 TB environment containing a mix of data typically found in an enterprise environment, including user files, VMware images, databases, and Microsoft Exchange data.

Figure 5 shows the combined benefits of progressive incremental backups combined with data deduplication technology in TSM 6. The red line at the top represents the total data protected, and the blue line indicates the data actually stored on disk after nightly deduplication and reclamation.

ESG Lab examined the capacity consumed in the file pool after running each backup, deduplication, and reclamation session. Immediately after the first backup had completed, the full data set was measured at just over 11 TB. After eleven backup jobs, with an average change rate of 2.1% per day, progressive incremental backups in combination with deduplication had reduced required capacity to only 7.8 TB. The capacity savings over time are also shown in Figure 5. In this example, the combination of TSM progressive incremental backups and data deduplication reduced disk capacity by a factor of 19:1 after just ten backups.

[2] Configuration details can be found in the Appendix.

FIGURE 5. CUMULATIVE DATA REDUCTION OVER TIME

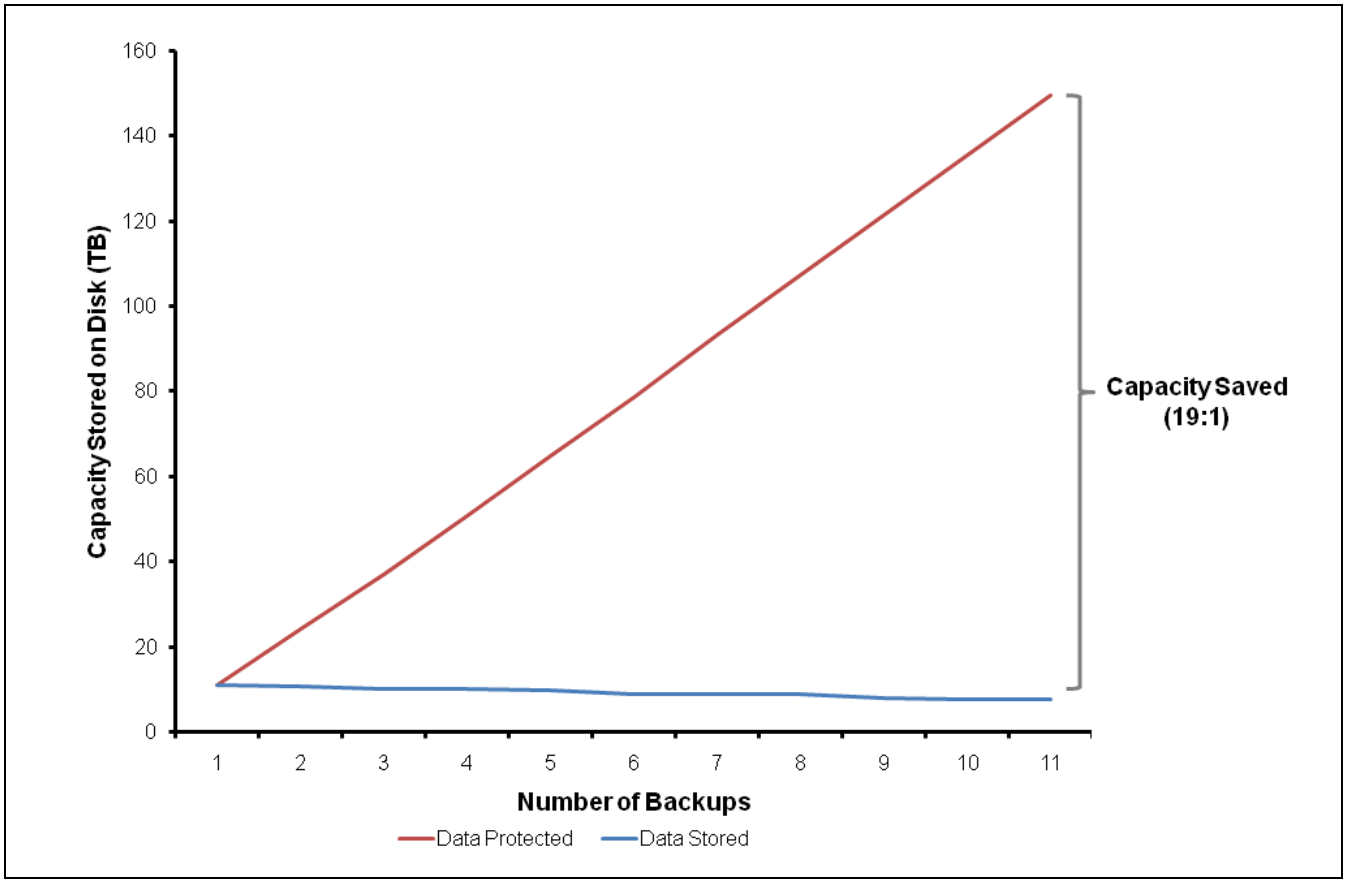

TABLE 1. DEDUPLICATION TESTING DETAILS

Day	Cumulative Data Protected (TB)	New or Changed Data (TB)	Data Stored on Disk (TB)	Data Reduction Factor
1	11.2	0.24	11.2	0.0%
2	24.2	0.24	11.0	54.5%
3	37.2	0.25	10.4	72.0%
4	50.9	0.25	10.3	79.8%
5	64.8	0.26	10.0	84.6%
6	78.8	0.26	9.0	88.5%
7	93.4	0.27	9.2	90.2%
8	107.5	0.27	9.0	91.6%
9	121.5	0.28	8.2	93.2%
10	135.5	0.28	7.8	94.3%
11	149.6	0.29	7.8	94.8%

At the end of the test period, ESG Lab compared these results to the data reduction achieved by progressive incremental backups alone and found that deduplication had shrunk the final data set by nearly 50%.

What the Numbers Mean

- IBM's data reduction technologies reduce capacity requirements dramatically while providing fast, single-pass restore capability.
- After just six incremental backups, data storage requirements were reduced by 90% as compared to daily full backups.
- Data deduplication enhanced data reduction in TSM by nearly 50% over progressive incremental backups alone.

Why This Matters

ESG research[3] indicates that cost is the leading obstacle to disk-based backup deployments. Data reduction technologies change the economics of backup to disk by reducing the cost of storage required to retain backups on disk.

ESG Lab has validated that progressive incremental backups and data deduplication in IBM TSM 6.1 can be used to significantly reduce disk capacity while enabling administrators to apply deduplication appropriately to address data type and retention needs. TSM 6.1 administrators can effectively provide high performance backup services, plus fast and reliable restores, using greatly reduced disk capacity. This lowers the cost per GB for backup data and enables companies to retain data exponentially longer for recovery purposes while minimizing the impact of deduplication on backup windows and recovery SLAs. The inclusion of deduplication in TSM EE at no additional cost adds to the value delivered to users.

Enhancing Scalability and Availability

Through version 5.5, TSM's architecture included a proprietary back-end relational database to maintain information about all clients, files, policies, and scheduling. A recovery log was also maintained to roll back any changes made if a storage transaction was interrupted before completion. The TSM server keeps the database updated and catalogued throughout the backup/restore process. The attributes of client data—number of versions, retention timeframes, and storage destination—are stored in the database. This enables TSM to define storage management policies for clients, satisfy recovery requests, and provide advanced data protection functionality, such as progressive incremental backups. Through version 5.5, the proprietary back-end database was one developed specifically for TSM. The database had a hard size limit of 530 GB, but in practice, users kept their databases much smaller than that for performance and manageability. TSM 6.1 replaces the proprietary database with IBM DB2 Enterprise Edition. DB2 offers several enhancements over the previous database, most notably the ability to manage a much larger volume of objects and improved backup performance.

ESG Lab Testing

ESG lab examined the installation and upgrade process for TSM using a Windows-based server running TSM 5.5 with a relatively small TSM database.[4] First, ESG installed the upgrade wizard on the TSM Server using the distribution CD. The next step was to uninstall the 5.5 server code and client using the Add/Remove software tool in the Windows Control Panel. At this point, ESG Lab launched the TSM 6.1 installation Wizard, seen in Figure 6. The installation wizard is consistent across all TSM components and products, providing a click-through interface that is straightforward and easy to follow.

[3] Source: ESG Research Report, *Data Protection Market Trends*, February 2008.
[4] Configuration details can be found in the Appendix.

FIGURE 6. THE TSM INSTALLATION WIZARD

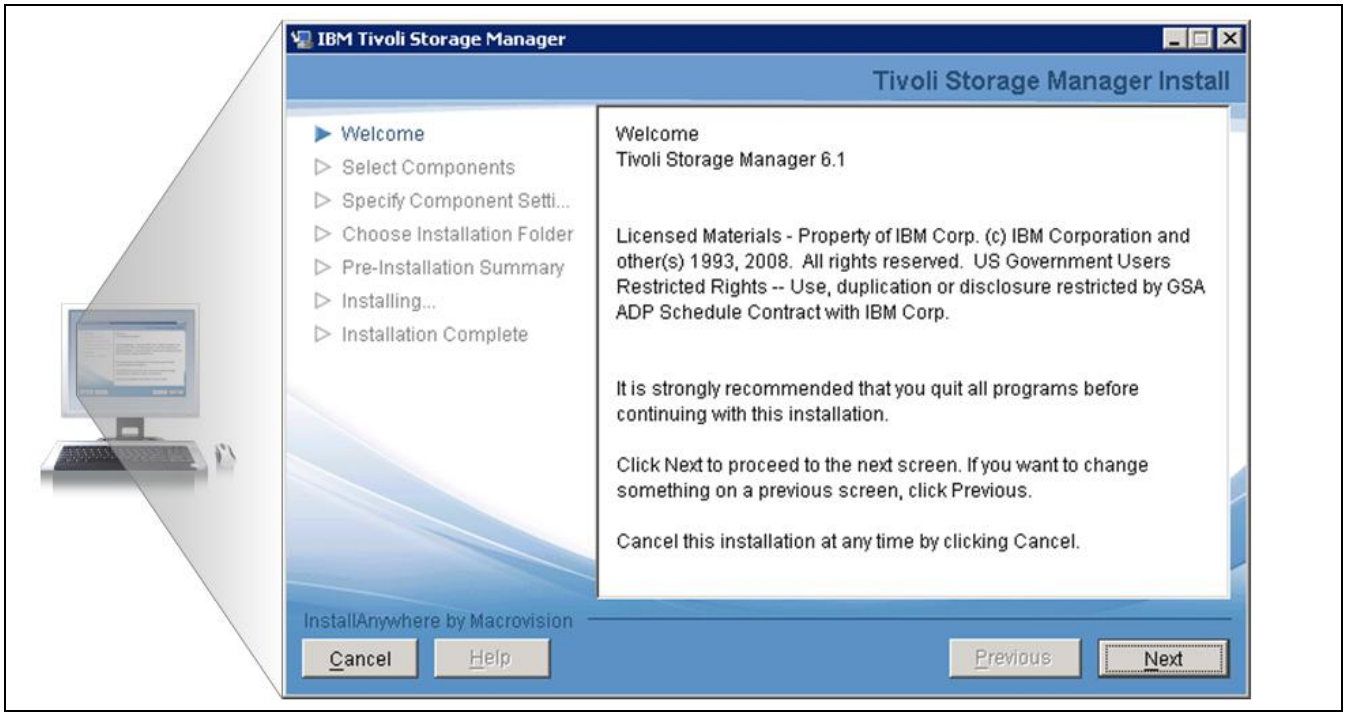

First, ESG Lab selected the components to install: the TSM 6.1 server, license, and driver. Shown in Figure 7, the wizard next prompted for the desired user name and password for the DB2 database that would be created in the final step. This is the only DB2 configuration required in the entire installation.

FIGURE 7. SETTING UP DB2

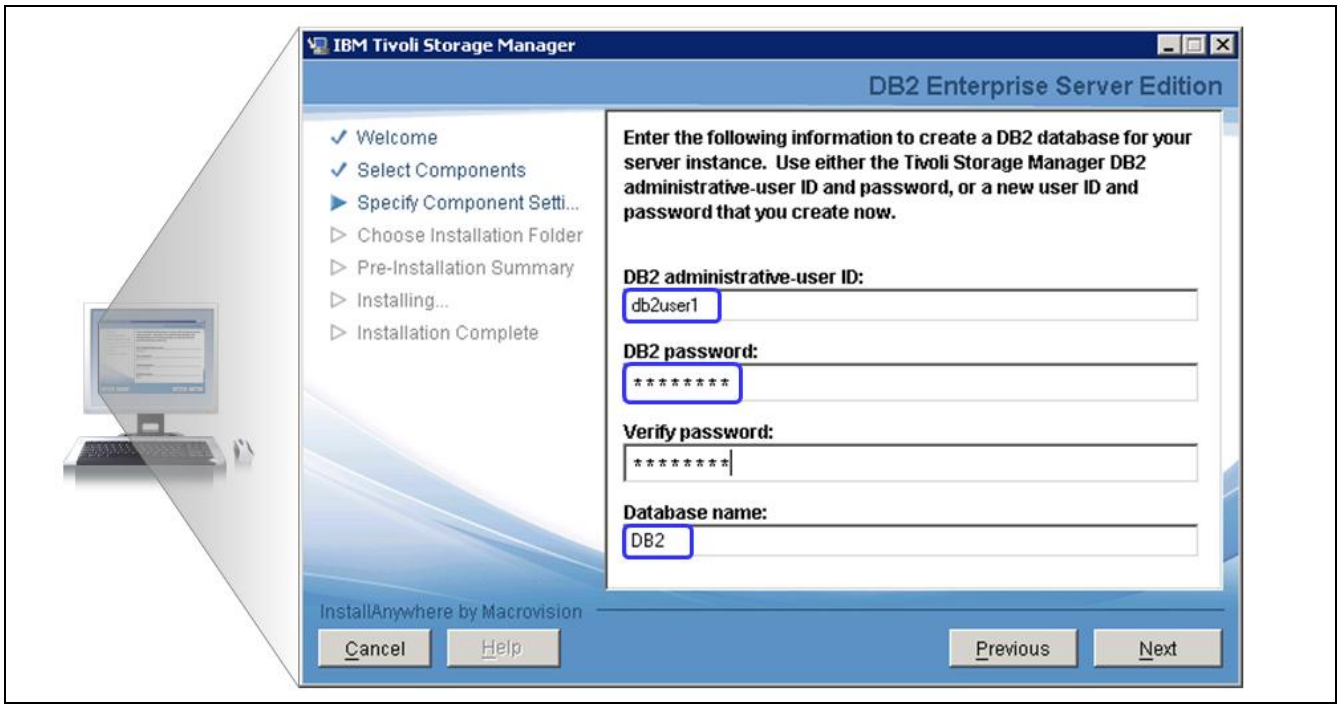

Once the destination folder had been selected and the pre-installation summary verified, the wizard installed TSM 6.1 and created the new DB2 database. Once the installation of TSM 6.1 was complete, ESG Lab launched the TSM Server Upgrade Wizard, seen in Figure 8.

FIGURE 8. THE TSM 6.1 UPGRADE WIZARD: SELECTING THE SERVER

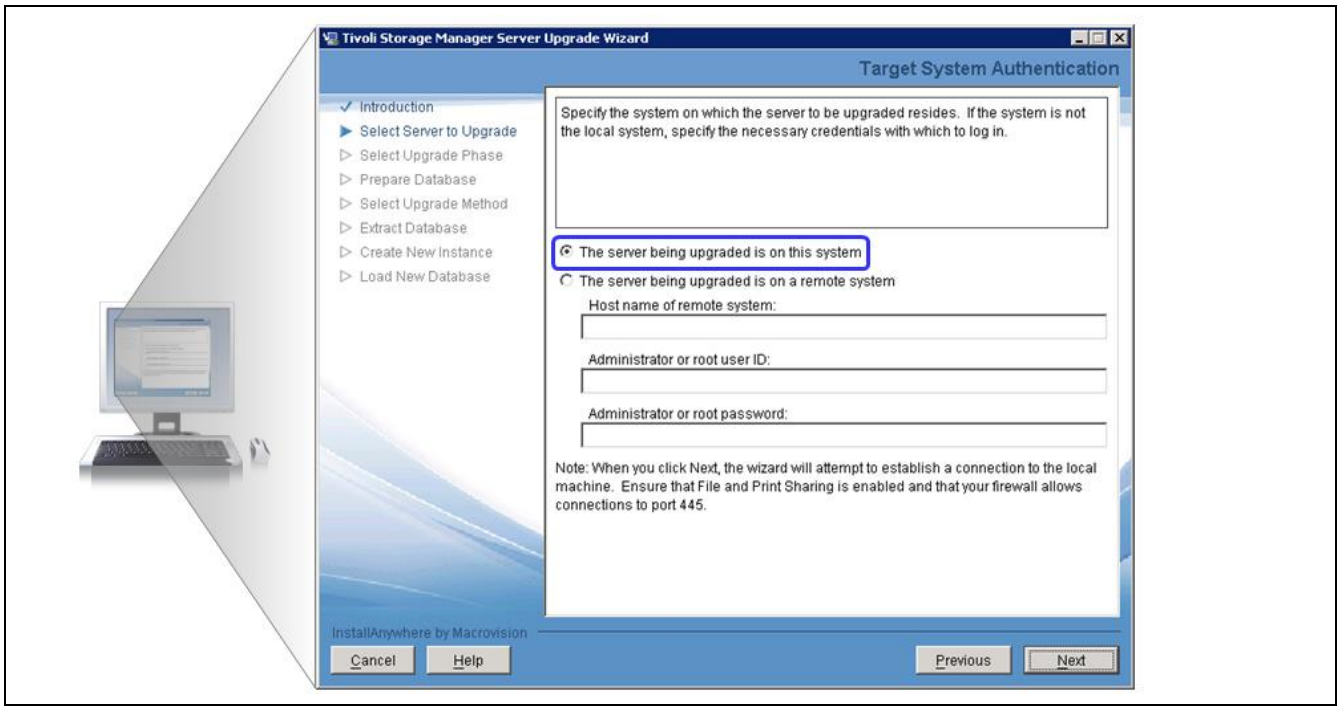

After prompting for the location of the server to be upgraded (local or remote) and the location of the database files, the upgrade wizard prepared the existing database and extracted the data from it. The wizard then created the new DB2 instance and loaded the data extracted in the previous step. After about 10 minutes of processing, the upgrade completed, as seen in Figure 9.

FIGURE 9. UPGRADE COMPLETE

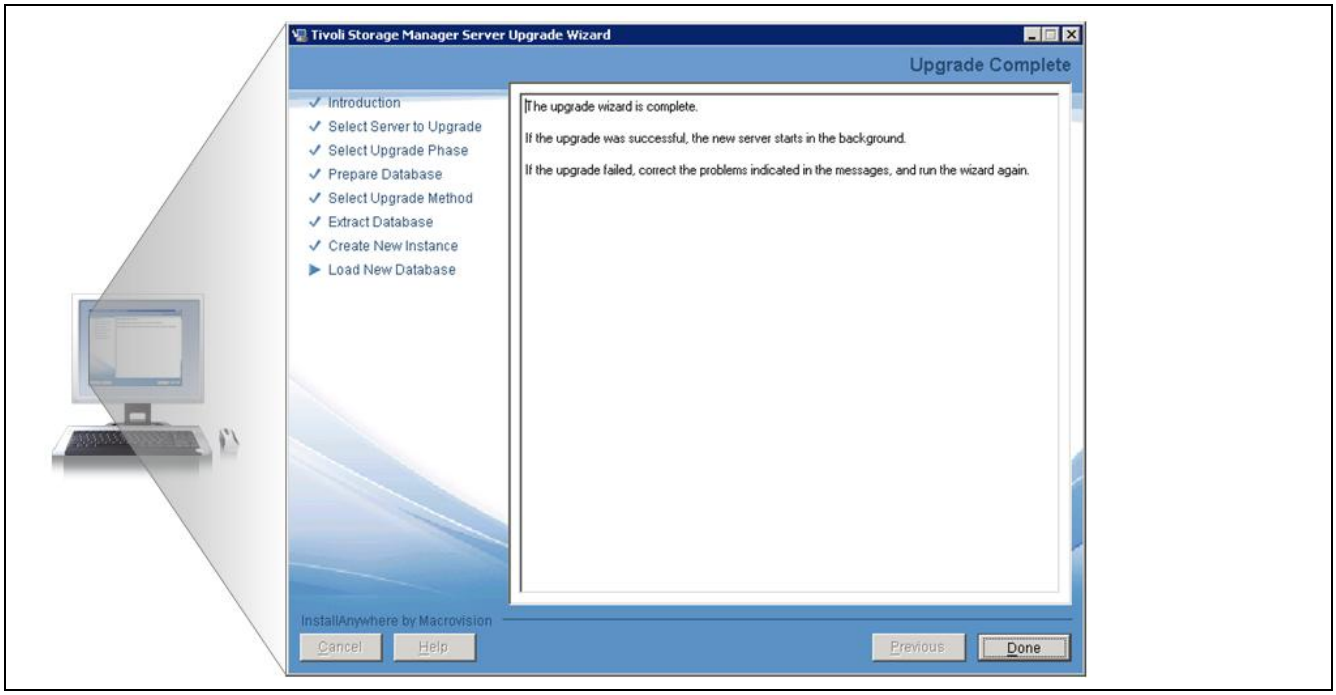

ESG lab then connected to the TSM database and verified that it was running and that all previous policies, clients, and objects were correctly represented in the new system.

ESG Lab audited IBM internal performance testing of concurrent backups using multiple clients and 10 KB files and observed TSM 6.1 backup throughput to be between two and four times greater than version 5.5 with reduced CPU utilization on the AIX platform.

Finally, ESG Lab connected to a running TSM 6.1 system in IBM's Client Environment Test (CET) lab. The CET lab is used to test TSM against large numbers of clients with large, dynamic data sets. This particular server, seen in Figure 10, was running a DB2 instance 728 GB in size (50% larger than the largest possible TSM Database in previous versions) and was managing protection of approximately 500 million objects for 1,400 TSM clients.

FIGURE 10. A 700 GB TSM DATABASE

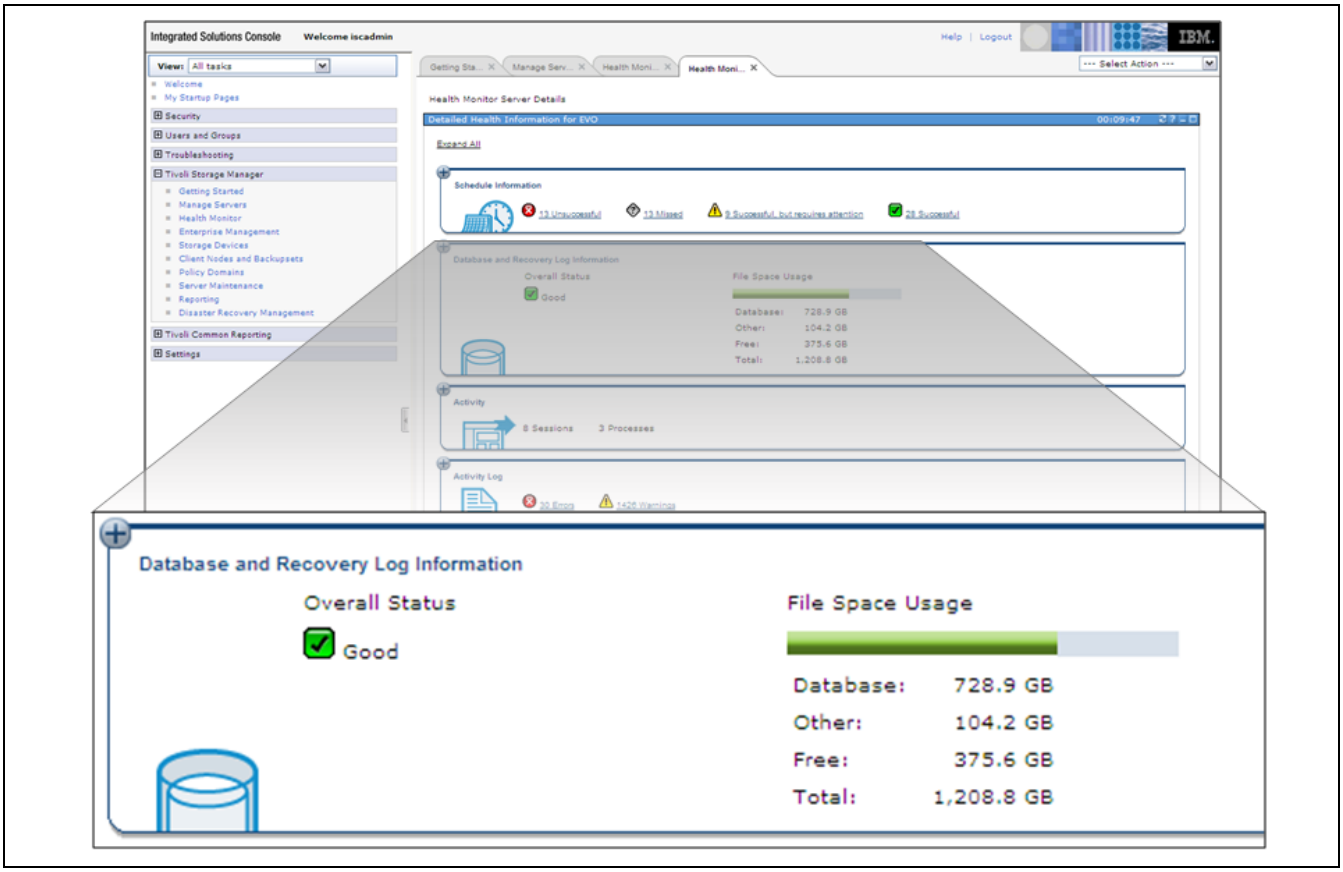

Why This Matters

ESG asked IT managers to identify major challenges with their data protection technologies and processes.[5] The top three responses were the need to reduce backup times, keeping pace with capacity increases, and the cost of storage systems. For years, backup administrators have been struggling to complete nightly backups before business resumes in the morning. As backup windows continue to shrink and data sets grow, IT managers need a solution that can scale to meet the challenge while reducing capacity stored on disk.

ESG Lab has confirmed that TSM 6.1 leverages the DB2 database to deliver significantly greater sustained backup throughput than previous versions while offering greater scalability and advanced functionality. Administrators can potentially consolidate their TSM environment onto fewer servers—even as they roll out new functionality like data deduplication.

[5] Source: ESG Research Report, *Data Protection Market Trends*, February 2008.

Enhancing Productivity with Monitoring and Reporting

TSM 6.1 plugs directly in to ITM (IBM Tivoli Monitoring, IBM's system monitoring software that manages operating systems, databases, and servers) to provide detailed, near real-time monitoring. Historical reporting is provided in TSM Admin Center via the ISC (Integrated Solutions Console). Both of these features are provided at no cost in TSM and TSM EE and users have the ability to create custom reports using the Tivoli Common Reporting framework and open-source BiRT (Business Intelligence and Reporting Tools).

ESG Lab Testing

ESG Lab walked through an installation of ITM and the TSM Monitor Agent to get a feel for the experience of a TSM customer who had not used Tivoli Monitoring before. ESG installed the ITM software and the TSM Monitor Agent on a running TSM Server using the distribution CD as seen in Figure 11.

FIGURE 11. INSTALLING THE TSM MONITOR AGENT

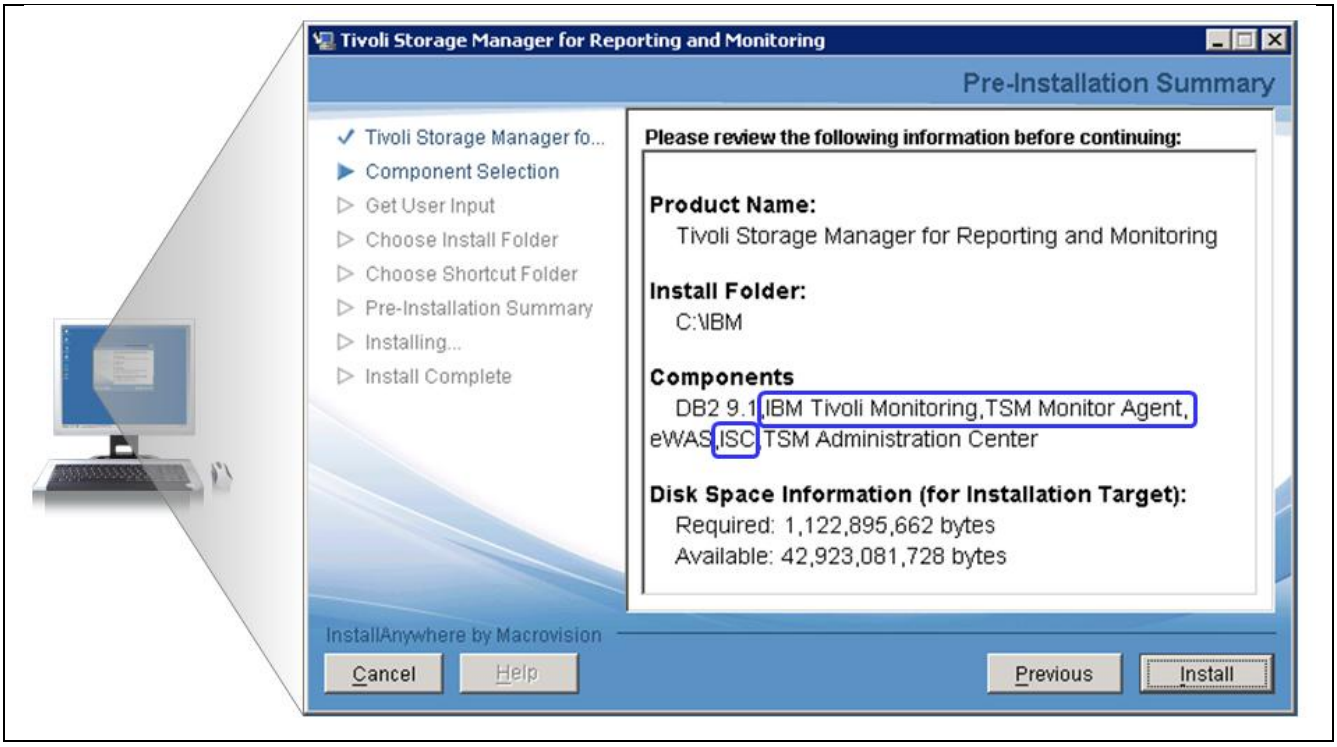

The installation completed in less than five minutes and ESG Lab was able to connect to IBM Tivoli Monitoring by pointing a web browser at the server where we installed the software. As seen in Figure 12, TSM shows up as an object in the ITM tree view.

FIGURE 12. MONITORING TSM SERVER ACTIVITY

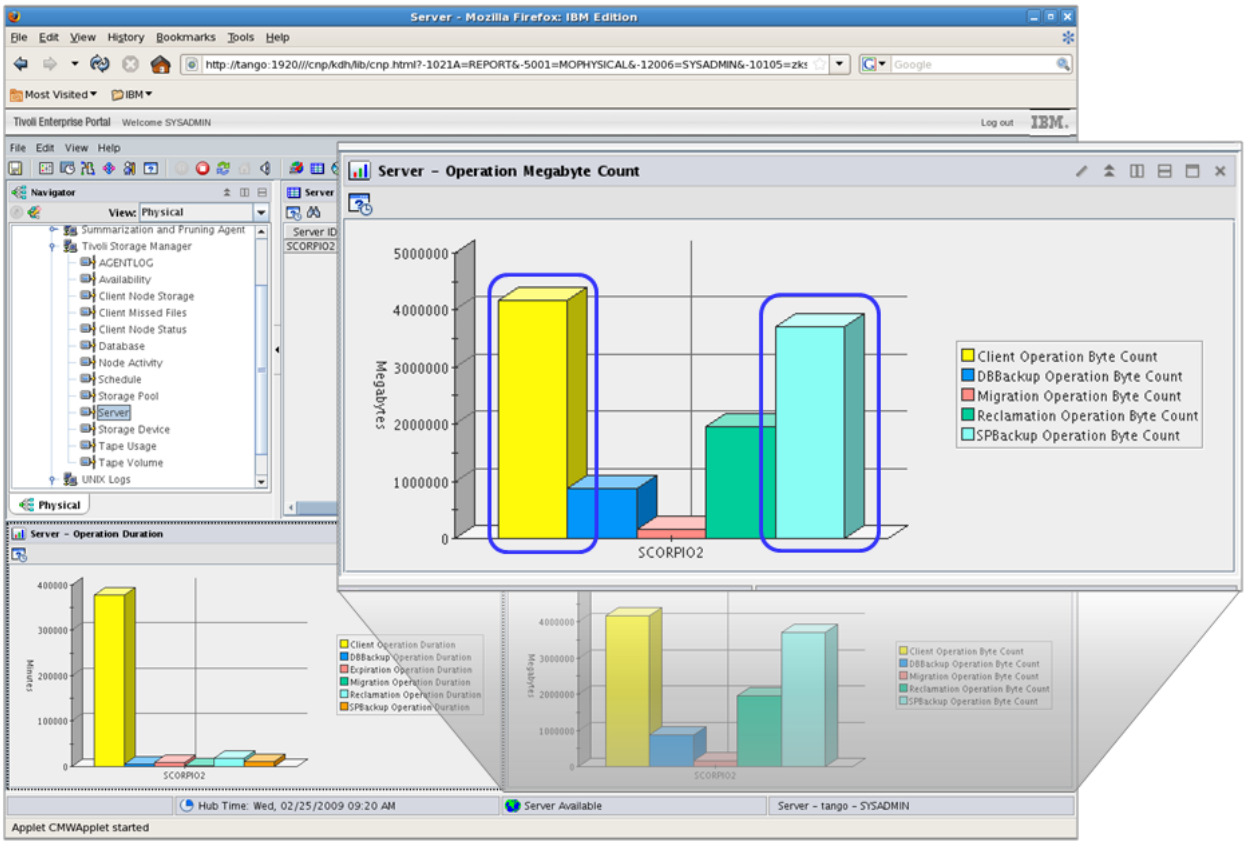

Under Tivoli Storage Manager on the left are more than a dozen categories ranging from client status all the way through storage device and tape status. All of the data in this view is for the previous 24 hours. ESG Lab clicked on 'Server' and examined the 'Operation Megabyte Count.' This view showed the total bytes processed by TSM in the past day, broken into several categories.

Next, ESG Lab clicked on 'Tape Volume' and was treated to a detailed look at the available capacity of every tape under management by TSM, shown in Figure 13.

FIGURE 13. MONITORING CAPACITY UTILIZATION

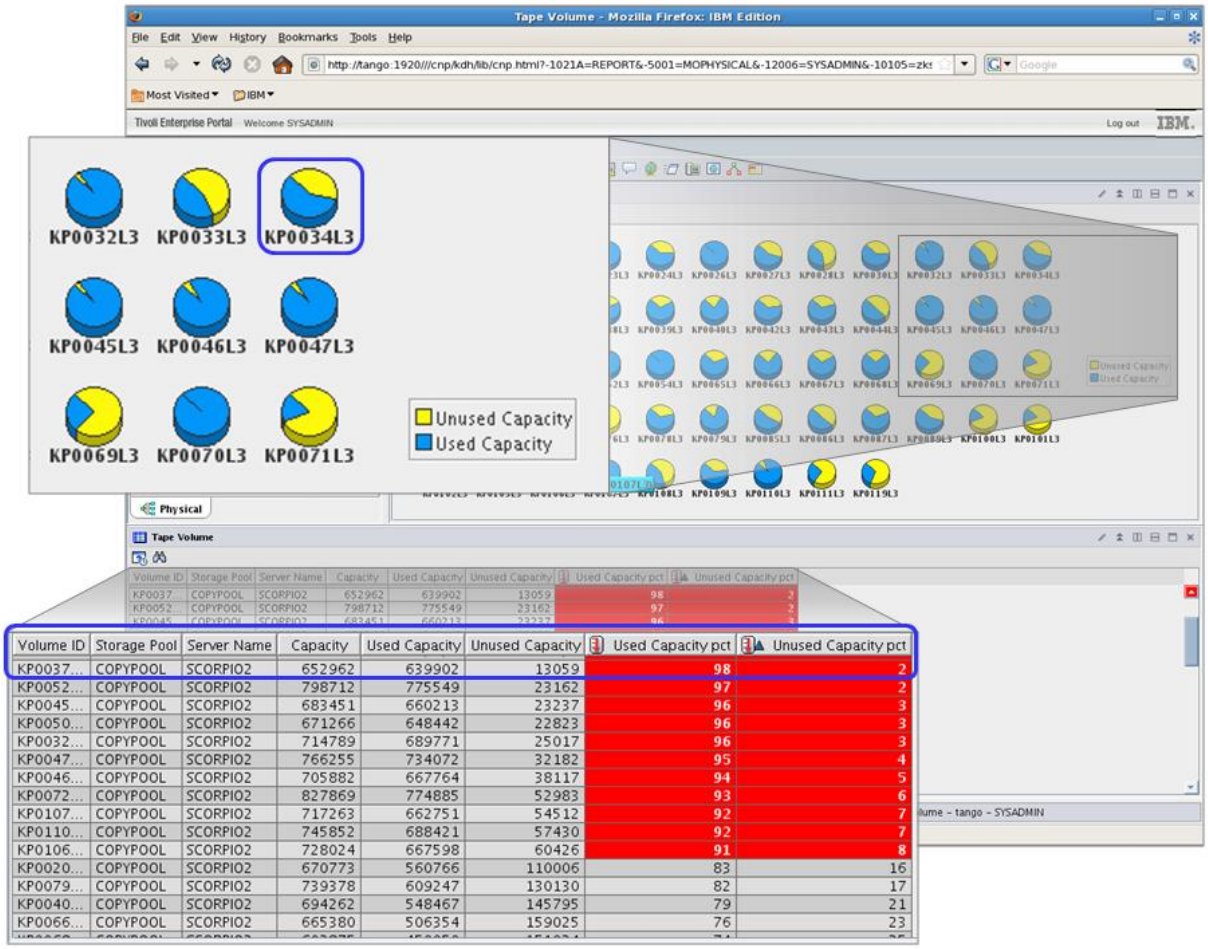

Volume ID	Storage Pool	Server Name	Capacity	Used Capacity	Unused Capacity	Used Capacity pct	Unused Capacity pct
KP0037...	COPYPOOL	SCORPIO2	652962	639902	13059	98	2
KP0052...	COPYPOOL	SCORPIO2	798712	775549	23162	97	2
KP0045...	COPYPOOL	SCORPIO2	683451	660213	23237	96	3
KP0050...	COPYPOOL	SCORPIO2	671266	648442	22823	96	3
KP0032...	COPYPOOL	SCORPIO2	714789	689771	25017	96	3
KP0047...	COPYPOOL	SCORPIO2	766255	734072	32182	95	4
KP0046...	COPYPOOL	SCORPIO2	705882	667764	38117	94	5
KP0072...	COPYPOOL	SCORPIO2	827869	774885	52983	93	6
KP0107...	COPYPOOL	SCORPIO2	717263	662751	54512	92	7
KP0110...	COPYPOOL	SCORPIO2	745852	688421	57430	92	7
KP0106...	COPYPOOL	SCORPIO2	728024	667598	60426	91	8
KP0020...	COPYPOOL	SCORPIO2	670773	560766	110006	83	16
KP0079...	COPYPOOL	SCORPIO2	739378	609247	130130	82	17
KP0040...	COPYPOOL	SCORPIO2	694262	548467	145795	79	21
KP0066...	COPYPOOL	SCORPIO2	665380	506354	159025	76	23

The view was presented both graphically and in a table, with the latter providing detailed data on each tape volume as well as a bright red highlight for volumes that were over 90% occupancy.

FIGURE 14. TSM REPORTING: MISSED FILES

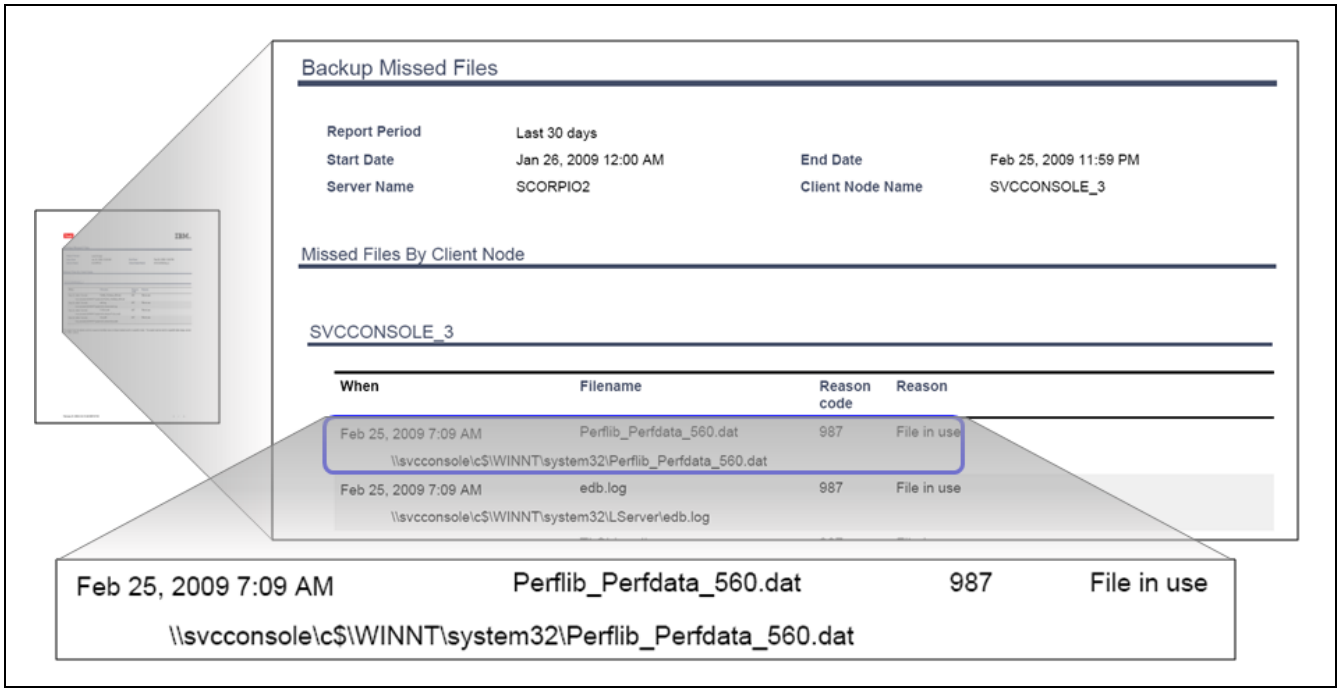

Finally, ESG Lab examined historical reporting. Users have the ability to collect, store, and manage a tremendous amount of detailed data about an organization's data protection environment. An example of this capability is the Backup Missed Files report, shown in Figure 14. Previous generations of TSM kept lists of all missed files for every backup in the main DB2 database, but extracting that data was a manual process. TSM 6.1's reporting capability provides user-friendly, detailed reporting on each missed file so administrators can easily target and correct any issues preventing backup of those files.

Why This Matters

The ability to know whether or not backup operations are actually successful has long been a vexing problem for backup administrators. In fact, in a recent ESG survey of 398 North American IT professionals responsible for their organization's data protection processes and technologies, 40% cited "difficulty validating backup/recovery success" as a major challenge in their current data protection environments.[6]

Not surprisingly, IT staffs leveraging some form of backup reporting and monitoring solution—whether included with a backup application or as part of a more comprehensive reporting/analytics or SRM tool—typically enjoy a noticeably higher rate of success for their backup operations than those with no formal tools or processes in place. In fact, ESG research indicates that users with some form of data protection reporting software had an 87% success rate for their backup jobs on average.

TSM 6.1's reporting and monitoring agent, included in TSM at no cost, plugs directly into IBM Tivoli Monitoring and can provide detailed, near real-time and historical reporting to keep administrators informed and ready to act on critical issues affecting their data protection environment, allowing them to keep backups running smoothly.

[6] Source: ESG Research Report, *Data Protection Market Trends*, February 2008.

Enhanced NAS Backup and VMware Support

TSM 6.1 incorporates two important enhancements to streamline backup and recovery in consolidated and virtualized environments. TSM leverages NetApp's Snapshot Difference and SnapMirror to Tape technologies, enabling customers to more efficiently back up IBM N series and NetApp FAS volumes. TSM 6.1 also offers enhanced integration with VMware Consolidated Backup (VCB), automating the backup and recovery processes within the GUI at the image level.

SnapMirror to Tape and Snapshot Difference API Integration

NDMP is an open standard protocol for backing up network-attached storage. The NDMP protocol was designed to enable NAS devices to directly communicate with tape and other backup devices. As seen in the top half of Figure 15, a TSM NDMP client running backup software sends control data only to the NAS device, instructing it to back up or restore. The NAS device (an IBM N Series or NetApp NAS appliance in this example) directly communicates with the backup media. This eliminates load from application servers and keeps backup traffic off the production LAN. In the traditional NDMP backup, the NAS device sends individual files and metadata to the backup media.

FIGURE 15. NDMP BACKUP AND TSM SNAPMIRROR INTEGRATION

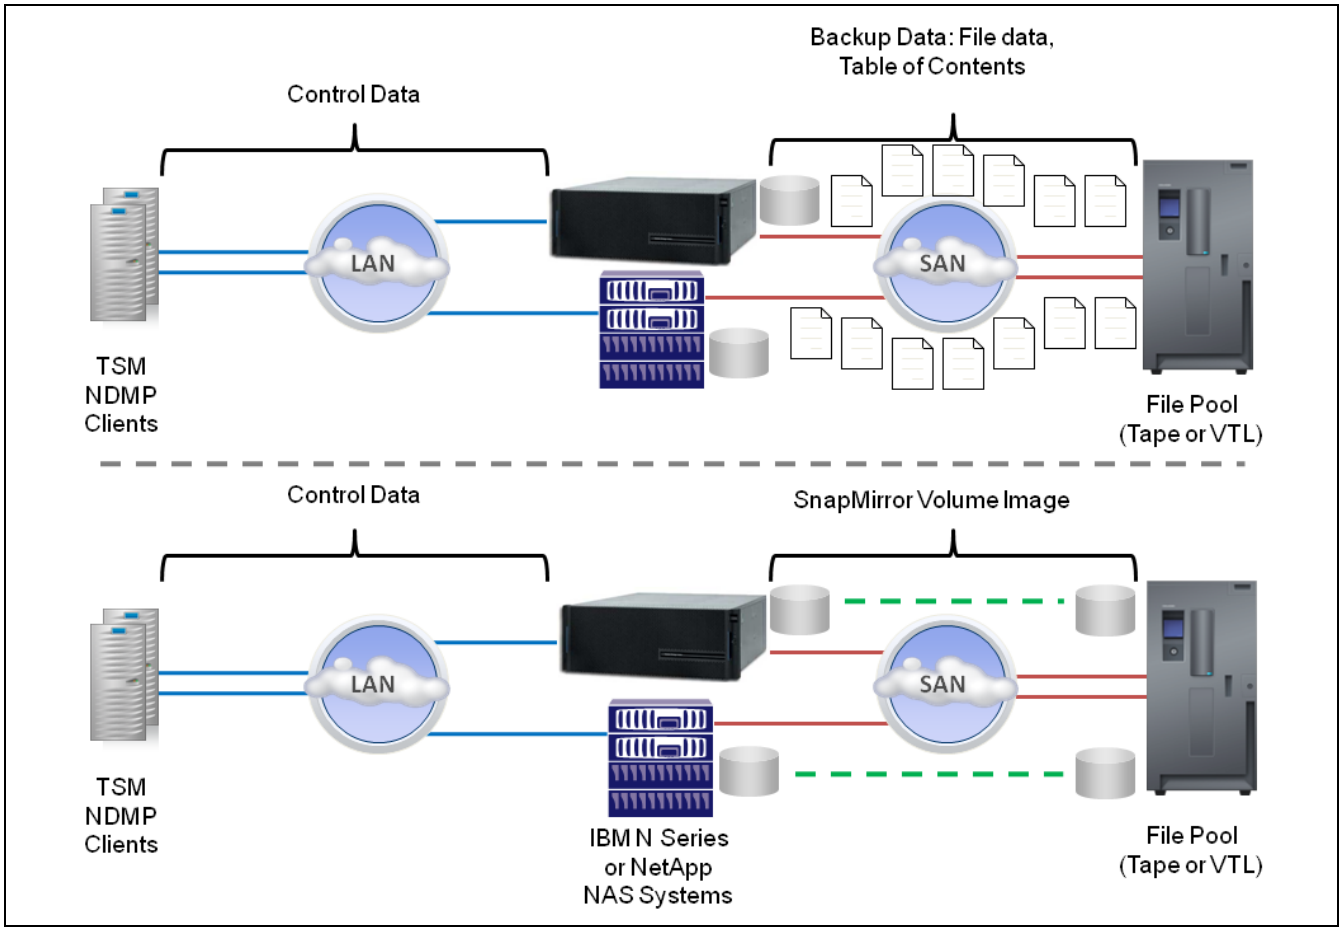

The bottom half of Figure 15 shows how TSM 6.1 takes this one step further using IBM N series and NetApp NAS appliances. The NDMP client sends control data to the NAS device and the NAS device uses SnapMirror functionality to send a full, point-in-time block-level image of the volume to the backup media. This reduces processing overhead and I/O on the NAS device while allowing TSM to manage backup and recovery.

Traditionally, scans of large file systems add hours to the incremental backup process because the entire file system (many millions of files and directories) must be traversed to check for changes (defined as new, modified, or deleted files) as shown in the top half of Figure 16.

FIGURE 16. TRADITIONAL INCREMENTAL BACKUP VS. SNAPSHOT DIFFERENCING INTEGRATION

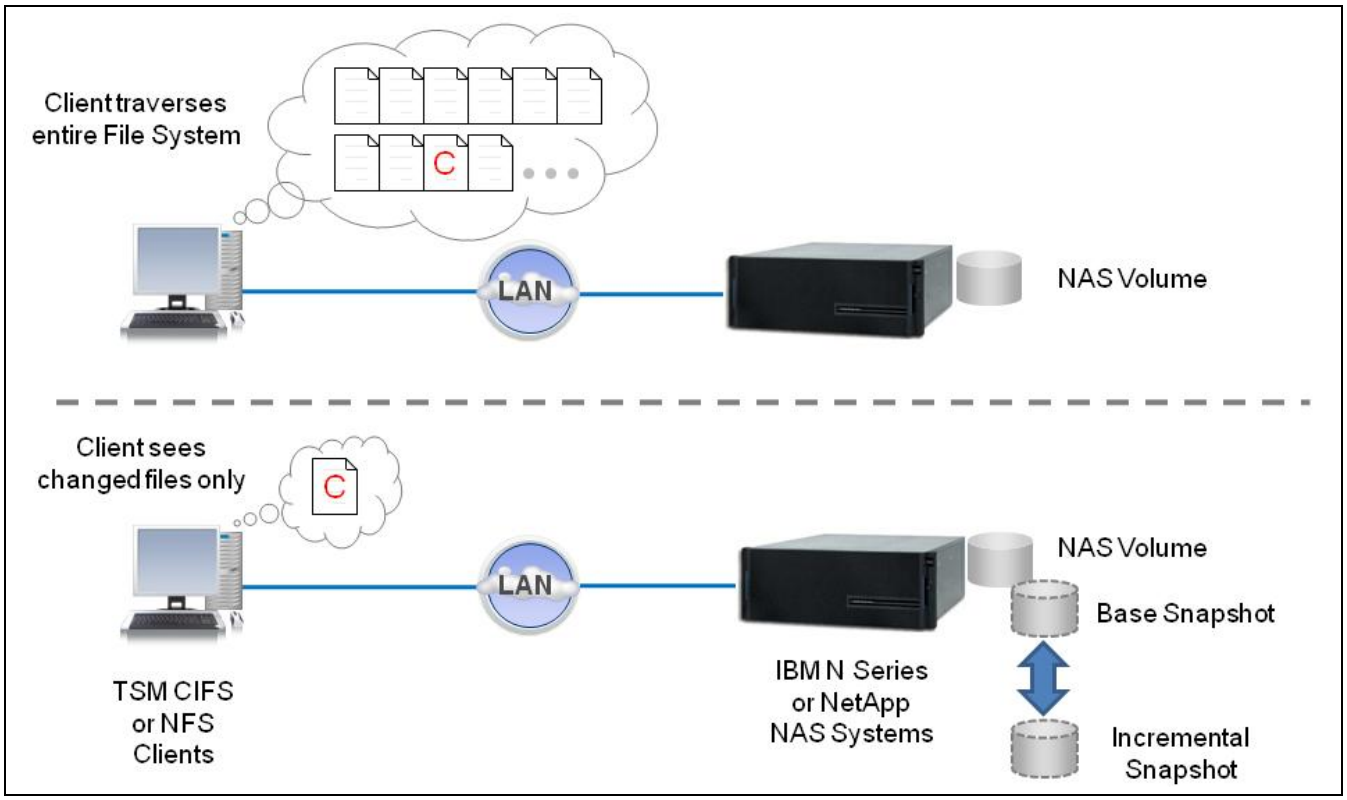

This process must complete before an incremental backup can begin. TSM 6 leverages NetApp Snapshot technology to sharply reduce the time required for progressive incremental file backup. When the TSM client performs an incremental backup using Snapshot Differencing, the IBM N Series or NetApp FAS system compares the current and previous snapshots and sends the TSM client only the changed files.

ESG Lab Testing

ESG Lab tested Snapshot Difference API integration by performing an incremental backup of a small file system shared over CIFS by a NetApp filer using both the traditional method and the Snapshot Difference API using the TSM command line. A full backup and base snapshot had been taken of this file system prior to beginning the test. First, several small files were added to a folder in the test file system. The traditional incremental was then performed by typing 'incremental z:' at the TSM command line. The backup process took several minutes while TSM traversed the file system, found the new files, and backed them up. To test Snapshot Differencing, ESG lab typed 'incremental z: -snapdiff' at the TSM command line and the filer instantly returned the list of changes and the backup completed in seconds, with no traverse of the file system. Overall, the incremental backup process was reduced from six minutes to six seconds on a small file system with several thousand files. Large production file systems with millions of files should see similarly dramatic results.

ESG Lab also performed a SnapMirror to tape backup of a file system exported by an IBM N series NAS appliance using the TSM NDMP client. The backup command was exactly the same as with a traditional NDMP backup, with the addition of a new command line option to enable SnapMirror to back up to tape. ESG Lab typed the following command: 'backup node nas1 type=snapmirror.' At this point, the NAS appliance took a point-in-time snapshot of the file system and began copying the image directly to tape. Once the backup completed, ESG

Lab performed a restore to a different volume and validated that the restore operation was complete by comparing the files in the restored volume to the originals in the backed up volume.

VMware Consolidated Backup Integration

TSM 6.1 can be used to protect VMware virtual guest systems via integration with VMware Consolidated Backup (VCB). TSM 6.1 works in concert with the VCB proxy server and VMware Converter to provide backup and recovery services for virtual machines while eliminating the complexity and overhead of managing agents running in virtual machines. As shown in Figure 17, administrators simply specify the VMs to back up on each Virtual Center server and the staging area where VM images will be stored on disk.

FIGURE 17. TSM VMWARE CONSOLIDATED BACKUP INTEGRATION

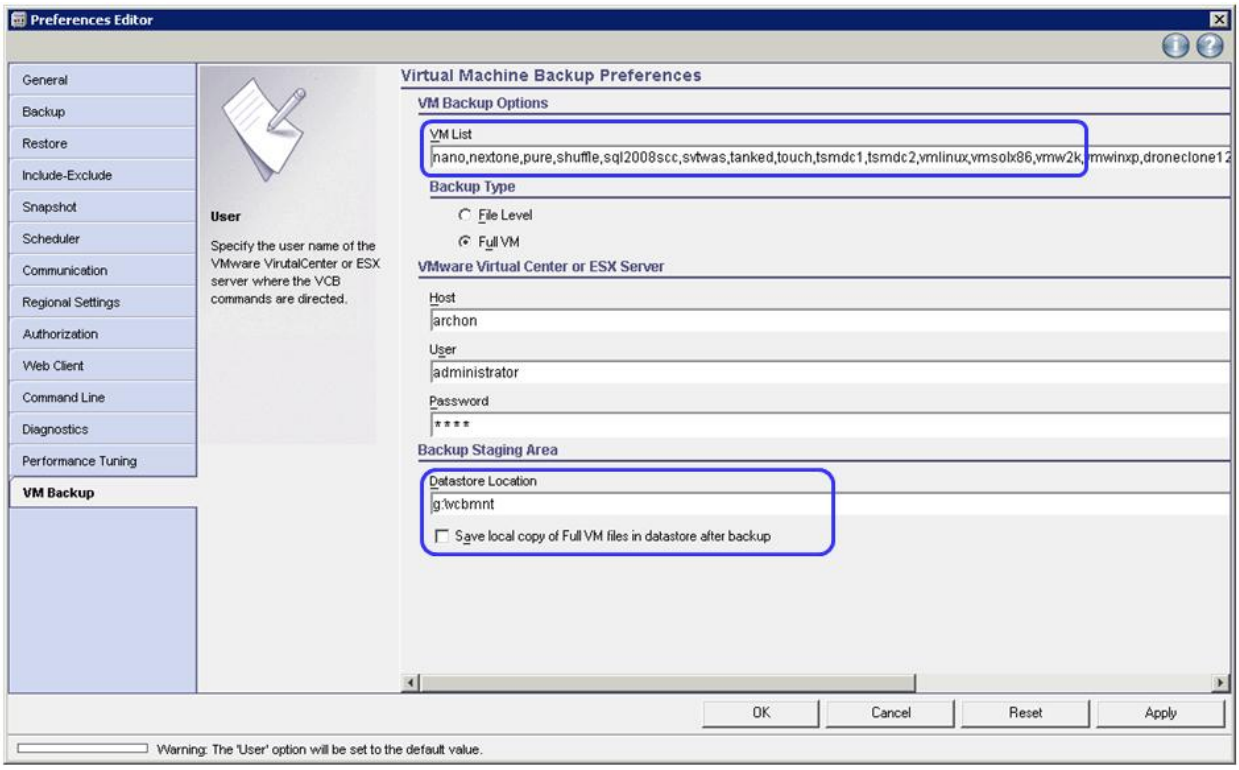

By default, the image is deleted from the staging area once the backup has been completed and the image is safely stored in a TSM pool. Administrators have the option to save the local copy of the VM in the staging area. This can speed recovery from recent backups. While copies can be deleted manually when no longer needed, most users prefer to let TSM automatically delete the files after backups complete.

ESG Lab Testing

ESG Lab performed a full restore of a virtual server on a VMware ESX server using TSM 6.1. As shown in Figure 18, ESG Lab selected the virtual server 'vmw2k' for restore.

FIGURE 18. RESTORING A VIRTUAL MACHINE

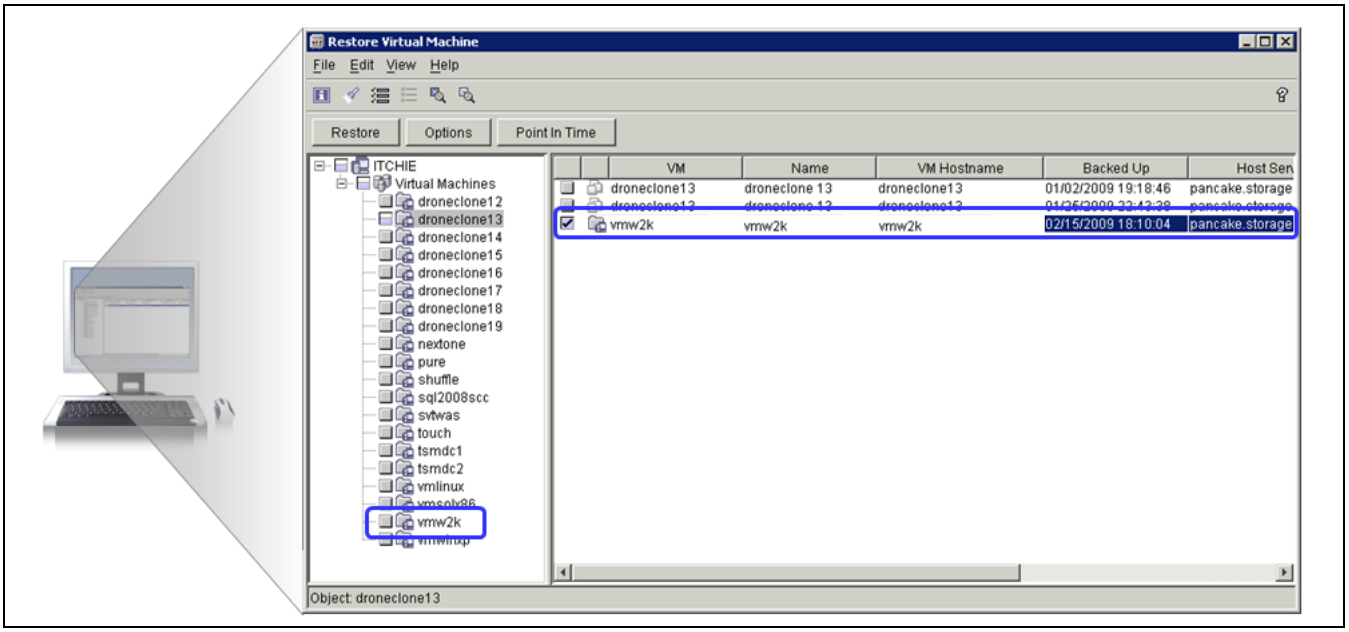

FIGURE 19. THE VMWARE CONVERSION WIZARD

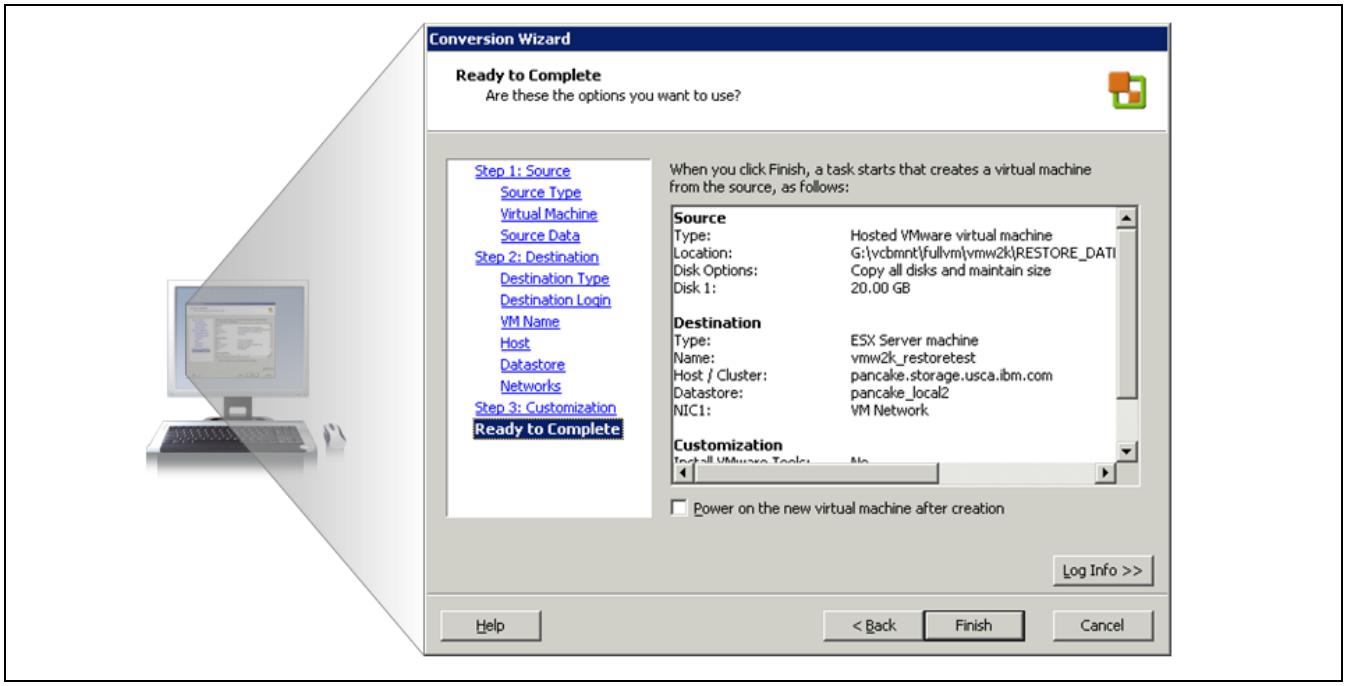

When ESG Lab clicked 'Restore,' the VM image was restored to the staging area and the VMware Conversion wizard was launched, as seen in Figure 19, to convert the VM image file into a live virtual machine. After specifying the source file and the destination ESX server, ESG Lab clicked 'Finish' and the virtual machine was restored.

Figure 20 shows the virtual machine powered up normally. ESG Lab was able to log on to the console of the VM and access all resources that belonged to this machine.

FIGURE 20. THE RESTORED VIRTUAL MACHINE

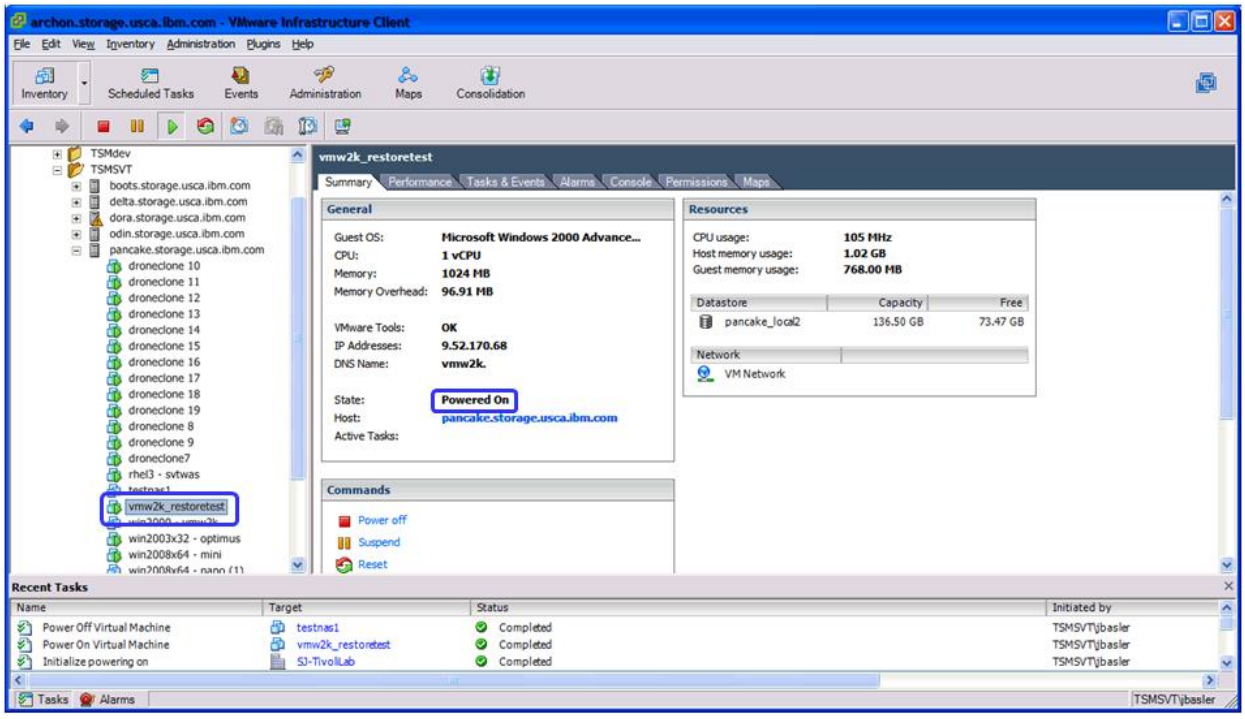

The entire restore took about ten minutes from start to finish, with TSM handling the entire process.

Why This Matters

Storage consolidation and server virtualization concentrate processing power and data storage in the data center on fewer servers and systems; protecting these assets presents a unique set of operational challenges thanks to growing volumes of data residing on NAS devices and on applications within virtual machines. In fact, ESG research found that more than one-third (37%) of current users report that the total amount of data they need to back up has increased with the deployment of server virtualization.[7]

ESG Lab has validated that TSM 6.1 provides quick and easy data protection for NAS volumes, operating systems, and application data residing in virtual machines. TSM integration with NetApp's Snapshot Difference and SnapMirror to tape APIs use familiar methodologies for administrators while providing powerful integration with volume snapshot and mirroring technology inside IBM N Series and NetApp NAS appliances. ESG Lab used TSM's VCB integration to perform a full recovery of a virtual machine directly from the TSM GUI. These features provide more efficient backups and faster time to recovery as they eliminate cost and overhead from virtual machines and NAS clients.

[7] Source: ESG Research Report, *The Impact of Server Virtualization on Storage*, December 2007.

ESG Lab Validation Highlights

- ☑ TSM 6.1 progressive incremental backup, combined with new data deduplication technology, provided an impressive 95% data reduction factor (more than 19:1) over just eleven days.
- ☑ DB2 provides an enterprise-class, scalable back-end for TSM 6.1. ESG Lab confirmed that TSM can scale to much larger capacities under management while providing higher performance and more functionality.
- ☑ TSM 6.1 offers excellent monitoring and reporting. ESG Lab examined near-real time statistics for servers, clients, and backup media, as well as historic reporting for an active TSM environment.
- ☑ TSM has enhanced NAS backup/recovery using integration with NetApp's Snapshot Difference and SnapMirror to tape APIs. ESG Lab was able to reduce the time required for an incremental backup of a file system by 98% using Snapshot Difference technology. A full SnapMirror-based backup of an IBM N series appliance was executed using a single, standard NDMP command.
- ☑ TSM 6.1 offers improved VMware integration. ESG Lab used TSM 6.1 to restore a full VMware virtual machine directly from the TSM interface, with no scripting or manual commands needed.

Issues to Consider

- ☑ Monitoring and reporting built into TSM 6.1 is limited to pre-defined reports. IBM has informed ESG that customers can use the Tivoli Common Reporting framework and BiRT (Business Intelligence and Reporting Tools) against the DB2 database to create custom reports, but ESG Lab believes it would be valuable to build custom reporting functionality directly into TSM.
- ☑ VCB Backup has the option to save a local copy of full VM backup files after they have been captured by TSM, but this feature requires a manual cleanup to remove these files. ESG Lab spoke with IBM and policy settings to perform this cleanup are planned for a future release of TSM.

ESG Lab's View

For over 15 years, TSM has been uniquely positioned for open systems data management—protecting, retaining, and ensuring accessibility to data. TSM's architecture and design principles were rooted in mainframe storage management concepts, distinguishing it from comparable solutions with features such as integrated archiving, disk-to-disk backup, and progressive incremental backup policies. These and other features are proof points of the level of optimization in the TSM platform. With TSM 6, IBM has responded further to the demands of organizations grappling with relentless growth in digital information through features that introduce even greater efficiencies.

ESG Lab performed hands-on testing of key TSM 6.1 features to demonstrate and audit capabilities. ESG Lab tested data deduplication to understand the value of this capacity optimization technique in combination with progressive incremental backups over time. ESG Lab found that the two technologies working in concert were able to achieve 90% data reduction after just six incremental backups and 95% reduction after ten backups.

With the inclusion of DB2 as the underlying TSM relational database, ESG sought to verify several related benefits, including the impact on scalability and performance, installation, upgrades, and database migration processes, as well as reporting and monitoring. ESG Lab verified active TSM 6.1 installations with databases in excess of 700 GB, managing hundreds of millions of objects. This is significantly larger than anything possible with previous versions of TSM and does not represent an upper limit. The installation/upgrade process was smooth and completely automated. The reporting options were extensive and detailed; administrators who may have relied on third party tools now have a robust built-in option.

ESG Lab also validated TSM's new NDMP SnapMirror to tape, NetApp Snapshot Difference API support and VMware VCB integration. All three delivered the optimization benefits desired by end-users. Block-level backup of a NAS volume was completed using SnapMirror to tape quickly and effortlessly using familiar tools and techniques. The time required for incremental backup of a filesystem was reduced by 98% using the Snapshot Difference API and TSM's VMware integration enabled simple and efficient protection and recovery of virtual machines without necessitating complex scripting or manual processes.

ESG Lab believes that the TSM platform achieves optimization in data management in multiple ways. The solution creates capacity efficiencies through integrated archiving, disk-based data protection, an "incremental forever" strategy, and now, data deduplication. Similarly, IBM has been and continues to be focused on backup and recovery performance with TSM. Optimizations, including integration of the DB2 database, can accelerate backup and recovery processes significantly while integrated image-based functionality simplifies and speeds backup and recovery in NAS and VMware environments. IBM has clearly put a lot of time and effort into TSM 6.1 and that hard work has paid off: IBM has reinforced TSM's position as a preeminent enterprise data protection platform with lots of enhancements for its large existing customer base—and plenty to offer potential new users.

Appendix

TABLE 2. TEST CONFIGURATION

Deduplication	
IBM TSM Server IBM p550, 16x 2.1GHz POWER5+ cores, 32 GB RAM 4x 4Gb/sec FC HBAs -QLA2340	IBM TSM 6.1 AIX 5.3
IBM XIV – 40 TB IBM DS4000 – 10 TB IBM TS3500 Tape Library – LTO3	
Installation and upgrade	
IBM TSM Server IBM x336, 2x 3.0GHz XEON CPU, 4 GB RAM	IBM TSM 6.1 Windows 2003 SP1
Windows, NAS and VMware	
IBM TSM Server IBM x3650, 8x 3.0GHz XEON CPU, 24 GB RAM 2x HBA -QLA2340	IBM TSM 6.1 Windows 2003 R2
IBM N3600 NAS Appliance IBM DS4800, DS6000 IBM TS3584 Tape Library – LTO4 IBM TS3592 VTL	

20 Asylum Street
Milford, MA 01757
Tel: 508-482-0188
Fax: 508-482-0218

www.enterprisestrategygroup.com