Enterprise Strategy Group | Getting to the bigger truth.™

ESG Lab Review

IBM Data Protection for Virtual Environments:

Tivoli Storage Manager for VMware Version 7.1 and FlashCopy Manager for VMware Version 4.1

Date: August 2014 **Author:** Vinny Choinski**,** Senior Lab Analyst, and Aviv Kaufmann, Lab Analyst

Abstract: *This ESG Lab Review documents remote hands-on testing of the IBM Tivoli Storage Manager data protection solution for virtual environments with a focus on ease of use, flexibility, and architecture efficiency.*

The Challenges

In spite of its many associated benefits, server virtualization can—and does—present numerous IT challenges. As an example, efficiently protecting and recovering business assets in highly virtualized environments often proves to be extremely difficult. Even organizations that are highly successful in recovering data in their current backup and restore environments constantly look for ways to improve RTO/RPO and overall data protection efficiency. In fact, ESG research indicates that only 8% of survey respondents rely solely on their backup applications as compared with a combined 73% who leverage snapshot technology at the datastore and virtual-machine level (see Figure 1).[1]

Figure 1. Mechanisms for Protecting VMware Virtual Machines

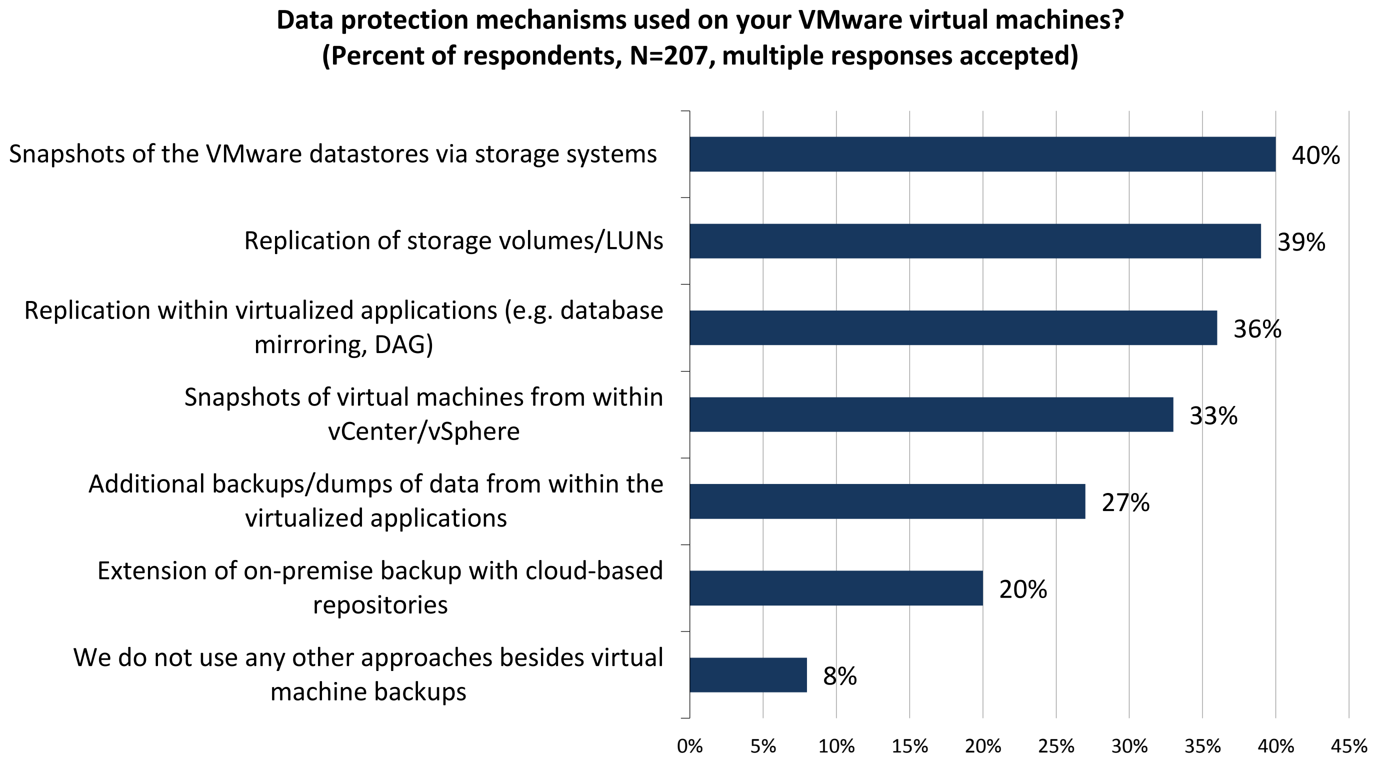

Source: Enterprise Strategy Group, 2014.

[1] Source: ESG Research Report, *Trends for Protecting Highly Virtualized and Private Cloud Environments*, June 2013.

The goal of ESG Lab reports is to educate IT professionals about data center technology products for companies of all types and sizes. 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.

The Solution: Tivoli Storage Manager – Data Protection for Virtual Environments

IBM Tivoli Storage Manager for Virtual Environments (TSM for VE) version 7.1 and Tivoli Storage FlashCopy Manager version 4.1 are software solutions designed to protect VMware environments by transferring block-level data from snapshots between the virtual environment and the Tivoli Storage Manager Server. TSM data protection for virtual environments relies on VMware snapshots and storage array-aware snapshots to facilitate data protection. It can be quickly and easily installed on virtual machines or physical hosts. The intuitive and unified user interface for virtual environment protection was designed with VMware administrators and backup administrators specifically in mind. It connects the enterprise-class, industry-recognized TSM server with the virtual environment, providing simple, easy-to-use wizards for configuration tasks. TSM for VE 7.1 enables TSM vStorage API for Data Protection (VADP) backups and disk hardware snapshots with FlashCopy Manager (FCM) for VMware 4.1 can be managed using the same interface. New features include:

- A single, easy-to-use interface for both TSM backups and FCM snapshot management.
- A management interface that can be launched from VMware vCenter or via any browser.
- A simplified, wizard-driven installation with all requisite components included on a single install media.
- The ability to deploy on both virtual machines and physical hosts.
- Instant Access/Restore enabling the TSM server to act as a virtual datastore that provides the ability to instantly power on and use a VM from the backup data.
- Object-level recovery from image-level backups (MS Exchange, MS SQL Server, and files).
- The ability to leverage all TSM server enterprise-class features including Operations Center enhancements.
- Small, Medium, and Large TSM server reference blueprints for improved ease of deployment.
- Enhanced TSM server deduplicated data ingest performance.

Figure 2 shows the TSM 7.1 for VE solution overview. On the left side of the figure is the VMware environment to be protected including FCM-managed snapshots. In the middle is the vStorage backup server and management interface. The right side of the figure shows the target TSM server and the new Operations Center interface.

Figure 2. Solution

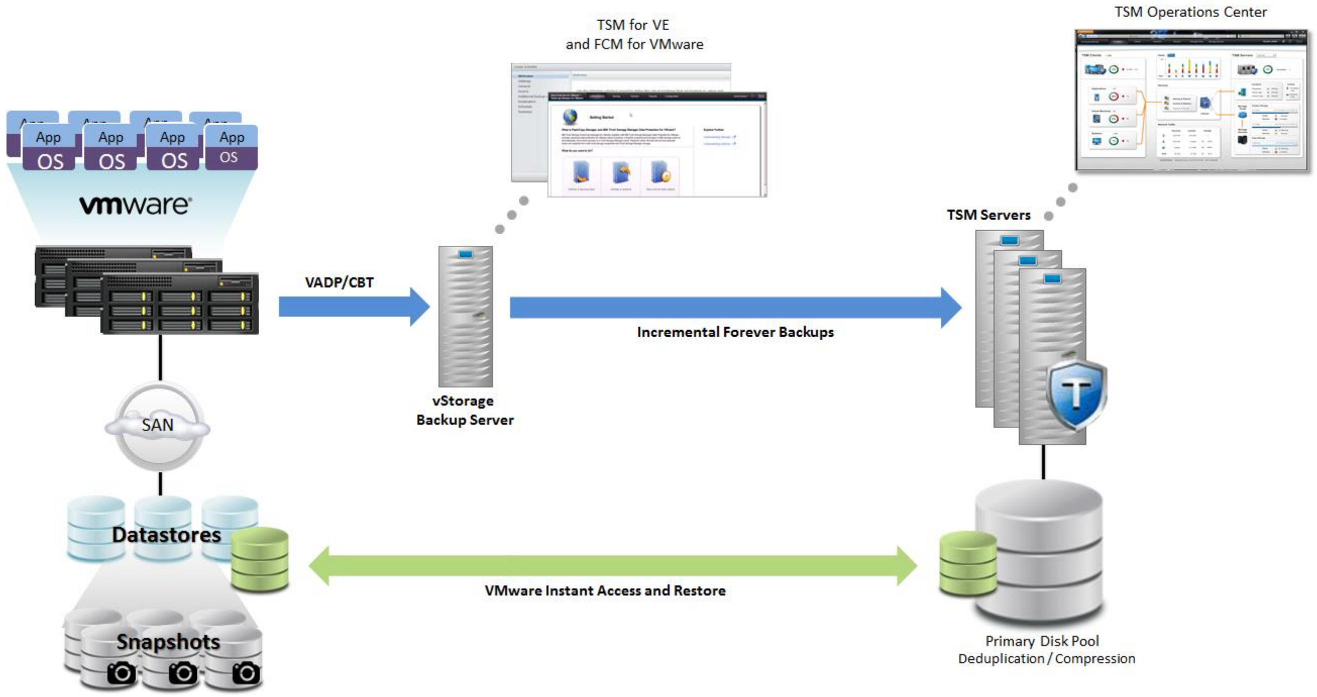

ESG Lab Tested

ESG Lab performed hands-on testing of IBM TSM for VE 7.1 with a focus on validating the ease of use and enhanced data protection flexibility offered to VMware administrators. Before testing, ESG Lab reviewed the test harness that consisted of an IBM TSM server, a TSM vStorage Backup Server, an IBM N-Series storage array, and several servers running a clustered ESXi environment managed by a vCenter instance running as a VM. Using this configuration, we next reviewed three high-level data protection options that were achievable using IBM TSM for VE. All three options could be initiated through the vCenter interface and leveraged VMware's VADP backup technology.

The first option, using TSM for Virtual Environments, enables individual virtual machines to be backed up to the TSM server using an initial full copy followed by incremental updates handled through VADP's change block tracking (CBT) feature. The second and third options leverage the integration of TSM for Virtual Environments with FCM for VMware to combine the performance benefits of volume-level hardware snapshots with the protection of an offloaded backup to the TSM server. FCM can be used to take snapshots of datastores that are stored locally on the storage array and simply catalogued by TSM for VE, or that are both stored on the local storage array and backed up to and catalogued in the TSM server. Using these FCM snapshots, an administrator can restore entire datastores with all of the VMs included, or simply restore one of the VMs stored within the datastore. The three high-level data protection choices for virtual environments are shown in Figure 3.

Figure 3. Overview

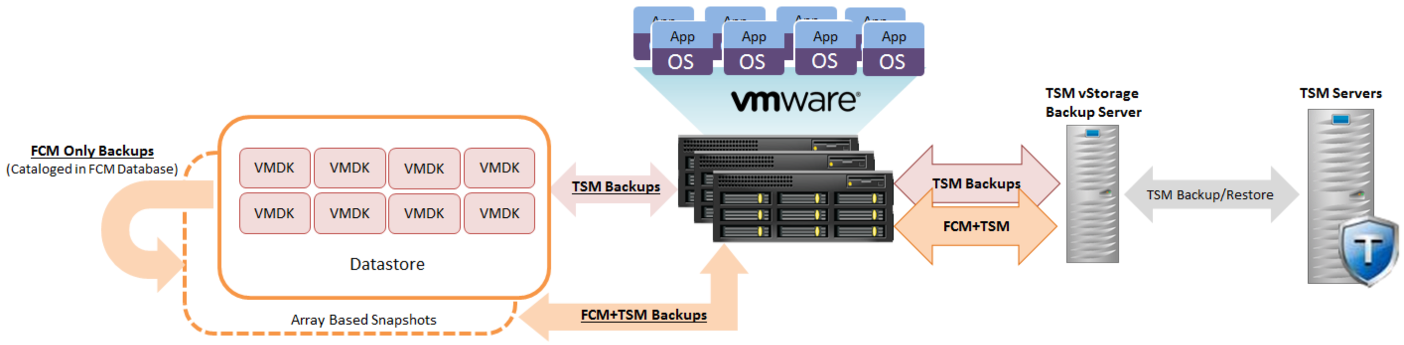

Flexible VMware Data Protection

To begin testing, ESG Lab first installed IBM TSM for VE 7.1 on the windows-based TSM vStorage Backup Server using a single install media, guided by a simplified configuration wizard. The installation process configured the server to be used as a vStorage Backup Server with a Data Mover for the virtual environment, as a vSphere plug-in server, and installed the recovery agent. After the quick and simple installation, we were able to use a browser interface to launch the TSM for VE GUI. Using the vSphere web client, we were able to easily access the same GUI through the vCenter plug-in. From the GUI, ESG Lab was able to quickly and easily define a backup task, initiate a restore, and monitor the status of the entire virtual data protection environment.

We were able to initiate the backup of a single VM by creating a backup task and simply following the wizard-driven prompts. Because this was the first time the VM was backed up to TSM, a full copy was sent to the TSM server, and a backup schedule was configured to carry forward the future incremental updates to the backup image. Figure 4 shows the TSM for VE management interface.

Figure 4. Management Interface

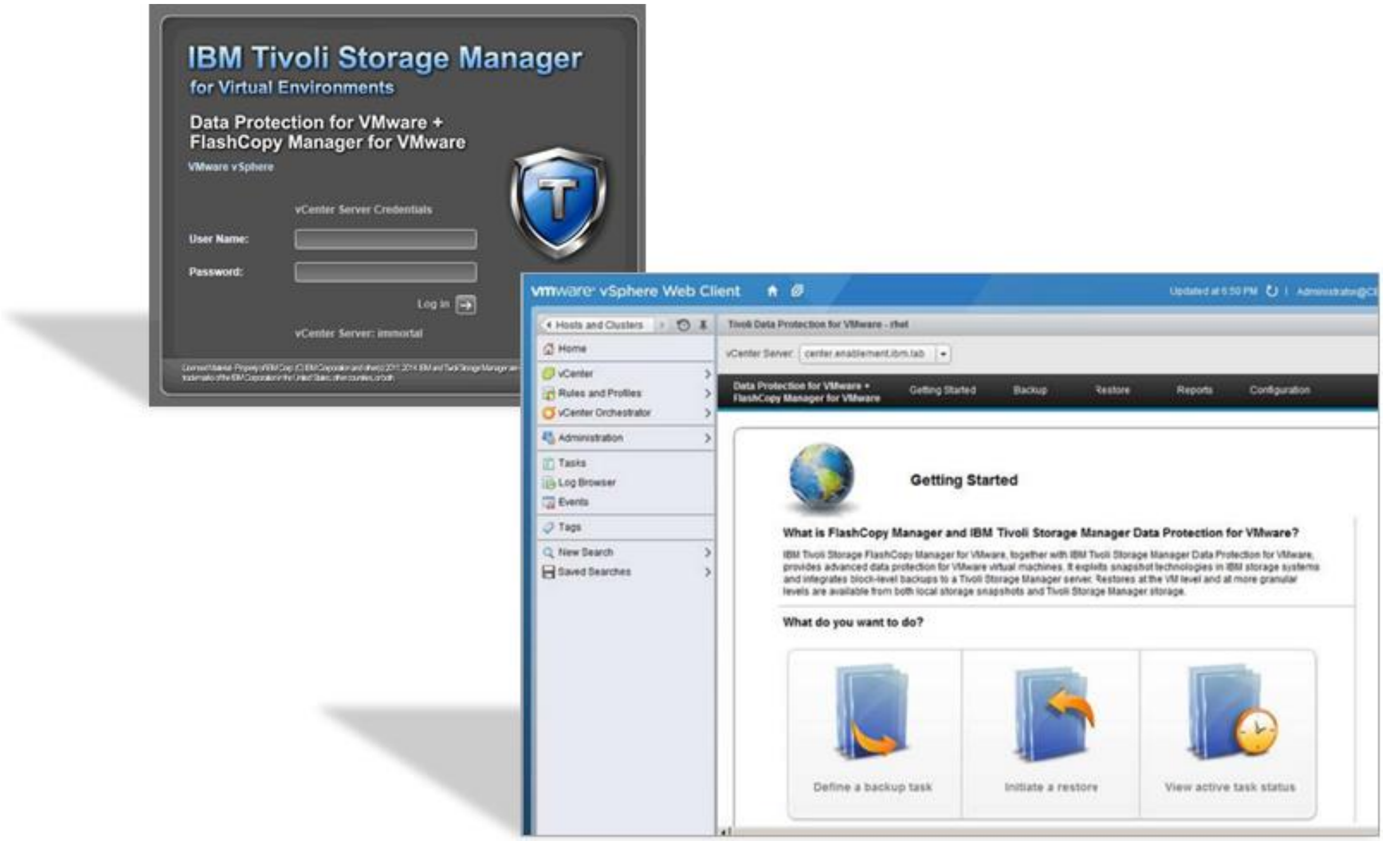

Getting Started

What is FlashCopy Manager and IBM Tivoli Storage Manager Data Protection for VMware?

IBM Tivoli Storage FlashCopy Manager for VMware, together with IBM Tivoli Storage Manager Data Protection for VMware, provides advanced data protection for VMware virtual machines. It exploits snapshot technologies in IBM storage systems and integrates block-level backups to a Tivoli Storage Manager server. Restores at the VM level and at more granular levels are available from both local storage snapshots and Tivoli Storage Manager storage.

What do you want to do?

Next, ESG Lab deleted the VM and used the restore wizard to recover the deleted VM. ESG Lab was able to accomplish this task using several wizard-guided restore options. **Full VM Instant Access** allowed us to power on the virtual machine directly from the TSM server backup image for our immediate use. This allowed us to quickly read and copy files from the VM without permanently applying the restore to the environment. **Instant Access and Restore** made the VM instantly accessible in much the same way, but continued to copy the data from the TSM server to the datastore in the background. Once all the data was copied, the VM was permanently recovered to the primary datastore using Storage vMotion, and any temporary writes were rolled forward into the VMDK. Finally, a **Full VM Restore** copied all of the data from the TSM backup to the datastore before allowing the machine to be powered on. Figure 5 highlights the functionality of performing an Instant VM Restore/Access.

Figure 5. Instant Virtual Machine Restore/Access

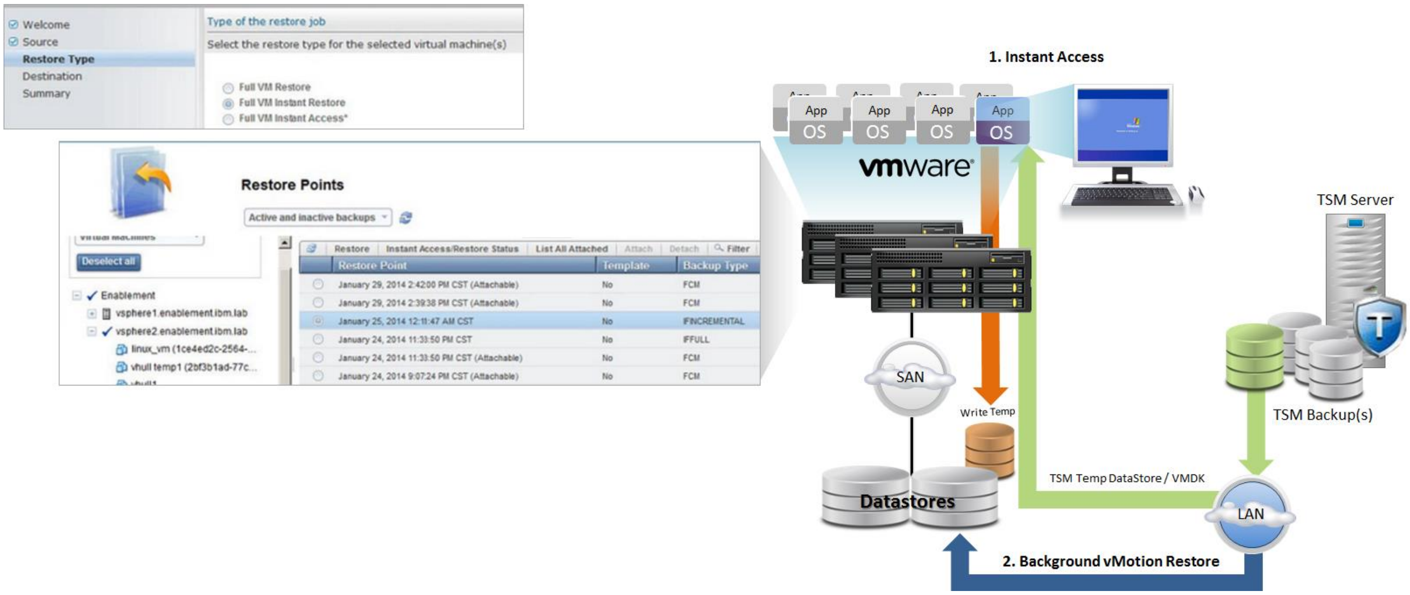

- ☑ Welcome
- ☑ Source
- **Restore Type**
- Destination
- Summary

Type of the restore job

Select the restore type for the selected virtual machine(s)

○ Full VM Restore
● Full VM Instant Restore
○ Full VM Instant Access*

Restore Points

Active and inactive backups

Deselect all

- ✓ Enablement
  - vsphere1.enablement.ibm.lab
  - ✓ vsphere2.enablement.ibm.lab
    - linux_vm (1ce4ed2c-2564-...
    - vhull temp1 (2bf3b1ad-77c...

	Restore Point	Template	Backup Type
○	January 29, 2014 2:42:00 PM CST (Attachable)	No	FCM
○	January 29, 2014 2:39:38 PM CST (Attachable)	No	FCM
●	January 25, 2014 12:11:47 AM CST	No	IFINCREMENTAL
○	January 24, 2014 11:33:50 PM CST	No	IFFULL
○	January 24, 2014 11:33:50 PM CST (Attachable)	No	FCM
○	January 24, 2014 9:07:24 PM CST (Attachable)	No	FCM

ESG Lab then validated the ability of TSM for VE to leverage storage array-based FCM snapshots to protect a virtual environment. Using the unified interface for TSM protection of VMware, an FCM backup of a datastore containing 11 active VMs was easily created by driving an array-based snapshot on the N-series storage array hosting the datastore. Next, the 11 VMs were deleted from the datastore. Using the TSM for VE recovery wizard, we were able to select the FCM datastore recovery point and perform an instant restore on the entire datastore, including automatically registering and starting the 11 VMs. It should be noted that the contents of the entire recovered datastore volume are replaced by the volume snapshot backup contents. Figure 6 highlights the key functionality of the Instant Datastore Recovery.

Figure 6. Instant Datastore Recovery

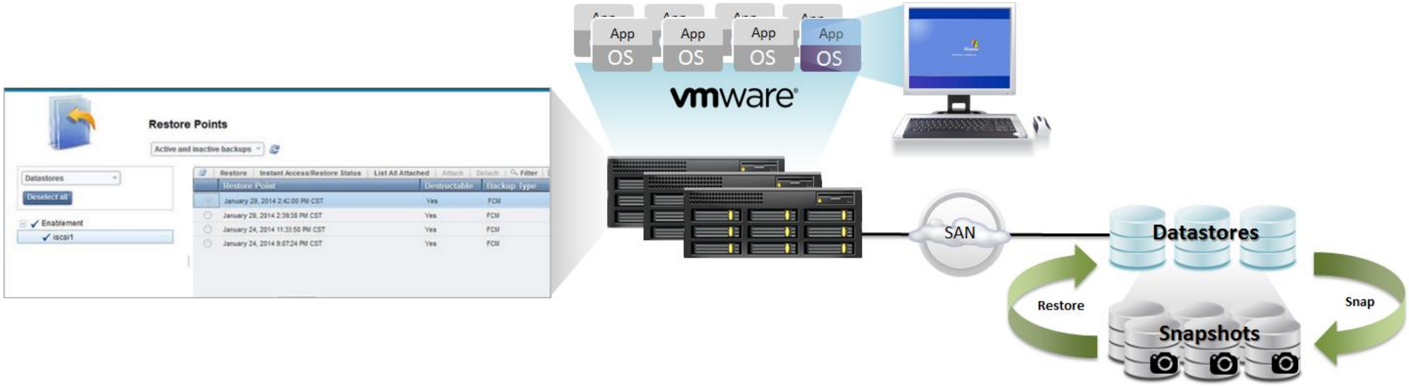

Restore Points

Active and inactive backups

Datastores

Deselect all

- ✓ Enablement
  - ✓ iscsi1

	Restore Point	Destructable	Backup Type
●	January 29, 2014 2:42:00 PM CST	Yes	FCM
○	January 29, 2014 2:39:38 PM CST	Yes	FCM
○	January 24, 2014 11:33:50 PM CST	Yes	FCM
○	January 24, 2014 9:07:24 PM CST	Yes	FCM

Finally, ESG Lab used the TSM recovery agent to quickly recover a deleted file directly from a VMDK TSM backup image. Using an active Windows VM that had been recently backed up to a TSM server, a Word document was permanently deleted from the C: drive. The recovery agent was started, and we selected the Virtual Machine backup image from the TSM server, and specified the disk to be mounted on our local VM as drive X:\. ESG Lab was then able to explore the mounted drive, select the file that we had previously deleted, and copy it to our local drive on the active VM. This is a quick and easy way to recover single files rather than restoring an entire VM from a backup image. Granular File-level Recovery functions are shown in Figure 7.

Figure 7. Granular File-level Recovery

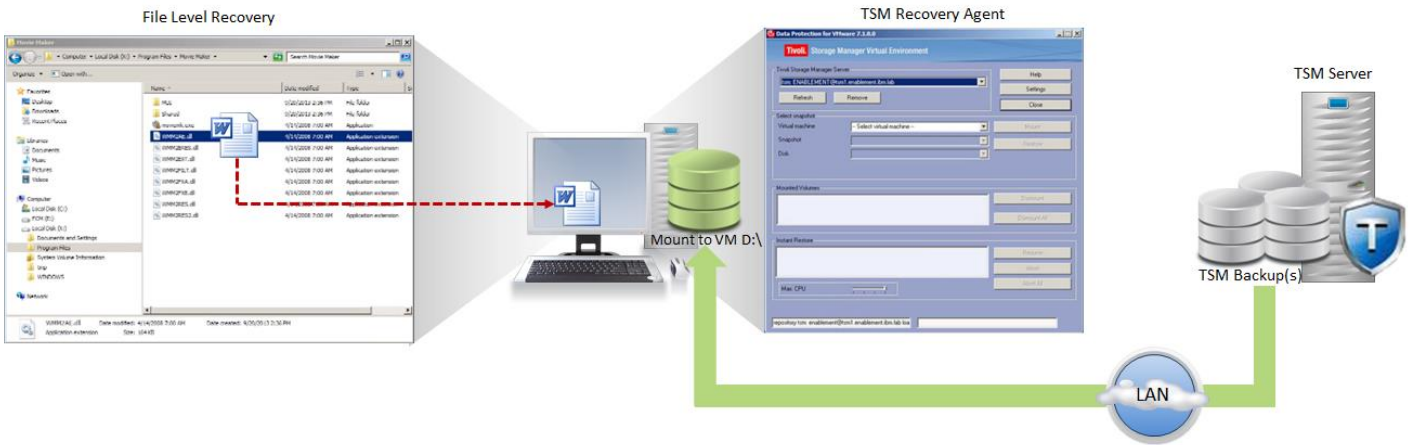

Why This Matters

Tackling the challenges associated with data protection in highly virtualized environments is difficult. Organizations are looking for easy-to-use tools delivered in a format familiar to their IT staff. They are looking for options that improve recovery time, allow quick validation of backups, and enable IT to restore only the data that matters to the business.

ESG Lab confirmed that TSM for VE 7.1 was easy to install and use. It can be launched by a VMware administrator or backup administrator right from vCenter or from any browser. It leverages VMware and volume-level hardware snapshots using FCM for quick access to recovery data; in addition, the TSM server can act as a virtual datastore when instant VM access or restore is required. TSM for VE 7.1 also supports granular object-/file-level restores from image-level backups.

Efficient, Flexible Architecture with Improved Ease of Management

ESG Lab validated the improved efficiencies built into the IBM TSM 7.1 release that benefit data protection for VMware. To begin, we took a look at the improvements to the TSM Operations Center GUI designed to help simplify the management and monitoring of an entire physical and virtual TSM environment. While TSM administrators will always be able to leverage their existing scripts to perform tasks through the command line interface, the new GUI does an excellent job of providing a single, consolidated view of the entire data protection environment including clients, services, network, backup servers, and storage.

Using the TSM Operations Center 7.1 GUI, ESG Lab was able to easily add and remove clients and select backup policies to perform backups of these clients to a TSM Server. We were able to easily identify and remediate issues with the configuration and drill down on the entire virtual infrastructure to see which machines were being protected, view and validate activity logs, and monitor the performance of scheduled backup jobs. The GUI was extremely intuitive to use and provided system configuration guidance and simple, integrated mouse-over descriptors. Figure 8 shows the TSM Operations Center GUI, including the data protection configuration overview, on the left and the TSM Server dashboard overview on the right. For a better understanding of the capabilities of TSM OC 7.1, readers are encouraged to try the online demo.[2]

Figure 8. Operations Center

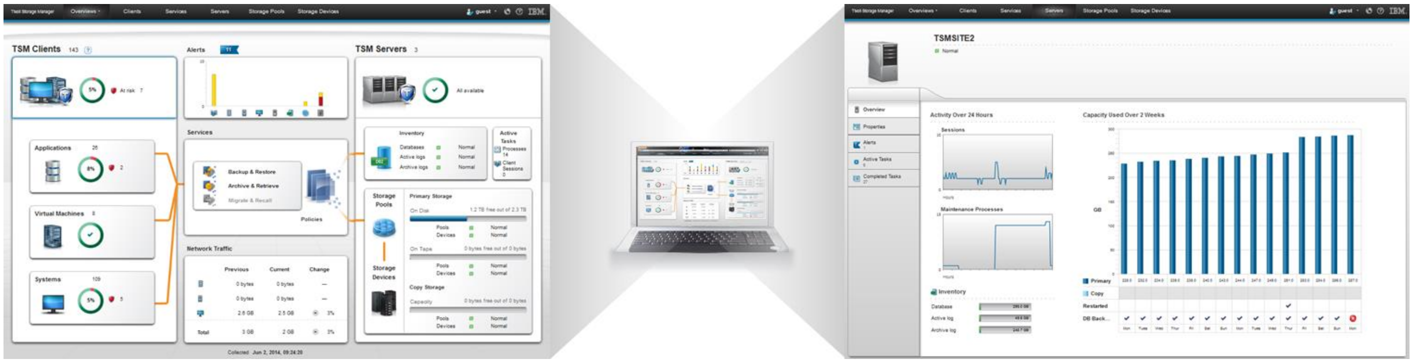

Next, ESG Lab audited some of the performance efficiencies included in the 7.1 release. Through suggested hardware configuration changes and software streamlining, IBM has been improving the achievable server-side deduplicated ingestion rates of TSM servers. While every configuration will perform differently, IBM has used a controlled test bed to quantify the improvements to the code with each major software release since v6.3.3. With each release, IBM has systematically targeted and improved each phase of the server-side ingestion, including the ingestion process and identification, deduplication dereference, and reclamation.

ESG Lab audited the test results of these improvements, and was impressed to see improvement of more than 600% in the ingestion rate of server-side deduplication on TSM servers from version 6.3.3 to 7.1.0. The results showed that IBM is committed to continually improving the efficiency of the TSM server with each release, rather than only adding new

[2] TSM Operations Center Live Demo: https://demo.tsm.ibmserviceengage.com:11090/TSMLiveDemo/

features and functionality. Figure 9 shows the improvement in the ingestion rate of a single TSM server for each tested release of code when compared with the ingestion rate measured for the 6.3.3 release.

Figure 9. Improved Data Ingestion Rate

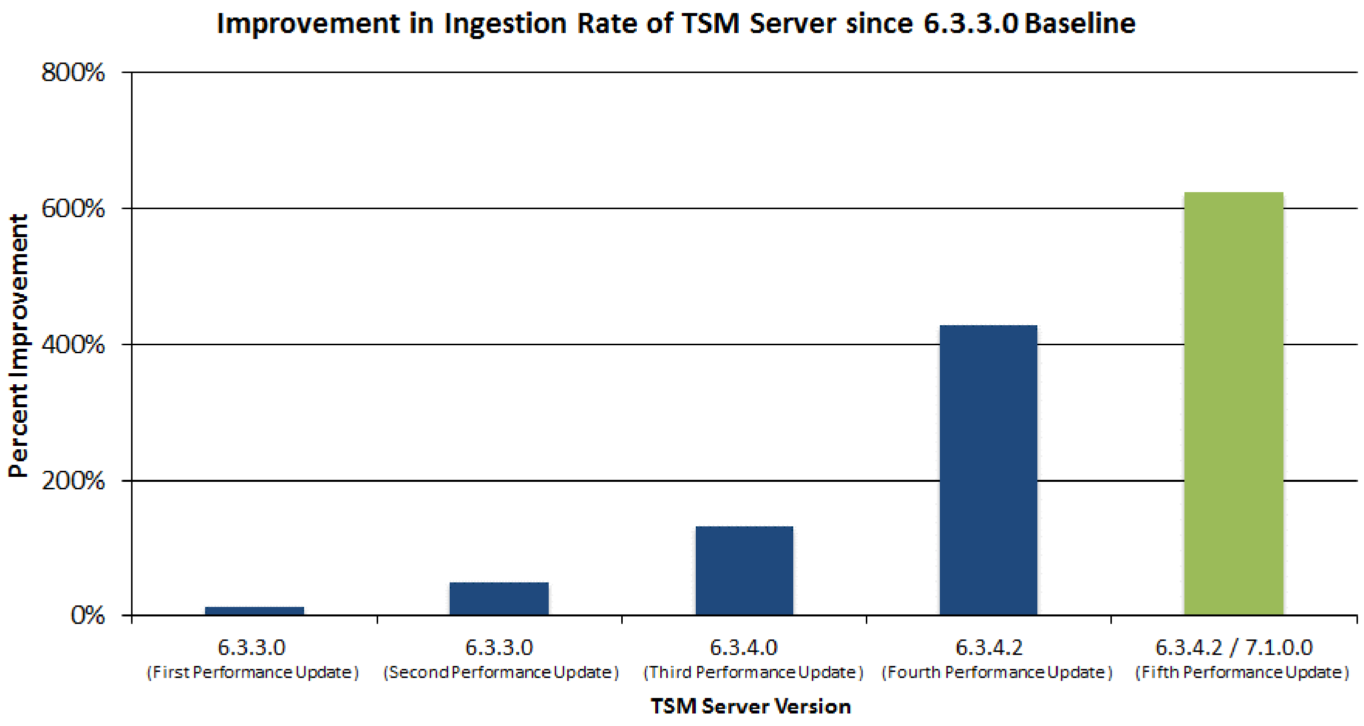

By streamlining the process and making the ingestion and deduplication processes more efficient, IBM enables customers to perform more backups in less time. Using the results measured on the controlled TSM server environment highlighted earlier, IBM demonstrated the time that was required to backup an environment with a 10TB daily change rate to the single TSM server. The same dataset used for version 6.3.3.0 was retested using version 7.1.0.0. The results demonstrate that the improvements included in release 7.1.0.0 enable shorter backup windows for existing data sets or the ability to protect larger environments (see Figure 10).

Figure 10. Improved Ingest Duration

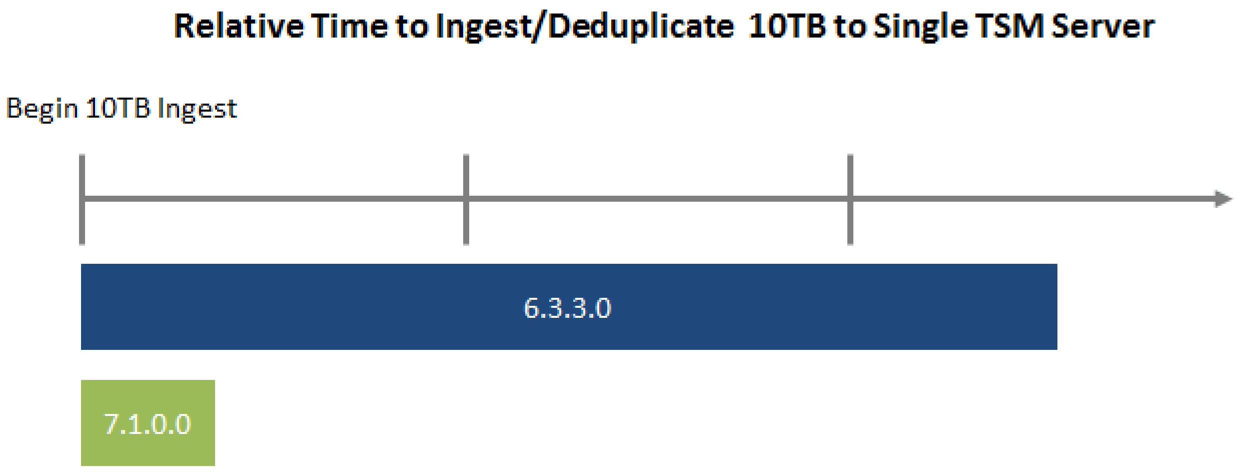

Finally, ESG Lab audited the improvements introduced in the 7.1 release to better handle the ingestion of larger objects. In previous releases, because objects were ingested in series, customers were encouraged to limit the size of objects to 50GB. Among other improvements, version 7.1 introduced server-side ingestion parallel processing of larger files. By breaking down larger files into smaller chunks on the server side and processing them in parallel streams, the data ingestion is completed faster. This enables a single TSM server to more quickly perform full initial backups of large objects. It should be noted that environments that are purely VMware may not benefit from this improvement because VM backups are presented in smaller 128MB chunks. However, VM backups in a mixed environment with larger non-VMware datasets will indirectly benefit due to the reduction in processing time for the larger objects. Figure 11 compares the data ingestion of 40 large objects between version 6.3.3.0 and 7.1.0.0. Version 7.1.0.0 was able to ingest a much larger dataset in less time, resulting in improved protection for large object environments.

Figure 11. Improved Large File Ingestion Duration

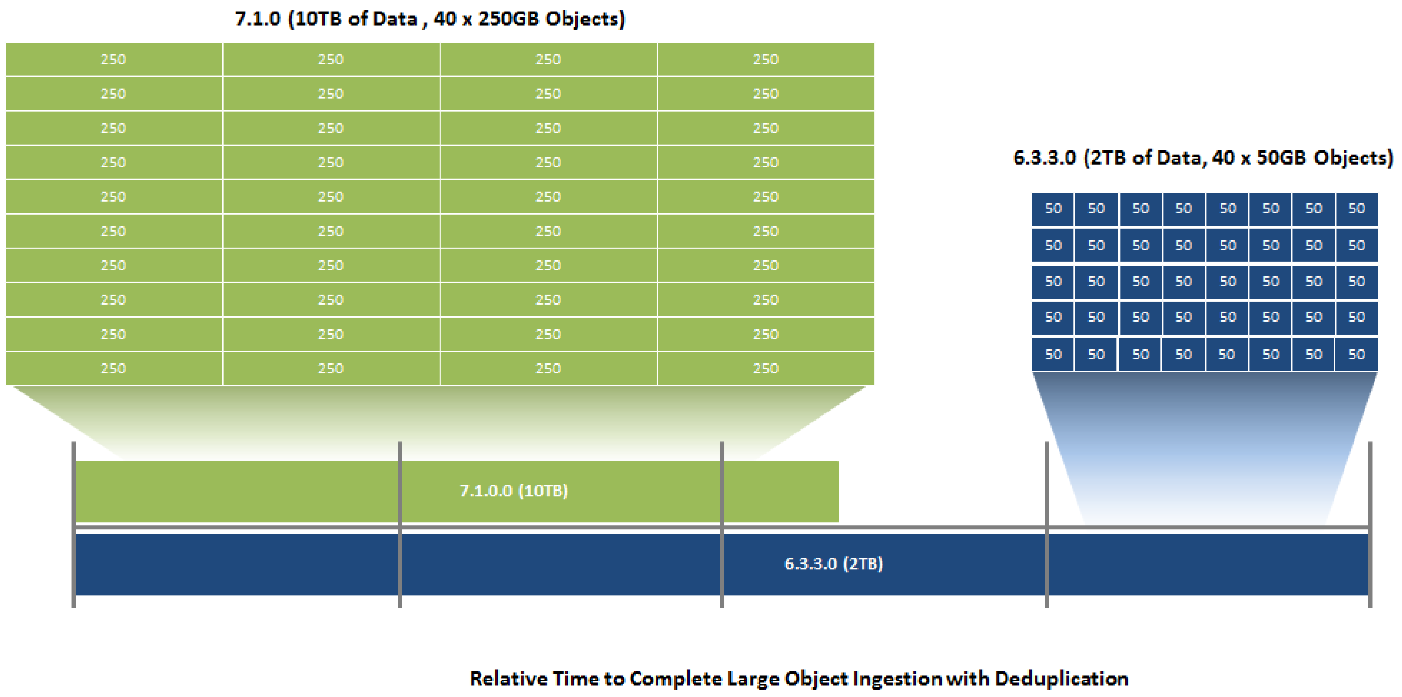

Why This Matters

As companies grow and invest in technologies that make it possible to scale to larger virtual environments, it is critical that the data protection technology keep pace. Companies can simply purchase additional data protection hardware to achieve this task, but the cost and complexity adds up very quickly. It is not enough to simply throw more hardware at the problem; the data protection environment must become more efficient.

IBM TSM for VE version 7.1 and FCM for VMware supports larger virtual environments and adds many new features, but IBM is also dedicated to improving the performance and management efficiency with each major release of code. ESG Lab validated that TSM Operation Center made it easy to manage and monitor a complete data protection environment. ESG Lab also audited the performance enhancements in version 7.1. Given a 24-hour backup window, a single TSM server running the 7.1 release was able to backup and deduplicate more than 30TB of data.

The Bigger Truth

IBM TSM has a long and prominent history in the enterprise data protection space. In fact, its origins start in the late 1980s with a product named Workstation Data Save Facility (WDSF), which evolved into ADSTAR Distributed Storage Manager (ADSM) and was most recently rebranded as Tivoli Storage Manager.

TSM was ahead of the game in its efforts to integrate information lifecycle management concepts into a data protection solution and to understand the value of a perpetual incremental backup schema. IBM developed these approaches to address the data protection challenges caused by the proliferation of distributed server environments. The product has long been feature-rich and builds on rock-solid infrastructure. However, its name has not been synonymous with intuitive ease of use and management.

Now, highly virtualized environments have taken the place of distributed server environments as among the top challenges that data protection professionals face, and IBM is meeting this challenge with its TSM for VE solution. ESG Lab tested TSM for VE 7.1 and was pleased to find a feature-rich and easy-to-use solution for VMware data protection. The Lab explored the intuitive, wizard-driven interface by configuring backup schedules and conducting multiple restores. We used both VM- and storage-level snapshots to validate the quick restore and access capabilities the solution provides. The Lab also confirmed the ability to do granular recoveries from image backups by restoring individual files and MSSQL objects.

However, the backup and restore process alone does not paint a complete picture. All the backed up data must be stored and managed. TSM for VE 7.1 accomplishes this through strong integration with the IBM TSM server. The TSM server brings all the mature data management capabilities (e.g., backup history, deduplication, migration, collocation, and replication) to the overall solution.

ESG Lab believes that IBM has done a good job making TSM for VE extremely user friendly while still offering all the bells and whistles that are required to meet today’s virtual data protection challenges. IBM is also working hard to make the TSM server more user friendly through the introduction and ongoing development of Operations Center along with the reference architecture blueprints for deploying TSM. If your organization is currently running TSM for physical server backups, ESG believes TSM for VE is worth a serious look. If your current solution is not meeting your data protection needs or the cost of maintaining it is eating into your IT budget and you are considering a replacement, you may want to engage IBM to conduct a free assessment of your data protection environment as part of your decision process.