



Tivoli Storage Manager 6 Launchpad



Transcript

Tivoli Storage Manager 6 Launchpad

Introduction: Managing and Protecting Data with TSM

THE TIDAL WAVE OF DATA CONTINUES

Businesses continue to face a tidal wave of data and information that is ever-increasing. The ability to successfully manage data has become a competitive necessity, and with increasingly stringent compliance demands and the growing complexity of the storage infrastructure, storage administrators are feeling the pressure.

Data loss can do serious harm to business activities: delaying or disrupting the revenue stream, undermining customer confidence, and even doing long-term damage to its reputation in the marketplace. Contrast that to the budget reality that most of us experience. IT budgets are being trimmed or at best, remaining flat. Storage infrastructures and the associated costs of administration are growing while the resources available to preserve and manage stored data are shrinking. We have to find ways to accomplish more at a lower cost.

HIGHER SERVICE LEVELS ARE REQUIRED

As more business is conducted online, the need for faster and more reliable data recovery increases. One underlying fact is that there is a lot more information currently online than ever before, often without an off-line option for conducting business. When something goes wrong, information and applications can become unavailable. And then you have the potential for loss of revenue, damage to your brand and regulatory non-compliance.

However, we're seeing a growing gap that's created by the increasing cost of information unavailability and the decreasing ability to deliver recoverability through traditional backup and restore technologies. It's widening because the costs of disruption continue to increase as more business is done digitally, and because the increasing complexity of systems make it more difficult to perform a timely recovery. The more data you have, in more places, the harder it is to recover when anything goes wrong.

THE PRESSURES ON STORAGE ADMINISTRATORS ARE MOUNTING

So the pressures are growing for storage administrators, who are constantly being asked to do more with less. They are trying to keep up with exploding data growth and longer retention requirements, higher service level requirements, and integration of heterogeneous systems from mergers and acquisitions. There are serious security issues – you need to protect your customers' data as well as your own. And you probably have a large amount of critical data distributed in remote and branch offices that also needs to be protected.

The bottom line: you can't afford any downtime, you can't afford to lose any data, and you can't afford to keep buying more storage capacity and hiring more administrators. You need a smarter way of getting the tidal wave of data under control.

INTRODUCING TIVOLI STORAGE MANAGER 6

IBM Tivoli Storage Manager, or TSM, has been a leader in helping organizations protect, retain and maintain availability of their most important asset, their data, for more than 15 years. The TSM family is an enterprise-wide data protection and recovery solution, with features including: Backup / Restore, Archive / Retrieve, Continuous Data Protection (CDP), Disaster Recovery, HSM / Information Lifecycle Management, Application Protection, Database Protection, Bare Machine Recovery and Robust Remote Office Recovery.

TSM also offers unified recovery across the broadest range of OS and HW platforms (from laptop to mainframe) and across all storage tiers and has been a leader in data reduction techniques since 1993. Introduced in early 2009, Tivoli Storage Manager 6 leapfrogs the competition to address your most challenging IT issue, the continuing tidal wave of data growth, through significant improvements in management, scalability, performance and availability.

Tivoli Storage Manager 6 is the core component of an enterprise-wide data protection and recovery management solution set. Tivoli Storage Manager works seamlessly with Tivoli Storage Manager FastBack™, which provides enhanced data protection and recovery of critical Microsoft® Windows servers in remote offices and in the data center and with Tivoli Continuous Data Protection for Files to protect desktop and laptop systems.

The net result is a unified recovery management platform for the entire enterprise.

SUMMARY

Tivoli Storage Manager 6 keeps IBM on the leading edge of enterprise-class data protection and recovery management, staying ahead of the ever-increasing growth in digital information, no matter where it resides.

You need Tivoli Storage Manager 6 if you want to reduce data protection infrastructure and operational costs; if you want to improve service levels by reducing backup and recovery times; and through improved integration with Microsoft Windows, VMware and NetApp/N series; and if you want to simplify administration and gain better visibility into your data management environment.

Reduce costs, improve service levels, mitigate risks – Tivoli Storage Manager 6 is a key component when building a dynamic storage infrastructure.

Demo 1: Restoring Active Directory Objects

Introduction

Active Directory has become a very popular identity and access management solution for Windows environments. Many customers are looking for a more granular recovery solution to help reduce downtime and ease the burden on administrators when something goes wrong in the Active Directory environment. This demonstration will show how to restore individual Active Directory objects in Tivoli Storage Manager (or TSM) version 6.1 or later.

Demonstration

We begin by going to the Active Directory users and computers screen. First we will delete a user and then, using TSM, we will restore that user.

Here we have created several organizational units: OU1, OU3, and so on. As an example, we'll take the role of an Active Directory Administrator and, going to "OU test", delete user "Test 2". Afterwards we realize that this user should not have been deleted – now what can we do?

In Tivoli Storage Manager we'll launch the backup archive GUI. We on "Restore", and you can see the Active Directory node in the tree. This is new to TSM version 6.1.

Expanding the Active Directory node shows us both "local deleted objects" and "server". The difference between the two is that Active Directory keeps track of the deleted objects, and they're not erased when you delete them; they're just marked for deletion. Although they are no longer shown to the administrator, it is possible to restore them. Active Directory does not preserve all of the attributes of the user, so you might lose some functionality when you restore from the local deleted container.

The best way to restore is from the server. If we have performed a system state backup in the past, which should be done every day, we are able to retrieve Active Directory information from the backup. So when we click on the server, it displays the list of existing system state backups ... including yhe one that we took only a few days ago.

We'll select this one, and push the "OK" button. This basically tells us that it's going to restore the Active Directory database to a temporary location for processing.

Now, looking at our organizational unit, OU test, we can see the user that we just deleted, Test 2. On the right hand side it shows all of the information about this user: the attributes, their names if they are stored in the Active Directory database, and their values. This information can be used for reference purposes.

We'll select this user and push the "restore" button, just as we do when we restore a file. Now that the restoration is complete, we'll go to the Active Directory users and computers screen and click F5 to refresh it ... and, as you can see, user Test 2 is restored back into the environment.

However, returning to the Restore screen, (in Windows Server 2008) we were able to see both the local deleted object and the server views. Using Windows Server 2003 only the local deleted objects are available.

Conclusion

As you have seen in this short demo, you now have the power in TSM 6 to restore individual user accounts and other data objects in Windows Active Directory environments.

Demo 2: Data Protection for MS Exchange

Introduction

Recovering from problems with your Microsoft Exchange server has always been a painful, time-consuming task. With the release of Tivoli Storage Manager 6 for Mail, we've significantly eased that pain by enabling fast and easy recovery of individual mailboxes and e-mail objects such as messages, calendars, contacts and attachments. We'll show you how fast and easy it is in this demo.

Demonstration

First, let's open up our TSM for Mail Version 6.1 product, and we can see that we have an Exchange server here called Rangoon VM2. This Exchange server is its Exchange 2007 Server has three storage groups.

We will now back up all three storage groups by clicking backup here. This is usually done through a scheduled operation, so in production there would be no graphical user interface; it would all be done automatically and run in the background. For demonstration purposes, we will use the graphical user interface here now.

So we can see the backup progress. We have already transferred our 30 megabyte Exchange Database. It has finished, and we can see the message right here. We have chosen to take a VSS backup; that's essentially a snapshot type backup which will provide a snapshot -- on local disk and also within the TSM server.

The return code is zero. We can see that it was successful. Now we can go to the Restore tab and see what's available for recovery. In this case, we have the first, second, and third storage groups available for recovery. We also have snapshots inside TSM server and we on local disk. So we now have a valid backup.

The next step is to bring up the mailbox of an individual user. The name of the user is Bob. Bob has already received some e mails here in his inbox. Let's assume that Bob accidentally deletes Project 4711 and the second e mail for Project 4711.

So Bob has delete the information and makes sure that's it's really been deleted. We go into the trash bin here and delete it from the trash bin as well. Synchronize really quickly. And now, let's assume a month later Bob wants those two e mails back.

This is very simple. The only thing we have to do is go into our TSM for Mail product -- in Version 6.1 we have a new utility called "Mailbox Restore". We click on this utility and we can look for Bob, actually perform a search here, and specify Bob's mailbox as the one that we want to restore.

We can add Bob's to the list of selected mailboxes, and pick a backup date to define how far back we want to go in time to retrieve the e mails or the mailbox that has been accidentally deleted.

In this case, we go with the default. So we'll take the last backup that has been performed. We can also specify filters. These filters select individual e mails which should be restored. Since we happen to know the subject, which was Project 4711, we will enter this as in the subject text.

We can also select where we want to restore the files to -- one option is to restore right back into the Exchange server. Another option is to restore into an alternate mailbox, somebody else's mailbox. Or, we can also restore into a PST file which can then be sent out to other people who might have an interest in that particular PST file.

In this case, we choose to restore right back into the Exchange server. So the only thing we have to do is click on Restore, and TSM for Mail will now use the Exchange Recovery Storage Group in the background. This process is 100 percent Microsoft compliant.

TSM for Mail will restore Bob's individual mailbox and we can see that this operation has completed. The backup was actually taken from a VSS snapshot type backup that is now being used for the recovery.

Now, let's open this up so we can see that the restore completed successfully and the individual items were recovered. We can see that two items were restored. Let's click okay, and go back to Bob's inbox.

We have to synchronize so we'll click on "get mail" and the e mails are being brought back. So we see Project 4711 here, and here again. This was a very simple one step procedure to restore individual e mails.

Let's switch over to another scenario. Another scenario would be if an administrator accidentally deletes Bob's complete mailbox. To simulate such a case, we open up the Microsoft Exchange Management Console, which is what an administrator typically would see. Within this console, the administrator will now remove this mailbox.

So we click on Remove, and Exchange asks us if we're sure. Yes, we are. We click on Refresh, and the mailbox is gone. If Bob now tries to log into this mailbox let's click on synchronize here we will get an error message simply because the mailbox doesn't exist anymore on the server.

Now, in this scenario, we open up TSM for Mail again, go to the utilities, specify a mailbox restore, we can see Bob right here. We select him. Now, we leave all of the filters empty, which means we want to restore the complete mailbox. Again, we want to restore to the existing Exchange server. The only thing we have to do is click on restore.

Essentially this is the same scenario as before. The mailbox is restored into the recovery storage group and we can see that the VSS operation completed. Now, we have to wait a little bit since all of the e mails are now being restored from the last backup right back into Bob's new mailbox. Okay, we have a success message here.

So here we can see that the mailbox restore completed successfully, and 24 items were restored. Let's close that. Let's go back to Bob's mailbox and synchronize again. Now all of the 24 messages are being downloaded from the Exchange server so they are all successfully restored. We have restored individual e mails, and we've also restored a complete mailbox, in both cases from Bob.

Let's look again at how we did this. It was a one step procedure, and the only thing we had to do was go into TSM for Mail 6.1 and select the mailbox we want to restore. Then, just click on Restore and the mailbox or the individual e mails are being restored automatically for us.

Conclusion

Recovering individual mail items or mailboxes is a one step procedure now with TSM 6 for Mail. It is also backward compatible, so you can use TSM 6 for Mail with the TSM 5.4 or TSM 5.5 server just as well.

Thank you.

Demo 3: New Features & Data De-duplication

Introduction

Tivoli Storage Manager has always been a leader in data reduction technologies – including progressive-incremental data capture that only backs up changed files forever; data compression; and automated migration of older files from primary storage to secondary storage. The data can be automatically deleted when its useful life has ended.

In the Extended Version of TSM 6, we have now added data deduplication to help further reduce capacity requirements and to help speed recovery times by enabling more data to be stored on the same amount of disk. In this short demo, we'll show you how the new data deduplication works.

Demonstration

As you can see here, we are logged into the TSM Administration Center. On the right side, you can see the Tivoli Storage Manager tab, and a new tab which is used for Tivoli Common Reporting. This is new in TSM Version 6.

First, let's have a quick look on the Health Monitor. The Health Monitor has been improved so we have now more functions to help you see how your systems are performing. We have one server here that we are managing, the template server, and you can see the health of this server on one page here.

You can see the storage device status, the activity log, and the overall status as well. You can also see if there were any missed job schedules or any type of warning and error here. Also new in Version 6, you can drill down on error messages by clicking on the hyperlink to bring you to the cause of the error, and then you can fix it right from the Health Monitor.

We also have a new tab for client nodes and backup sets. In the past, you had to navigate through policy domains and policy sets in order to get to client nodes, but now we can directly access the client nodes and backup sets, and we can perform all of the node related actions right from here. You can sort this by server, and we can even search for individual nodes across multiple TSM servers. This is also a new function in Version 6.

Well, let's go back to our Storage Devices tab, and you can see that we have five storage pools defined here. And at this point in time, they're all empty. You can see this by the percentage that's been utilized; they all show zero percentage utilized. So all of the five storage pools which are all on disk are empty at this point in time.

So let's perform a backup and see how this reflects a change in the storage pool. We are using our new backup archive client here, which is now written in Java to allow for portability and gives us flexibility in terms of how we can deploy this new client.

Let's back up one directory. We want to back up the critical directory. We select a full backup, which means that we are not leveraging the Incremental Forever function in TSM; we're just backing up this directory. And you can see the backup report that has been completed. We have actually backed up 10 objects.

Now, let's refresh our storage pools here and we can see that our backup pool where the backup data is going is now 13.1 percent utilized. So this is approximately about three megabytes that we have now backed up into this particular storage pool.

Let's do it again. We'll back up the same directory, the critical directory. We will again choose an Always Backup, which will force the client to back up everything that's in there again. We see Backup completed, 10 objects. Let's refresh the storage pool, and we can see that this storage pool is now utilized by 26.2 percent. That will be approximately six megabytes.

Now, let's open the TSM Administrative Command Line client, and what we want to do now is migrate this storage pool. We type in the migrate storage pool command, forcing the migration of our backup pool. And then we go back here and refresh the table.

We can see that the backup pool is now empty, and instead we have moved the data into the sequential deduplication pool, the SEQ DUP pool. This one has a capacity of 40 megabytes, and we have only used our six megabytes, which shows 13.1 percent utilized. So all of the data is now in a different disk pool.

The next thing we would like to do is to back up this storage pool. The reason why we're backing up this storage pool is because we have defined that data deduplication can only take place when a backup has been performed on this exact same storage pool.

Now we have backed up the storage pool. Let's do it again, and see how this reflects in our table right here. We have to refresh it, and can see that we now have the same data in the copy storage pool that we have in the sequential deduplication pool. So we also have this one pool being utilized 13.1 percent with six megabytes sitting in both pools at this point in time.

Okay, now let's take a look at the new TSM 6 Data Deduplication function. We would like to perform this function on our sequential data deduplication pool here. So we select the radio button, and can actually scroll down here and see a new function called Identify Duplicates.

We click on this function and define the number of processes which are being used to identify duplicate data, We only want to go with one process at this point in time. Click okay. This is essentially all we have to do, which means we have now marked this storage pool as a data deduplication pool. It has to be a sequential disk pool, otherwise data deduplication will not be supported.

At this point in time nothing has been de duped; we have just enabled the storage pool to be de duped. And again, it's 40 megabytes and 30.1 percent is being utilized right now.

Let's go back to our TSM administration command line, and reclaim space on that storage pool, and on the sequential de dup pool, the SEQ DUP pool.

In a production environment this would be a scheduled task that runs on the TSM server in predefined time slots. So this is typically an action that takes place overnight. Let's do it.

Okay, we can see that the reclaimed storage pool command has been started, and we can go back in here and refresh the table again. Almost magically we can see that the sequential deduplication pool is now only half the size of what it was before.

Remember we had about 13.1 percent utilized before with about six megabyte? Now we only have 6.6 percent, which is half of that. And why is that? Because we backed up the same data twice and now have de duped it so the TSM server has basically eliminated all the duplicates, saving about 50 percent here. Okay. Proof is that our copy pool, which is up here, is still 13.1 percent utilized.

So as you can see the only thing for TSM has to do in order to perform data deduplication is mark the storage pool as a to be de duped pool. Then you just perform a reclaim storage pool function and data de duplication takes place immediately.

Conclusion

The design goal of TSM 6 is to help you manage more data while reducing costs. Through the new data deduplication feature in TSM 6 Extended Edition, you can reduce current and future storage capacity requirements while keeping more backup data on disk to help speed recovery times.

Demo 4: New Reporting and Monitoring in TSM 6

Introduction

With the release of Tivoli Storage Manager 6, we have made significant improvements in ease-of-management to help you reduce the time and the complexity of administering your TSM environment. We've simplified administration by incorporating all new reporting and near-real-time monitoring to help you gain better visibility into current operations, and see historical trending. This will help you make smarter decisions about future capacity requirements and other management tasks. In this demo, you will see some of the power of these new capabilities that are built into TSM 6 with no additional licensing required.

Demonstration

First, let's see how the information for monitoring and reporting is being captured. We now have so called monitoring agents. In this case, we have to find three monitoring agents on our local server. These

monitoring agents connect to a Tivoli Storage Manager and extract the information that is needed for reporting and monitoring from the Tivoli Storage Manager.

We have three agents here - one connects to Tucson, one connects to a local TSM server, and one connects to a TSM server in our local area network. You can create as many agents as you like. You basically click on a predefined template, and in this way you can create your own agent, which then reports to a specific TSM server.

Now, let's have a look at what the monitors look like. Once you have this agent running, information will be collected and you can display that information within the Tivoli Enterprise Portal, or TEP.

This is the TEP graphical user interface here, and we can open TEP and navigate to the Tivoli Storage Manager entry here. This is a new entry, and it shows all the TSM servers which are available to be monitored.

For example, our local TSM server is right here and we can open this up to see all the monitors that are available. There's an availability monitor. This one basically shows which monitors are available at this point in time or which managers are active. And as you can see here, we have some active and inactive ones.

Now as we scroll down a little bit, you can see all the predefined monitors. We have monitors for client node storage, client missed files, clients node status, a database monitor, a TSM client node activity monitor, a schedule monitor, a storage pool monitor, a TSM server monitor, storage device monitor, tape usage monitor and tape volume monitor.

Let's connect to our server in Tucson and see what's available there. These are the same monitors, and we can look at the first one, which shows client node storage. We can see all of the client nodes, their disk usage, their tape volume count, their maximum file space capacity and some other parameters that are available in this monitoring.

The Client Missed Files monitor will show files that were missed in the last backup by client node name when the files were missed, and, the actual file name of the missed file. So in this case, Patrick has missed backing up his ntuser.dat file on February 1st at 6:43.

Bar and pie charts help us present these statistics in a more visual format. There's a client node status, ... and a database monitor. Here we can see the actual TSM server database, the current space, the allocated size, and how much total space has been used.

We also have a node activity here so we can see when backups were being performed, when they were started, and when they ended. We can see which schedule they belong to. We can see how the backups performed, how many files or objects were inspected, how many were processed, and also how many failed. And again, this can be visualized in the graphs at the bottom of the screen. These windows are adjustable, so you can modify them to your individual needs and store these settings as a separate TEP workspace.

Now let's have a look at the schedules option. We can see all of the schedules that were being defined on that server, their specific status, and storage pools on that server. This information is displayed in gigabytes. We can see the total space, the total usage, how many volumes are actually being used, and we can see some trending information here as well.

Let's look at another monitor, the server monitor. This monitor shows us how our TSM server is doing, how many client backup operations are being performed, how many times the TSM server database has been backed up, what the migration status is, what the reclamation was doing, and also how the backups were being done.

Let's look at the storage device monitor. So here we can see all the storage devices on that server. So we have device class file here, and we also have device class disk.

The Tape Usage monitor tracks the actual tape usage so we can view the particular clients, the storage pools, and the server name.

By looking at the Tape Volumes monitor we can actually see how much capacity has been used and how much remains unused. And we can see that our server is pretty busy and all of the tape volumes which are pretty full are being shown in red.

You can move your mouse over the tapes and you will get additional information for that particular tape volume. So you can see we have one, two, three, four, five, six, seven, eight, nine, 10, 11 monitors that come with TSM reporting monitoring. And, of course, you can define your own.

Let's exit TEP and look at some reports. Reports are a little bit different from monitors. They're also being collected from the TSM servers, and the information is stored within a DB2 database.

You can create those reports through the TSM Administration Center. Here you see the TSM Administration Center, and you can click on Work With Reports. This will bring up our new screen for reporting. I'm going to maximize this.

You can see Tivoli products show up and Tivoli Storage Manager will appear, showing you that it has two sections: client reports and server reports. Let's look at the client reports. We can do this by simply double clicking, and we can see that we have six client reports already defined here.

Now let's just take a look at the client activity history which provides a set of charts showing key client metrics over a period of time. We can click here on the HTML option, which will format this particular report in HTML format. We then get an On-Demand Report parameter screen where we can pick what we would like to see.

In this case, we want to see the last 30 days and would like to report on our Tucson server. You can actually pick the server that you want to report on. Remember, we're collecting data in this example for three TSM servers. So let's run this.

Now the report has been created, and we can see the client activity history here. We can see all the nodes that are a part of this particular report. We can see date, and we can see a very nice chart here, too. We can also see the bytes and number of objects that are being moved. And of course, we see the creation date and time.

Let's look at another report - client backup currency. This time we would like to see this report in PDF format. Again, we pick the server name here and run the report.

Now, this one is in PDF format, and it can be printed, of course. We can see client backup currency here, so we can see all of the clients and we can see the dates of their last backups. For example, Vista Test Client was backed up on January 26th. We can also scroll down here and see this report for all of the clients that are defined on that particular server.

The same principle applies to the server reports. Let's check them out here. We will select the server activity detailed report. Click on HTML, and here we get a more sophisticated report parameter screen. We want to select the last 90 days. We want to select our Tucson server, and we would like to see how the database server backup went on that particular server.

Now, click on run, and the report is being formatted so we can see the duration. We can also see the date, how many objects were being evaluated, and how many objects were being processed.

Let's take a look at the tape volume capacity analysis. We will select PDF format this time. This report here is created in real time, and the data is not in real time, so it's not being pulled from the TSM server, but from a DB2 database which holds all of the collected data over time. So this is independent from the TSM server.. The TSM server does not even have to run in order to create these reports. They are mentioned as historical reports.

So let's put date and time in here, so we would like to start from February 1st, again, on our Tucson server. This is now formatted in PDF format; which is why it takes a little bit longer to run.

So here we have a tape volume capacity analysis. You see the number of volumes and we can see what percentages of these volumes are being used. We can see a pie chart that shows used capacity and unused capacity. And here we can view by volume and see in a detailed fashion what's happening on all of these tape volumes.

We have now looked at four reports and as you can see, we have five server reports and we have also six client reports available - that's a total of 11 reports. These reports are available right out of the box, or you can define them on your own. A development tool called BIRT is available to do that, and using that tool you can essentially create your own reports.

Conclusion

In the past, you would have had to purchase third-party reporting and monitoring tools to get this level of visibility into your TSM environment. We're now providing it as an integral part of the TSM 6 product, and it should help you to react faster to changes or problems in your environment, help you to make better decisions, and ultimately make you more productive and successful.