

IBM TIVOLI STORAGE MANAGER V6.3 FUNDAMENTALS COURSE CONTENT

❖ SECTION 1: IBM TIVOLI STORAGE MANAGER FAMILY:

A. Given the online resources available, identify which IBM Tivoli Storage Manager (Tivoli Storage Manager) data protection product to use for a given solution.

With emphasis on performing the following tasks:

- a. Tivoli Storage Manager for Databases - Perform online, consistent ,and centralized backups to avoid downtime, protect vital enterprise data infrastructure and minimize operation costs.
 - 1) Includes Data Protection for Oracle, which interfaces with Oracle Recovery Manager to support Oracle backup and restore utilities.
 - 2) Includes Data Protection for Microsoft SQL Server, which enables users to perform on-line backups of SQL databases, to Tivoli Storage Manager storage.
- b. Tivoli Storage Manager for Mail - Utilizes the application program interfaces (APIs) provided by e-mail application vendors to perform online "hot" backups and improve restores without shutting down the e-mail server.
 - 1) Includes Data Protection for Lotus Domino, which exploits the "transaction logging" feature of Domino, enabling the capture of just the database changes for logged databases, thus resulting in less-frequent full backups.
 - 2) Includes Data Protection for Microsoft Exchange, which can produce the different types of backups specified by Microsoft backup APIs: Full Backups, Incremental Backups, Differential Backups, Copy Backups and Database Copy Backups.
- c. Tivoli Storage Flashcopy Manager - Provides fast application-aware backups and restores leveraging advanced snapshot technologies in IBM storage systems.
- d. Tivoli Storage Manager for Microsoft Sharepoint - Provides backup and recovery of Microsoft SharePoint Portal 2003 and Microsoft Office SharePoint Server 2007 environments.
 - 1) Include ability to do both a granular and platform-level backup and restore.
- e. Tivoli Storage Manager data protection for virtual environments - Provides backup and recovery for virtual environments (Vmware, HyperV).
 - 1) Include the ability for recover VM's to original and alternate environments.

B. Given online resources available, describe Tivoli Storage Manager deduplication functionality so that knowledge of where the deduplication action takes place is identified.

With emphasis on performing the following tasks:

- a. Client-side deduplication -Deduplication performed at the data source (e.g., by a backup client), before data is transferred to a target location.
 - 1) Advantages:
 - Deduplication before transmission conserves network bandwidth.
 - Awareness of data usage and format may allow more effective data reduction.
 - Processing at the source may facilitate scale-out.
 - Immediate data reduction, minimize disk storage requirement.
 - No post-processing (identify, reclamation) to regain duplicate data.
 - 2) Disadvantages:
 - Deduplication consumes CPU cycles on the file/application server.
 - Requires Tivoli Storage Manager V6.2+ software deployment at source and target endpoints.
 - Might be bottleneck for data ingestion (e.g., longer backup times).
 - Only one deduplication process for each I/O stream.
 - No deduplication of legacy data on the target server if only method used.
- b. Server side deduplication - Deduplication performed at the target (e.g. by backup server or storage appliance).
 - 1) Advantages:

- No deployment of client software required.
 - Possible use of direct comparison to confirm duplicates.
 - No impact to data ingestion.
 - Potential for deduplication of legacy data.
 - Possibility for parallel data deduplication processing.
- 2) Disadvantages:
- Deduplication consumes CPU cycles on the target server.
 - Data must be processed twice (during ingestion and subsequent deduplication).
 - Storage needed to retain data until deduplication occurs.

❖ **SECTION 2: IBM TIVOLI STORAGE MANAGER SERVER ARCHITECTURE CONCEPTS:**

A. Given the use of IBM Redbooks, identify IBM Tivoli Storage Manager (Tivoli Storage Manager) "wheel of life" data protection products so that knowledge how/when Tivoli Storage Manager processing takes place is obtained.

With emphasis on performing the following tasks:

- a. The daily schedule has a period where clients perform their backups.
- b. Once the clients are finished, the server performs housekeeping.
 - 1) The server makes copies of the primary storage pools for off-site storage.
 - 2) The server backs up its database, deletes volume history, saves the device configuration, and creates a list of tapes for.
 - 3) The server then migrates the data from disk storage pools to on-site tape pools, reclaims blank space from tape pools.
 - 4) Finally the expiration process runs.
- c. A new client backup begins for the next night.

B. Given online resources, describe the role Tivoli Storage Manager policies play in the enforcement of data retention so that customer data is retained per business or legal requirement.

With emphasis on performing the following tasks:

- a. Identify business data retention requirements.
- b. The Tivoli Storage Manager policy structure is built on the server by using either the GUI or CLI.
- c. Tivoli Storage Manager policy hierarchy structure is as follows:

1) Policy Domain:

A policy domain is a way to group Tivoli Storage Manager clients depending on how to treat their data.

2) Policy Set:

A policy set is a group of management classes. There can be multiple policy sets within a policy domain, but only one of them is active at a time.

3) Management Class:

The management class is a tier in the policy management that essentially serves as an interface between the client's data and the copy groups whose rules govern the versioning and/or the retention of data. A management class usually contains both a backup and an archive copy group, or it can house either group. A management class can

even be empty, that is, without a backup or archive copy group, but in such a case, the management class is useless, as there are no rules that would govern the data.

4) Copy Group:

Copy groups consist of rules used to govern the retention of data. There are two types of copy groups: a backup copy group, which holds the rules for backup data, and an archive copy group which holds the rules for archive data.

C. Given online resources, identify the different types of backups available so that the methodology that Tivoli Storage Manager uses for backups are understood.

With emphasis on performing the following tasks:

- a. Identify backup methodology types.
 - 1) Full - All files are backed up.
 - 2) Incremental - Only files that have changed since last backup occurred are backed up.
 - 3) Selective - Overrides Incremental backup and saves selected files regardless if they have changed or not since the last backup.
 - 4) Image or Logical volume backup - Back up a file system, or raw logical volume as a single object from the client machine.
 - 5) Adaptive subfile backup - This is when only changed portions of files are backed up.
 - 6) Journal-based backup - This is when the client keeps a running list of files as they change. This list is what is used to determine what files get saved when the backup occurs.
- b. Determine which appropriate methodology should be used in any given customer situation.

D. Given book/documentation, describe the Tivoli Storage Manager backup versioning and data retention so that the Tivoli Storage Manager data storage rules are defined.

With emphasis on performing the following tasks:

- a. Tivoli Storage Manager backup versioning - is where you save multiple versions of the same object - There are multiple parameters to set Tivoli Storage Manager backup versioning.

VEREXISTS: Number of object copies or versions to keep.

VERDELETE: Number of versions to keep when object has been deleted on the client.

RETEXTRA: Number of days to keep inactive versions.

REONLY: How many days to keep the last object copy in Tivoli Storage Manager when the object deleted on the client.

E. Given knowledge of Tivoli Storage Manager, describe how the Tivoli Storage Manager server and client work together so that Tivoli Storage Manager can provide backup and restore data.

With emphasis on performing the following tasks:

- a. Progressive Backup Methodology:

Progressive backup can be thought of as combining the backup benefits of the incremental approach with the restore benefits of the differential approach. Files are backed up incrementally to reduce network traffic, while recovery media is consolidated to provide better restore performance.

- b. Data Management Policy:

- 1) Policy Domain:

A policy domain is a way to group Tivoli Storage Manager clients depending on how to treat their data.

- 2) Policy Set:

A policy set is a group of management classes. There can be multiple policy sets within a policy domain, but only one of them is active at a time.

- 3) Management Class:

The management class is a tier in the policy management that essentially serves as an interface between the client's data and the copy groups whose rules govern the versioning and/or the retention of data. A management class usually contains both a backup and an archive copy group, or it can house either group. A management class can even be empty, that is, without a backup or archive copy group.

- 4) Copy Group:

Copy groups consist of rules used to govern the retention of data. There are two types of copy groups: a backup copy group, which holds the rules for backup data, and an archive copy group which holds the rules for archive data.

- c. Storage pool:

A storage pool is a collection of storage pool volumes; each storage pool represents one type of media.

Tivoli Storage Manager has two types of storage pools: Primary storage pools and copy storage pools.

- 1) Primary storage pools:
When a client node backs up, archives, or migrates data, the data is stored in a primary storage pool.
 - 2) Copy storage pools:
A copy storage pool provides an additional level of protection for client data and is created for the express purpose of backing up a primary storage pool.
 - 3) Active-data pool:
An active-data pool is a special type of copy storage pool that contains only the active version of data.
- d. Register Tivoli Storage Manager client, associated with policy domain.
 - e. Tivoli Storage Manager client backup/archive objects store in Tivoli Storage Manager data storage rules.
 - f. Tivoli Storage Manager client back up/store data by using Tivoli Storage Manager scheduler or on demand (manual).
 - g. Tivoli Storage Manager client back up/restore protocol (TCP/IP, LAN-Free or share memory).

F. Given knowledge of Tivoli Storage Manager, describe the Tivoli Storage Manager Disaster Recovery Manager feature so that the complete backup of all data and ability to recover it in a timely fashion is understood.

With emphasis on performing the following tasks:

- a. Tivoli Storage Manager server database backup:
The heart of every Tivoli Storage Manager is the database, and the database backup is vital for server recovery.
- b. Copy storage pool data:
When Tivoli Storage Manager backs up clients, the new client data is stored in primary pools. Copy storage pools contain backups of the primary storage pools, and are intended for shipment off-site. Copy all new primary storage pool files to a copy storage pool, ensuring that the copy storage pool is up-to-date with the most recent backup. Each time the primary pool is backed up to the copy storage pool, the newly generated tapes should be sent off-site.
- 1) Active-data pool data:
Active-data pool data can be handled like copy storage pool data. You will need less storage, but you will only have copies of your latest version.
- c. Send copy storage data and Tivoli Storage Manager database backup to off-site location.
- d. Volumes tracking:
Tivoli Storage Manager disaster recovery manager (DRM) provides several levels of volume tracking. DRM volume tracking includes:
 - 1) Identify which off-site volumes are needed for a given recovery.
 - 2) Integrate with tape management systems: Tivoli Storage Manager DRM is fully integrated with tape management. Every time a new tape is created in the corresponding copy storage pools or Tivoli Storage Manager database backup, it is automatically eligible for off-site movement.
 - 3) Recycle partial filled volumes: Off-site volumes are reclaimed in the same way as on-site volumes. By using DRM, you can query which volumes have reached an empty state because of reclamation, and request them to be returned on-site.
 - 4) Tracking off-site volumes: DRM manages media (volumes) by assigning a special state to each volume. The state is one of a number of predefined states used by DRM. There are two possible directions for a volume; move from on-site to off-site and move from off-site to on-site.
- e. Server recovery plan:
The special DRM "prepare" command generates a recovery plan that contains critical information needed for recovery. It contains several scripts that can be exploded into

executable files that can be used for recovery.

- f. Disaster recovery methods - There are several methods for disaster recovery.
 - 1) Mirror site:
A mirror site has an up-to-date replica of all required data, on line and ready to go. Updates from the primary site are transmitted and applied in real time to the mirror site.
 - 2) Hot site:
Generally, a hot site has all of the infrastructure hardware and software in place, A hot site can also contain copies of the data, but possibly not up-to-date.
 - 3) Cold site:
With a cold site, it may have all the required hardware and software available, but not running (thus the term "cold"). It may not even be properly configured or ready for usage. The installation of Tivoli Storage Manager and all required software are needed in order to restore Tivoli Storage Manager server database and Tivoli Storage Manager clients' data. DRM will help to rebuild the Tivoli Storage Manager server, so that that Tivoli Storage Manager clients' data can be restored.
 - 4) Reconstruction:
Reconstruction is rebuilding the previous working environment at the same place and to the same state as it was before. DRM will help to rebuild the Tivoli Storage Manager server, so that that Tivoli Storage Manager clients' data can be restored.

G. Given at least 2 Tivoli Storage Manager V6.3 servers connected via server-to-server and a properly set up storage pool and policy environment, perform node replication so that specified nodes are replicated to the server.

With emphasis on performing the following tasks:

- a. Verify proper server-to-server connection.
- b. Verify storage requirements.
- c. Define target server.
- d. Back up node on source server.
- e. Set up replication rules.
- f. Replicate node:
 - 1) Command
 - 2) Administrative Schedule
- g. Verify successful replication via query occupancy.

❖ **SECTION 3: IBM TIVOLI STORAGE MANAGER STORAGE AND DEVICE CONCEPTS:**

A. Given knowledge of IBM Tivoli Storage Manager (Tivoli Storage Manager), describe the storage hierarchy in Tivoli Storage Manager so that the way Tivoli Storage Manager organizes storage pools is understood.

With emphasis on performing the following tasks:

- a. Disk storage pool
A disk storage pool is a random access device (referred as hard disk). The disk storage pool can be configured as primary storage pool.
- b. Sequential storage pool
This can be a hard disk device but configured as FILE device class, or a tape or optical devices. The sequential storage pool can be configured as primary or copy storage pool.
- c. Primary storage pools
When client node data are backed up, archived, or migrated, the data are stored in a primary storage pool.
- d. Copy storage pool
A copy storage pool provides an additional level of protection for client data and is created for the express purpose of backing up a primary storage pool. Copy storage pool volumes are intended for shipment off-site, to provide recoverability of the Tivoli Storage Manager server

environment. The copy storage pool contains all current versions of all files, active and exactly as they appear in the primary storage pool.

1) Active-data pool

Active-data pool provides faster recovery. It can reside on file device or tape. For faster recovery, it should reside on disk.

e. Migration thresholds

They specify when Tivoli Storage Manager server starts or stops migrating data from a storage pool to the next storage pool in the Tivoli Storage Manager storage pool hierarchy.

B. Given knowledge of Tivoli Storage Manager, describe Tivoli Storage Manager reclamation so that the Tivoli Storage Manager reclamation process, which minimizes the use of sequential volumes is understood.

With emphasis on performing the following tasks:

a. Reclamation threshold:

Reclamation threshold is an optional parameter when defining or modifying a sequential storage pool. Reclamation makes the fragmented space on a sequential volume usable again by moving any remaining unexpired objects from one volume to another volume within the same storage pool, thus making the original volume available for reuse.

b. Data expiration:

When the backup/archive data is no longer needed (exceed versioning or data retention expired), Tivoli Storage Manager server will delete (expire) that object, thus freeing up the storage media.

c. Off-site volumes reclamation:

Tivoli Storage Manager cannot physically move the data from one of these volumes to another because they are in an off-site vault, not available in the library. Tivoli Storage Manager manages reclamation for an off-site copy pool by obtaining the active files from a primary storage pool or from an on-site volume of a copy pool. These files are then written to a new volume in the copy pool, and the database is updated. The new volume will be moved to the off-site location, and the off-site volume, will be moved back to the on-site scratch pool for reuse.

C. Given knowledge of Tivoli Storage Manager, describe the data movement between the Tivoli Storage Manager storage pools so that how data moves between the Tivoli Storage Manager storage pools is understood.

With emphasis on performing the following tasks:

a. Disk storage pool:

Disk storage pool is a random access device (referred as hard disk). The disk storage pool can be configured as primary storage pool.

b. Sequential storage pool:

This can be hard disk device but configured as FILE device class, or a tape or optical devices. The sequential storage pool can be configured as primary, active data, or copy storage pool.

c. Automatic Migration:

By setting up storage pool migration thresholds, a storage pool will automatically migrate data from one storage to next storage pool in the Tivoli Storage Manager storage pool hierarchy.

d. Migration processes:

A process of migrating data from one storage pool to the next storage pool in the Tivoli Storage Manager storage pool hierarchy automatically or manually.

e. Restore storage pool:

In the event of lost or damaged volume(s) in the primary storage pool, restore storage pool will copy data from copy storage pool back to primary storage pool.

f. Reclamation processes:

Reclamation is a server process that consolidates data and free space on tape (or optical) volumes in sequential storage pools.

❖ **SECTION 4: IBM TIVOLI STORAGE MANAGER CLIENT CONCEPTS:**

A. Given detailed information about data protection strategies described in IBM Tivoli Storage Manager (Tivoli Storage Manager) Documentation, explain backups and archives and how they differ so that the concepts have been explained.

With emphasis on performing the following tasks:

- a. Explain how backup uses versioning.
 - 1) Backup is, when data will be stored in versions on the Tivoli Storage Manager server, controlled and managed by policies in versions and retention time.
 - 2) The backup versions and retention time will be managed in the Tivoli Storage Manager server policies by the appropriate management class.
- b. Explain how archive is a powerful and extremely flexible mechanism for storing long term Data.
 - 1) Archive is the function, to store long term data on a basis of retention time bound to a description.
 - 2) Archived objects will be stored with no version limit and they will be retained for the defined time period regardless of whether they are deleted on the client.

B. Given Tivoli Storage Manager documentation, explain the difference between restoring and retrieving data so that the concepts have been explained.

With emphasis on performing the following tasks:

- a. Restore is used for the recovery of active or inactive files that have been saved by using the "backup" function of Tivoli Storage Manager.
- b. Retrieve is used for the recovery of data that has been archived for long term retention.
- c. Explain how active and inactive Versions from backup can be restored.
 - 1) Active version can be restored by BA-Client Command Line or GUI. By default the active version will be used.
 - 2) Inactive versions can be restored from back up by Backup-Archive client CLI or GUI. The inactive version has to be selected explicitly.
- d. Use the package description to search for archives and to determine which files to retrieve.

C. Given Tivoli Storage Manager documentation, describe journal-based backups (JBB) so that the concept has been explained.

With emphasis on performing the following tasks:

- a. Describe the advantages of JBB.
 - 1) JBB can help to significantly reduce the time to compare the changes in the file system.
- b. Set up the journal engine.
 - 1) The journal engine has to be set up by a wizard on Windows or by configuring and starting of the journal daemon in AIX.
- c. Dependencies for JBB
 - 1) If there is heavy load in the file system, the journal service needs to be supported by a frequently full incremental backup without journal (no journal) to recreate the Journal database.
 - 2) During its lifecycle the journal service needs additional CPU power and storage space for the journal database.
- d. Describe how the journal backup works and what restrictions exist.
 - 1) During the incremental backup the journal database was queried about the changes in the file system and the comparison process will be eliminated.
 - 2) The idle time of the Backup-Archive client will be significantly reduced.
 - 3) JBB only works on AIX, Linux ,and Windows, and is limited to local file systems.

D. Given knowledge of Tivoli Storage Manager and storage Area Network Infrastructure and the basics of LAN-free data movement, explain the principles of LAN-free and LAN-based backup so that the configuration and principles of operation is understood.

With emphasis on performing the following tasks:

- a. Explain the principles of operation for LAN-free data movement.
- b. Explain how to install and customize the storage agent.
- c. Understand the intercommunication between Tivoli Storage Manager server, storage agent and Tivoli Storage Manager client.

❖ **SECTION 5: IBM TIVOLI STORAGE MANAGER CLIENT IMPLEMENTATION:**

- A. Given a correctly set up system and network which meet the requirements, and access to the IBM Tivoli Storage Manager (Tivoli Storage Manager) code, determine whether the installation is new and follow the documented installation steps for any supported operating system so that the Tivoli Storage Manager client is installed on the system.**

With emphasis on performing the following tasks:

- a. Determine the OS platform.
- b. Determine if this is a new installation or an upgrade.
- c. Read appropriate documentation and Support Flashes on the IBM Tivoli support Website.
- d. Identify location of most current Tivoli Storage Manager code, patches, and fixes.
- e. Select the necessary product parts and features.
- f. Install the Tivoli Storage Manager client code to any supported operating system from the appropriate media /CD or electronic image.
- g. Install the language packs from the appropriate media (CD or electronic image).
- h. Verify the code installed correctly.

- B. Given the need to register a client to Tivoli Storage Manager, log on to Tivoli Storage Manager as admin, determine policy information based on SLA, run commands to register node, schedule the client for automatic backup so that a user can perform backups and restore the user's data.**

With emphasis on performing the following tasks:

- a. Open admin CLI.
- b. At the CLI prompt; log in to Tivoli Storage Manager with admin account.
- c. Perform commands to register node.
- d. Authenticate node with Tivoli Storage Manager server.
- e. Associate the node with a client schedule.
- f. Verify the scheduler is started.
- g. Start the Tivoli Storage Manager Backup-Archive GUI to perform backup (if necessary).

- C. Given the need to perform a client backup manually as opposed to automated scheduling, log on to the Backup-Archive GUI client, select the backup icon, determine files/folders identified requiring backup so that the files identified as needing a manual back up have been copied to the Tivoli Storage Manager server.**

With emphasis on performing the following tasks:

- a. Open the Backup-Archive GUI client.
- b. Prepare for backup by selecting the backup icon.
- c. Identify and select files identified requiring backup.
- d. Perform backup of selected files.
- e. Verify files have been successfully backed up.

- D. Given the need to perform a client restore: log on to the B/A GUI client, select the restore icon, select versioning criteria, identify files/folders that are needing to be restored, so that files saved on the Tivoli Storage Manager server have been copied back to the local system.**

With emphasis on performing the following tasks:

- a. Open the Backup-Archive GUI client.
- b. Prepare for restore by selecting the Restore icon.
- c. Select versioning options.
- d. Identify and select files identified as requiring restore.
- e. Perform restore of selected files.

f. Verify files have been successfully restored on the system.

- E. Given the need to perform a client restore manually to a different location on the system: log on to the Backup- Archive GUI client, select the restore icon, select versioning criteria, determine files/folders that are identified as needing restore, so that files saved on the Tivoli Storage Manager server have been copied back onto the local system to an alternate location.**

With emphasis on performing the following tasks:

- a. Open the Backup-Archive GUI client.
- b. Prepare for restore by selecting the restore icon.
- c. Select versioning criteria.
- d. Identify and select files identified as requiring restore.
- e. Select the new location.
- f. Perform restore of selected files.
- g. Verify files have been successfully restored to desired location.

- F. Given a working Tivoli Storage Manager environment where one node restores the data backed up from another node, create a node access so that one node now has been granted the authority to restore files for another specified node.**

With emphasis on performing the following tasks:

- a. Create a node access list to grant a given node the ability to restore ALL data backed up from another node.
 - 1) Initiate a client through:
 - CLI
 - GUI
 - Web client
 - 2) Determine objects to be restored.
 - 3) Perform restore.
 - 4) Monitor progress and confirm completion.

- G. Given the need to verify a client backup completed: log on to Tivoli Storage Manager, query the events to determine the status of the backup, determine course of action based on results so that the status of the backup has been identified.**

With emphasis on performing the following tasks:

- a. Open admin CLI.
- b. At the CLI prompt; log in to Tivoli Storage Manager with admin account.
- c. Perform query to determine the status of the backup.
- d. Determine next steps based on the output. ie failure, success, missed.
- e. Query the Tivoli Storage Manager activity log to gather further information if necessary.
- f. Recommend course of action based on findings.

- H. Given a properly installed and customized Tivoli Storage Manager client environment and an error has occurred, review the various logs so that the Tivoli Storage Manager client issue has been identified.**

With emphasis on performing the following tasks:

- a. Check the client error log for error messages.
- b. Check the schedule log and webcl log for additional messages.
- c. Check the Tivoli Storage Manager server activity log for further troubleshooting.
- d. Search for Error conditions on the operating system level.
- e. Check the error description in the Tivoli Storage Manager documentation for clarification.
- f. Search for assistance at IBM Support Websites.

- I. Given a properly set up scheduling for backup on client and server side, set up and review the schedule and event logs, so that it has been confirmed that the client backups are running as scheduled.**

With emphasis on performing the following tasks:

- a. Set up client schedules.
- b. Verify the schedule log on client side for successful operation.
- c. Check the event log on server side for schedule event conditions and statistics.
- d. Check the date and time when the schedule was planned for execution and compare it with the actual values.

❖ **SECTION 6: IBM TIVOLI STORAGE MANAGER ADMINISTRATION:**

A. Given a working IBM Tivoli Storage Manager (Tivoli Storage Manager) V6.3 system with the latest PVU table in the Tivoli Storage Manager server instance directory, perform estimation of processor value units so that the license model is validated.

With emphasis on performing the following tasks:

- a. Use the QUERY PVUESTIMATE command to generate reports that estimate the number of server devices and client devices managed by the Tivoli Storage Manager server.
- b. Adjust the role classification in the nodes definition.
- c. For detailed report issue the SQL SELECT * from PVUESTIMATE_DETAILS.

