Enterprise Backup and Restore technology and solutions

LESSON V

HP Data Protector
Overview and Concepts
Part I

Veselin Petrunov
Backup and Restore team / Deep Technical Support

HP Bulgaria Global Delivery Hub
Global Operations Center
November, 2013

HP Data Protector

Overview and Concepts

Part I

Veselin Petrunov / November 2013

Course overview

Part I

1. About Data Protector

Part II

2. Media management and devices
3. Security / Users and groups

Part III

3. The Data Protector internal database
4. Service management

Part IV

7. How Data Protector operates
8. Integration with applications

Part V

7. Disk backup
8. Synthetic backup
9. Split mirror
10. Snapshot concepts
11. Microsoft VSS Service
12. Deduplication
12. VM environments integration
13. Backup scenarios (labs)
14. Further information

Contents

Part I

HP Data Protector

Overview

- Scalable and Highly Flexible Architecture
- Easy Central Administration
- High Performance Backup
- Data security
- Supporting Mixed Environments
- Easy Installation for Mixed Environments
- High Availability Support
  - Integrating with clusters to ensure fail-safe operation with the ability to back up virtual nodes
  - Enabling the Data Protector Cell Manager itself to run on a cluster
  - Supporting all popular online database Application Programming Interfaces
  - Integrating with advanced high-availability solutions (HP, EMC, etc.)
  - Providing various disaster recovery methods for Windows and UNIX platforms
  - Offering methods of duplicating backed up data during and after the backup
- Backup Object Operations
- Easy Restore
- Automated or Unattended Operation
- Service Management
- Monitoring, Reporting and Notification
- Integration with Online Applications
- Integration with Other Products

HP Data Protector

Overview

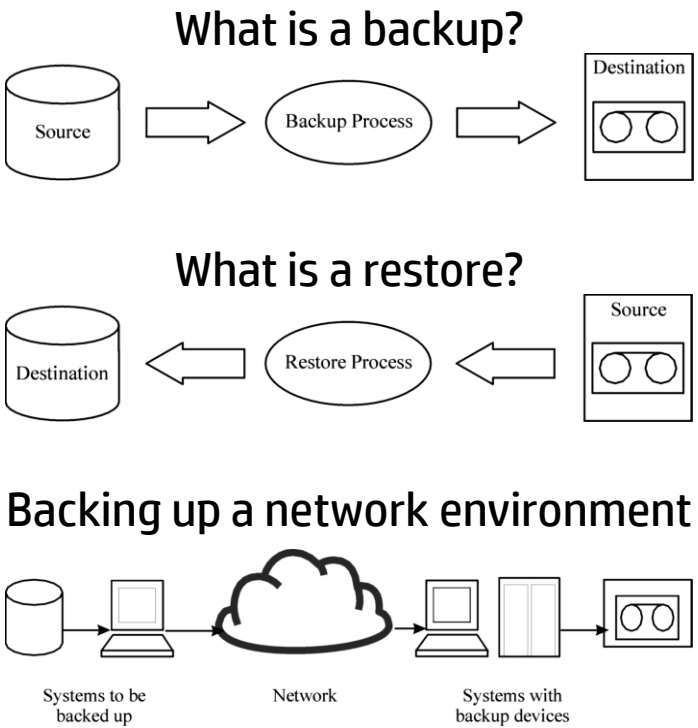

HP Data Protector

Data Protector architecture

Data Protector cell (physical view and logical view)

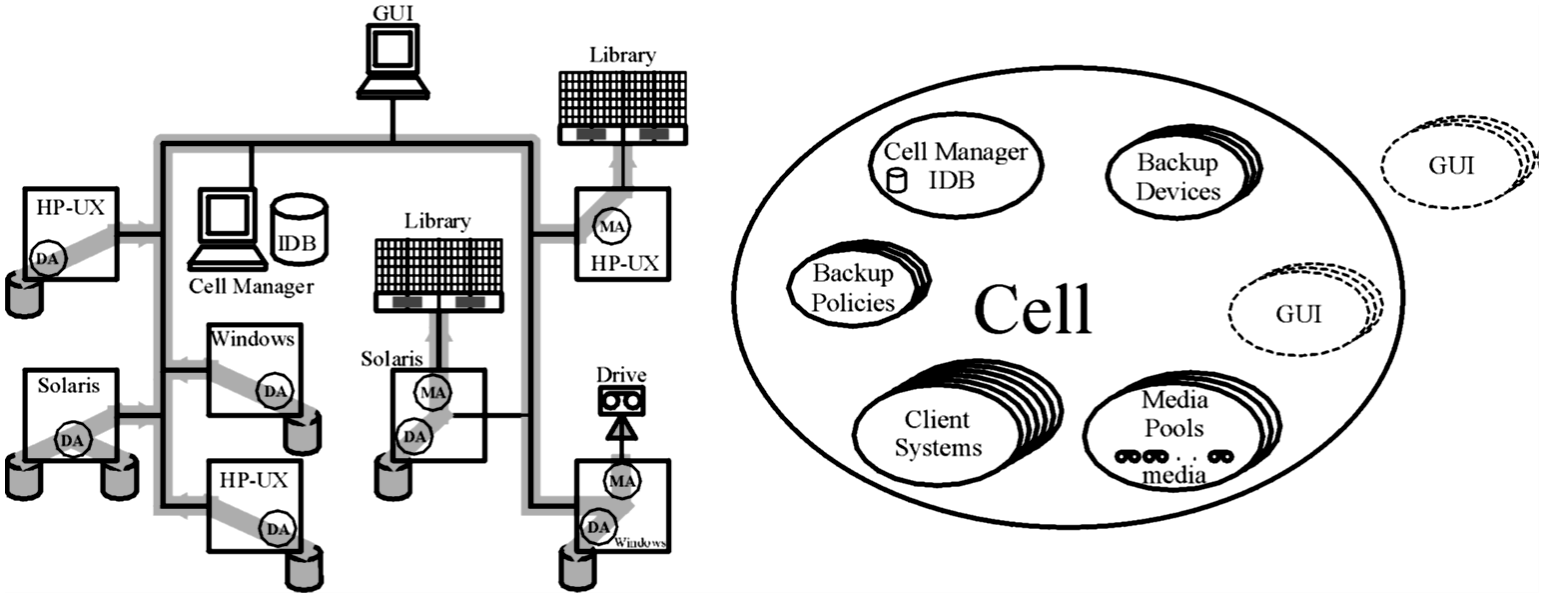

HP Data Protector

Data Protector architecture

Cell Manager

- Manages the cell from a central point
- Contains the IDB (The IDB contains information about backup details such as, backup durations, media IDs, and session IDs)
- Runs core Data Protector software
- Runs Session Managers that start and stop backup and restore sessions and write session information to the IDB

HP Data Protector

Data Protector architecture

Other cell components

- Clients (Systems to be backed up)
- Media servers (Systems with backup devices)
- Systems with a user interface
- Installation Server

HP Data Protector

Operations in the cell

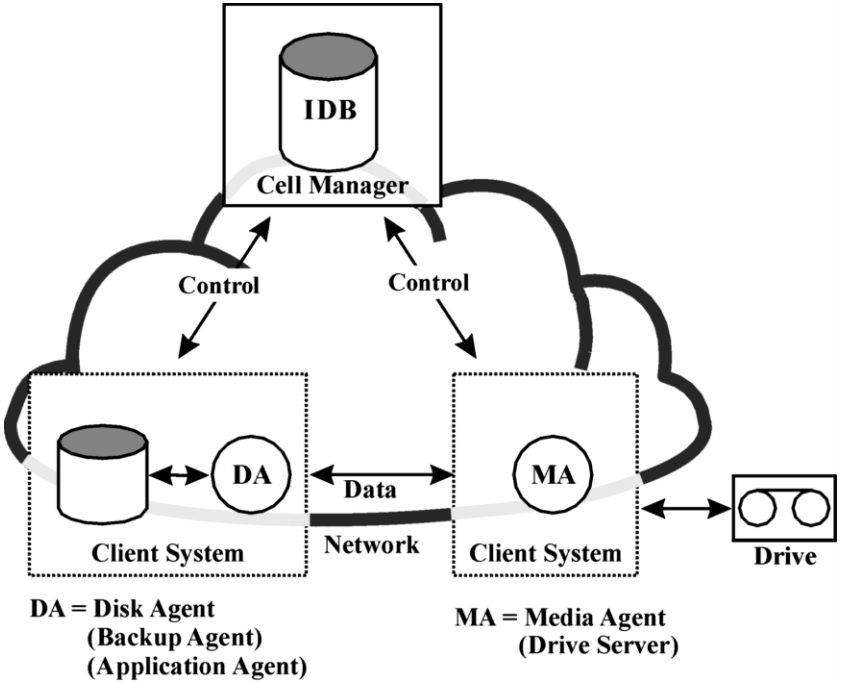

HP Data Protector

Backup and Restore sessions

Backup sessions

Restore sessions

HP Data Protector

Enterprise environments

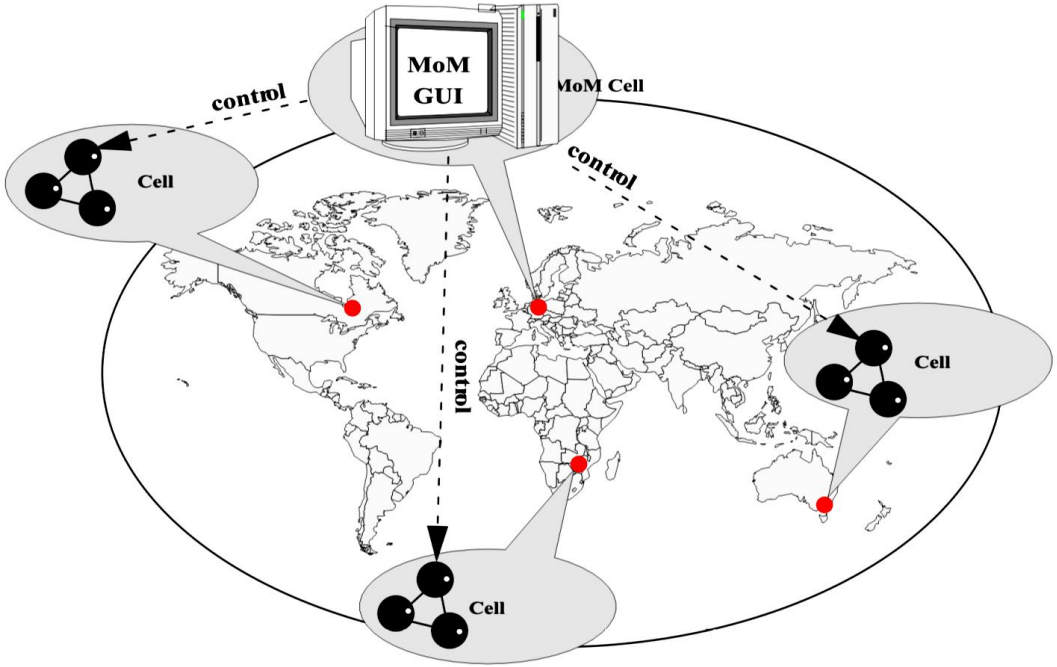

HP Data Protector

Splitting an environment into multiple cells

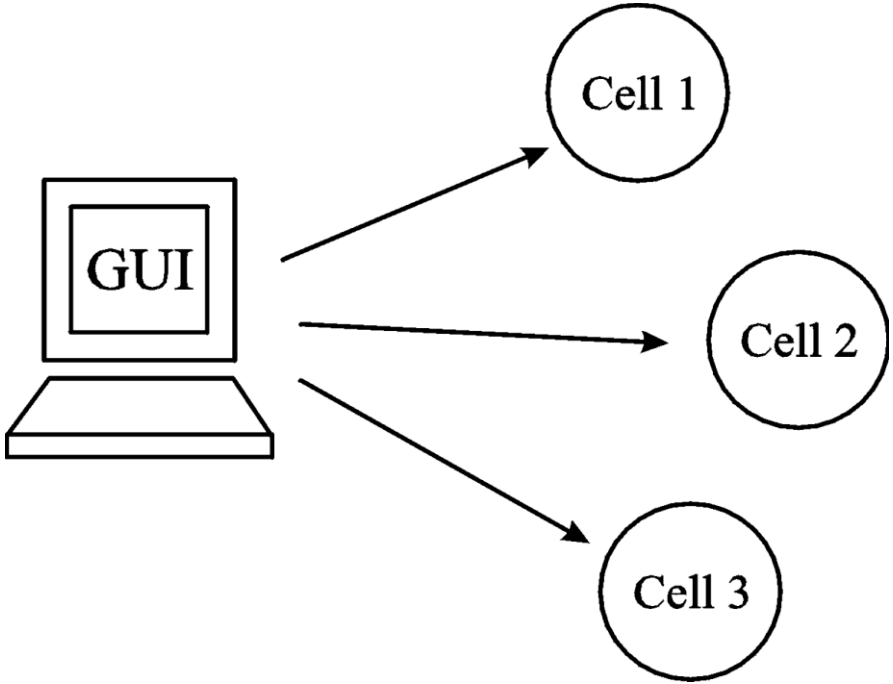

HP Data Protector

MoM - Manager-of-Managers

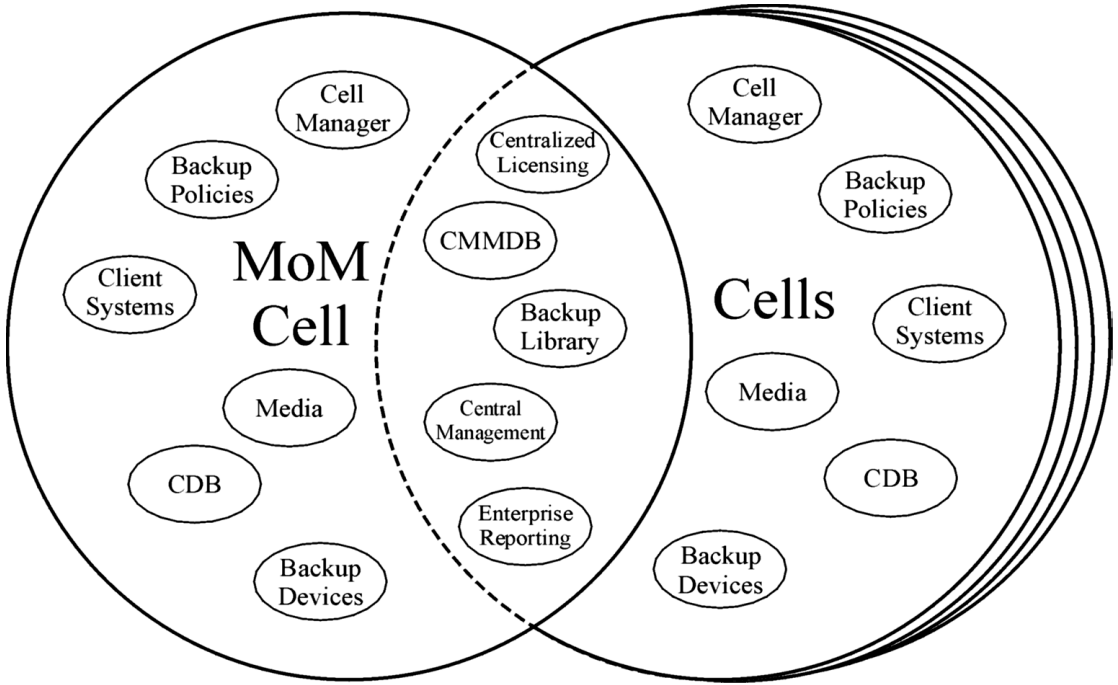

HP Data Protector

Media management

- Grouping media into logical groups, called **media pools**, which allows you to think about large sets of media without having to worry about each medium individually.
- Data Protector keeps track of all media and the status of each medium, data protection expiration time, availability of media for backup, and a catalog of what has been backed up to each medium.
- Fully automated operation. If Data Protector controls enough media in the library devices, the media management functionality lets you run the backup sessions without operator intervention.
- Automated media rotation policies that allow media selection for backups to be performed automatically.
- Recognition and support of barcodes on large library devices and silo devices with barcode support.
- Recognition, tracking, viewing, and handling of media used by Data Protector in large library devices and silo devices.
- The possibility of having information about the media in a central place and the sharing of this information among several Data Protector cells.
- Interactive or automated creation of additional copies of the data on the media.
- Support for media vaulting.

HP Data Protector

Backup devices

How backup specifications, devices, and media pools are related:

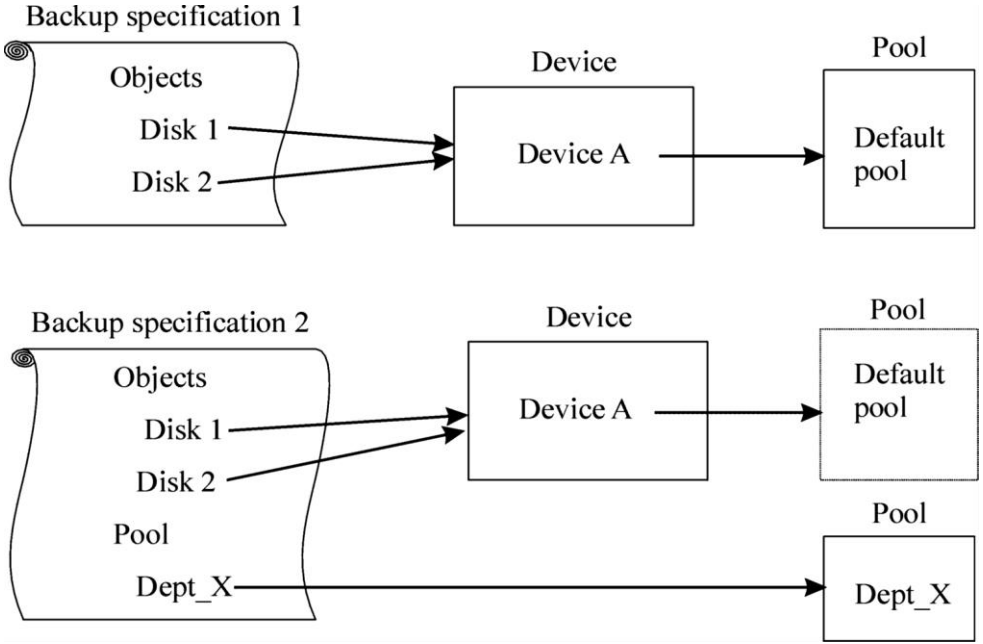

HP Data Protector

User interfaces

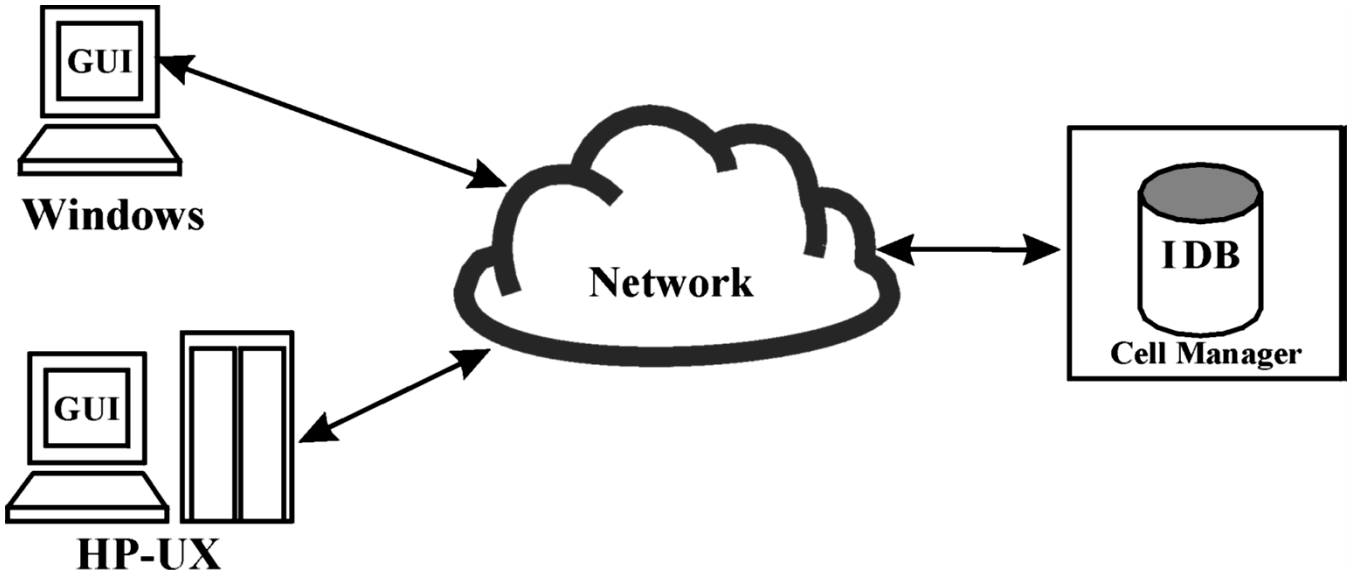

HP Data Protector

Data Protector GUI

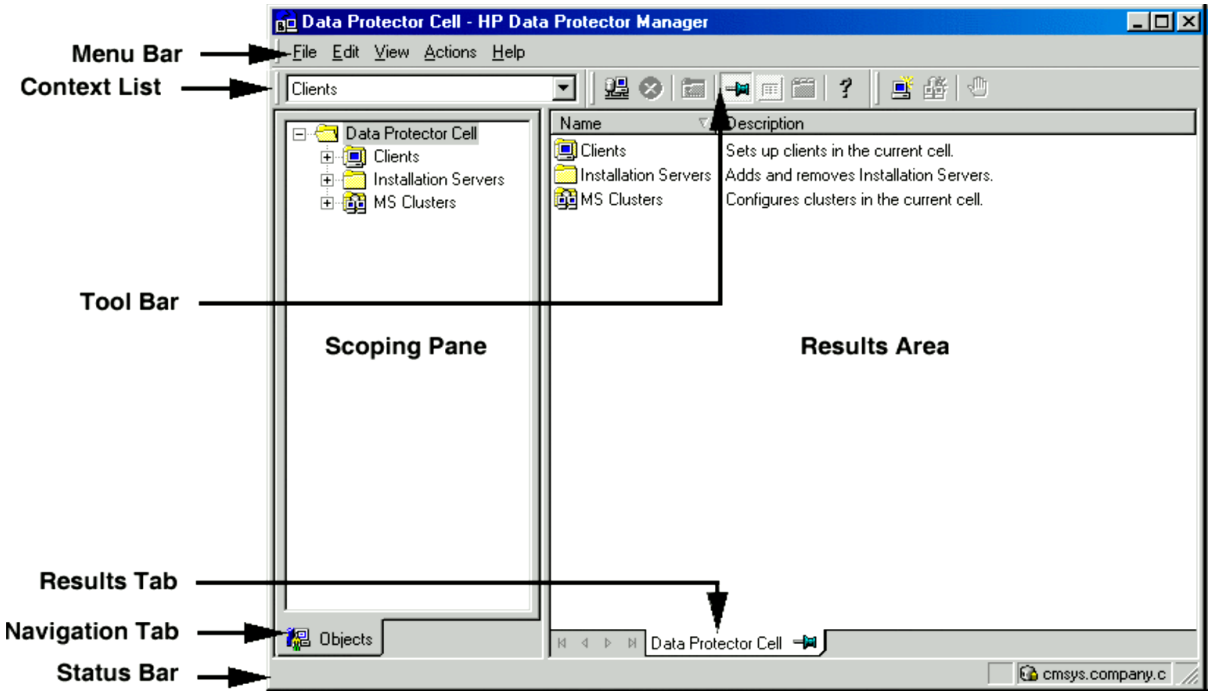

HP Data Protector

Overview of tasks to set up Data Protector

1. Analyze your network and organizational structure. Decide which systems need to be backed up.
2. Check if there are any special applications and databases which you want to back up, such as Microsoft Exchange, Oracle, IBM DB2 UDB, SAP R/3, or others. Data Protector provides specific integrations with these products.
3. Decide on the configuration of your Data Protector cell.
4. Purchase the required Data Protector licenses for your setup. This way you obtain the passwords you will need to install.
5. Consider security aspects.
6. Decide how you want to structure your backups.
7. Install and configure your Data Protector environment.
8. Become familiar with tasks such as Handling failed backups, Performing restores, Maintaining the IDB, etc.

HP Data Protector

Creating cells

1. Creating cells in the UNIX environment

2. Creating cells in the Windows environment
   - Windows domains/Mapping a Data Protector cell into a Windows domain
   - Windows workgroups

3. Creating cells in a mixed environment

HP Data Protector

Clustering

Cluster concepts

What is a cluster?

- Ensure that mission-critical applications and resources are as highly-available as possible
- Tolerate component failures
- Support either the addition or subtraction of components

Components:

- Cluster nodes (two or more)
- Local disks
- Shared disks (shared between nodes)

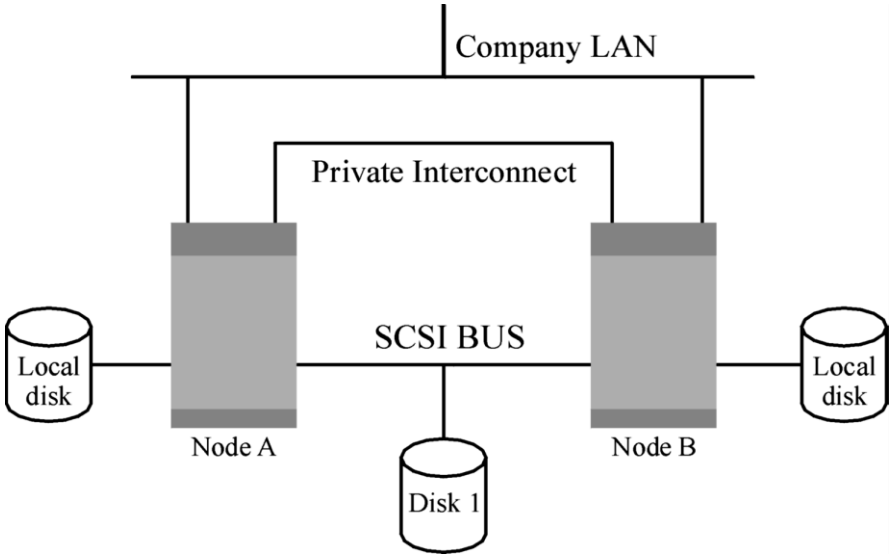

HP Data Protector

Cluster support

The Data Protector cluster support means the following:

- The Data Protector Cell Manager is installed in a cluster. Such a Cell Manager is fault tolerant and can restart operations in the cell automatically after the failover.
- The Data Protector client is installed in a cluster. The Cell Manager (if not installed in the cluster) in such a case is not fault tolerant; the operations in the cell must be restarted manually.

The behavior of the Cell Manager after the failover is configurable as far as the backup session (failed due to the failover) is concerned - the failed session can be:

- restarted as a whole
- restarted only for the failed objects
- not restarted at all

HP Data Protector

Example cluster environments

Cell Manager installed outside a cluster

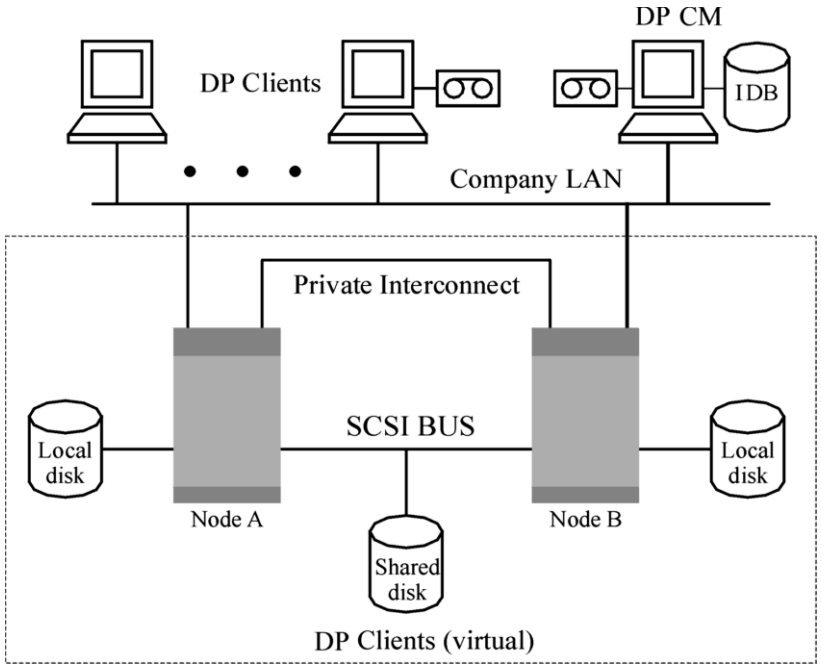

Backup behavior

Condition	Result
Failover of the node before a backup starts	Successful backup
Failover of the node during backup activity	**Filesystem/disk image backup**: The backup session fails. The completed objects from the session can be used for restore, the failed (running and pending) objects need to be backed up again by restarting the session manually.
	Application backup: The backup session fails. The session needs to be restarted manually.

HP Data Protector

Example cluster environments

Cell Manager installed outside a cluster, devices connected to the cluster nodes

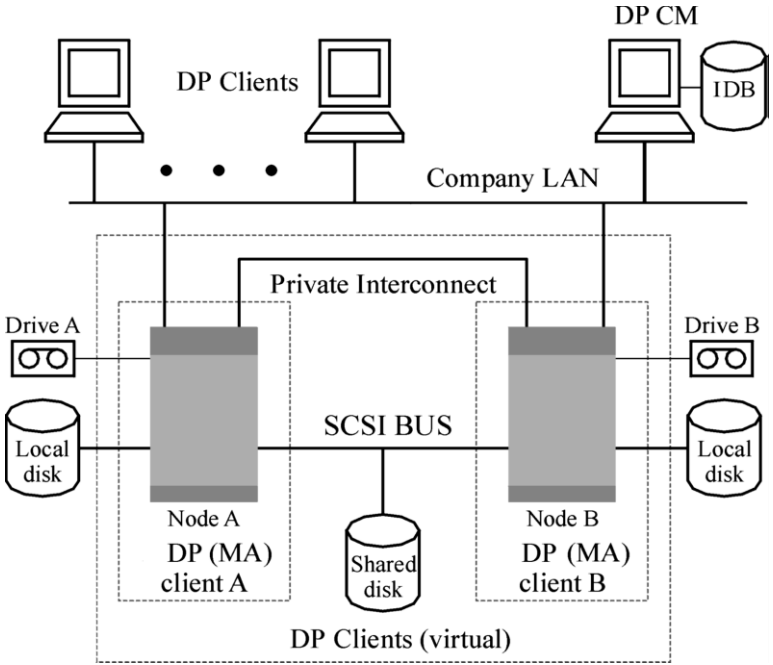

Backup behavior

Condition	Result
Failover of the node before a backup starts	Successful backup due to automatic device switching (load balancing)
Failover of the node during backup activity	**Filesystem/disk image backup**: The backup session fails. The completed objects from the session can be used for restore, the failed (running and pending) objects need to be backed up again by restarting the session manually.
	Application backup: The backup session fails. The session needs to be restarted manually.

HP Data Protector

Example cluster environments

Cell Manager installed in a cluster, devices connected to the cluster nodes

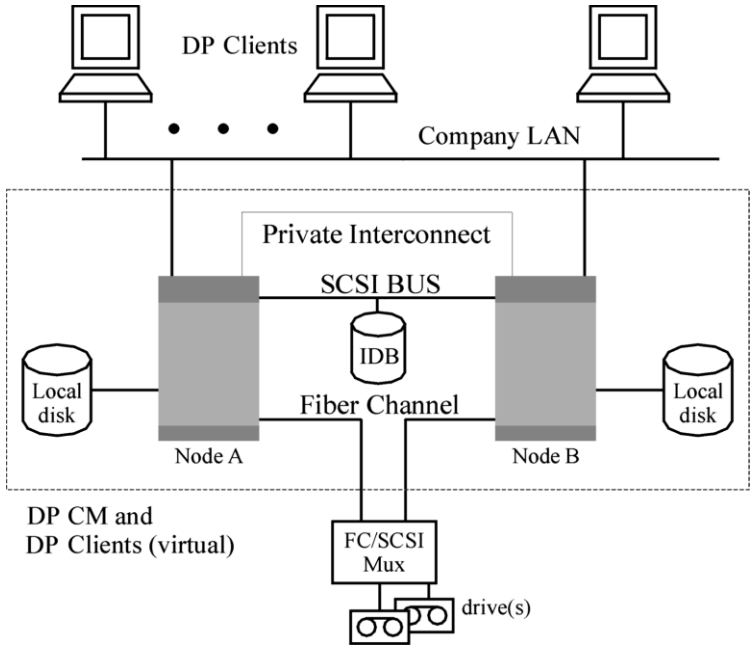

HP Data Protector

Example cluster environments

Cell Manager installed in a cluster, devices connected to the cluster nodes - **Backup behavior**

Condition	Result	
Failover of the node before a backup starts	Successful backup	
Failover of the application and the Cell Manager during backup activity (Cell Manager runs on the same node as the application).	**Filesystem/disk image backup** The backup session fails. The completed objects from the session can be used for restore, the failed (running and pending) objects are backed up again by restarting the session automatically.	**IMPORTANT** To restart the session, the appropriate Data Protector option must be selected. For information on defining all possible Data Protector actions in case of failover of the Cell Manager, see the online Help index: "cluster, managing backups".
	Application backup The backup session fails. The session is restarted automatically.	
Failover of the application during backup activity without Cell Manager failover (Cell Manager runs on other node than the application).	**Filesystem/disk image backup** The backup session fails at failover of the node where the filesystem is installed. The completed objects from the session can be used for restore, the failed (running and pending) objects need to be backed up again by restarting the session manually.	
	Application backup The backup session fails. The session needs to be restarted manually.	

HP Data Protector

Full and incremental backups

- Full backups
- Incremental backups
- Synthetic backup
- Conventional incremental backup
- Enhanced incremental backup
- Incremental backup using Change Log Provider

	Full backup	Incremental backup
Resources	Takes more time to complete than incremental backup and requires more media space.	Backs up only changes made since a previous backup, which requires less time and media space.
Device handling	If you use a standalone device with a single drive, you need to change the media manually if a backup does not fit on a single medium.	It is less likely that the backup will require additional media.
Restore	Enables simple and quick restore.	A restore takes more time because of the number of media needed.
IDB impact	Occupies more space in the IDB.	Occupies less space in the IDB.

HP Data Protector

Types of incremental backups

Incremental backups

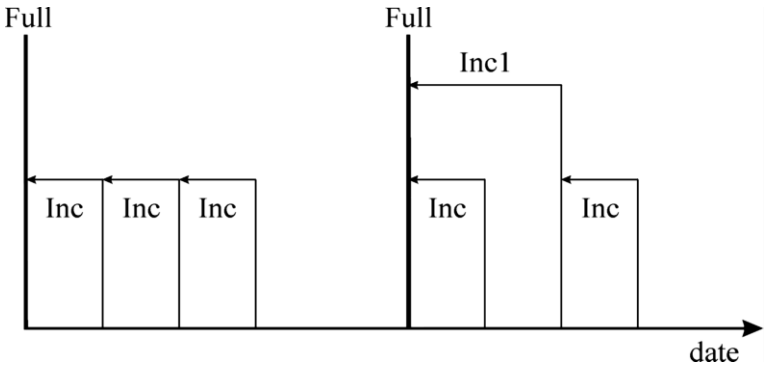

Leveled incremental backups

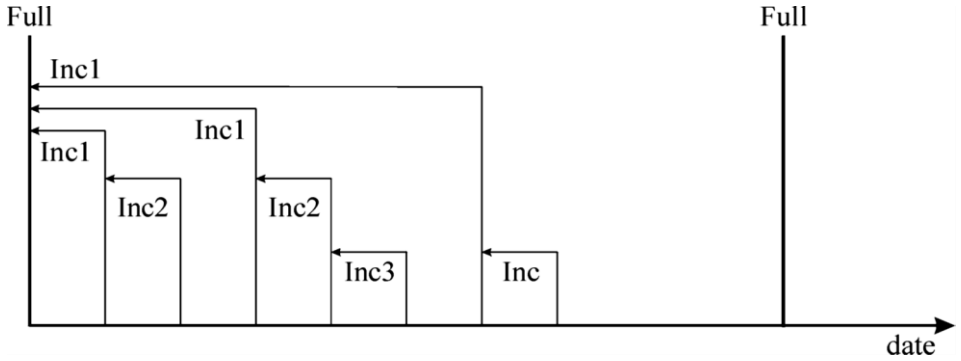

HP Data Protector

Considering restore

Media needed to restore from simple and leveled incremental backups

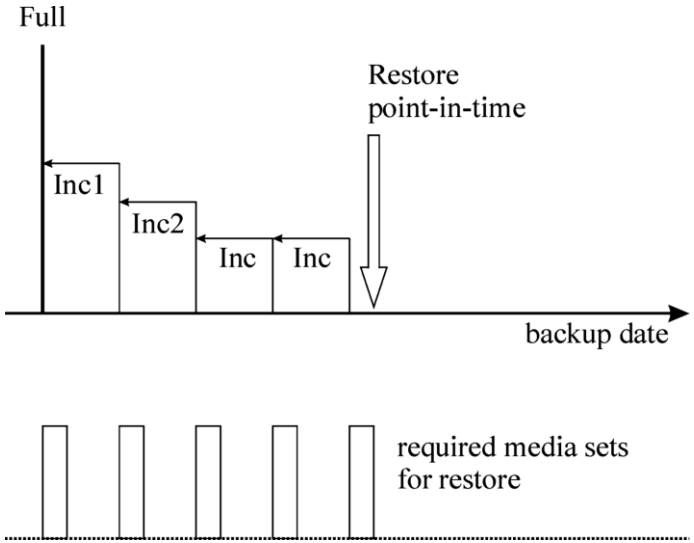

Media needed to restore from leveled incremental backups

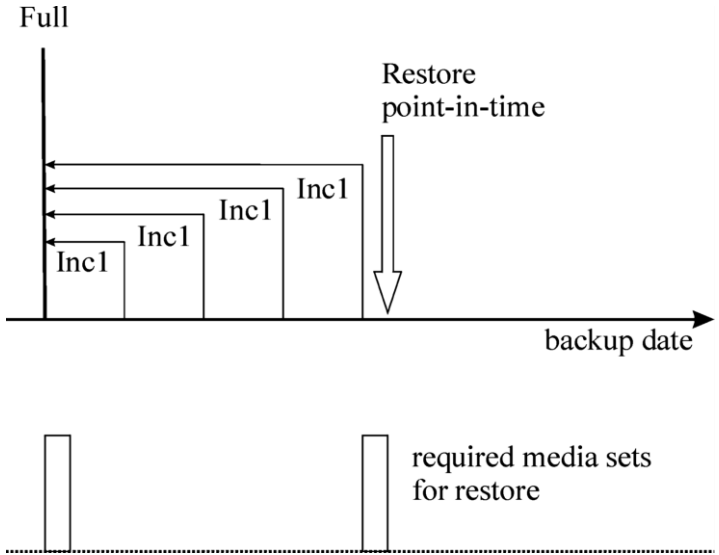

HP Data Protector

Keeping backed up data and information about the data

- Data Protector Internal Database
- Data protection
- Catalog protection
- Logging level

Browsing files for restore

- Enabling the browsing of files and quick restore
- Enabling the restore of files, but not browsing
- Overwriting backed up files with new data
- Exporting media from a cell

HP Data Protector

Backing up data

- Selecting what to back up, from which client system - the source of data.
- Selecting where to back up - the destination.
- Selecting to write the same data to additional media sets - mirroring.
- Selecting how to back up - backup options.
- Scheduling a backup for automated operation.

HP Data Protector

Backup specification

Creating a backup specification

- What is a backup specification?
- How to create a backup specification

Selecting backup objects. What is a backup object?

- Client name: a hostname of the Data Protector client where the backup object resides.
- Mount point: an access point in a directory structure (drive on Windows and mount point on UNIX) on the client where the backup object is located.
- Description: uniquely defines the backup objects with identical client name and mount point.
- Type: backup object type, for example filesystem or Oracle.

HP Data Protector

Backup session

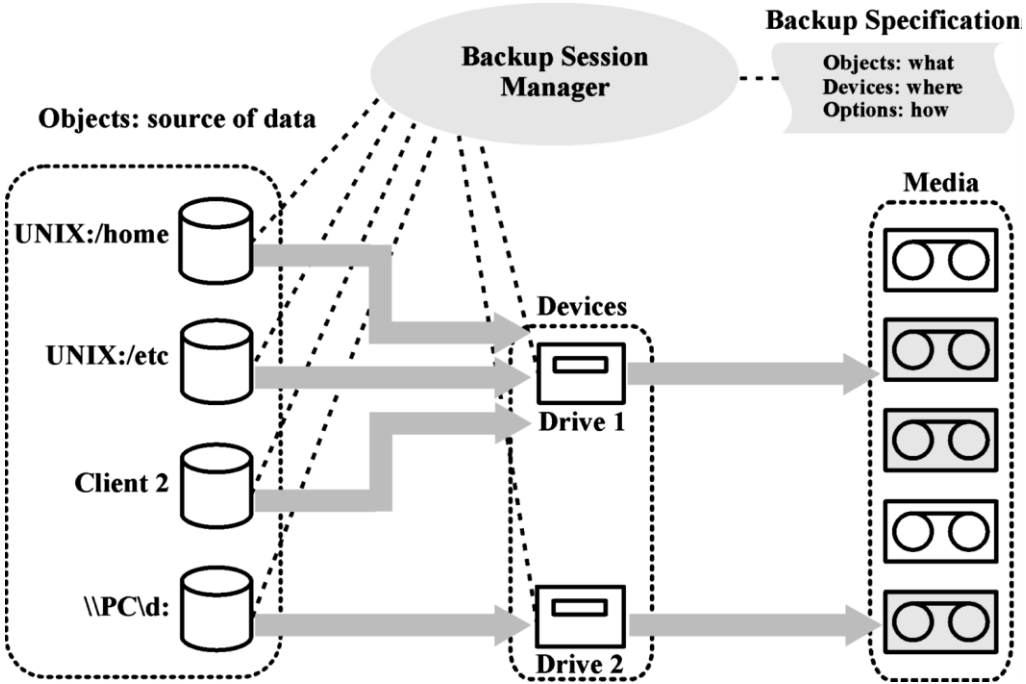

HP Data Protector

Scheduling, backup configurations and sessions

- Backup configuration
- Backup session
- Optimizing backup performance

Scheduling tips

- When to schedule backups
- Staggering full backups:

	Mon	Tue	Wed	…
system_grp_a	Full	Incr1	Incr1	…
system_grp_b	Incr1	Full	Incr1	…
system_grp_c	Incr1	Incr1	Full	…

HP Data Protector

Scheduling tips

- Optimizing for restore

Full backup with daily simple incremental backups

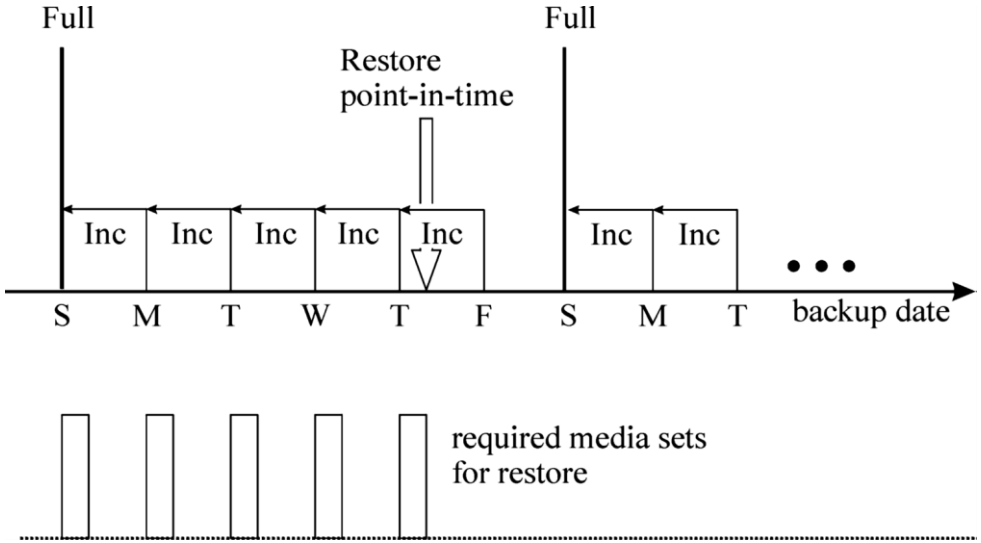

HP Data Protector

Scheduling tips

- Optimizing for restore

Full backup with daily level 1 incremental backups

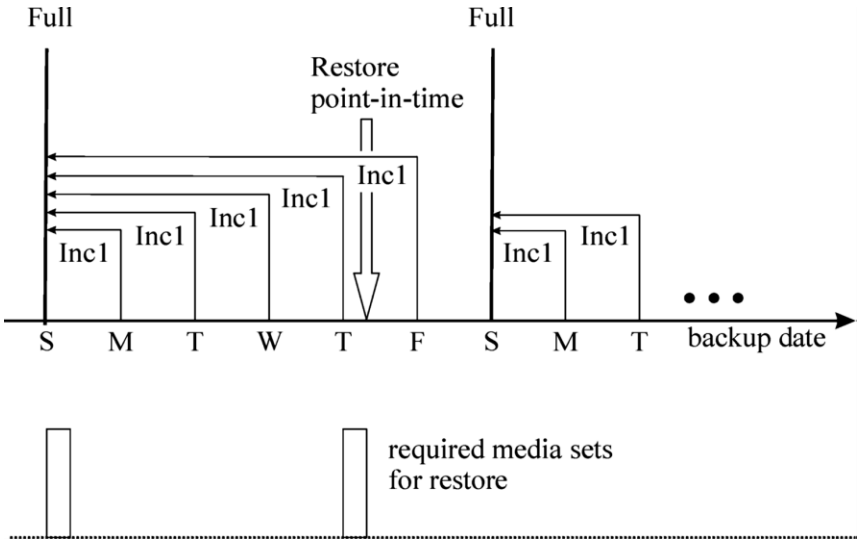

HP Data Protector

Scheduling tips and tricks

- Optimizing for restore

Full backup with mixed incremental backups

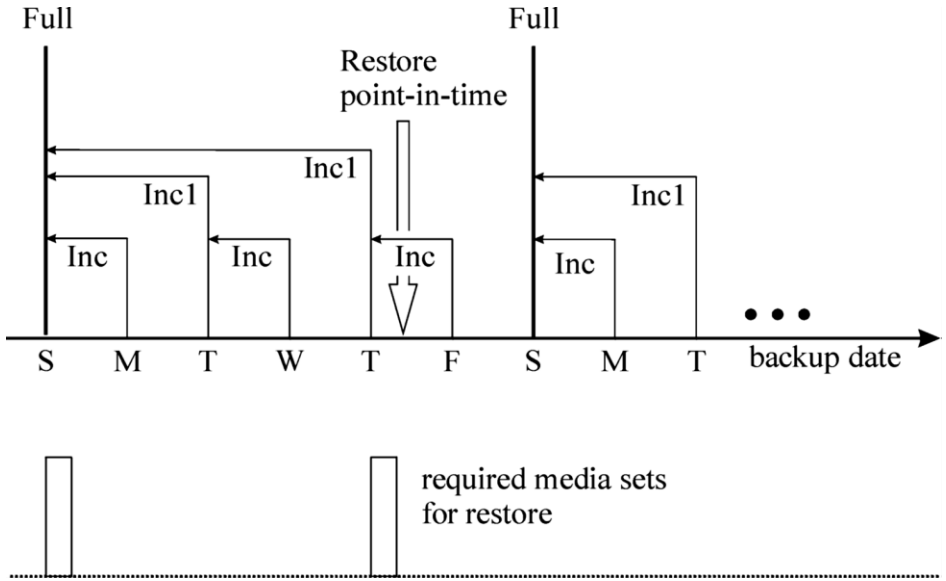

HP Data Protector

Duplicating backed up data – Data Protector data duplication methods

Object copy	Object copy	Object mirror	Media copy	Smart Media Copy
What is duplicated	Any combination of object versions from one or several backup, object copy, or object consolidation sessions	A set of objects from a backup session	An entire medium	An entire medium
Time of duplication	Any time after the completion of a backup	During backup	Any time after the completion of a backup	Any time after the completion of a backup
Media type of source and target media	Can be different	Can be different	Must be the same	Are different as disk-based storage is combined with tape-based storage
Size of source and target media	Can be different	Can be different	Must be the same	Must be the same
Appendability of target media	Yes	Yes	No	No
Result of the operation	Media containing the selected object versions	Media containing the selected object versions	Media identical to the source media	Media identical to the source media

HP Data Protector

Copying objects

What is object copy?

Start of object copy session

- Post-backup object copying
- Scheduled object copying

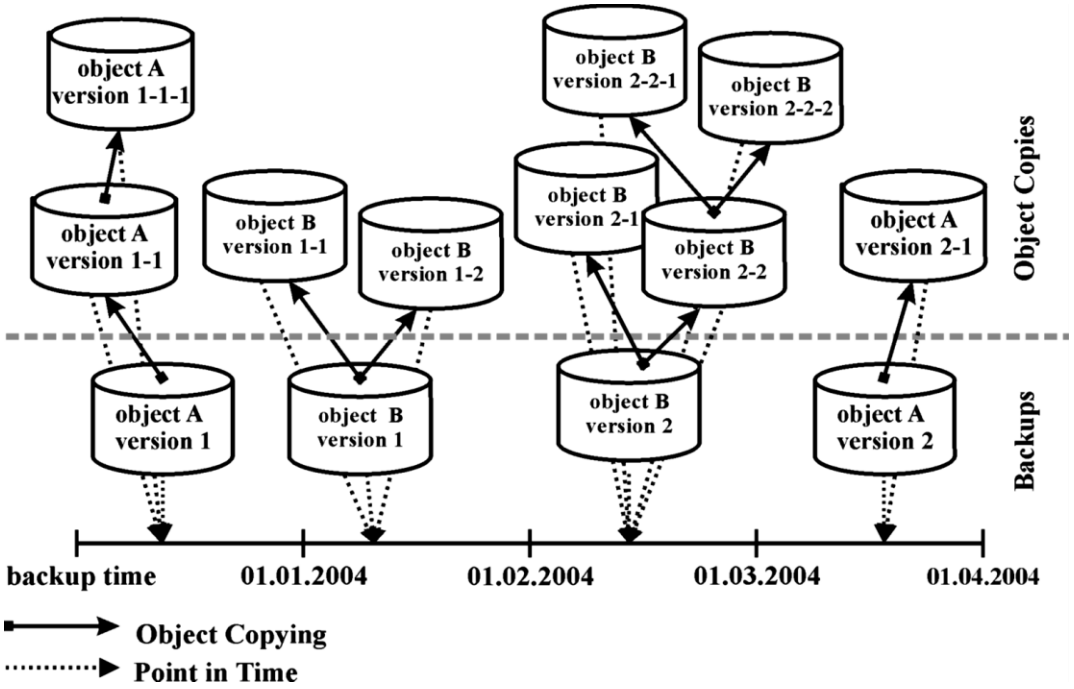

HP Data Protector

Copying objects

Why use object copy?

- Vaulting
- Freeing media
- Demultiplexing of media
- Consolidating a restore chain
- Migration to another media type
- Support of advanced backup concepts

HP Data Protector

Copying objects

Freeing media

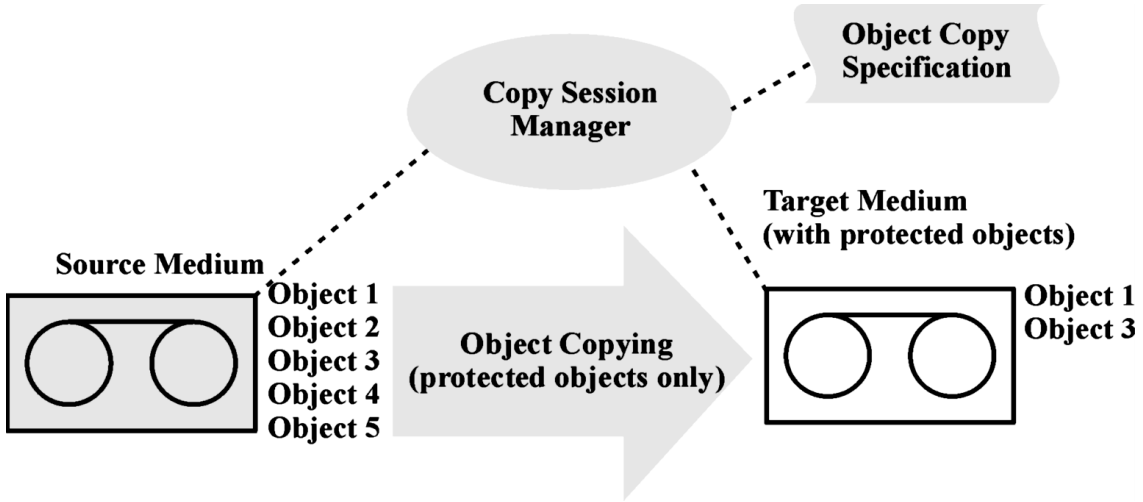

HP Data Protector

Copying objects

Demultiplexing of media

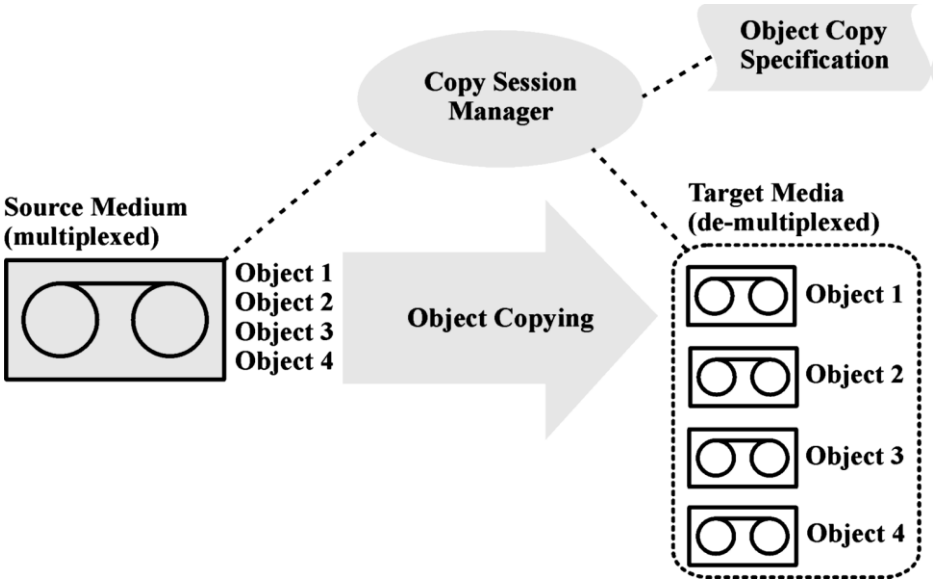

HP Data Protector

Copying objects

Disk staging concept

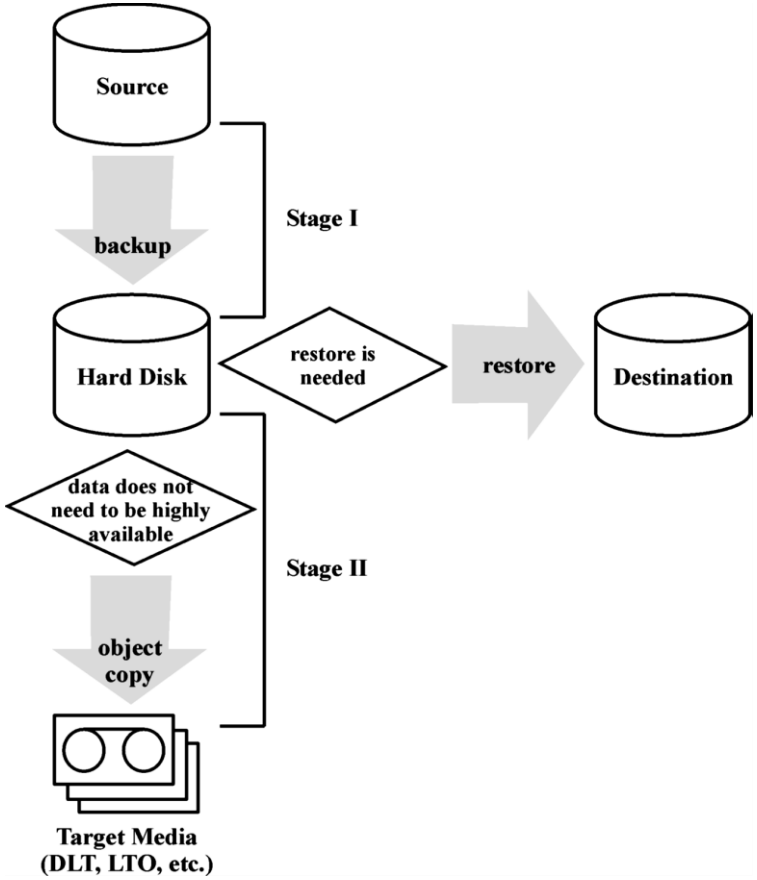

HP Data Protector

Object mirroring

What is object mirroring?

Benefits of object mirroring

- It increases the availability of backed up data due to the existence of multiple copies.
- It enables easy multi-site vaulting, as the backed up data can be mirrored to remote sites.
- It improves the fault tolerance of backups, as the same data is written to several media. A media failure on one medium does not affect the creation of the other mirrors.

HP Data Protector

Object mirroring

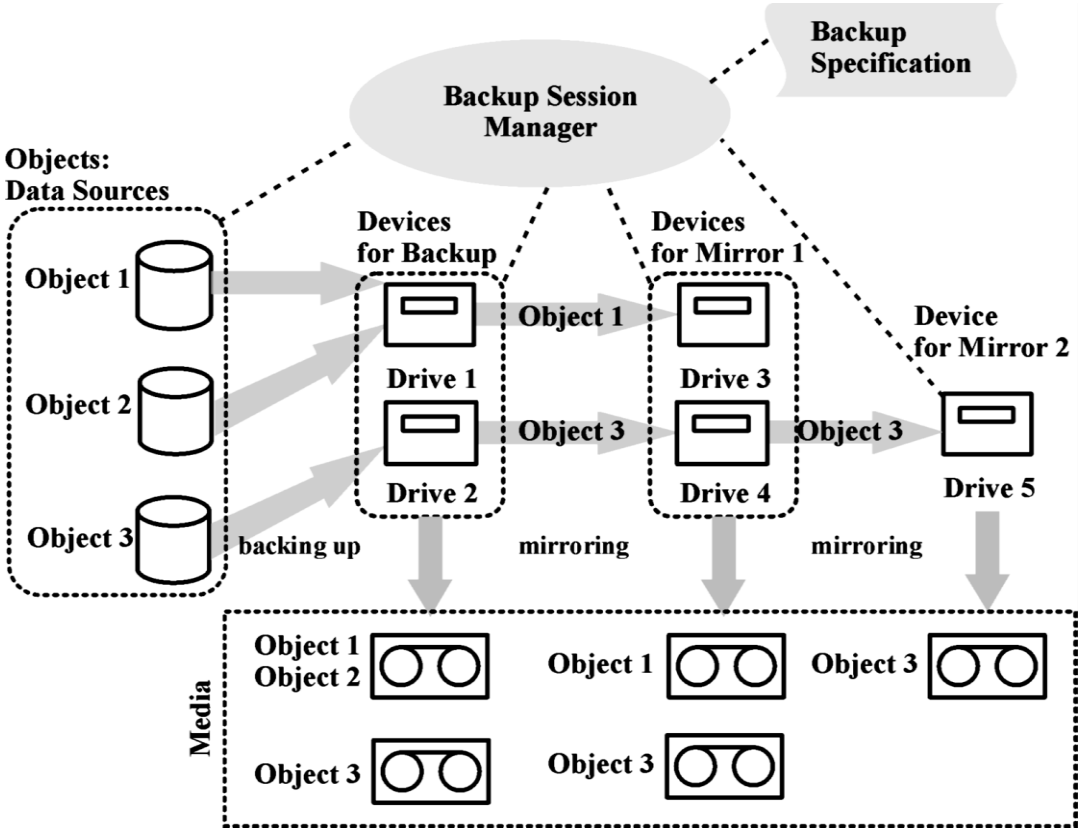

HP Data Protector

Copying media

What is media copying?

How to copy media?

What is the result?

Automated media copying

Smart media copying using VLS

HP Data Protector

Verifying backup media and backup objects

What is media verification?

- You've copied a medium for archive purposes and you want to check the validity of the copy before placing it in a vault.
- A backup medium has become full and you would like to check all the objects on it before sending for long-term storage.

HP Data Protector

Verifying backup media and backup objects

What does media verification do for you?

- Checks the media identification, description, and location information in the Data Protector headers
- Reads all blocks on the medium and verifies block format
- If a cyclic redundancy check (CRC) was performed during backup, recalculates the CRC and compares it with the one stored on the medium

HP Data Protector

Verifying backup media and backup objects

What is object verification?

Data Protector object verification allows you to check the validity of backup objects, as opposed to backup media. You can use it to check:

- single or multiple objects
- on single or multiple media
- interactively, or in scheduled or post-operation sessions

You might want to use object verification:

- after an object copy to a different medium
- after performing object consolidation on the restore chain of an object backed up incrementally
- to check all backup objects produced within a specified time-frame after a backup device change

HP Data Protector

Verifying backup media and backup objects

What does object verification do for you?

- a single backup object, without having to check the complete medium, potentially saving a
- lot of time with large backup media
- large objects that span more than one medium
- several objects on several media
- a specific object version (interactive only)
- the media agent host, avoiding any network traffic
- another host, factoring in network effects

HP Data Protector

Restoring data

Policies for restoring data are an essential part of the overall backup strategy in the company. Keep the following in mind:

- Backing up and restoring files is essentially the same as copying files. Therefore, ensure that only authorized people have the rights to restore confidential data.
- Ensure that unauthorized people cannot restore files of other people.

This section describes some possible implementations of the restore policy using Data Protector. You can restore your filesystem data by browsing through restore objects or restore sessions. By default, data is restored to its original location. However, you can specify any location to be the destination of restored data.

HP Data Protector

Restoring data

The restore duration depends on a number of factors, such as:

- The amount of data to be restored. This also directly influences all the following items.
- A combination of full and incremental backups.
- Media and devices used for backup.
- Speed of networks and systems.
- The application you are recovering, for example, Oracle database files.
- The use of parallel restore. Several objects can be restored with a single read operation, depending on how the data was backed up.
- Speed and ease of selecting the data to be restored, which depends on the logging level settings used during the backup and on catalog protection time.

HP Data Protector

Restoring data

Selection of the media set:

- Media set selection algorithm
- Selection of restore chain
- Media location priority

Selection of devices:

- Automatic device selection
- Original device selection

Operators are allowed to restore?

Use this policy in the following cases:

- In a large network environment where it is best to have a dedicated person to do such jobs.
- In an environment where end users do not have the necessary computer knowledge to restore files, operators can be trusted to restore sensitive data.

End users are allowed to restore?

Use this policy in the following cases:

- When the end users have sufficient knowledge to handle restores. You may need to provide some training for the users on basic backup concepts and restore operations.
- You use library backup devices with media of most recent backups. The end user Data Protector user group, by default, does not allow end users to handle mount requests for needed media. The end users will still need the assistance of the backup operator in case of mount requests. This can be avoided by using large libraries.

HP Data Protector

Disaster recovery

It is too late to prepare for a disaster recovery once a disaster has occurred!

The disaster recovery process consists of 4 phases:

1. Phase 0 (planning/preparation) is the prerequisite for a successful disaster recovery.
2. In Phase 1, DR OS is installed and configured, which usually includes repartitioning and reformatting of the boot partition, since the boot or system partition of the system are not always available and the environment needs to be recovered before normal restore operations can resume.
3. In Phase 2, the operating system with all the configuration information that defines the environment with Data Protector (as it was) is restored.
4. Only after phase 2 is completed, is the restore of applications and user data possible (Phase 3). A well-defined, step-by-step process has to be followed to ensure a fast and efficient restore.

HP Data Protector

Disaster recovery

Disaster recovery methods

- Manual disaster recovery
- Automated disaster recovery
- Disk delivery Disaster recovery
- Enhanced Automated Disaster Recovery (EADR)
- One Button Disaster Recovery (OBDR)
- Alternative disaster recovery methods
- Recovery using third-party tools (for Windows)

Q&A

Thank you