Enterprise Backup and Restore technology and solutions

LESSON VI

HP Data Protector
Overview and Concepts
Part II

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 II

Veselin Petrunov / November 2013

Contents

Part II

2. Media management and devices
    1. Media life cycle
    2. Media pools
    3. Media management before backups begin
    4. Media management during backup sessions
    5. Media management after backup sessions
    6. Devices
    7. Standalone devices
    8. Small magazine devices
    9. Large libraries
    10. Data Protector and Storage Area Networks (SAN)
3. Security / Users and user groups
    1. Security
        1. Cells
        2. Data Protector users accounts/groups/rights
        3. Visibility of backed up data/what is backup ownership?
        4. Data encryption and encrypted control communication
    2. Increased security for Data Protector users/Access to backed up data
    3. Users and user groups - Using predefined user groups/Data Protector user rights

HP Data Protector

Media management

Data Protector provides the following media management functionality that allows simple and efficient management of a large number of media:

- Grouping media into logical groups, media pools, that enable 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 state of each medium, the data protection expiration time, the availability of media for backups, and a catalog of what has been backed up to each medium.
- The capability to transfer all media-related catalog data from one Data Protector Cell Manager to another one without physically accessing the media.
- Automated media rotation policies so that you do not need to take care of tape rotation manually.
- The possibility to explicitly define which media and which devices you want to use for backup.
- Optimized media management for specific device types, such as standalone, magazine, library devices and large silo devices.

HP Data Protector

Media management (continues)

- Fully automated operation. If Data Protector has control of enough media in the library devices, the media management functionality enables the running of backups without the need for an operator to handle media for weeks.
- Recognition and support of barcodes on large libraries with barcode support and silo devices.
- Automatic recognition of Data Protector media format and other popular tape formats.
- Data Protector only writes to blank media initialized (formatted) by Data Protector. You cannot force Data Protector to overwrite foreign tape formats during a backup, thus you avoid accidental overwrites of media that belong to other applications.
- Recognition, tracking, viewing, and handling of media used by Data Protector and separating it from media used by other applications in library and silo devices.
- Keeping information about the media used in a central place and sharing this information among several Data Protector cells.
- Support for media vaulting.
- Interactive or automated creation of additional copies of the data on the media.

HP Data Protector

Media life cycle

A typical media life cycle consists of the following steps:

1. Preparing media for backup.

   This includes initializing (formatting) media for use with Data Protector and assigning media to media pools, which are used to track the media.

2. Using media for backup.

   This defines how media are selected for backup, how the condition of the media is checked, how new backups are added to the media, and when data on the media is overwritten.

3. Vaulting media for long-term data storage. You can use one of Data Protector's data duplication methods to make copies of the backed up data for vaulting purposes.
4. Recycling media for new backups once the data on the media is no longer needed.
5. Retiring media - Once a medium has expired, it is marked poor and will no longer be used by Data Protector.

HP Data Protector

Media pools

What is a media pool?

A pool is a logical set, or group, of media with a common usage pattern and media properties. It can only have media of the same physical type. DLT and DAT/DDS media cannot be in the same pool for instance. The current location of a medium has no influence on its relation to the pool. Whether the medium is in a drive, in a repository slot of a library, in the vault or somewhere else, does not matter; it always belongs to its pool until it is recycled and exported from the cell. Several devices can use media from the same pool.

Media pool property examples

- Appendable - This allows Data Protector to append data to the media in this pool when performing subsequent backup sessions. If this option is not selected, then the media will contain data from a single session only.
- append incrementals only - A backup session appends to a medium only if an incremental backup is performed. This allows you to have a complete set of full and incremental backups on the same medium, if there is enough space.
- media allocation policy - There are several levels of strictness as to which media can be used for backup. They range from strict, where Data Protector requires a specific medium, to loose, where Data Protector accepts any suitable medium in the pool, including new (blank) media.

Every device is linked to a default pool. This pool can be changed in the backup specification.

HP Data Protector

Media pools - Free pools

What is a free pool?

A free pool is an auxiliary source of media of the same type (for example, DLT) for use when all free media in a regular pool run out. It helps to avoid failed backups due to missing (free) media.

When is a free pool used?

Media are moved between regular and free pools on two events:

- Allocation. Media are moved from a free pool to a regular pool
- Deallocation. Media are moved from a regular pool to a free pool. You can specify in the GUI whether deallocation is done automatically. Media from the PC backup pool in “Free pools”, for example, are not automatically deallocated.

A free pool has the following benefits:

- Sharing of free media between pools - All free (unprotected, empty) media can be grouped in a free pool and shared between all media pools that support free pool usage.
- Reduced operator intervention for backup - Assuming that all free media are shared, the need for mount requests is reduced.

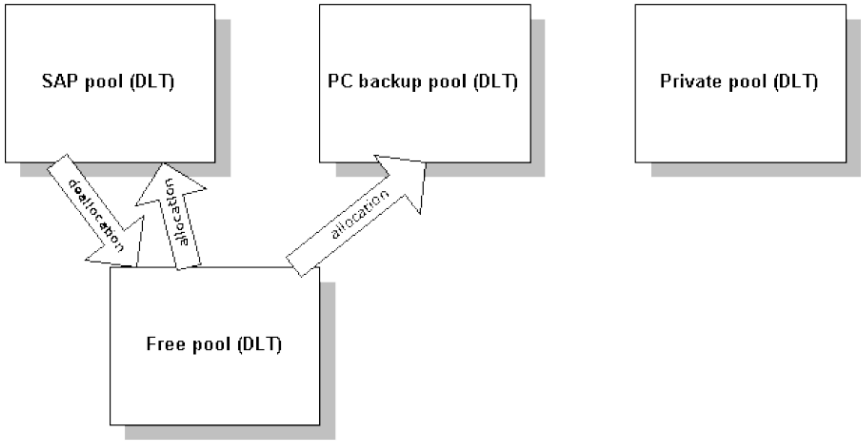

HP Data Protector

Media pool usage examples

Example 1 - A simple one device/one media pool relation

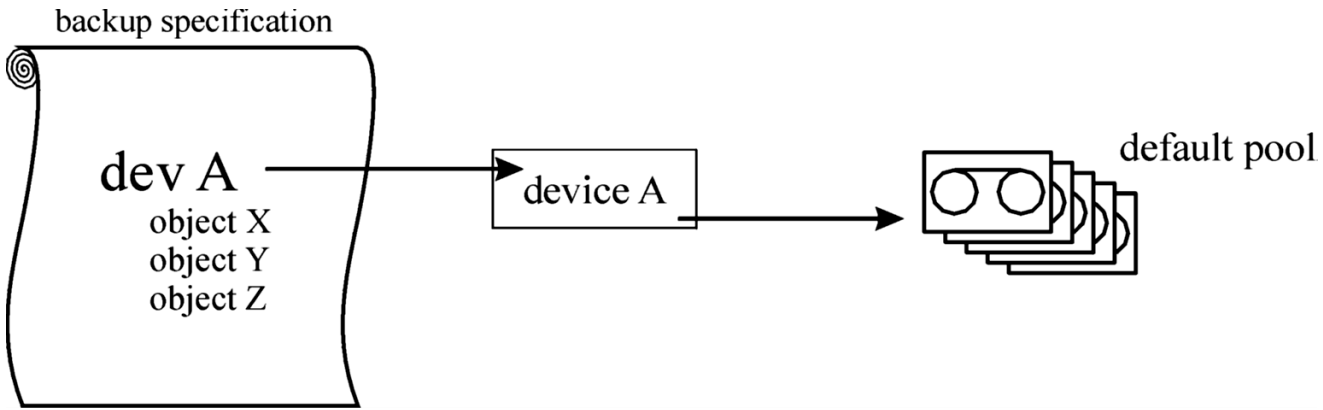

HP Data Protector

Media pool usage examples

Example 2 - Configuration of media pools for large libraries

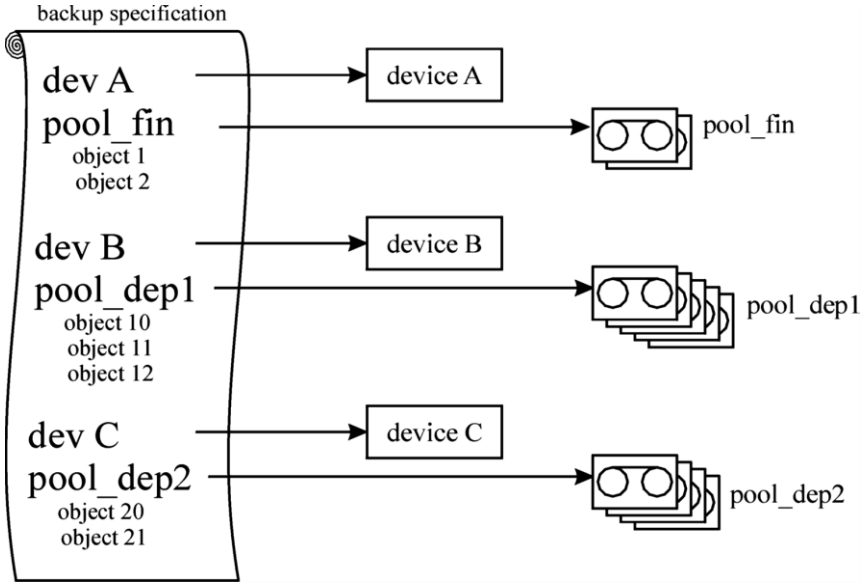

HP Data Protector

Media pool usage examples

Example 3 - Multiple devices, single media pool

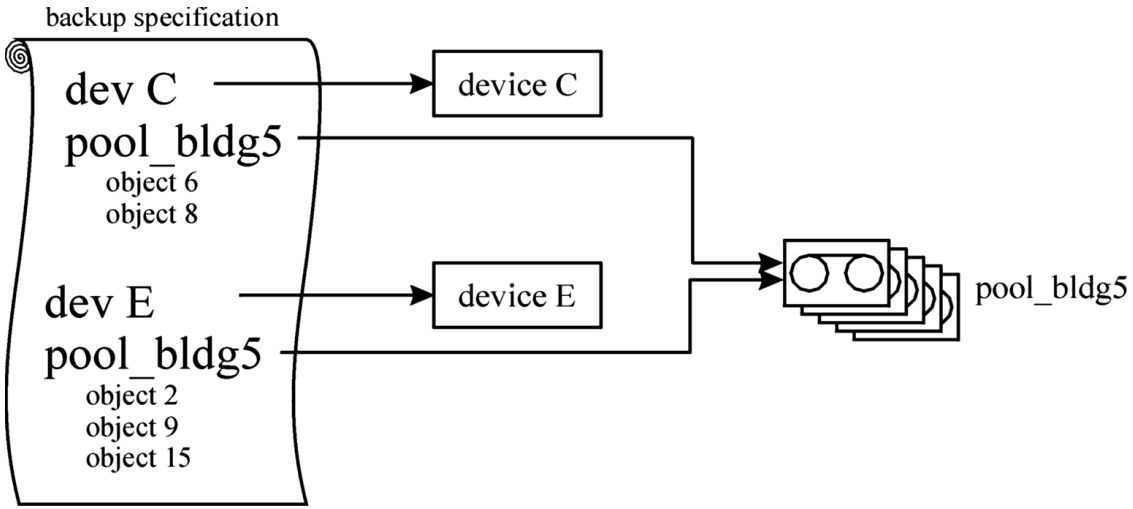

HP Data Protector

Media pool usage examples

Example 4 - Multiple devices, multiple media pools

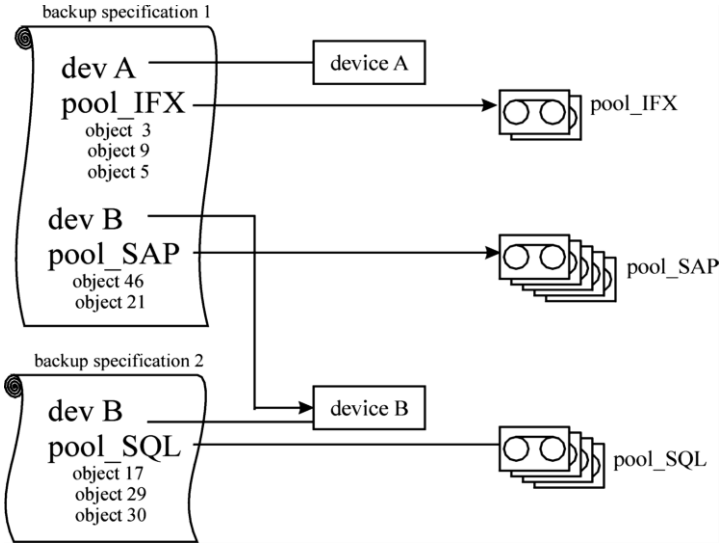

HP Data Protector

Media management before backups begin

- Initializing or formatting media
- Labeling Data Protector media
- Location field

HP Data Protector

Media management during backup sessions

- Selecting media for backups
  - Media allocation policy
  - Pre-allocating media
  - Media condition

- Adding data to media during backup sessions
  - Media usage policy
  - Distributing objects over media

HP Data Protector

Media management during backup sessions

Distributing objects over media

Multiple objects and sessions per medium, sequential writes

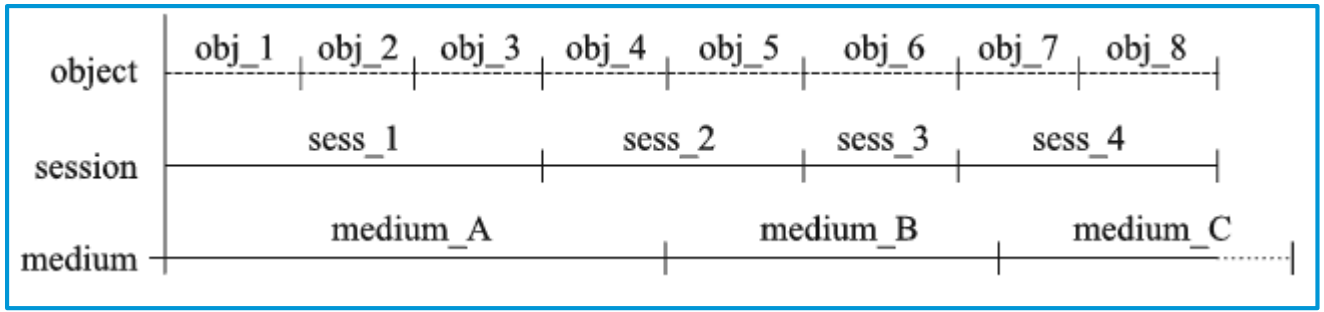

HP Data Protector

Media management during backup sessions

Distributing objects over media

Multiple objects and sessions per medium, concurrent writes

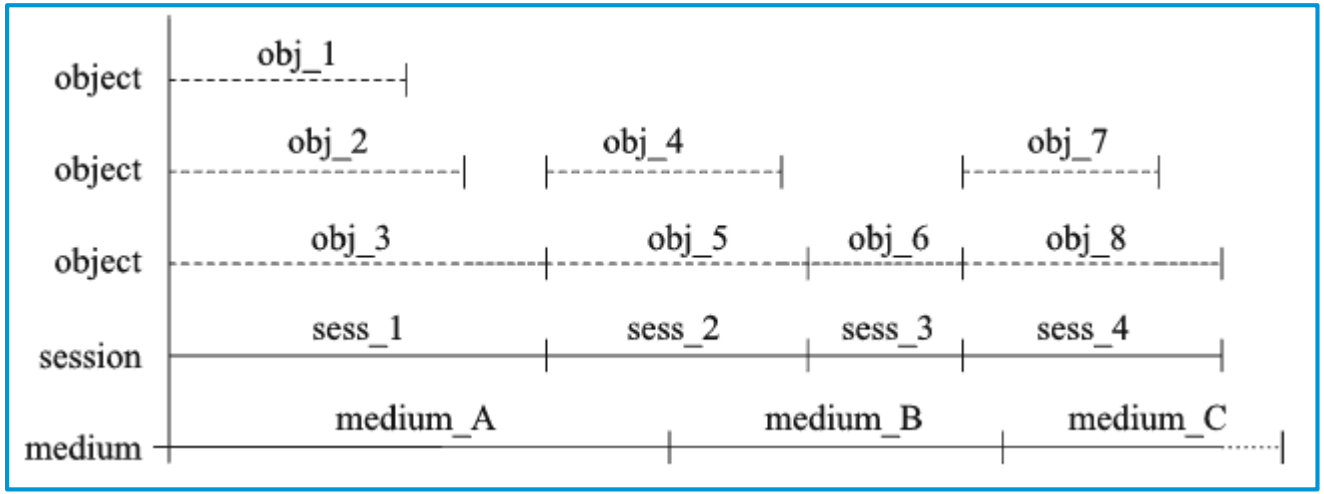

HP Data Protector

Media management during backup sessions

Distributing objects over media

Multiple media per session, multiple media per object

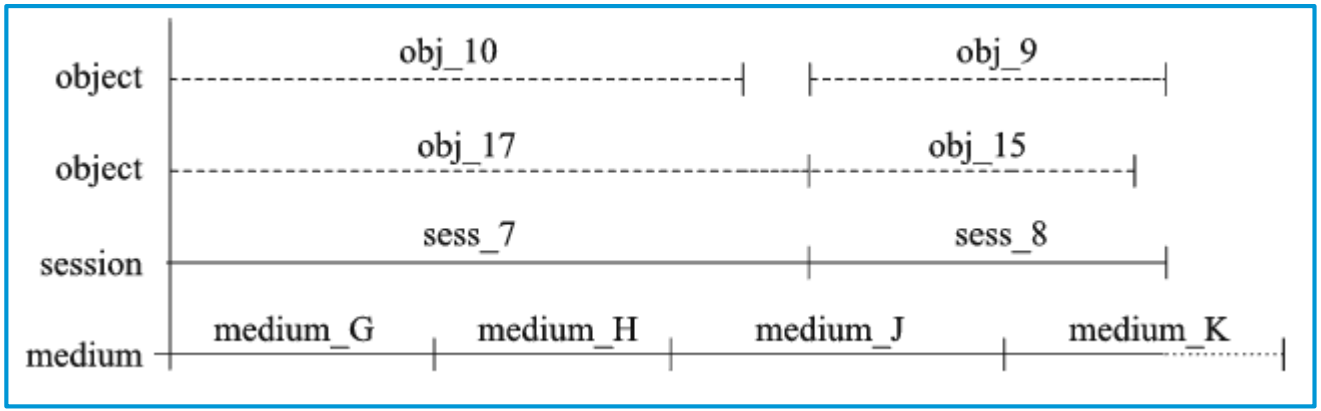

HP Data Protector

Media management during backup sessions

Distributing objects over media

Each object written on a separate medium

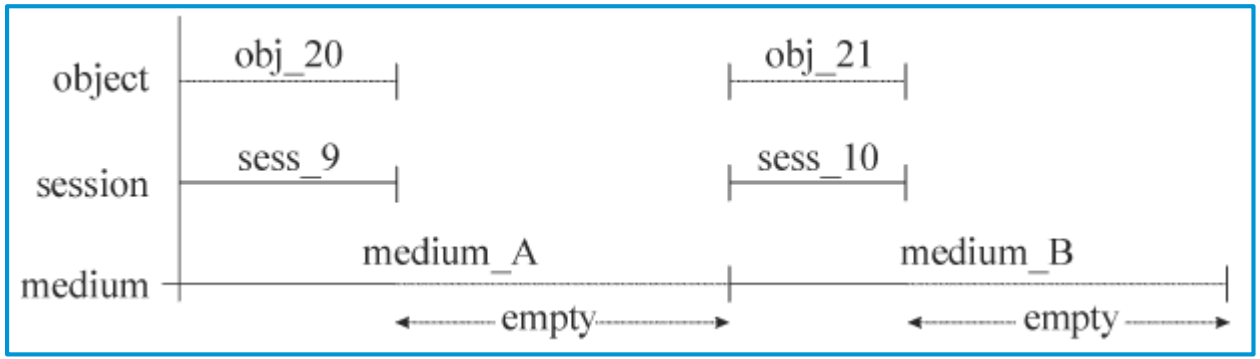

HP Data Protector

Media management during backup sessions

Calculating media condition

Media can have three states: **good**, **fair**, or **poor**.

On a per-medium basis, the following is used for calculating the condition:

- number of overwrites

  The usage of a medium is defined as the number of overwrites from the beginning of the medium. Once the medium has more than the threshold number of overwrites, it is marked as poor.

- media age

  The age of a medium is calculated as the number of months that have elapsed since you formatted, or initialized, the medium. Once a medium is older than the threshold number of months, it is marked as poor.

- device errors

  Some device errors result in the medium being marked as poor. If a device fails during a backup, the medium used for the backup in this device is marked as poor.

HP Data Protector

Media management after backup sessions

Once the data is stored on the media, you must take the right precautions to protect the media and the data on the media. Consider the following:

- **Protecting media from overwrites.**

  You have specified this when you configured a backup of data, but you can change this after the backup is done. For more information on data and catalog protection.

- **Protecting media from physical damage.**

  Media with permanent data may be stored to a safe place.

- **Copying backed up data and keeping the copies at a safe place.**

HP Data Protector

Media management after backup sessions

Vaulting

Vaulting is a process of storing media with important information to a safe place, where they are kept for a specific period of time. The safe place for media is often called a vault.

Data Protector supports vaulting with the following features:

- Data protection and catalog protection policies.
- Easy selecting and ejecting of media from a library.
- The field media location tells you the physical location where the media are stored.
- A report showing media used for backup within a specified time-frame.
- A report showing which backup specifications have used specified media during the backup.
- A report showing media stored at a specific location with data protection expiring in a specific time.
- Displaying a list of media needed for a restore and the physical locations where the media are stored.
- Filtering of media from the media view based on specific criteria.

HP Data Protector

Media management after backup sessions

Vaulting

Implementing vaulting

The implementation of vaulting depends on your company's backup strategy and policies for handling data and media. Generally, it consists of the following steps:

1. Specifying the desired data protection and catalog protection policies when configuring backup specifications.
2. Configuring a vault in Data Protector. Essentially, this means specifying a name for the vault you will use for media, for example: Vault_1.
3. Establishing the appropriate media maintenance policy for media in the vault.
4. Optionally, creating additional copies of the backed up data for vaulting purposes, using the object mirror functionality during backup, or the object copy or media copy functionality after backup.
5. Selecting the media you want to store in a vault, ejecting the media and storing it in the vault.
6. Selecting the media with expired data which is in a vault and inserting the media in a library.

HP Data Protector

Media management after backup sessions

Vaulting

Restoring from media in a vault

Restoring media from a vault is no different than restoring from any other media. Depending on how your data and catalog protection policies are defined, you may need to do some additional steps:

1. Bring media from a vault and insert the media into a device.
2. If the catalog protection for the media is still valid, restore data simply by selecting what you want to restore using the Data Protector user interface.
   If the catalog protection for the media has expired, Data Protector does not have detailed information about the backed up data. You must restore by manually specifying the files or directories you want to restore. You can also restore the complete object to a spare disk and then search for files and directories in the restored filesystem.

HP Data Protector

Devices

- **Using devices with Data Protector**
  Data Protector supports a number of devices available on the market. To use a device with Data Protector, you must configure the device in the Data Protector cell. When you configure a device, you specify a name for the device, some device specific options, such as barcode or cleaning tape support, and a media pool. The process of configuring devices is simplified with a wizard that leads you through all the steps and can even detect and configure devices automatically.
- **Library management console support**
  Many modern tape libraries provide a management console that allows libraries to be configured, managed, or monitored from a remote system. The scope of tasks that can be performed remotely depends on the management console implementation, which is independent of Data Protector. Data Protector eases access to the library management console interface. The URL (web address) of the management console can be specified during the library configuration or re-configuration process.
- **TapeAlert**
  TapeAlert is a tape device status monitoring and messaging utility that makes it easy to detect problems that could have an impact on backup quality. From the use of worn-out tapes to defects in the device hardware TapeAlert provides easy-to-understand warnings or errors as they arise, and suggests a course of action to remedy the problem.
  Data Protector fully supports TapeAlert 2.0, as long as the connected device also provides this functionality.

HP Data Protector

Device lists and load balancing

- **Multiple devices for backup**
  When configuring a backup specification, you can specify several standalone devices or multiple drives in a library device that will be used for the operation. In this case, the operation is faster because data is backed up in parallel to multiple devices (drives).
- **Balancing the use of devices**
  By default, Data Protector automatically balances the load (the usage) of devices so that they are used evenly. This is called **load balancing**. Load balancing optimizes the usage by balancing the number of the objects backed up to each device. Since load balancing is done automatically during backup time, you do not have to manage the allocation of objects to devices used in the session; you just specify the devices to be used.
- **Device chaining**
  Data Protector allows you to configure several standalone devices of the same type, connected to the same system, as a device chain. When a medium in one device gets full, the backup automatically continues on the medium in the next device in the device chain.

HP Data Protector

Device lists and load balancing

When to use load balancing? Use load balancing when:

- You back up a large number of objects.
- You use library (autochanger) devices with several drives.
- You do not need to know on which media objects will be backed up.
- You have a good network connection.
- You want to increase the robustness of the backup. Data Protector automatically redirects the backup operation from failed devices to other devices in a device list.

When not to use load balancing? Do not use load balancing when:

- You want to back up a small number of large objects. In this case Data Protector often cannot effectively balance the load among devices.
- You want to explicitly select to which device each object will be backed up.

HP Data Protector

Device streaming and concurrency

- **What is device streaming?**
  To maximize a device performance, it must be kept streaming. A device is streaming if it can feed enough data to the medium to keep the medium moving forward continuously.
- **How to configure device streaming?**
  To allow the device to stream, a sufficient amount of data must be sent to the device. Data Protector accomplishes this by starting multiple Disk Agents for each Media Agent that writes data to the device.
- **Disk agent concurrency**
  The number of Disk Agents started for each Media Agent is called Disk Agent (backup) concurrency and can be modified using the Advanced options for the device or when configuring a backup.
- **Increased performance**
  If properly set, backup concurrency increases backup performance. Device streaming is also dependent on other factors, such as network load and the block size of the data written to the device.
- **Multiple data streams**
  This feature is useful for backing up very large and fast disks to relatively slow devices. Multiple Disk Agents read data from the disk in parallel and send the data to multiple Media Agents.

HP Data Protector

Segment size

A medium is divided into data segments, catalog segments and a header segment. Header information is stored in the header segment, which is the same size as the block size. Data is stored in data blocks of data segments. Information about each data segment is stored in the corresponding catalog segment. This information is first stored in the Media Agent memory and then written to a catalog segment on the medium as well as to the IDB.

Some tape technologies place limitations on the number of file marks per medium. Ensure that your segment size is not too low.

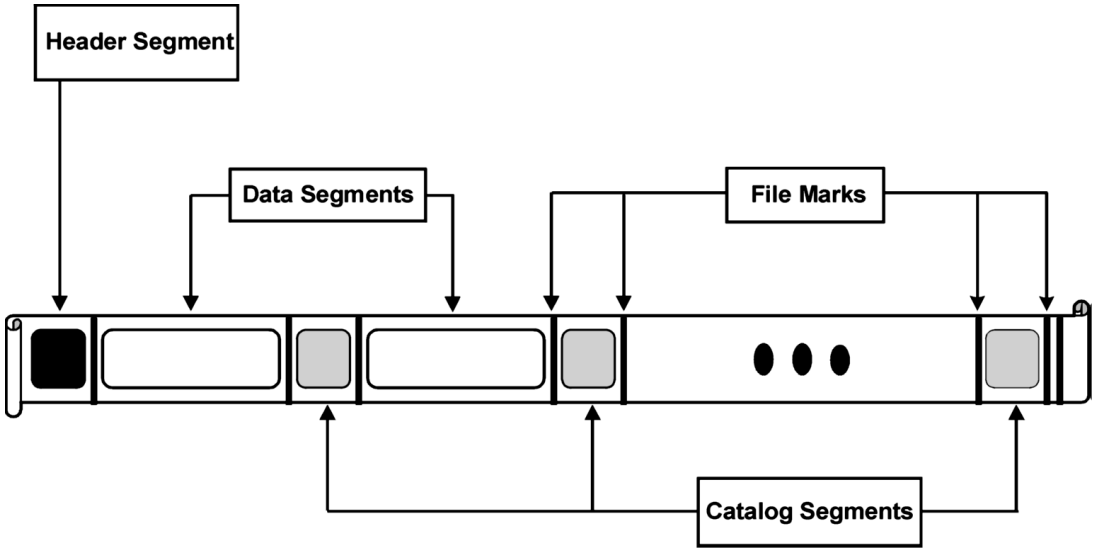

HP Data Protector

Block size

Segments are not written as a whole unit, but rather in smaller subunits called blocks. The hardware of a device processes data in units of a device-type specific block size. Data Protector allows you to adjust the size of the blocks it sends to the device. The default block size value for all devices is 64 kB.

Increasing the block size can improve performance. Changing the block size should be done before formatting tapes. For example, a tape written with the default block size cannot be appended to using a different block size.

HP Data Protector

Device locking and lock names

- **Device names**
  When configuring devices for use with Data Protector, you can configure the same physical device many times with different characteristics simply by configuring the same physical device in Data Protector with different device names.
- **Physical device collision**
  When specifying a device used for backup, you may specify one device name in one backup specification and another device name of the same physical device in a different backup specification. Depending on the backup schedule, this may result in Data Protector trying to use the same physical device in several backup sessions at the same time, thus creating a collision.
- **Preventing collision – Device lock name**
  To prevent this collision, specify a virtual lockname in both device configurations. Data Protector checks if the devices have the same lockname and prevents collision.

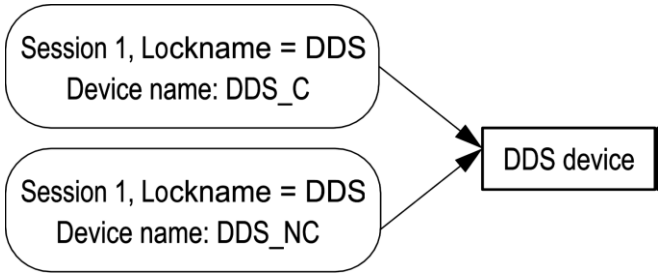

HP Data Protector

Standalone devices

What are standalone devices?

Standalone devices are devices with one drive that reads/writes to one medium at time. Standalone devices are used for small scale backups or special backups. When the medium is full, the operator must manually replace it with a new medium for the backup to proceed.

Data Protector and standalone devices

Once you have connected a device to the system, you use the Data Protector user interface to configure the device for use with Data Protector. To do this, you must first install a Data Protector Media Agent on the system with the device connected. Data Protector can detect and automatically configure most standalone devices. During a backup, Data Protector issues a mount request when the medium in a device is full. The operator must replace the medium for the backup to continue.

What are device chains?

Data Protector allows you to configure multiple standalone devices to a device chain. When a medium in one device gets full, the backup automatically continues on the medium in the next device in the device chain. Device chains allow running unattended backups using several standalone devices without having to manually insert/eject media when the media are full.

Stacker devices

Stacker devices, similar to device chains, contain a number of media that are used in a sequential order. When a medium gets full, the next medium is loaded and used for backup.

HP Data Protector

Small magazine devices

What are magazine devices?

Magazine devices group a number of media into a single unit called a magazine. Data Protector treats the magazine as if it were a single medium. A magazine has a larger capacity than a single medium and is easier to handle than several single media.

Data Protector and magazine devices

Data Protector allows you to perform media management tasks on magazines as sets, emulating single media by providing magazine and media views, or on a single medium. You can alternatively use magazine devices as normal libraries without using Data Protector magazine support. Data Protector can detect and automatically configure magazine devices.

Cleaning dirty drives

Using cleaning tapes, Data Protector can automatically clean magazines and other devices when they get dirty.

HP Data Protector

Large libraries

What are library devices?

Library devices are automated devices, also called autoloaders, exchangers or jukeboxes. In Data Protector, most libraries are configured as SCSI libraries. They contain a number of media cartridges in a device's repository and can have multiple drives writing to multiple media at a time. A typical library device has a SCSI ID for each drive in the device and one for the library robotic mechanism that moves media from slots to drives and back. For example, a library with four drives has five SCSI IDs, four for the drives and one for the robotic mechanism.

Data Protector also supports silo libraries, such as HP Libraries, StorageTek/ACSLS and ADIC/GRAU AML.

Handling of media

The Data Protector user interface provides a special library view, which simplifies managing library devices.

Media in a large library device can all belong to one Data Protector media pool, or they can be split into several pools.

Configuring a library

When configuring a device, you configure the slot range you want to assign to Data Protector. This allows sharing of the library with the other application. The assigned slots may contain blank (new) media, Data Protector or non-Data Protector media. Data Protector checks the media in the slots and displays the information about the media in the library view. This allows you to view all kinds of media, not just the media used by Data Protector.

HP Data Protector

Large libraries

Sharing a library with other applications

A library device can be shared with other applications storing data to media in the device. You can decide which drives from the library you want to use with Data Protector. For example, out of a four-drive library you may choose to use only two drives with Data Protector. You can decide which slots in the library you want to manage with Data Protector. For example, out of the 60 slots library you might use slots 1-40 with Data Protector. The remaining slots would then be used and controlled by a different application. Sharing of the library with other applications is especially important with large HP libraries and silo libraries, such as StorageTek/ACSLS or ADIC/GRAU AML devices.

Enter / eject mail slots

Library devices provide special enter/eject mail slots an operator uses to enter or eject media to or from the device. Depending on the device, more than one enter/eject slot can be provided. In case of a single mail slot, media are inserted one by one, while in case of multiple mail slots, a particular number of slots can be used in one enter/eject operation. Data Protector allows you to enter/eject several media in one step. For example, you can select 50 slots in the device and eject all media in one action. Data Protector will automatically eject media in the correct order for the operator to remove the media from the enter/eject mail slot.

Barcode support

Data Protector supports library devices with a barcode reader. In these devices, each medium has a barcode that uniquely identifies media.

HP Data Protector

Large libraries

Cleaning tape support

HP Data Protector provides automatic cleaning for most devices using a cleaning tape. This medium will be used automatically by Data Protector if a dirty drive event from the device is detected.

- For SCSI libraries it is possible to define which slot holds a cleaning tape.
- For devices with a barcode reader, Data Protector recognizes cleaning tape barcodes automatically if they have the CLN prefix.
- For devices without a cleaning tape, a dirty drive detection will cause a cleaning request to be displayed on the session monitor window. The operator must clean the device manually.
  You cannot continue your backup without cleaning the drive, since the backup may fail because data may not be correctly written and stored on the media.

HP Data Protector

Sharing a library with multiple systems

What is library sharing?

Device sharing allows you to connect different drives of a physical library to different systems. These systems can then perform local backups to the library. The result is significantly higher backup performance and less network traffic. To enable library sharing, the drives in the library must have the possibility to connect to separate SCSI buses. This is useful with high performance libraries to allow the drive to receive data in a continuous stream from multiple systems, further enhancing performance. Data Protector internally redirects the robotic commands to the system that manages the robotics.

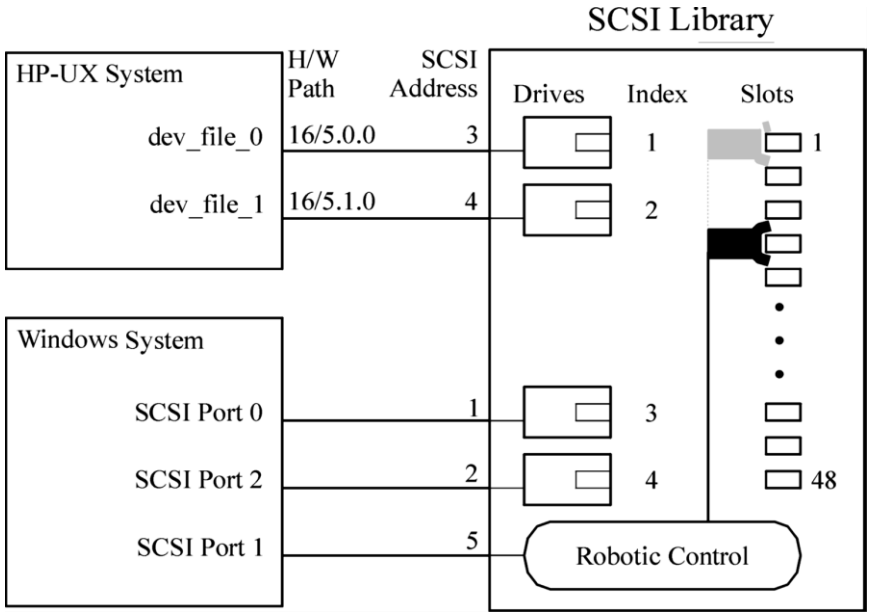

HP Data Protector

Large libraries

Control protocols and Data Protector Media Agents

The drives in the library must be able to physically connect to different systems that have a Data Protector Media Agent (the General Media Agent or the NDMP Media Agent) installed. With Data Protector, there are two types of protocols used for drive control:

- SCSI – for SCSI or Fibre Channel connected drives.
  This protocol is implemented in both the General Media Agent and in the NDMP Media Agent.
- NDMP – for NDMP dedicated drives.
  This protocol is implemented in the NDMP Media Agent only.

On the other hand, there are four types of protocols used for library robotic control:

- ADIC/GRAU – for ADIC/GRAU library robotics
- StorageTek ACS – for StorageTek ACS library robotics
- SCSI – for robotics other libraries
- NDMP – for NDMP robotics

All four library robotic control protocols are implemented in both the General Media Agent and in the NDMP Media Agent.

HP Data Protector

Large libraries

Drive control

Any Data Protector client system configured to control a drive in a library (regardless of the drive control protocol and platform used) can communicate with any Data Protector client system configured to control the robotics in the library (regardless of the robotics control protocol and platform used). Thus, it is possible to share drives in any supported library among Data Protector clients systems on various platforms using various robotic and drive protocols. The NDMP Media Agent is needed only on client systems controlling the backup of an NDMP server (on client systems configured for NDMP dedicated drives). In all other cases the two Data Protector Media Agents are interchangeable.

Robotic control

A Data Protector client system controlling the library robotics can have either the General Media Agent or the NDMP Media Agent installed, regardless of the type of drive protocol (NDMP or SCSI) used with the drives in the library.

	Robotic control protocol			
	ADIC/GRAU	StorageTek ACS	SCSI	NDMP
Drive control protocol (NDMP or SCSI)	NDMP Media Agent or General Media Agent	NDMP Media Agent or General Media Agent	NDMP Media Agent or General Media Agent	NDMP Media Agent or General Media Agent

HP Data Protector

Sharing a SCSI library (robotics attached to a Data Protector Client System)

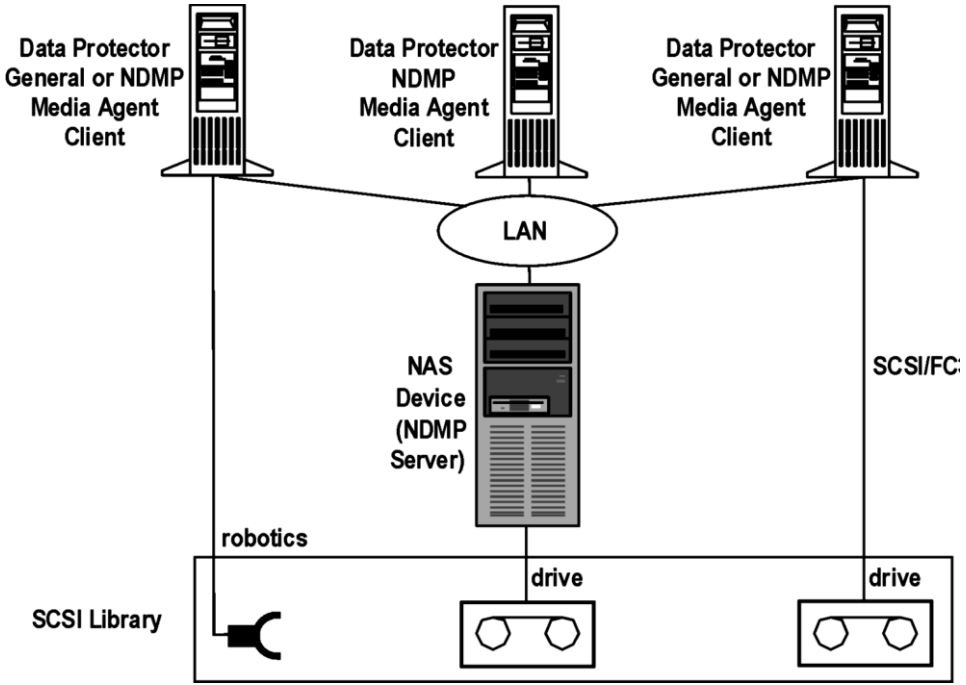

HP Data Protector

Sharing a SCSI library (robotics attached to an NDMP Server)

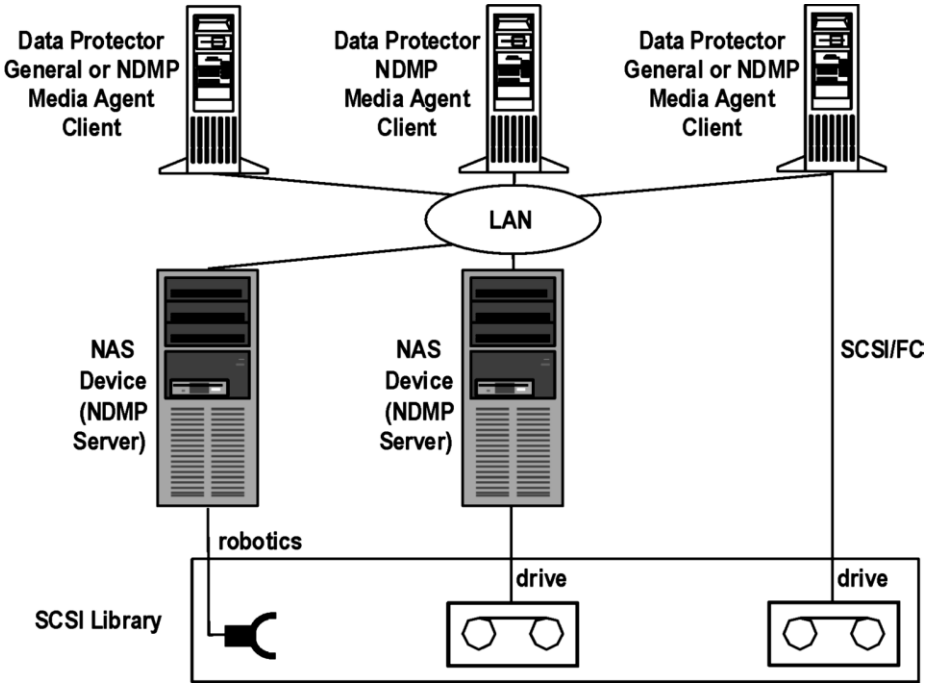

HP Data Protector

Data Protector and Storage Area Networks

Where and how you store data in your enterprise may have a serious impact on your business. Information is becoming increasingly mission-critical to most companies. Today, terabytes of data must be accessible to users across the network. The Data Protector implementation of SAN-based Fibre Channel technology provides you with the data storage solution you need.

Storage Area Networks

A Storage Area Network (SAN) is a new approach to network storage that separates storage management from server management with a network devoted to storage.

A SAN provides any-to-any connectivity for all network resources, thus enabling device sharing between multiple client systems and increasing data traffic performance as well as the availability of devices.

The SAN concept allows the exchange of information between multiple data storage devices and servers. The servers can access data directly from any device and do not need to transfer data over the conventional LAN. A SAN consists of servers, backup devices, disk arrays, and other nodes, all connected with a fast network connection, typically Fibre Channel. This additional network provides off-loading storage operations from the conventional LAN to a separate network.

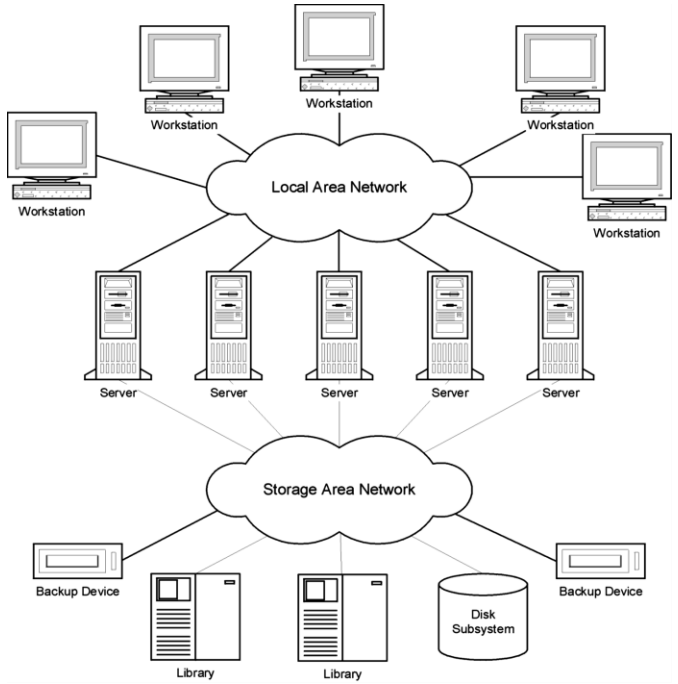

HP Data Protector

Device sharing in SAN

Data Protector supports the SAN concept by enabling multiple systems to share backup devices in the SAN environment. The same physical device can be accessed from multiple systems. Thus, any system can perform a local backup on some device or any other device. Because data is transferred over the SAN, backups do not need any bandwidth on your conventional LAN. This type of backup is sometimes referred to as a “LAN-free” backup. Backup performance is also improved, because SAN-based Fibre Channel technology typically provides an order of magnitude higher throughput than LAN technologies.

You need to prevent several computer-systems from writing to the same device at the same time. This can become even more complex when devices are used from several applications. Access to the devices needs to be synchronized between all systems involved. This is done using locking mechanisms.

SAN technology provides an excellent way to manage the robotics of a library from multiple systems. This allows the option to manage the robotics from one system (classic) or allow each system that uses the library to access the robotics directly, provided the requests to the robotics are synchronized between all the systems involved.

HP Data Protector

Multipath

A device in a SAN environment is usually connected to several clients and can thus be accessed through several paths, that is client names and SCSI addresses (device files on UNIX). Data Protector can use any of these paths. You can configure all paths to a physical device as a single logical device - **multipath device**.

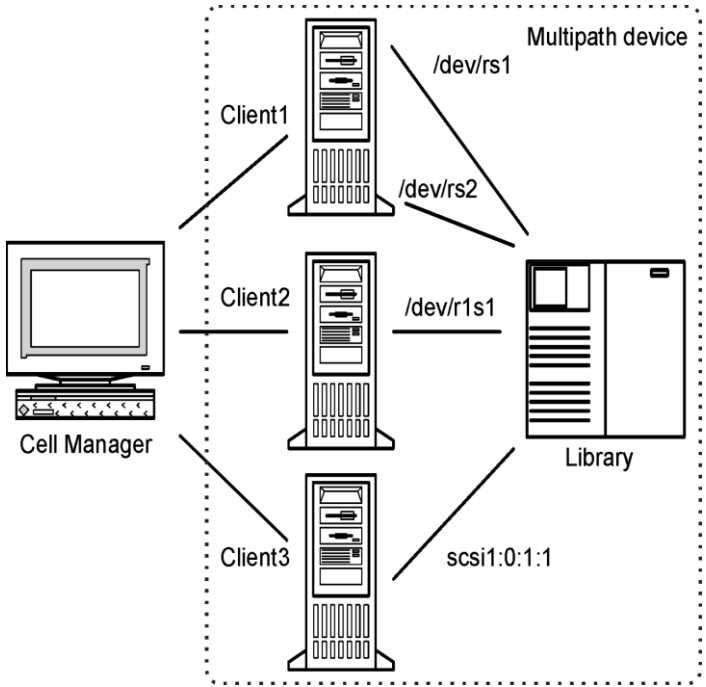

HP Data Protector

Multipath

Why use multiple paths

Data Protector provides easy way to configure a single multipath device for all paths. Multipath devices increase system resilience. Data Protector will try to use the first defined path. If all paths on a client are inaccessible, Data Protector will try to use paths on the next client. Only when none of the listed paths is available, the session aborts.

Path selection

During a backup session, the device paths are selected in the order defined during the device configuration, except if a preferred client is selected in the backup specification. In this case, the preferred client is used first.

During a restore session, the device paths are selected in the following order:

1. Paths that are on the client to which the objects are restored, if all objects are restored to the same target client
2. Paths that were used for backup
3. Other available paths

HP Data Protector

Device locking

Locking devices must cover the possibility of several applications using the same device, as well as only Data Protector using a device by sending data and commands to it from several systems. The purpose of locking is to ensure that only one system at a time communicates with a device that is shared between several systems.

If Data Protector is the only application that uses a drive, but that same drive needs to be used by several systems, Device Locking has to be used.

If Data Protector is the only application that uses a robotics control from several systems, Data Protector handles this internally, provided the library control is in the same cell as all the systems that need to control it. In such a case, all synchronization of access to the device is managed by Data Protector internal control.

HP Data Protector

Device sharing in clusters

Clustering, which is often used in combination with the SAN concept, is based on sharing network resources (for example network names, disks, and tapes devices) between nodes.

Cluster-aware applications can at any time run on any node in a cluster (they run on virtual hosts). To perform a local backup of such an application, you need to configure devices with virtual hostnames instead of real node names. Configure as many devices for each physical device as you need, using the Lock Name device locking mechanism.

Static drives

Static drives are devices that are configured on a real node in a cluster. They can be used to back up data from systems with disks that are not shared. However, they are not useful for backing up cluster-aware applications, because such application can run on any node in the cluster.

Floating drives

Floating drives are device that are configured on a virtual host, using virtual system names. Floating drives should be configured for the backup of cluster-aware applications. This ensures that no matter on which node in the cluster the application is currently running, Data Protector always starts a Media Agent on that same node.

HP Data Protector

Media management and devices

HP Data Protector

Security

What is security?

Security in the backup context typically refers to:

- Who can administer or operate a backup application (Data Protector).
- Who can physically access client systems and backup media.
- Who can restore data.
- Who can view information about backed up data.

Data Protector provides security solutions on all these levels.

Data Protector security features

The following features allow and restrict access to Data Protector and the backed up data. The items in this list are described in detail in the following sections.

- Cells
- Data Protector user accounts
- Data Protector user groups
- Data Protector user rights
- Visibility and access to backed up data
- Data encryption
- Encrypted control communication

HP Data Protector

Security

Cells

- **Starting sessions**
  Data Protector security is based on cells. Backup and restore sessions can only be started from the Cell Manager unless you have the Data Protector Manager-of-Managers functionality. This ensures that users from other cells cannot back up and restore data from systems in your local cell.
- **Access from a specific Cell Manager**
  Additionally, Data Protector allows you to explicitly configure from which Cell Manager a client system can be accessed, that is, configuring a trusted peer.
- **Restrict pre- and post-execution**
  For security reasons, various levels of restrictions can be configured for pre-exec and post-exec scripts. These optional scripts allow a client system to be prepared for the backup by, for example, shutting down an application to obtain a consistent backup.

HP Data Protector

Security

Data Protector users accounts

Anyone using any Data Protector functionality, administering Data Protector, or restoring personal data, must have a Data Protector user account. This restricts unauthorized access to Data Protector and backed up data.

Who defines user accounts?

An administrator creates this account specifying a user login name, systems from which a user can log in, and the Data Protector user group membership that defines the user rights.

When is the account checked?

When a user starts the Data Protector user interface, Data Protector checks user rights. User rights are also checked when specific tasks are performed by a user.

HP Data Protector

Security

Data Protector user groups

What are user groups?

When a new user account is created, the user becomes a member of the specified user group. Each user group contains defined Data Protector user rights. All the members of the group have the user rights set for the group.

Why use user groups?

Data Protector user groups simplify user configuration. The administrator groups users according to the access they need. For example, an end-user group could allow members to restore personal data to a local system only, while the operator group allows the starting and monitoring of backups, but not the creating of backups.

HP Data Protector

Users and user groups

Increased security for Data Protector users

Data Protector provides advanced security functionality that prevents unauthorized backing up or restoring of data. Data Protector security involves hiding data from unauthorized users, data encoding, and restricted grouping of users according to their responsibilities.

Access to backed up data

Backing up and then restoring data is essentially the same as copying data. Therefore, it is important to restrict access to this data to authorized users only.

Data Protector provides the user-related security - all users intent on using any of the Data Protector functionality must be configured as Data Protector users.

Visibility of backed up data

- Backed up data is hidden from other users, except the backup owner. Other users do not even see that data was backed up. For example, if the backup operator has configured a backup, only the backup operator or the system administrator can see and restore the backed up data. You can make data visible to other users using the Data Protector Public option.

HP Data Protector

Users and user groups

To use Data Protector, you must be added to the Data Protector configuration as a Data Protector user with certain privileges. Note that adding a new user is not a prerequisite for backing up the system this user is using.

Users are grouped into user groups with specific user rights, for example, to monitor sessions in the cell, configure backups, and restore files.

Predefined user groups

To simplify the configuration of your backup, Data Protector provides predefined user groups with specific rights to access Data Protector functionality. For example, only members of the admin user group can access all Data Protector functionality. Operators can, by default, start and monitor backups. Depending on your environment, you may decide to use the default Data Protector user groups, modify them, or create new ones.

Default administrators

During installation, the following users are automatically added to the Data Protector admin user group:

- UNIX root user on the UNIX Cell Manager system
- User installing Data Protector on the Windows Cell Manager system

This allows them to configure and use the complete Data Protector functionality.

HP Data Protector

Using predefined user groups

The following default groups are provided by Data Protector:

User group	Access rights
Admin	Allowed to configure Data Protector and perform backup, restore, and all other available operations.
Operator	Allowed to start backups and respond to mount requests.
End-user	Allowed to perform restore of their own objects. In addition, users can monitor and respond to mount requests for their own restore sessions.

Data Protector user rights

Data Protector users have the Data Protector user rights of the user group they belong to. For example, all members of the admin user group have the rights of the Data Protector admin user group.

When configuring a user from the Windows domain in Data Protector running on the UNIX Cell Manager, the user must be configured with the Domain Name or the wildcard group "*". Additionally, you can complement the user security layer provided by Data Protector user groups with restrictions of user actions to certain systems of the cell. Such restrictions can be configured in the user_restrictions file that is located on the Cell Manager. They apply only to members of the Data Protector user groups other than admin and operator.

HP Data Protector

Security

Data Protector user rights

What are user rights?

Data Protector user rights define the actions that a user can perform with Data Protector. They are applied on the Data Protector user group level and not to each user individually. Users added to a user group automatically gain the user rights assigned to this user group.

Why use user rights?

Data Protector provides flexible user and user group functionality, which allows the administrator to selectively define who can use a particular Data Protector functionality. It is important to carefully apply the Data Protector user rights: backing up and restoring data is essentially the same as copying data.

Visibility of backed up data

Backing up data means creating a new copy. Therefore, when you deal with confidential information, it is important to restrict access to both the original data and to the backup copy itself.

Hiding data from other users

When configuring a backup, you can decide whether during a restore the data is visible to everyone (public) or only to the owner of the backup (private).

HP Data Protector

Security

What is backup ownership?

Who owns a backup session?

Each backup session and all the data backed up within it is assigned an owner. The owner can be the user who starts an interactive backup, the account under which the CRS process is running, or the user specified as the owner in the backup specification options.

Backup ownership and restore

Backup ownership affects the ability of users to see and restore data. Unless the object is marked as Public, only the owner of the media set or an administrator can see the data saved in the media set. The right to see and restore private objects can be granted to groups other than admin as well.

HP Data Protector

Data encryption

Open systems and public networking make data security in large enterprises essential. Data Protector lets you encrypt backed-up data so that it becomes protected from others. Data Protector offers two data encryption techniques: software-based and drive-based.

Data Protector software encryption, referred to as AES 256-bit encryption, is based on the AES-CTR (Advanced Encryption Standard in Counter Mode) encryption algorithm that uses random keys of 256-bit length. The same key is used for both encryption and decryption. With AES 256-bit encryption, data is encrypted before it is transferred over a network and before it is written to media.

Data Protector drive-based encryption uses the encryption functionality of the drive. The actual implementation and encryption strength depend on the drive's firmware. Data Protector only turns on the feature and manages encryption keys.

The key management functionality is provided by the Key Management Server (KMS), which is located on the Cell Manager. All encryption keys are stored centrally in the keystore file on the Cell Manager and administered by the KMS.

You can encrypt all or selected objects in a backup specification and also combine encrypted and unencrypted sessions on the same medium.

In addition to the encryption functionality, Data Protector also offers the encoding functionality that uses a keyless, built-in algorithm for this purpose.

HP Data Protector

Backup session with AES 256-bit encryption

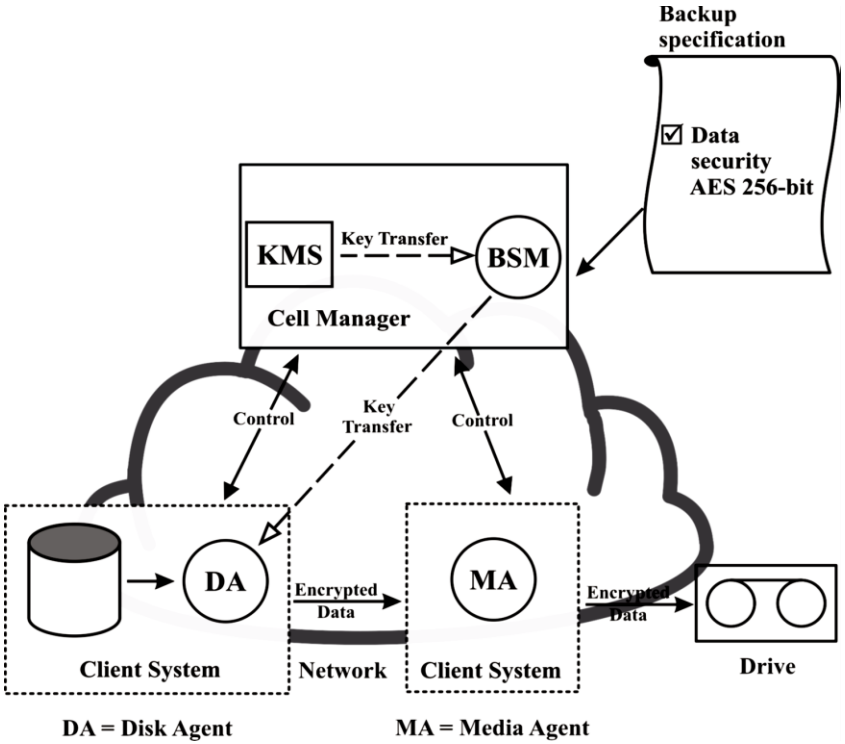

HP Data Protector

Data Protector drive-based encryption

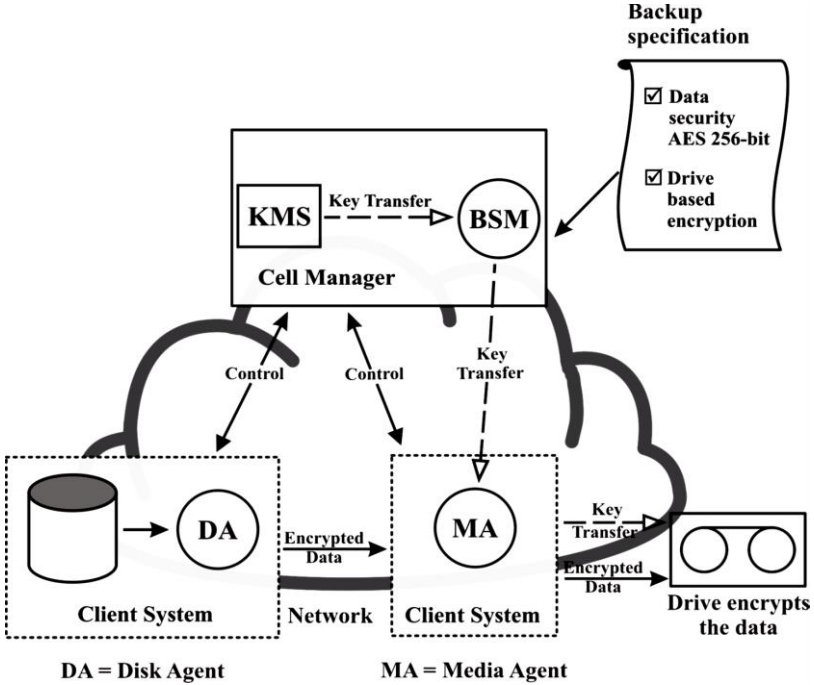

HP Data Protector

Restore from encrypted backups

No additional encryption related preparations are needed for restore of encrypted backups, as Data Protector automatically obtains the appropriate decryption keys.

HP Data Protector

Data Protector encrypted control communication

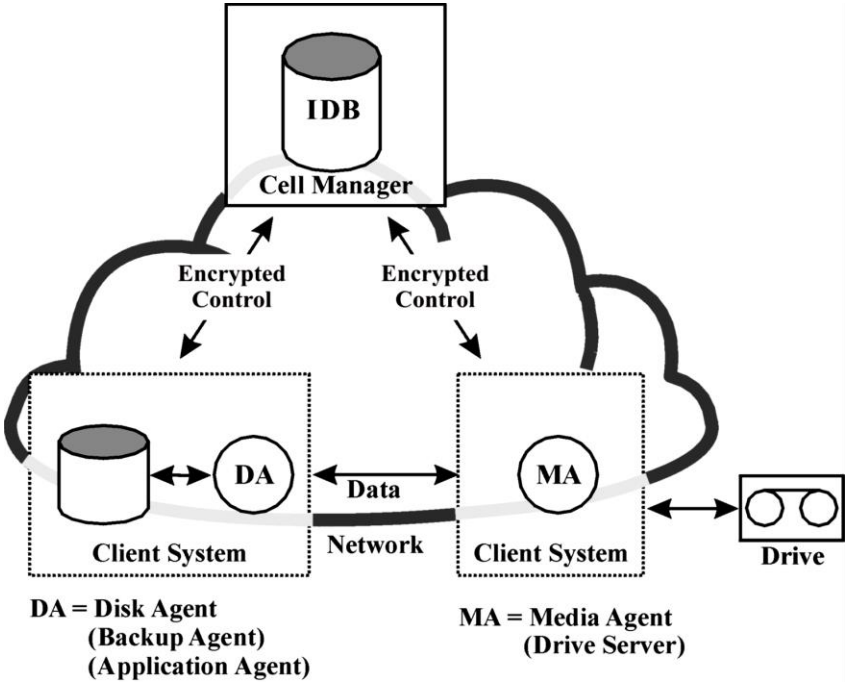

HP Data Protector

Data encryption and encrypted control communication

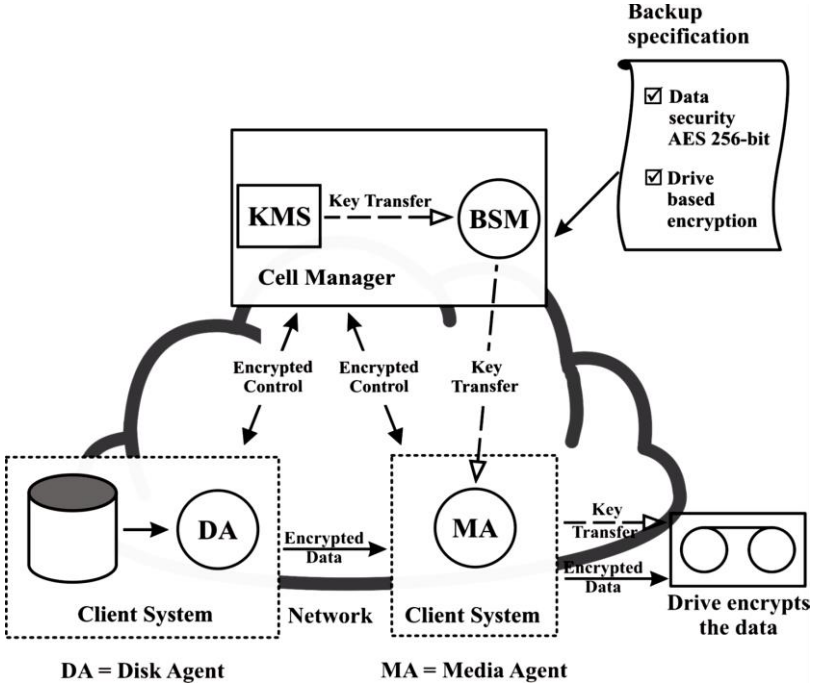

Q&A

Thank you