



HP StoreOnce Catalyst and HP Data Protector 7 – Implementation and Best Practice Guide – Release 2

Executive Summary

This guide is intended to enable the reader to understand the basic technology of HP StoreOnce Catalyst and to design a Data Protector solution. It is not intended to be a full guide to HP Data Protector 7 as there already exists extensive documentation on this software. However this guide will provide the extra information concerning best practice for a StoreOnce B6200 implementation using the StoreOnce Catalyst technology.

Contents

Executive Summary	1
Introduction to HP StoreOnce catalyst	1
HP StoreOnce Catalyst – the basics	2
Introduction to the StoreOnce B6200 Architecture	3
Terminology StoreOnce B6200	6
HP StoreOnce B6200 Catalyst Store specifications	7
HP StoreOnce Catalyst best practice – What information is required for planning?	8
HP StoreOnce Catalyst and Data Protector 7 integration	11
Terminology HP Data Protector	12
Data Protector Gateways	14
Creating a Data Protector 7 specification for backup to a StoreOnce Catalyst Store	17
Sizing the Backup Server	34
Reclaiming disk space from ‘expired’ backups	39
Using Data Protector 7 ‘Object Copy’ with StoreOnce Catalyst	40
Restoring Data from a StoreOnce Catalyst backup & replication	45
Recovering from total loss – ‘Disaster Recovery’ planning (DR)	48

Networking best practice	49
Housekeeping	51
HP StoreOnce Catalyst Licensing	52
HP StoreOnce B6200 Autonomic Restart	55
Conclusion and key learning points	56

Introduction to HP StoreOnce Catalyst

HP StoreOnce Catalyst brings the HP StoreOnce vision of a single, integrated enterprise-wide deduplication algorithm a step closer. It allows the seamless movement of deduplicated data across the enterprise to other StoreOnce Catalyst systems without rehydration. This means that you can benefit from:

Simplified management of data movement from a single pane of glass: tighter integration with your backup application to centrally manage file replication across the enterprise.

Seamless control across complex environments: supporting a range of flexible configurations that enable the concurrent movement of data from one site to multiple sites, and the ability to cascade data around the enterprise (sometimes referred to as multi-hop).

Enhance performance: distributed deduplication processing using StoreOnce Catalyst stores on the B6200 and on multiple servers can optimize loading and utilization of backup hardware, network links and backup servers for faster deduplication and backup performance.

Faster time to backup to meet shrinking backup windows: up to 100TB/hour *aggregate throughput, 4x faster than backup to a NAS target

*Actual performance is dependent upon configuration data set type, compression levels, number of data streams, number of devices emulated and number of concurrent tasks, such as housekeeping or replication.

HP StoreOnce Catalyst is currently available on the HP B6200 Backup System and also as a software component of HP Data Protector 7. In addition to HP Data Protector 7, HP StoreOnce Catalyst is also supported by Symantec NetBackup 7.x and Backup Exec 12. The HP B6200 can support Catalyst Stores, Virtual Tape and NAS (CIFS/NFS) on the same system and so is ideal for customers who have legacy requirements for VTL and NAS but wish to move to StoreOnce Catalyst.

HP StoreOnce Catalyst does require a separate license. VTL/NAS emulations do not require licenses except if they are used as replication targets devices. If VTL/NAS replication is used in addition to StoreOnce Catalyst then both licenses are required.

HP StoreOnce Catalyst - the basics

HP StoreOnce Catalyst is a new type of storage and is more closely integrated with the data protection software. In the case of Data Protector 7 the application programming interface is embedded within the Data Protector media agent (Fig 1). Data transfers and commands are transferred by standard IP connection and the HP B6200 offers both 1GbE and 10GbE connections. (10GbE recommended for performance). HP StoreOnce catalyst offers advanced features such as deduplication at the backup server and movement of backups between systems under the command of the HP Data Protector. HP StoreOnce catalyst also very importantly allows the data Protector 7 software to release disk space occupied by 'expired' backups. This feature is not available in virtual tape. Normally customers develop a scheme where backups are kept for varying periods of time. For example: A full backup is made on say a weekly basis with incremental backup performed every day. The incremental backups are expired when the next full backup is taken as they are no longer required. The weekly full backups could be kept for 4 weeks and then a monthly full backup is created. The weekly backups can then be expired and so on. This really is customer dependent and varies according to the data. HP StoreOnce Catalyst has the additional advantage that backups can then be moved offsite to another catalyst store all under control of the software. The data is moved without rehydration i.e. only new data 'chunks' are moved between stores. It is possible to move to multiple stores. Data duplication uses the HP Data Protector 'object copy' functionality and can replicate data to multiple HP StoreOnce Catalyst stores.

Figure 1 shows the data paths between the B6200 and the backup server equipped with Data Protector Software (media agent). The B6200 is shown as a 2 node/single couplet system. Only catalyst stores are shown but VTL and NAS can co-exist. The network connection is shown as a WAN or LAN because HP StoreOnce Catalyst Protocol is designed from the outset to accommodate possible latency differences between a local network and a wide area network. Only a Data Protector media server is shown and could either backup data which is contained on directly attached disks or from clients which only have a media agent. The networking will be covered later in a separate section. The HP StoreOnce Catalyst API is embedded in the Data Protector media agent.

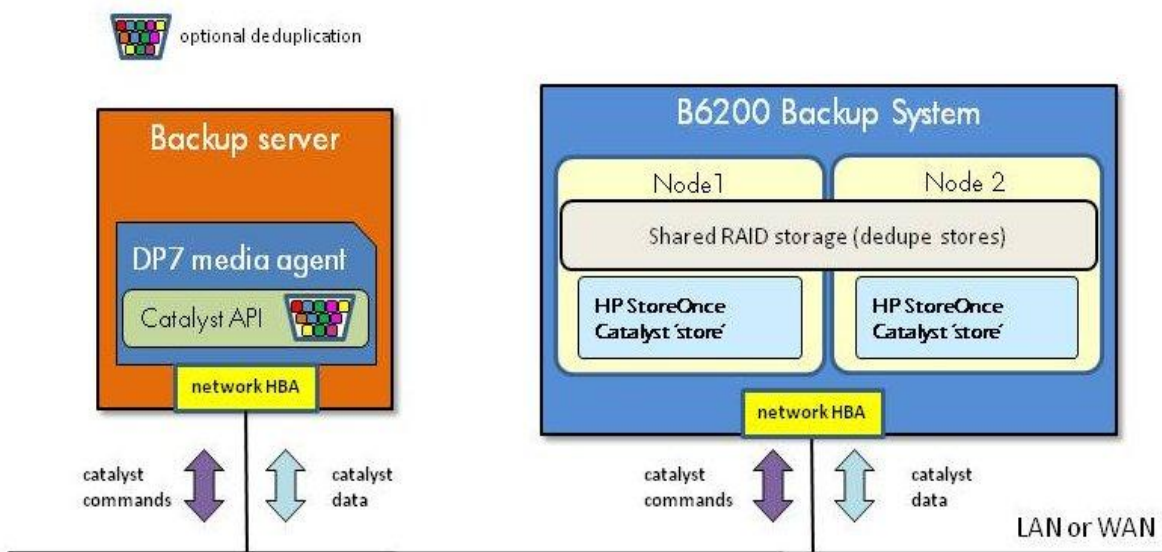


Fig.1: StoreOnce B6200 data paths to Data Protector backup server

Key Points:

- HP StoreOnce Catalyst is a unique interface and is fundamentally different from virtual tape or NAS.
- Optional deduplication at the backup server enables greater performance and reduced bandwidth requirements. This can be controlled at backup session/job level.
- Enables advanced features such as duplication of backups between appliances in a network efficient manner under control of the backup application.
- HP StoreOnce catalyst protocol runs on a standard IP network.
- Enables space occupied by 'expired' backups to be returned for re-use.
- Enables asymmetric expiry of data
- Enables store creation if required from within Data Protector.
- Backup jobs can restart automatically if the B6200 has a node failover condition. (Requires a restart script).
- Scalable licensing – pay as you grow

Introduction to the StoreOnce HP B6200 Backup System Architecture

The B6200 Backup System is designed for high availability and consists of up to 8 'nodes' arranged in pairs known as 'couplets'. Each node consists of a HP Proliant server with dual hex core Intel Xeon CPUs. These are linked with disk arrays for the data storage. The 'master disk' unit (one per node) contains dual 6Gbps SAS interfaces and each array is connected to the 'partner' node in the 'couplet'. This means that in the case of node failure the other node can access all the data stores. A basic architecture layout is shown in Fig 2. There is an internal 10GbE network for internal system management complete with dual ProCurve network switches. There is a 1GbE network which manages the network switches and connects to the iLO port on each server. The iLO port is used to shut down nodes in the case of malfunction. There is no shared storage between couplets. Each node has separate dual 8Gb fibre channel connections, dual 10GbE ports and dual 1GbE ports for connection to the customer's SAN and network infrastructure. All network connections are 'bonded' and are accessed by a 'virtual' IP address or VIF.

Each node has a service set which consists of the software modules which run the virtual tape, NAS and StoreOnce Catalyst deduplication devices. In failover the whole service set runs on the partner node in addition to its own service set. The shared storage is accessed by the additional 6Gbps SAS connection. (Disk controllers have dual 6Gbps SAS interfaces and the master disk unit has 2 controllers for resilience. The virtual IP addresses of the service sets remain the same so no reconfiguration is necessary. The management console will always be active on one node. If that node fails another node will activate its management console to take over. The management console network connection will be maintained throughout failover.

There is an interruption of service during failover and arrangements have to be made to restart backup jobs. This procedure varies according to ISV applications. As StoreOnce Catalyst protocol had regular 'checkpoints' backups can normally resume from the last checkpoint. This feature is known as "Autonomic Restart".

Power supplies to the B6200 are all dual (n+1) and it is highly recommended to arrange for dual mains power supply.

Each node has a usable capacity of 64TB. If this is exceeded a node can use storage from its partner node in the cluster. However performance is compromised if this overflow mode is used so best practice is to avoid using it. And of course there may not be space available on the partner node.

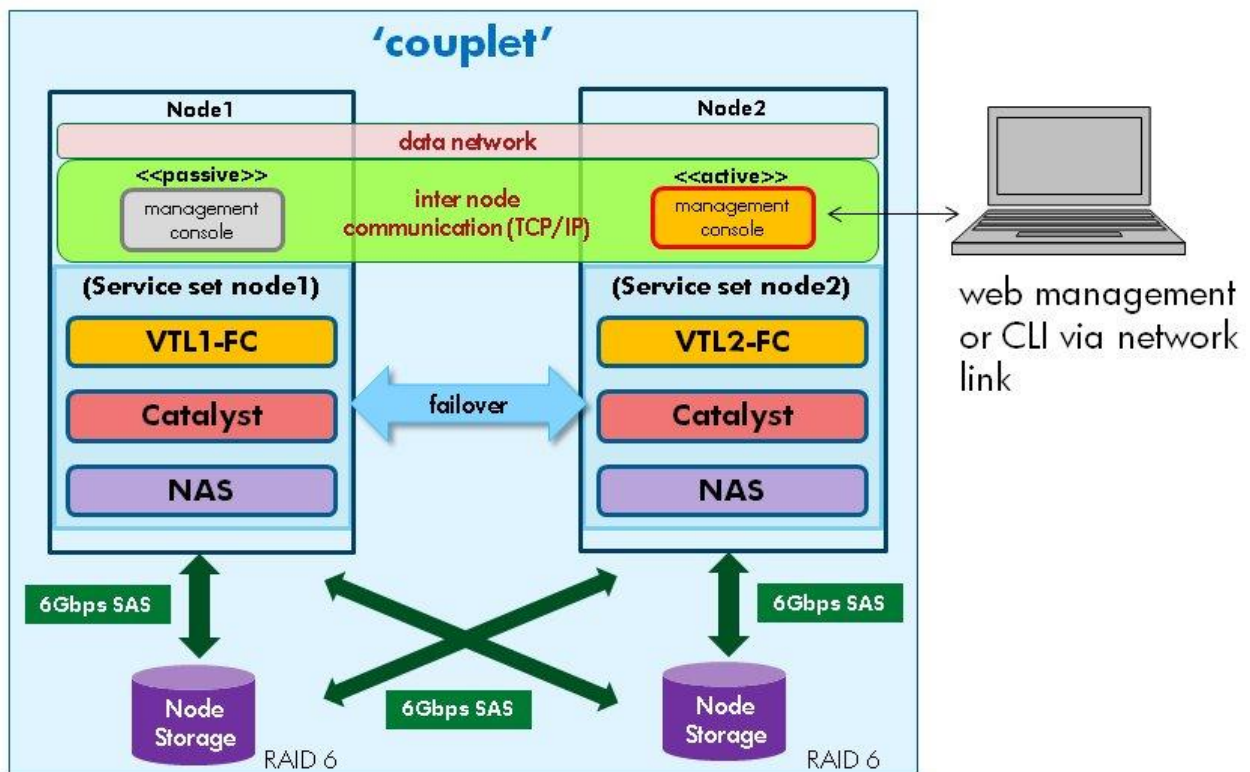


Fig 2. B6200 Architecture

Key Points:

- Maximum configuration is 64TB per node up to a maximum of 8 nodes.
- Nodes operate in pairs known as 'couplets'.
- Autonomic restart operates at a couplet level.
- There is NO user access to the internal networking or iLO ports.
- User networks connect to the node network ports.

- The ProCurve switches are for the internal network only.
- Customer data and management access is directly to a node.
- All network connections are 'bonded'. No special switch configuration is required.
- 1GbE and 10GbE network connections available for user management and/or data.
- HP StoreOnce Catalyst uses the network connections only.
- Fibre Channel is for Virtual Tape emulation only.
- Dual mains supplies are required and there are 4 power connectors per rack.
- Use the HP B6200 planning guide to select the correct power connection and to plan the networking.
- 10GbE is essential for large configurations which require maximum performance. 10GbE supports copper or fibre connections (10GbE SFPs are NOT supplied and need to be ordered separately)

Terminology – HP StoreOnce B6200

Terminology is important to understand and can be confusing. This section will cover the terms used in the HP B6200.

Term:	Details:
Cluster	Generally used for the whole HP B6300 appliance
Couplet	Consists of 2 interconnected HP Proliant servers each with disk storage. Failover occurs within a couplet.
Node	HP Proliant server hardware with SAS attached disk arrays
Management console	Software residing on all nodes managing the cluster and monitoring for node failure. Accessed via virtual ip address. Based on Ibrix fusion manager.
Bonded network ports	A method of combining network ports for either resilience or performance. The ports have each a physical IP address but a common virtual IP address. The MAC address appears the same. Sometimes this is known as link aggregation or in Windows 'teaming'.
VIF	Virtual Interface. VIF has an IP address for access to the management GUI.
Service set	HP term for the group of software modules which provide the Virtual Tape, NAS or StoreOnce Catalyst storage devices.
Autonomic Failover	The process in which a node is shut down (or fails) and the software modules all move over to the partner node. Failover only can occur within a 'couplet'.
Failback	Manual process which is the reverse of failover. Used following a failover.
Chunk	The unit of data in which the HP StoreOnce deduplication process is divided up a data stream. The average chunk size is 4KB.
Low-bandwidth backup	This refers to a StoreOnce Catalyst Backup where deduplication takes place at the client – in the case of data protector the media agent. After the first backup only new data is sent over the network to the B6200 catalyst store.

High-bandwidth backup	This refers to a StoreOnce Catalyst Backup where there is no deduplication at the client (data protector media agent).
-----------------------	--

HP B6200 Catalyst Store Configuration Specifications

	Per node***	Per Couplet	Per Cluster (max config)
Maximum number of catalyst stores*	48	96	384
Number of 'clients' supported per store*	no limit	no limit	no limit
Maximum number of inbound backup/restore/ copy jobs**.	192	384	1536
Maximum number of outbound copy jobs	48	96	384

*This maximum includes and NAS shares or VTLs configured.

** If Data Protector backups are using multiple streams each stream counts as a job.

*** More accurately per service set as node in failover will run 2 service sets and could support 96 catalyst stores. Of course performance is reduced in failover.

HP StoreOnce Catalyst best practice – What information is required for planning?

In order to achieve success with HP StoreOnce Catalyst it is important to ask the correct questions at the planning stage. There are many different data protection scenarios deployed by customers. To understand some of the important variable it is best described in several sections. This paper will not go in to depth but regard it as a starting point. However it will highlight the next steps and what tools are available.

Topology and geographical location of data

Firstly it is important to determine the number of servers which require data protection and where they are located. The tendency in recent years has been to centralize servers in order to simplify management and maintain much higher service levels. This has been made possible by increased WAN speeds. But some business models particularly in the retail sector still maintain servers at multiple sites for greater resilience but keeping central data warehouses. Any server located on a remote site

still needs backup and the preference is not to require staff on that site to be involved. A table of sites and the size of the local data storage is most useful as a starting point. Try and build a table similar to the example below. For enterprise customer this will of course be much larger and there may be multiple central locations.

Site location	Datacenter/ROBO	#servers	Data volume	Retention period	Data type	WAN links
Mainsite #1	main datacenter	100	80TB	6 months	database/file data	1GbE
Mainsite #2	main datacenter	100	80TB	1 week	database/file data	1GbE
remote #1	ROBO	5	6TB	1 week	file data	10mbps
remote #2	ROBO	5	5TB	1 week	file data	2Mbps
remote #3	ROBO	2	5TB	1 week	file data	2Mbps
remote #4	ROBO	5	5TB	1 week	file data	2Mbps
remote #5	ROBO	3	4TB	1 week	file data	2Mbps
remote #6	ROBO	2	4TB	2 weeks	database/file data	2Mbps
remote #7	ROBO	3	3TB	2 weeks	database/file data	2Mbps

Working from this data start to plan the solution. In the example above the customer has 2 data centers and a number of remote sites. The WAN links may already be in place in which case the sizing will dictate the backup window. It will also be necessary to determine how many server will have their own media agent and backup directly to the HP StoreOnce B6200.

The key decision is whether the local sites need to keep data onsite for faster recovery and data is stored locally or whether all backups are held in the main data centers. Before StoreOnce Catalyst it was necessary to install a local D2D system and replicate back to the main data centre. This is still an option with catalyst but now HP Data Protector 7 with catalyst support in software changes the options. Any server with a media agent loaded can perform a low-bandwidth backup over a WAN link. Only new data is transferred at each subsequent backup. There is no additional charge for Data Protector media agents. It is also possible to backup locally using software deduplication and object copy the data to the central data centers.

Key Point:

- If fast local recovery is required then the last backup should be held locally as restore are NOT 'low-bandwidth'.
- Although supported the HP catalyst protocol for backup is best restricted to national WAN connections but replication has been designed for international WAN links with higher latency values.

Type of data

Deduplication performance varies with the type of data and how it is being stored. Data falls into 2 general classes: structured and unstructured. Structured data would be database files which are

application specific. Unstructured data is normally stored in a standard filesystem and can vary in content. Some data such as files which have a degree of compression and encrypted data cause poor deduplication performance. Most common database applications are Microsoft SQL server and Oracle. Data Protector has agents for both these products. Although not a database as such Microsoft Exchange is structured data and has a dedicated Data protector 'agent'. Note that the data type is unimportant to HP StoreOnce Catalyst technology but has performance implications and implications for Data Protector Software.

Key Points:

- Best practice is to keep similar data in the same catalyst store. E.g. dedicate a B6200 catalyst 'store' to Oracle backups and a different store for SQL server.
- The number of Data Protector 'client' systems which need to be backed up per data protector cell is dependent on the number of unique file names. Server with large complex file systems put a greater load on the cell manager internal database. Guidelines are around 300 'clients' per cell manager. Multiple cell manager can be controlled by data Protector Manager of Managers option (MoM)
- Incremental backup ideally should be in a separate store from full backups. (This is not always possible in certain customer rotational schemes.)

Virtual Machines

It is likely that customer will have extensive virtual machines to back up. These normally achieve very high deduplication ratios. Data Protector 7 is well equipped to back up virtual machines. Keep these in a similar store.

How the data is stored also has implications because later in this document we will discuss how to set up multiple streams for performance.

Data Volumes, change rate and retention periods.

These are all essential parameters in designing the best solution. The B6200 has a maximum practical capacity of 64TB per node. Do NOT exceed this limit in the design phase. The system can use storage if available on its partner node but this could lead to loss of performance. The retention period is important because deduplication works best with retention of multiple backups. General best practice is not to retain data on disk for longer than 12 months. Data Protector can easily migrate data to tape for longer storage periods. Change rate is more difficult to estimate. The incremental backup sizes give some idea but for accurate information a trial run with an actual D2D would be recommended. Customers typically increase their data volumes on an annual basis and it is recommended to estimate this. Having accumulated this data then the HP Sizer tool will make the appropriate hardware recommendation.

Replication over the WAN link can be sized by the Sizer Tool and either the replication time window for a fixed WAN speed can be given, or a link speed for a given replication 'window' requirement.

In regards to sizing HP StoreOnce catalyst is very little different from VTL and NAS. The only additional feature is backup from clients directly to the catalyst store over the WAN.

Mixed StoreOnce Catalyst, VTL and NAS environments

It is highly likely that customers will be using VTL and NAS at the same time as StoreOnce catalyst unless this is a completely new project. This is perfectly possible as it causes minimum disruption. The number of VTLs, NAS shares and catalyst stores per service set (= a node) is limited to 48. The different devices can exist in any combination. The replication and limits for VTL and NAS devices are separate from catalyst job limits.

Key Points:

- For VTL and NAS replication it is necessary to purchase a replication license for each node (or target D2D system) which will be hosting a replication 'target'. Separate catalyst licenses are required and are covered later.
- If specifying a combined NAS/VTL/StoreOnce Catalyst system pay attention to the limits set per service set.

Optimizing Data Streams

As part of the planning stage and using some information about the customer environment it is important to realize that to achieve decent performance it is recommended that the backups should use multiple streams to a StoreOnce catalyst device. To achieve multiple streams the source data must be either from separate mount points or can be specified as directory 'trees'. However if server end user can modify the directory structure this is not always possible. For Oracle/SQL Server and Exchange this is achieved either automatically or in the case of Oracle by a RMAN parameter.

For customer from a tape background please note that one tape drive can normally handle one stream of data. However several sources of data could be "multiplexed into one tape drive. This is not possible with a StoreOnce Catalyst store. Please note that for VTL users multiplexing results in poor deduplication ratios. Multiplexing was popular for real tape drives because it kept the tape drive 'streaming'.

Key Points:

- Multiple streams are recommended for best performance with HP StoreOnce catalyst.
- Multiplexing cannot be configured within Data Protector 7 with catalyst devices. (Known as 'Concurrency' within Data Protector)
- It is necessary to select source data correctly for multiple streams. E.g. for filesystem backup separate mount points, drive letters or directory selections are required.
- Backup servers running multiple streams and deduplication need to be sized appropriately. Later in this document there are guidelines.
- Allocate backup server to different nodes in order to balance the load across the multi-node system. This will maximize the throughputs.
- Use the sizer tool – this is calibrated with the latest test results from HP R&D and will take into consideration data retention, data change and growth. It will also size any WAN links.

HP StoreOnce Catalyst and Data Protector 7 integration

HP Data Protector 7 has built in support for the HP StoreOnce Catalyst API. Minor version 7.01 of Data Protector is required in order that data copied between catalyst stores does not require rehydration before previous ongoing transmission. To upgrade top this version it is necessary to install the patch DPWINBDL_00701 (for Windows environments) or patch DPLNXBDL_00701 for Linux. As the HP StoreOnce Catalyst API is built in there is no requirement for a 'plug-in' software module as in Symantec NetBackup. HP Data Protector has the ability to select whether to perform deduplication at the backup server. This paper will cover the configuration and best practice for HP StoreOnce Catalyst stores using HP Data Protector as well as some of the advanced features such as duplication of backups using object copy.

Configuration of a StoreOnce Catalyst device in HP Data Protector:

Introduction

HP Data Protector classes StoreOnce Catalyst data storage as a 'backup to disk' device type. This device type consists of 2 different devices:

StoreOnce Backup System - this refers to the hardware based StoreOnce appliance based storage.

StoreOnce Software Deduplication – this refers to the software based deduplication store with local disk storage referred to in Data Protector 7 as a StoreOnce Library

(This paper will only cover the use of the B6200 StoreOnce appliance based catalyst device.)

The data paths from clients to the deduplication stores are similar for both type of devices and introduce the 'gateway' concept.

HP Data Protector 7 has the ability to perform deduplication at backup server using StoreOnce software embedded in the media agent.

Terminology – HP Data Protector 7

Earlier in this document the HP B6200 terminology was explained. The following terms apply to HP Data Protector 7 and StoreOnce Catalyst. For a more comprehensive glossary please consult the HP data Protector 7 concepts guide.

Term:	Details:
Cell Manager	The main server in a Data Protector Cell where the Data Protector software is installed. The cell manager controls all backup and recovery procedures. The cell manager also contains the internal database (IDB). Each DP cell has one cell manager.
Media Agent	The software module which enables access to storage devices – tape/ANS/StoreOnce catalyst. Receives data from disk agents either local to the server or via a network connection. The StoreOnce Catalyst API is embedded within the media agent software from Data Protector version 7. StoreOnce deduplication can be enabled.

Disk Agent	The software module which reads and writes data to or from a disk (or disks) and then sends it to a media agent.
Gateway	The gateway is used as the destination 'device' for backups to the B6200 catalyst store or a software-based StoreOnce Library. (Software-based deduplication store). Each data protector media server can have access directly to the StoreOnce catalyst store via a gateway.
Manager of Managers (MoM)	A method of grouping multiple Data Protector cells and managing from a central cell. Essential for very large installations.
IDB	The internal Data Protector database which maintains a record of data backed up and the media it resides on as well as all library/NAS share/catalyst store configuration.
Backup to disk device (B2D)	A disk based storage device target for Data Protector Backups. Can be either a software based deduplication store or B6200 appliance.
Source side deduplication	Deduplication takes place within the backup server. The StoreOnce deduplication code is embedded within the Data Protector media agent. Source side deduplication uses the 'implicit' gateway. Can only backup data which is located on the same server.
Server side deduplication	Deduplication takes place within the backup server as in source side deduplication. Can only use the explicit gateway. Can be selected for backing up data from other servers with DP7 disk agents as well as data held itself. Data is transferred via network connection to backup server and then will use the "explicit" gateway.
Target side deduplication	Deduplication all takes place in the B6200 system. Use the 'explicit' gateway only but with server side deduplication not selected.
Backup server	A dedicated server running the HP Data Protector media agent software. The server can send data to the backup to disk device via a 'gateway'. (This is then equivalent to a Symantec NetBackup 'media server')
Application server	A server running software applications for end-users. May also have a media agent installed or just a disk agent. Normal practice is to install media agents on large application servers so backups can be made directly to the backup appliance or tape drive instead of data passing by the network to a backup server.
Backup object	Defined as a backup unit which contains all the items from one disk or volume (logical disk or mount point). Can also be a raw disk image or a database entity.
Synthetic full backup	Supported with StoreOnce Catalyst. This is a technique where after one full backup is performed only incremental backups are performed but Data Protector merges these into a 'synthetic full' backup.
Virtual Synthetic full backup	Not supported with StoreOnce Catalyst. A more efficient version of 'synthetic full backups'. Use a special distributed filesystem for faster recovery of a full backup.

Access to the catalyst store is via a 'gateway' concept. The 'gateway' is roughly comparable to the tape drive device as used for real or virtual tape configurations. Hosts which are required to backup or restore data from a B6200 Catalyst Store require a gateway. The gateway really defines the nature of

the host to B6200 catalyst store access via the Data Protector media agent. The block diagram in Fig. 3 below illustrates the basic Data Protector 7 usage with the B6200 with a catalyst store.

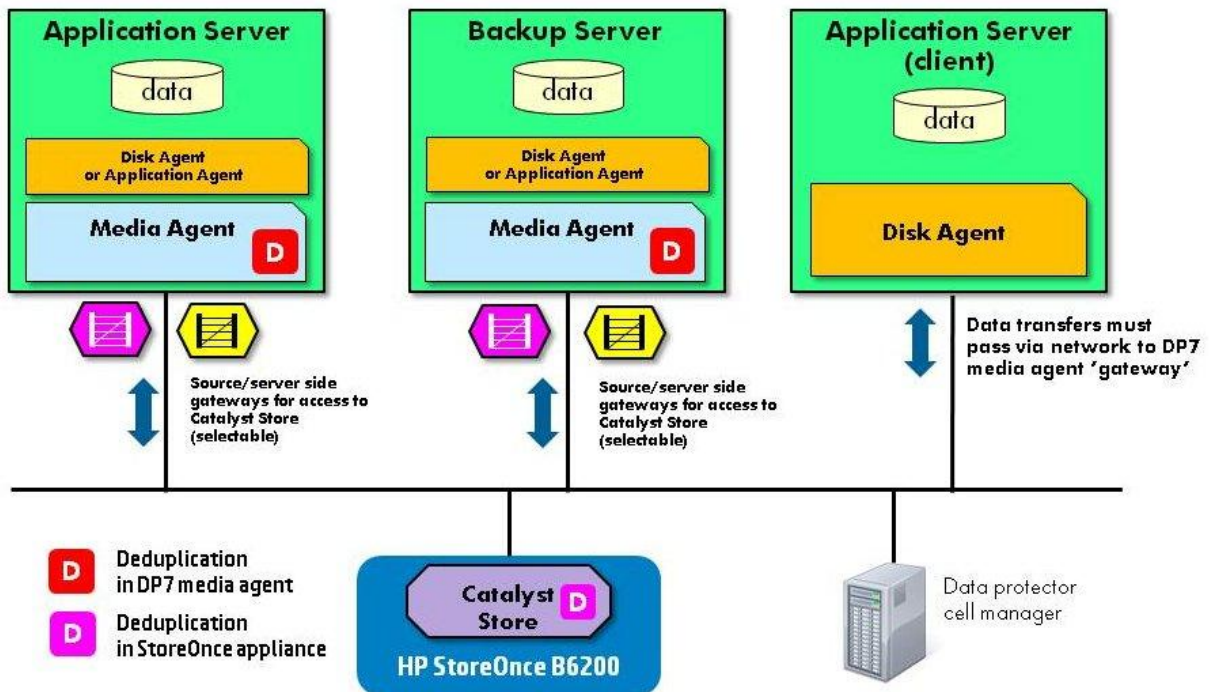


Fig. 3 Data Protector 7 'Gateways'

In Fig.3 it shows 3 servers and a B6200 system. Two of the servers have both media agents and disk agents installed. The 3rd application server has only a disk agent so does not have any gateways. All its backups and restores must go over the network to a server with a media agent.

Key points:

- Data Protector 7 gateways give access to the StoreOnce Catalyst store for backups and recovery.
- The gateway controls whether part of the deduplication process runs on the backup server.
- If deduplication occurs within the backup server then the data transfer is said to be a low bandwidth backup.
- Use the gateway configuration to control loading of the backup/application server.
- The gateway configuration controls the maximum number of data streams which can be sent to the catalyst store simultaneously.
- Data Protector 7 clients can backup over LAN or WAN directly to the B6200 StoreOnce Catalyst appliance if required. There is no additional license charge for additional client servers running media agents.
- Application/File server running just the disk agent can only backup up to a backup server running a DP7 media agent

- The DP7 cell manager should always run on a separate server.

The Data Protector Gateway types:

Within data protector there are 2 types of gateway, implicit and explicit. The configuration/selection of the gateway can specify whether part of the deduplication process (hashing/chunking and data compression) is performed by the server hosting the Data Protector media agent or within the HP StoreOnce B6200 'appliance'. The number of media agent processes (and hence data streams) is also part of the advanced gateway configuration. Fig 4. shows gateways configured for a backup to disk device.(StoreOnce B6200 appliance).




Fig. 4 Data Protector 7 devices configuration showing configured gateways

The source side or '**implicit**' gateway is configure once only but can be used by any of the clients in the cell which have a media agent installed. The implicit gateway does not have an assigned media agent but will start media agents as required on any server equipped with media agent software. In effect it like a 'virtual' gateway for every media agent equipped server in the cell where 'source-side' deduplication is specified for backup. The gateway will have the same configuration parameters on every media agent equipped backup server. This gateway is designed so that only files or data resident on the server can be backed via this gateway. Files or data resident on an application server with only the disk agent or application agent installed cannot be backed up or restored via an 'implicit' gateway.

The implicit gateway always invokes deduplication in the media agent and is referred to as 'source-side' deduplication. Object copy is not available using the implicit gateway. (Object copies of data which used the implicit gateway in the backup process can be 'remapped' to use an explicit or server-side gateway). Configuration of the 'source-side' implicit gateway is optional.

The server-side or "**explicit**" gateway is assigned individually to any server running the Data Protector 7 media agent. The 'explicit' gateway configuration can specify the maximum number of streams and whether part of the deduplication process is performed in the media agent which results in a low-bandwidth transfer of data. Selection of deduplication in this case still takes place in the media agent but the backup job specification must specify 'server-side' deduplication. Backup data can reside on other servers and be directed to a 'server-side' gateway by network transfer.

Deduplication Selection with HP Data Protector 7

Data Protector 7 offers 3 different deduplication options:

Source side deduplication – deduplication of data is performed within the server (or client hosting a media agent). After the first backup only new data 'chunks' are sent across the network to the B6200 hardware based HP StoreOnce Catalyst store. Source side deduplication is selected when creating a backup job specification by ticking the 'source side' deduplication box in the initial screen. Can **only** backup data stored locally to the backup server. Only the 'implicit' or source side gateways can be selected if source side deduplication is specified in a backup job. All other gateway destinations will be 'grayed out'. This is used for multiple clients (equipped with media agent and disk agent) because they can all be configured globally instead of setting each individual gateway as in server side deduplication. This is used for direct backup to the catalyst store from the 'client'. Hence the name 'implicit gateway'.

Server side deduplication – deduplication of data is performed within the dedicated backup server. Server side deduplication can be used for data held locally on the backup server **and** from other servers which have a disk agent installed. In this case data is transferred over the network to the backup server and then processed by the media agent and sent on to the catalyst store. Selecting server side deduplication in a backup specification requires the use of the 'explicit' gateway for the backup destination.

Target side deduplication- data is held on client with only a disk or application agent installed. This system can be remote from the backup server. Source or server side deduplication is not selected. All data is transferred at high bandwidth across the LAN or WAN to a backup server hosting a gateway to the StoreOnce Catalyst appliance. This may be necessary when Data Protector 7 has only application/disk agent support for a particular data type (E.g. OpenVMS backup).

Key Points:

- The implicit gateway is used for source side deduplication on any server in the cell running a media agent. A server running just a disk or application agent cannot select the implicit gateway as it is restricted to data held on the server running the media agent.
- The parameters (max. stream etc.) are the same for every server using the implicit gateway. Useful for limiting server loading.

- At least one explicit gateway must be configured. You cannot configure just an implicit gateway.
- For files or data held on application server with only a DP7 disk agent installed backups must be directed to an explicit gateway using server side deduplication.
- Target side deduplication is useful when the extra load of deduplication is not wanted on the backup server and can only use an explicit gateway.
- Only 64-bit servers can be configured for a gateway.
- The deduplication process is the exactly the same for server-side and source-side deduplication.
- The StoreOnce Deduplication and catalyst client binaries are built in to the media agent code. There is no requirement for a 'plug-in' software module as required for Symantec NetBackup and Backup Exec integration.

Configuration of StoreOnce Catalyst Stores in Data Protector 7

This section covers the configuration of a catalyst store and the provision of gateways to the media agent equipped servers.

Data Protector 7 has the unique ability to create StoreOnce Catalyst stores by itself from the Data Protector management GUI. This is optional and of course stores can be created via the B6200 management GUI and then integrate into Data Protector. This guide will cover both methods.

For the purpose of these notes there are 3 servers in the configuration and a B6200 system. The server 'Zen' is the cell manager and the servers Bill and Ben are clients with the DP 7 media agent installed. Server 'Zip' just has a disk agent loaded. The B6200 is a 2 couplet system and uses a 10GbE network for data and a 1GbE for management (known as template 1).

Creating a catalyst store using the B6200 management GUI

1. Login to the management GUI. From the left hand navigator section expand B6200 StoreOnce menu. Then expand the StoreOnce Catalyst Store menu and click on 'stores'. The menu screen shown in Fig.5 below displays all the existing stores. The 'Create' button is located at the top right of the screen.

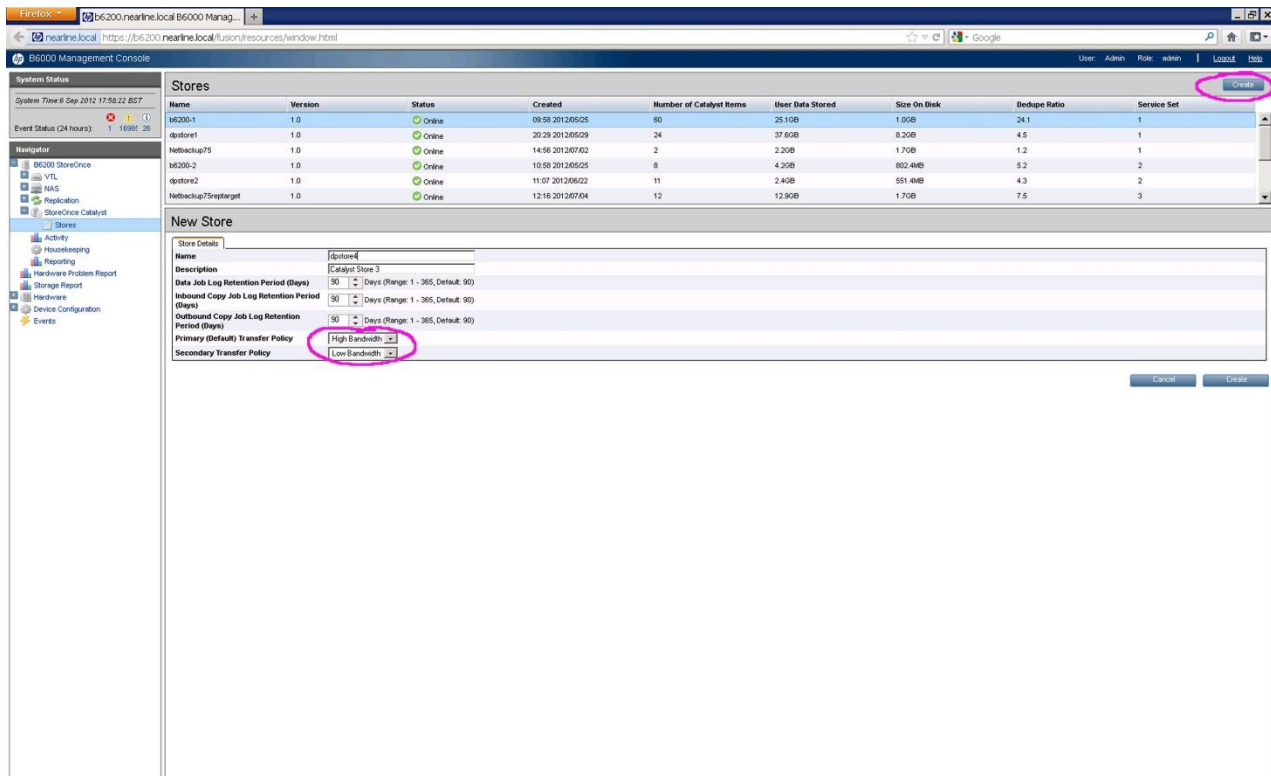


Fig. 5 Creating a StoreOnce Catalyst Store.

2. Select the service set you wish to create the store on. The 'pop-up' window will display each service set and how many devices are remaining. Each service set supports up to 48 stores, VTLs or NAS shares in any combination. (Note in this example it is a 4 node system, therefore there are 4 service sets available.)
3. The high/low bandwidth selection can be left at default as Data Protector 7 will control this setting via the catalyst protocol.
4. Allocate a store name and description as desired and select create on the left of the screen.
5. The store is ready for use. However it may be desirable to control access by setting permissions.



Fig. 6 Creating a StoreOnce Catalyst Store

Controlling Access to the catalyst store:

The access permissions control access to each store. This is desirable for reasons of security. Firstly highlight a 'service set' and then select StoreOnce Catalyst from the left hand navigation panel of the B6200 management GUI. Select the edit function located on the right hand side of the menu (not shown here and then check client access checking. Note that in this screen there are setting for maximum concurrent inbound copy jobs and outbound copy jobs and the port numbers used for catalyst command and data protocol. Select the clients tab.

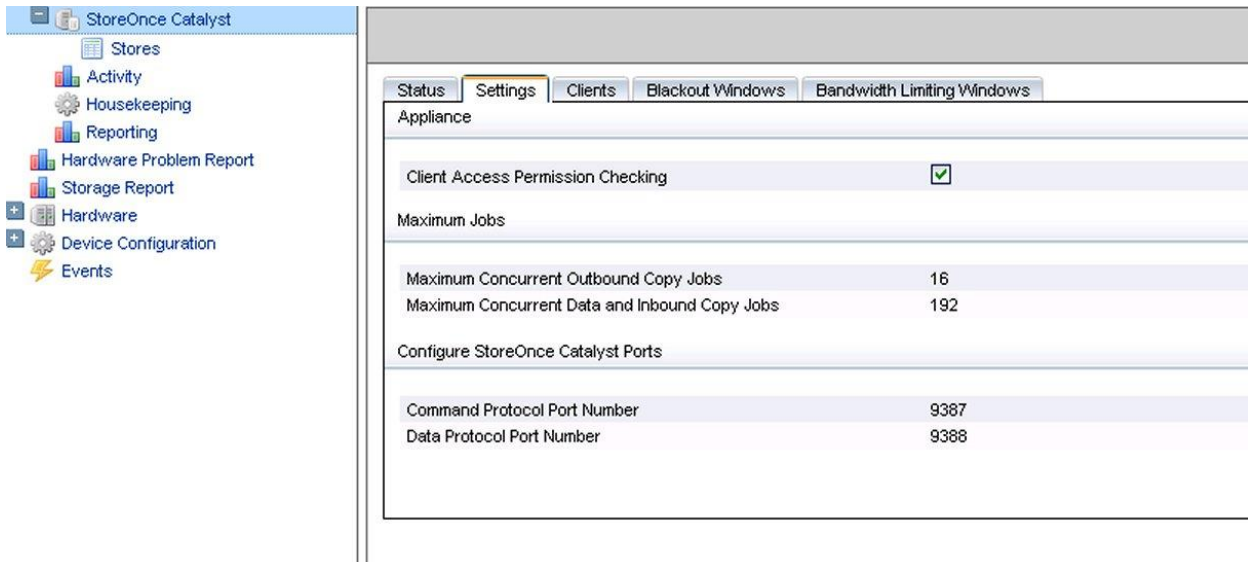


Fig.7 Client Access – settings.

Once the clients tab is selected clients can be added as shown in Fig. 8

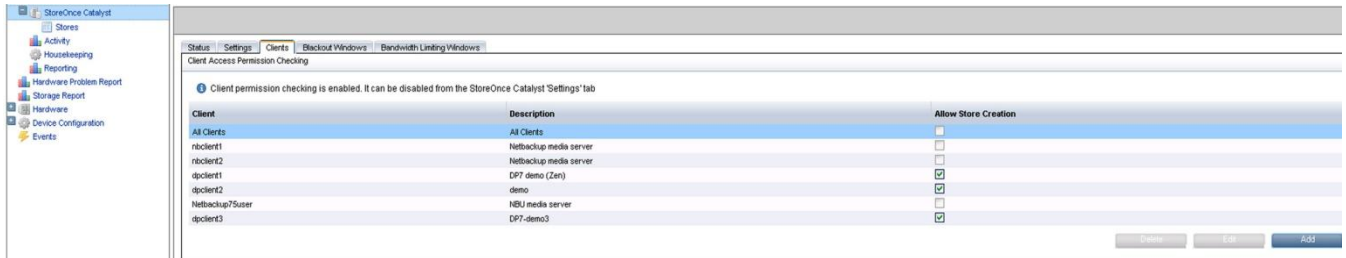


Fig. 8 Client selection

Select the 'Add' button and then add the client name and descriptions. Note the box to enable store creation. If required that HP DP7 creates its own store and client access is required it will be necessary to add a client in advance.

Once the the client is created then at the store level access level can be set. Select the store from the left-hand navigation menu and then select the permissions tab. The screen is as in Fig.9 below.

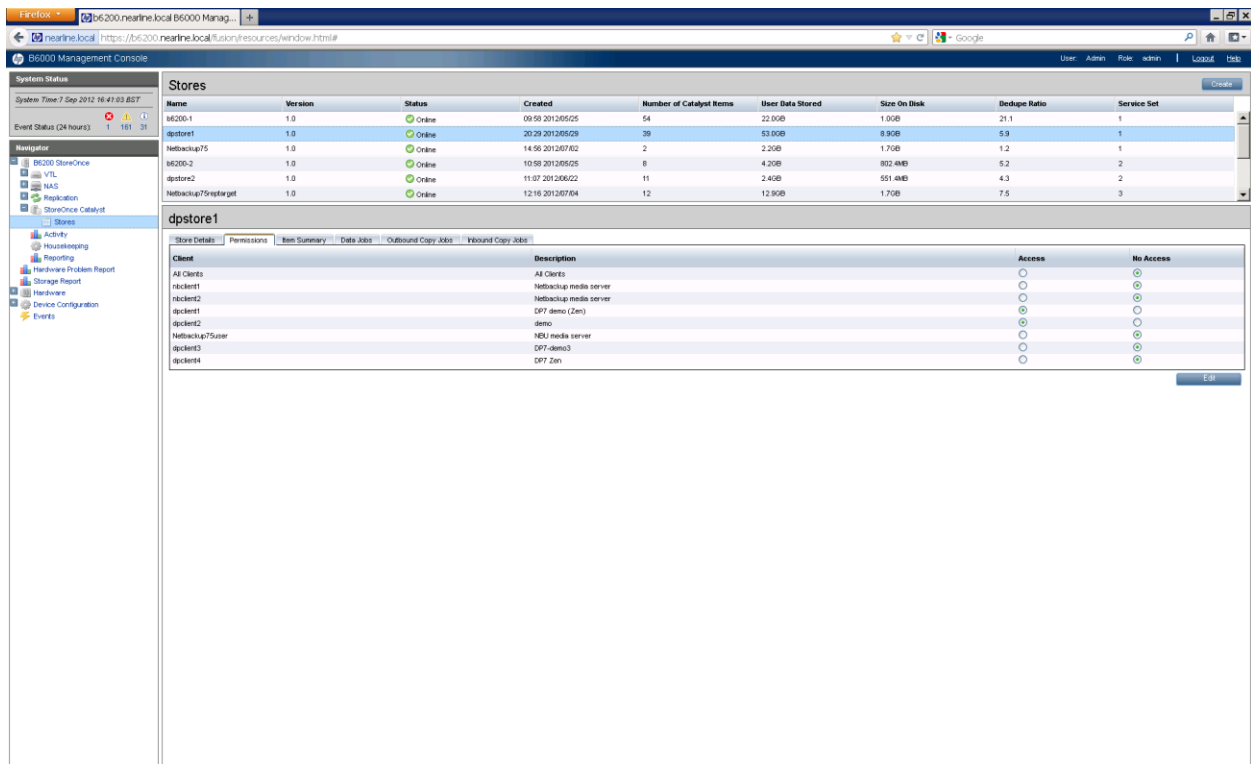


Fig. 9 StoreOnce Catalyst Store access permissions.

Key Points:

- Client Access permissions checking is set for each service set.
- Access to individual store is set at store level.

Configuration of HP Data Protector 7

Having configured a store on the B6200 system now proceed to the Data Protector management client. This guide will cover configuration of both gateway types. The cell manager is on server 'Zen' and following HP data Protector best practice this is a separate server. The cell manager creates quite a loading which is not desirable on a media server. There are 3 other servers: Bill, Ben and Zip. The server 'Zip' has only a Data Protector 7 disk agent installed. This guide will show the creation of both 'implicit' and 'explicit' gateways. The B6200 is configured for 'template1' network configuration which uses 10GbE for data and 1GbE for B6200 management. DNS is in use and the B6200 VIFs (virtual ip addresses) will be referenced by their fully qualified domain name. For the purpose of this exercise there will be 4 catalyst stores configured: dpstore1 - 4. Fig.10 shows the example layout. As dpstore1 has been created in the steps outlined above the configuration of Data Protector can now proceed. The catalysts store will be called 'B2D1' and be configured as a backup to disk device. Start the Data Protector management GUI and select 'devices & media' from the drop down box at the top left of the page. Right click on 'Device' and select 'Add Device'. On the screen displayed add the chosen device name (B2D1), description (optional), the device type is Backup to Disk and the interface type is 'StoreOnce Backup System'. Select next to continue the configuration.

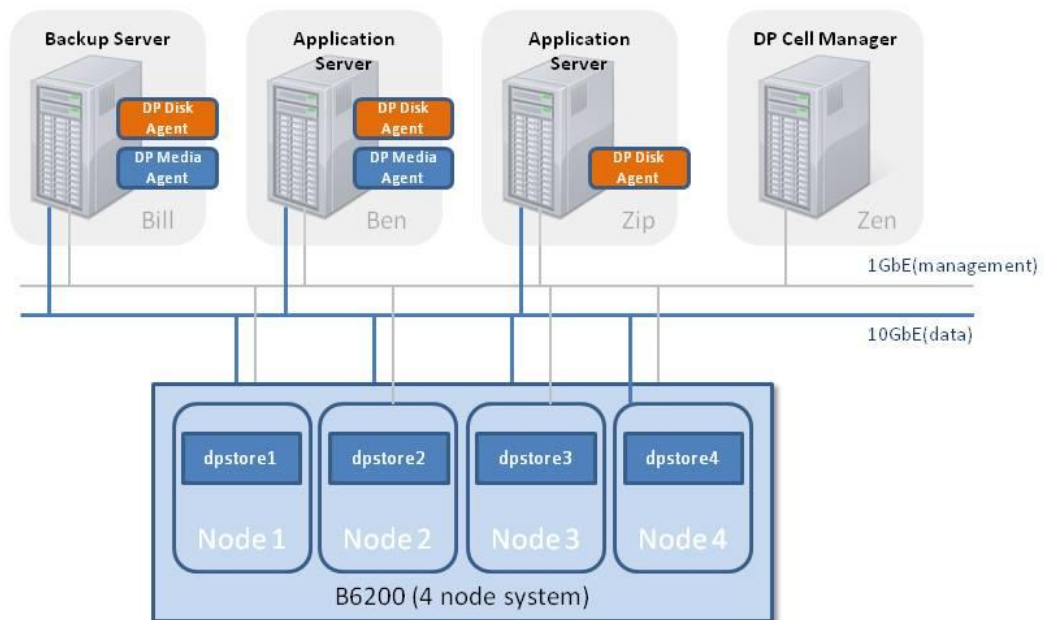


Fig. 10: Example B6200/Data Protector configuration

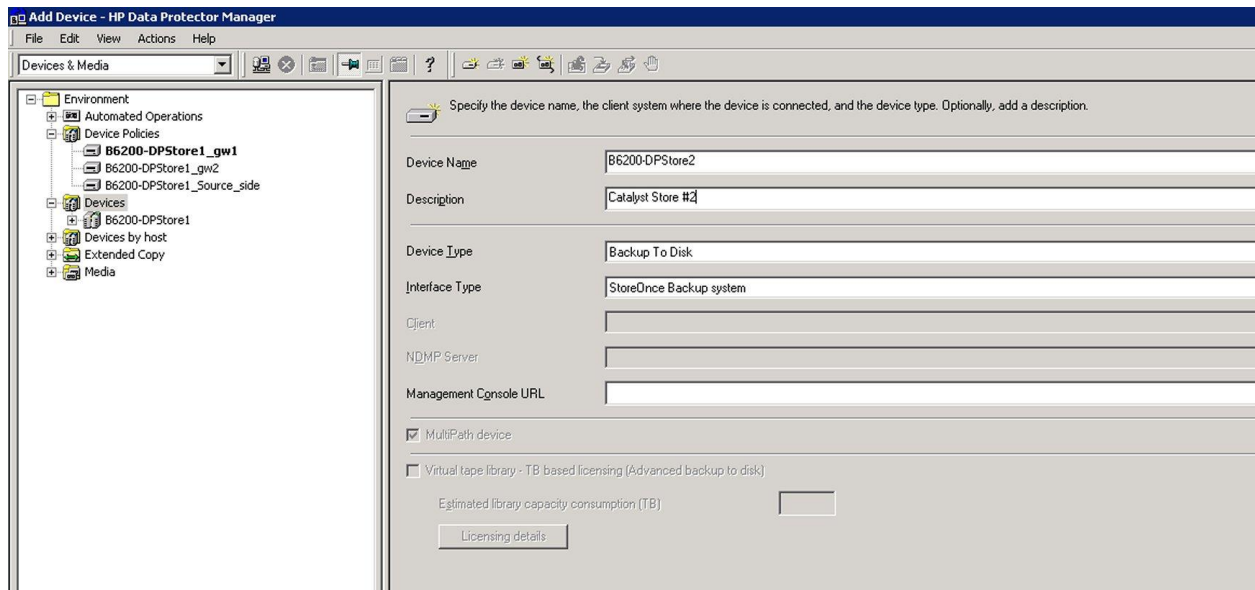


Fig. 11 HP Data Protector – creating catalyst device (1)

The next screen is used to select (or create) the HP StoreOnce Catalyst store located on the appliance specifying the VIF address. (Use ip address or fqdn) If client access permissions have been selected and a client name added to the B6200 store, then it is necessary to enter the client name to browse or create a store. If the client name is not entered any pre-made stores will not be accessible. The screen is shown below in Fig. 12. If the store has not been pre-configured on the B6200 a new store will be created providing the client ID is specified correctly.

The next part of the screen is used to configure an ‘implicit’ gateway (optional) by selecting ‘Source-side’ deduplication. The ‘Properties’ settings allow access to the advanced settings. The advanced settings allow the maximum number of streams per client to be specified. This has a default setting of 2 and with the implicit gateway will be applied by any media server when using source side deduplication. This is an important setting when optimizing the number of streams of data per catalyst store. The blocksize setting is not used for StoreOnce Catalyst but is still used by disk and application agents. It is recommended to increase this to at least 256KB.

The section below this adds ‘explicit’ gateways. At least one ‘explicit’ gateway must be configured. These gateways are applied individually to each data protector ‘client’ which has media agent software installed. The client server names are shown in the drop-down box (servers must be added as clients and the media agent installed prior to gateway configuration). Each server required to have a gateway is selected and then added. The ‘properties’ setting allows selection of ‘server-side’ deduplication if required. (You can alternatively right-click in the main window). If this box is not ticked then all deduplication will take place on the B6200 appliance and the backup will be in that case ‘high-bandwidth’ and is therefore ‘target-side’ deduplication. If the server-side deduplication box is ticked the deduplication occurs on the server hosting this gateway. Fig. 13 shows the ‘Advanced Settings’ options. The advanced setting also specifies the maximum number of streams per client. Default is the maximum available but the setting in the backup specification will set the limit. Optimization of streams per store will be discussed in a later section.

Once each gateway is configured use the 'Check' button on the right hand side (above 'properties') to check the communication with the backup server. This is important when the servers have dual networks (10GbE and 1GbE) as the data path to the B6200 is used for all catalyst data and commands. If using DNS both subnets must be capable of resolving their respective service set VIFs for data and the management VIF on the 1GbE network.

Select 'Finish' and the stores are now ready for use.

Key Points:

- The optional implicit gateway when selected will start media agents on any (media agent equipped) server but only for local data on that server. Uses the same settings for every server. Data for backup must reside on the same server. Used for source-side deduplication only. Useful for providing an overall limit on data streams to match backup server specification.
- The explicit gateways can be configured individually on each media agent equipped server which is registered as a client in the cell. Can be used for server-side or target-side deduplication. Can backup data which is resident on other servers via the network.
- For each data stream a media agent process is started. For each mount point a disk agent is started.
- The maximum number of connections per store can be set by Data Protector. By default the host appliance is to the physical limit of the store (192 per service set).

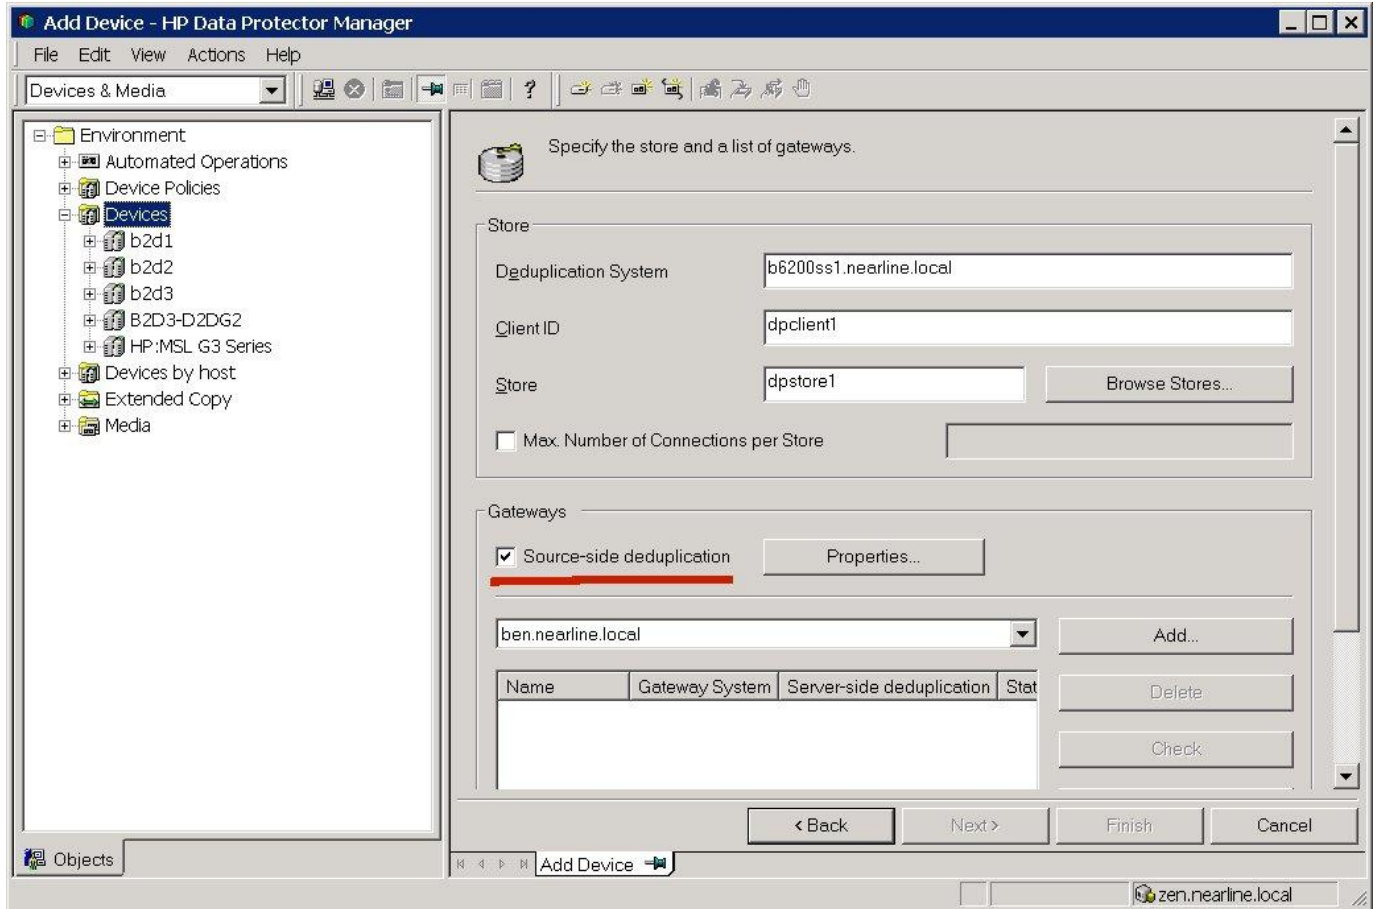


Fig. 12: Store & Gateway configuration.

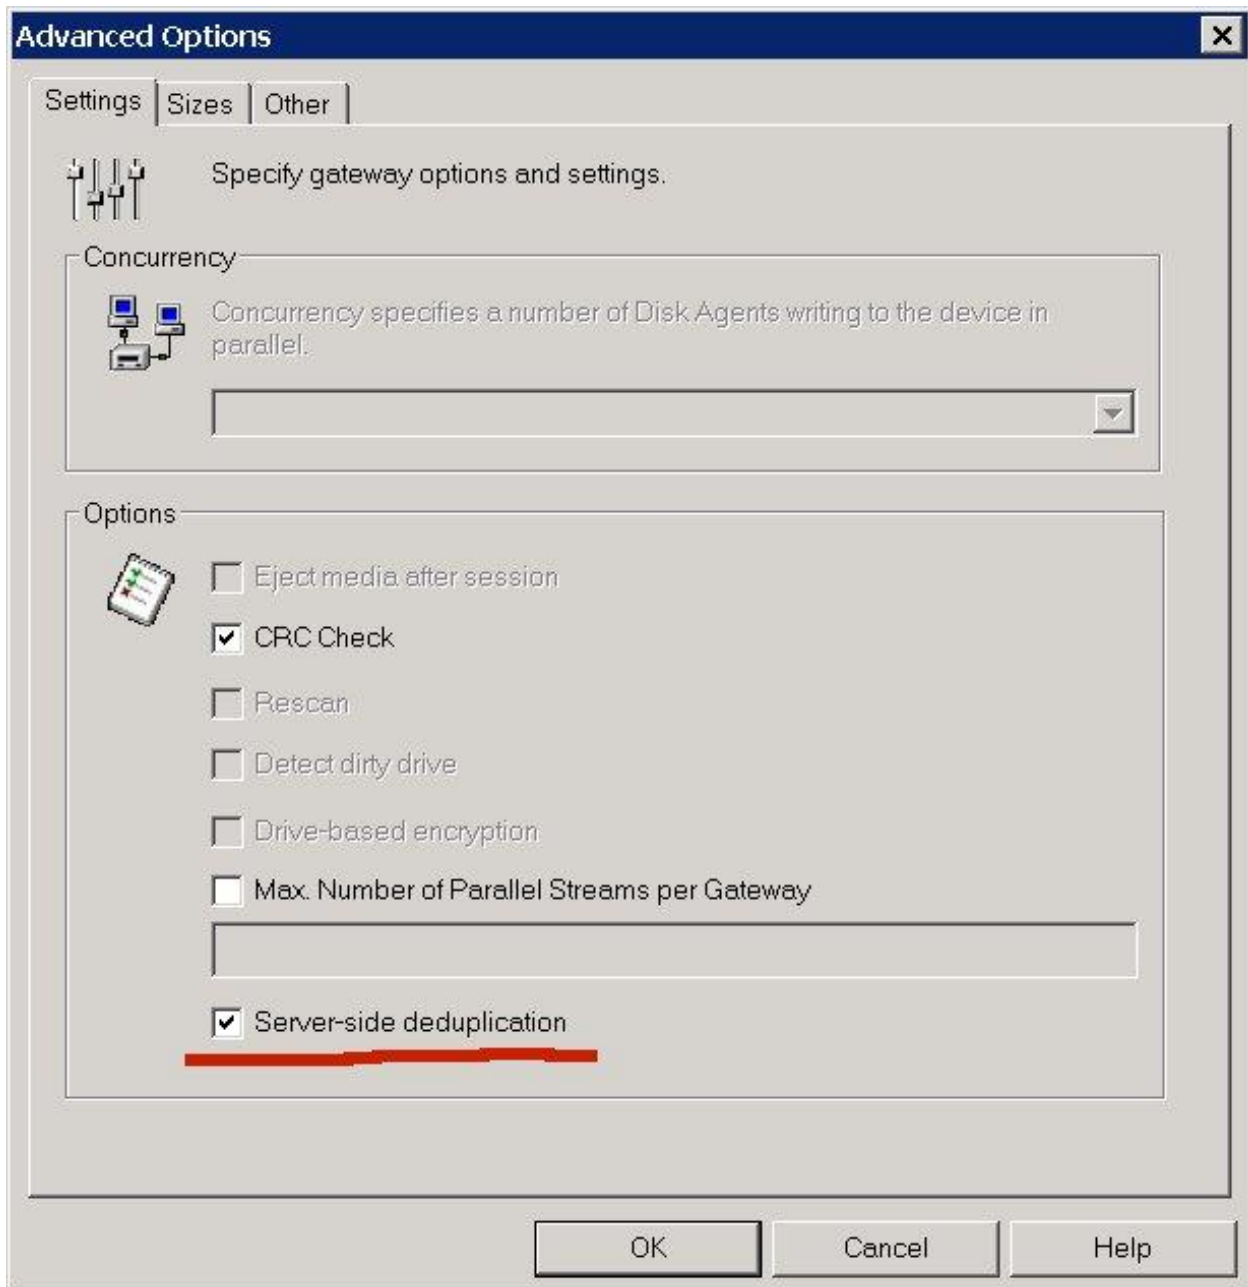


Fig. 13: Explicit Gateway configuration – Advanced Options.

Creating a Data Protector Specification for backup to a StoreOnce Catalyst Store

Using the demonstration configuration shown in Fig. 10, a backup specification will be created to perform a backup of data using source-side, server-side or target-side deduplication.

Backup using Source-side deduplication (implicit gateway)

Deduplication occurs at the backup server and the backup is often called low-bandwidth.

From the Data Protector Management screen select backup from the drop-down box at the top of the page. Right click on 'Filesystem' and select 'Add Backup'. Select 'Blank Filesystem' and check the 'Source-side' deduplication box. Click 'next'.

Select some files for backup and click 'next'. The software will now show the destination devices as shown in Fig. 14. below. The destination is the device 'B2D1'. Note that the explicit gateway is not available and is 'grayed' out. Select the source-side gateway. Remember this gateway is selectable only if 'source-side' deduplication is specified. Highlighting the gateway will allow the properties button to be selected. This is used to specify a media pool. There is a default media pool created for the 'backup to disk' device but additional media pools can be created if desired.

Click 'next' and specify the required options for retention. The backup specification options and schedule can be modified if desired. It is often useful to tick the 'display statistical information' box.

The backup specification can then be saved.

Backup using server-side deduplication (explicit gateway)

Deduplication occurs at the backup server but data can be backed up which is not stored on the server running the media agent. This is used for backing up clients which only have a disk agent loaded; in this case the server called 'Zip'.

From the Data Protector Management screen select backup from the drop-down box at the top of the page. Right click on 'Filesystem' and select 'Add Backup'. Select 'Blank Filesystem' . DO NOT check the 'Source-side' deduplication box. Click 'next'.

Select some files for backup from the server 'Zip'. Select next and the screen will show the destination devices. Expand B2D1 and now the 'Source-side' gateway is 'grayed' out and the explicit gateway on the backup server 'bill' is available. This gateway has 'server-side' deduplication selected in the advanced options. Select the options required and save the backup specification. Fig 15. shows the gateway selected.

Click 'next' and specify the required options for retention. The backup specification options and schedule can be modified if desired. It is often useful to tick the 'display statistical information' box.

The backup specification can then be saved.

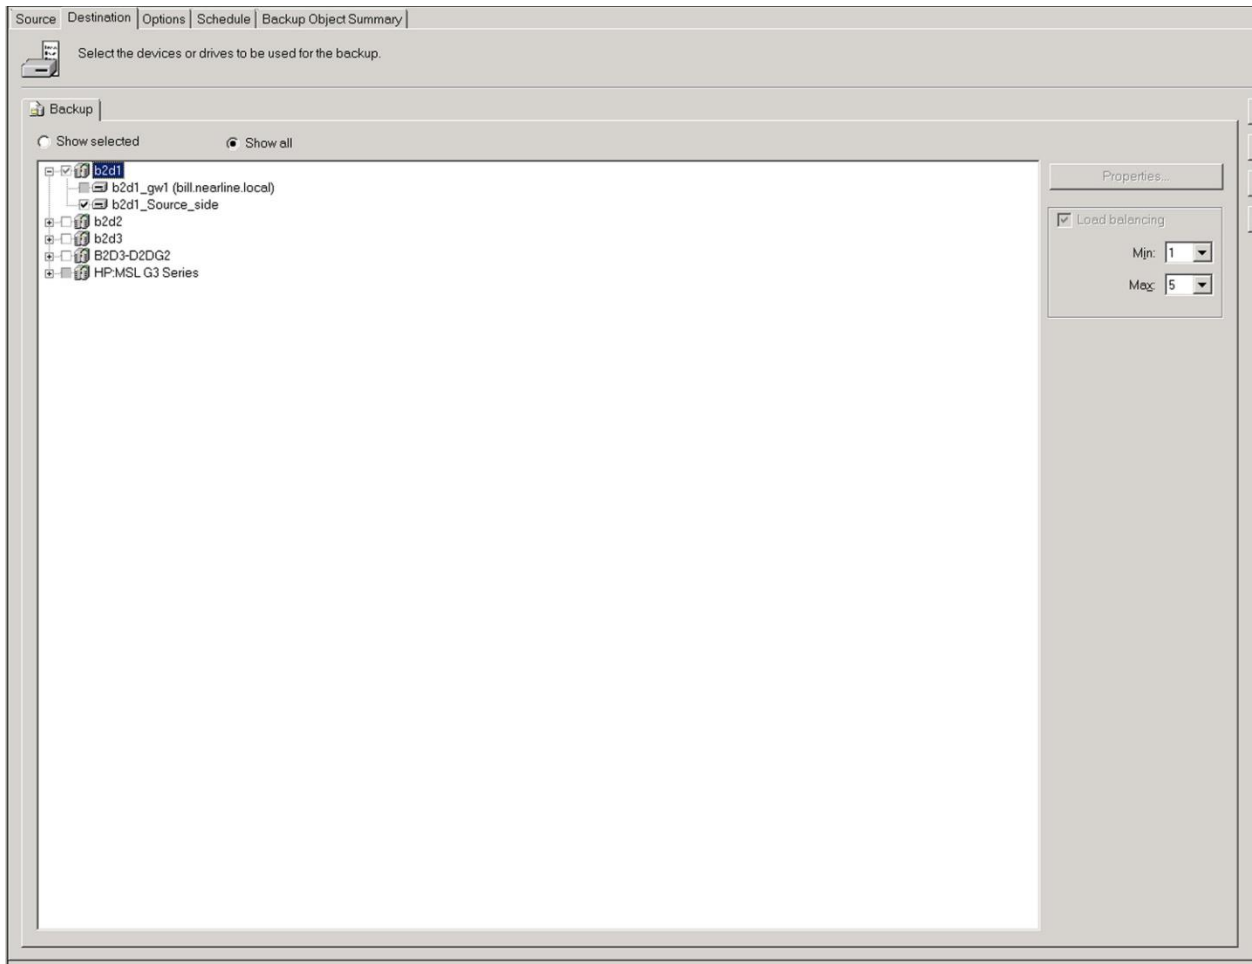


Fig. 14: Selecting the 'implicit' gateway.

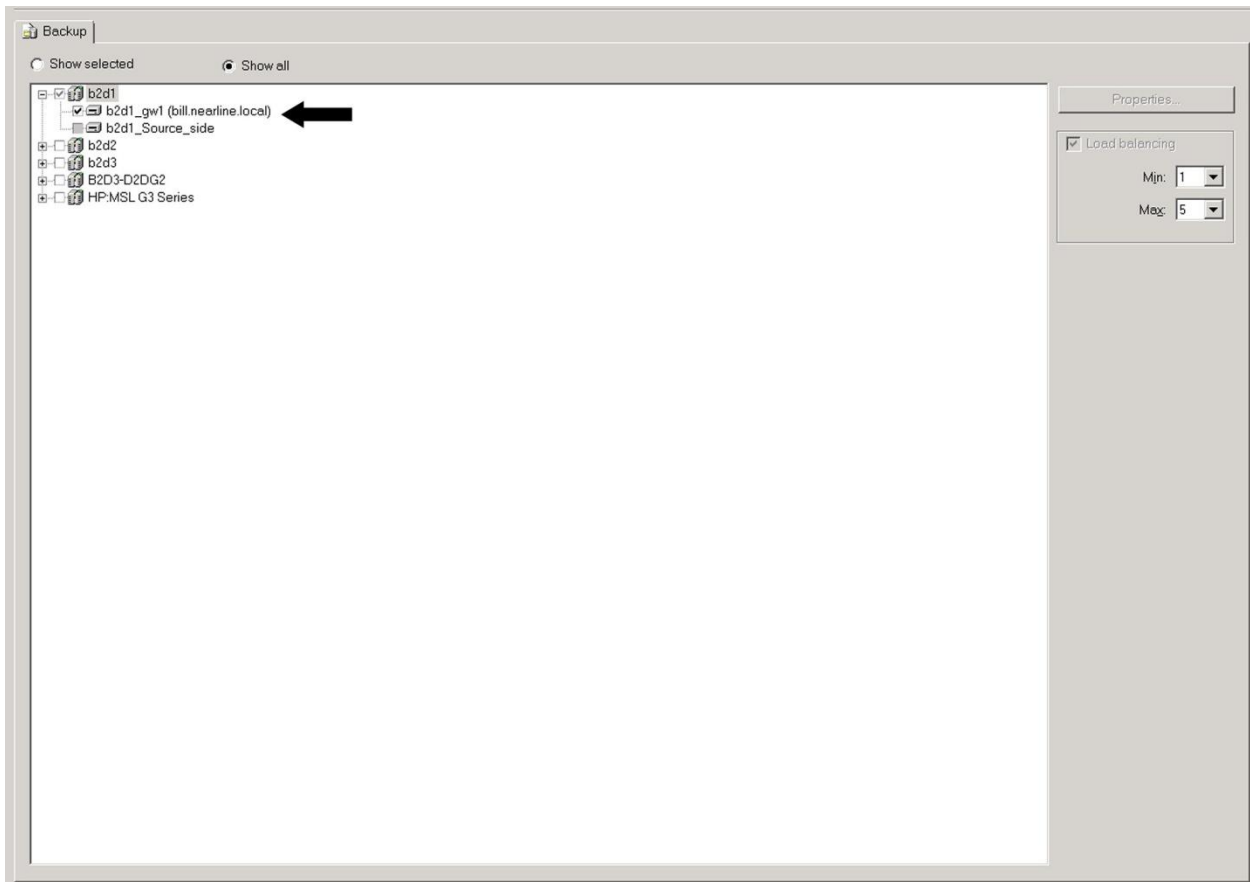


Fig. 15: Selecting the 'explicit' gateway for server-side backup (see arrow)

Backup using 'target-side deduplication'

Deduplication will occur at the B6200 'appliance' and source data can be on the server hosting the gateway or another server with a disk agent installed. If data is hosted by another server the disk agent will move the data via the network to the backup server.

It is necessary for 'target-side' deduplication to create a gateway with 'server-side' deduplication not selected in the 'Advanced' properties. In this example a gateway has been configured for the server 'Ben'. (Gateways can be modified from the devices menu. Only ONE gateway per backup to disk store is permitted per backup server/application server.)

From the Data Protector Management screen select backup from the drop-down box at the top of the page. Right click on 'Filesystem' and select 'Add Backup'. Select 'Blank Filesystem'. DO NOT check the 'Source-side' deduplication box. Click 'next'.

Select some files for backup from the server 'Zip'. Select next and the screen will show the destination devices. The gateways on 'Bill' and 'Ben' are available and the source-side gateway is 'grayed' out. Note that as you can name the explicit gateways it is good practice to include indication of whether the gateway performs high or low-bandwidth backups. (e.g. b2d1_gw2(high-bandwidth)). Note that the default name includes the backup server name. Fig. 16 shows the gateway configuration and Fig 17.

shows the gateway selection showing gateways on both 'Bill' and 'Ben' are available. For target-side deduplication select the high-bandwidth gateway.

Click 'next' and specify the required options for retention. The backup specification options and schedule can be modified if desired. It is often useful to tick the 'display statistical information' box.

The backup specification can then be saved.




Fig. 16: Explicit Gateway configuration

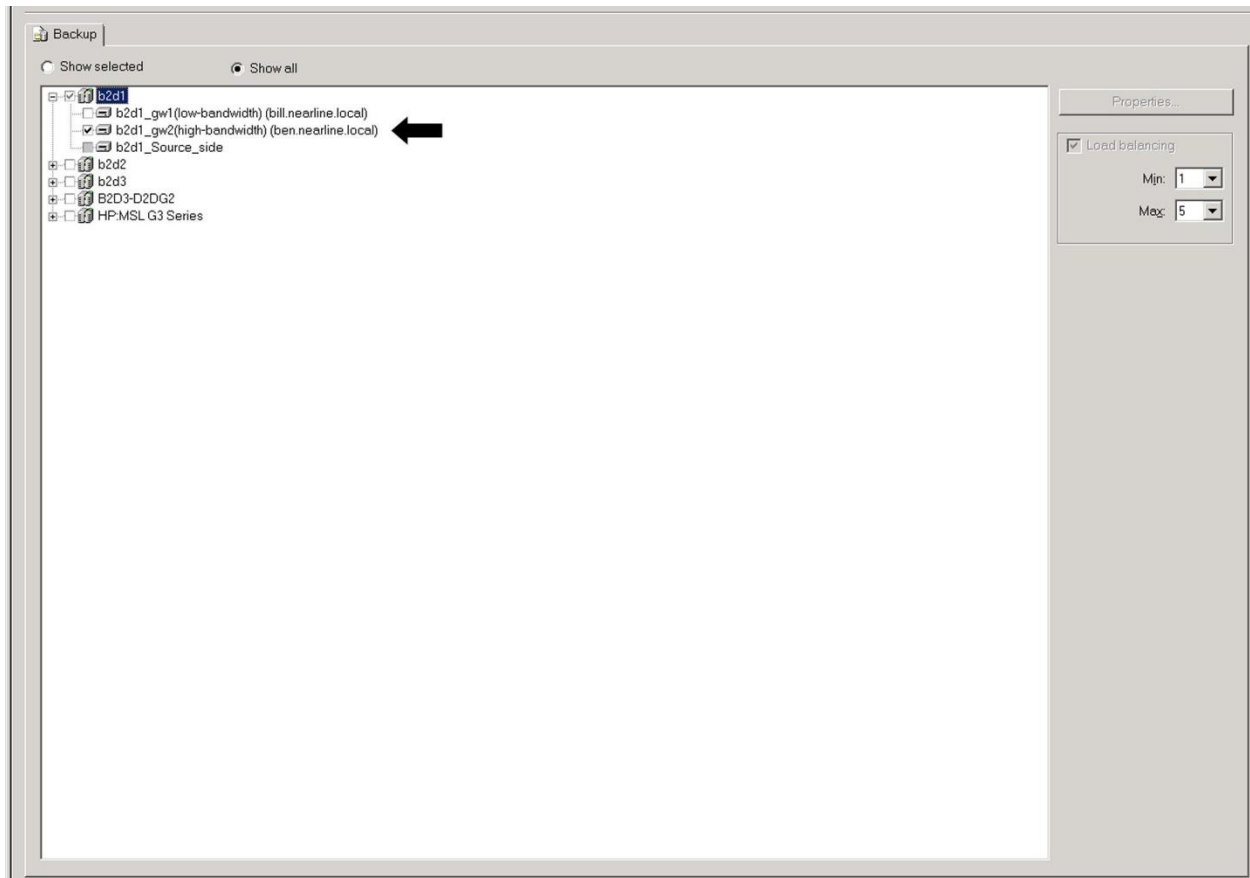


Fig. 17: High-bandwidth gateway selection for 'Target-side' deduplication.

Running a backup job interactively is a good way of checking where the deduplication is taking place. The following screenshots are from running the backup specification created previously for 'source-side', 'server-side' and target-side' deduplication. Fig. 18 'Source-side' deduplication, Fig. 19 'Server-side' deduplication and Fig. 20 'Target-side' deduplication.

Key Points:

1. Source-side and server-side both can perform deduplication on the backup server but are used in very different ways. The StoreOnce deduplication is contained within the media agent.
2. Deduplication places a loading on the backup server and this does need to be considered.
3. The implicit gateway for source-side deduplication is optional. It cannot be used to backup (or restore) data from other servers. Data must use the gateway on the server it resides on.
4. Use the explicit gateway for server-side deduplication which can backup data residing on any servers running the DP7 disk agent (or application agent).
5. Explicit gateways can be configured without server-side deduplication. Deduplication then takes place on the StoreOnce Catalyst appliance. This is known as 'target-side' deduplication.

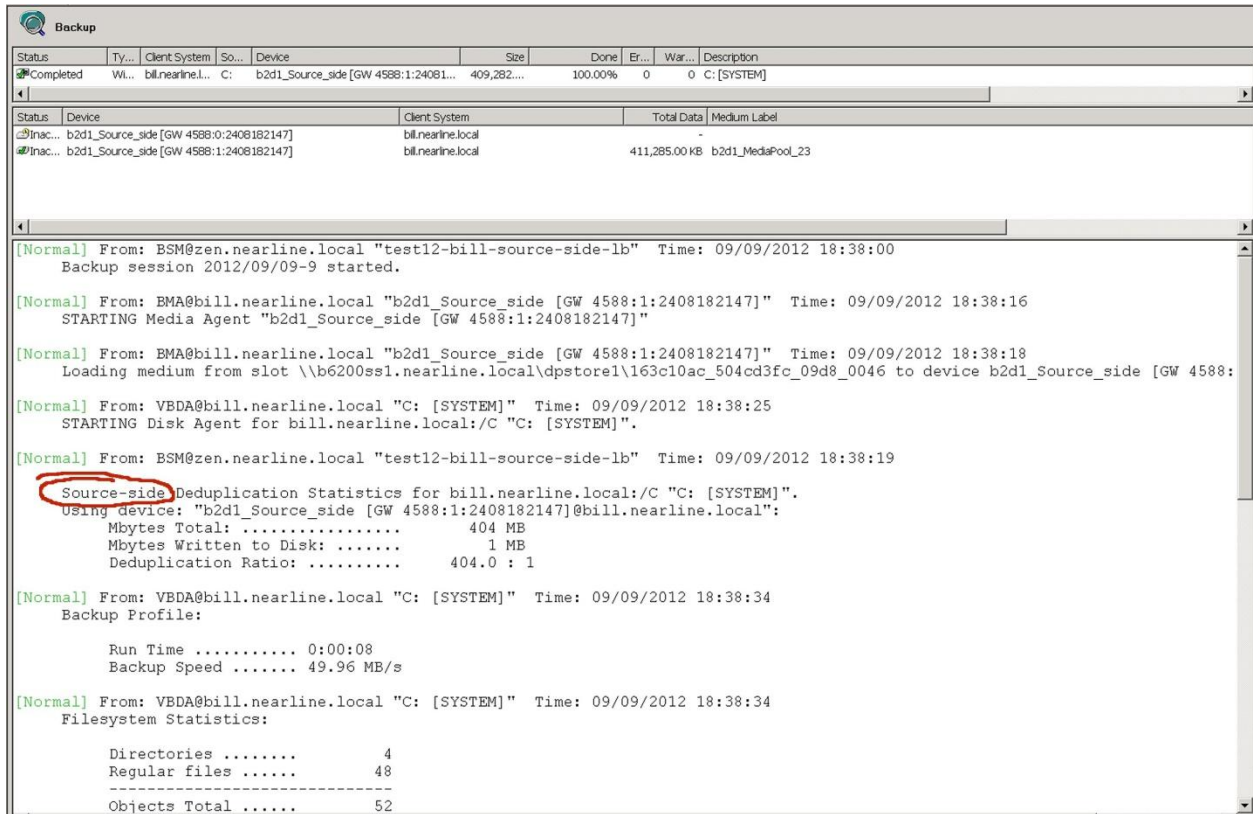


Fig. 18: Backup job specification running using 'source-side' deduplication

This screen shows that the backup job used source-side deduplication. Note that the top box of the windows shows the data being backed up by the disk agent and the window below shows the media agents in use. As this is a default source-side deduplication you will not that only 2 media agents are available and only one in use as there was only one mount point in the backup selection. the deduplication ratio reported is that reported by the B6200 appliance regardless of where the deduplication actually took place.

There is more detail in the next section on how to configure a backup to use multiple streams.

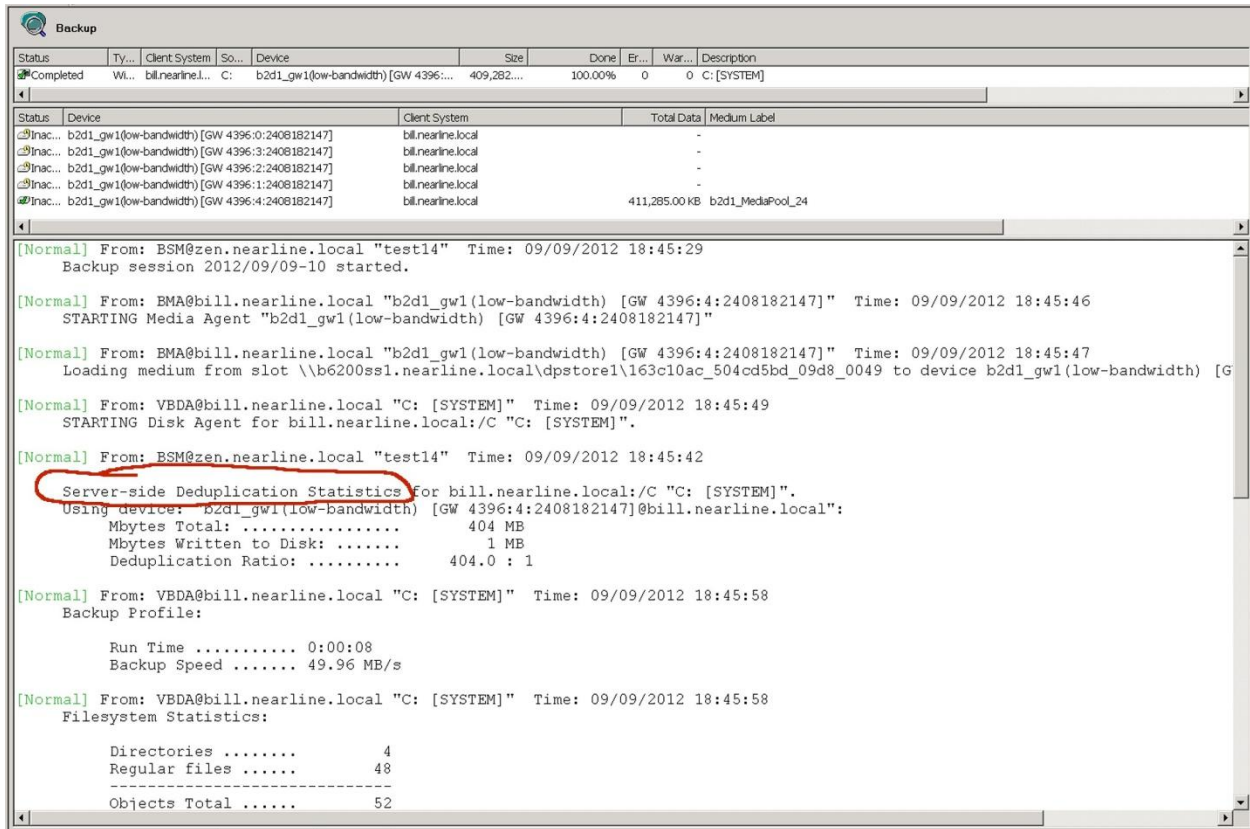


Fig. 19 Backup using 'server-side' deduplication via explicit gateway

Note here that the default streams setting for the explicit gateway are 5. Although 5 media gents are available, once again the file selection determined that only one media agent could be used.

Status	Ty...	Client System	So...	Device	Size	Done	Er...	War...	Description
Completed	Wl...	zip.nearline.l...	C:	b2d1_gw2(high-bandwidth) [GW 6016...	1,083,60...	100.00%	0	0	C:[SYSTEM]

Status	Device	Client System	Total Data	Medium Label
lnac...	b2d1_gw2(high-bandwidth) [GW 6016:4:2408157491]	ben.nearline.local	-	-
lnac...	b2d1_gw2(high-bandwidth) [GW 6016:3:2408157491]	ben.nearline.local	-	-
lnac...	b2d1_gw2(high-bandwidth) [GW 6016:2:2408157491]	ben.nearline.local	-	-
lnac...	b2d1_gw2(high-bandwidth) [GW 6016:1:2408157491]	ben.nearline.local	-	-
lnac...	b2d1_gw2(high-bandwidth) [GW 6016:0:2408157491]	ben.nearline.local	1,089,988.0...	b2d1_MedaPool_25


```

[Normal] From: BSM@zen.nearline.local "test11-zip-target-side" Time: 09/09/2012 18:50:34
Backup session 2012/09/09-11 started.

[Normal] From: BMA@ben.nearline.local "b2d1_gw2(high-bandwidth) [GW 6016:0:2408157491]" Time: 09/09/2012 18:50:50
STARTING Media Agent "b2d1_gw2(high-bandwidth) [GW 6016:0:2408157491]"

[Normal] From: BMA@ben.nearline.local "b2d1_gw2(high-bandwidth) [GW 6016:0:2408157491]" Time: 09/09/2012 18:50:52
Loading medium from slot \\b6200ssl.nearline.local\dpstore1\163c10ac_504cd6ed_09d8_004c to device b2d1_gw2(high-bandwidth) [

[Normal] From: VBDA@zip.nearline.local "C: [SYSTEM]" Time: 09/09/2012 18:50:51
STARTING Disk Agent for zip.nearline.local:/C "C: [SYSTEM]".

[Normal] From: BSM@zen.nearline.local "test11-zip-target-side" Time: 09/09/2012 18:51:20
Target-side Deduplication Statistics for zip.nearline.local:/C "C: [SYSTEM]".
Using device: "b2d1_gw2(high-bandwidth) [GW 6016:0:2408157491]@ben.nearline.local":
Mbytes Total: ..... 1076 MB
Mbytes Written to Disk: ..... 323 MB
Deduplication Ratio: ..... 3.3 : 1

[Normal] From: VBDA@zip.nearline.local "C: [SYSTEM]" Time: 09/09/2012 18:51:33
Backup Profile:

Run Time ..... 0:00:40
Backup Speed ..... 26.46 MB/s

[Normal] From: VBDA@zip.nearline.local "C: [SYSTEM]" Time: 09/09/2012 18:51:33
Filesystem Statistics:

Directories ..... 239
Regular files ..... 1116
-----
Objects Total ..... 1355

```

Fig. 20: Backup using 'target-side' deduplication.

In this backup a gateway was selected with 'server-side' deduplication not selected. All deduplication takes place at the StoreOnce B6200 appliance and the network transfer of data is referred to as 'high-bandwidth'. This means that all data is sent over the network and not just the new 'chunks' as in deduplication at the backup server. This setting may be necessary if the user does not wish to load up the backup server.

Sizing the Backup Server

HP StoreOnce Catalyst and Data Protector 7 now have the ability to perform deduplication at the backup server there is now extra load on that server. In the previous sections it has been shown how the gateway settings are used to specify where deduplication takes place. (i.e. server-side and source-side deduplication). A simple 'rule of thumb' can be used to determine the capability of the backup server.

Key Points:

- Allow 50MB/s of stream data per GHz of CPU core and 30MB RAM (allow 2 –cores for the DP media agent software)
- Allow at least 16GB of RAM overall
- Ignore hyperthreading (E.g. 12 cores=24 with hyperthreading).

Example:

Dual Hex-core CPU running at 3.4GHz (12 cores).

10 cores x 3.4GHz = 34GHz (remember 2 cores allocated to the media agent).

34 streams @ 50MB/s = 1700 MB/s

Optimizing performance

In order to obtain the maximum performance with HP StoreOnce Catalyst it is essential that the backup servers keep the store supplied with the optimal number of data streams. Some basics:

Key Points:

- One stream ingest rate is around 330 MB/s (4 disk shelves in low-bandwidth backup)
- Optimum performance per node has been achieved with 6 streams to each of 8 catalyst stores. (in house test results with 2:1 compression achieved 4200MB/s total for 1 node).
- You can direct 2 backup servers to one store or 1 backup server to multiple stores. (This is useful to separate different data types).
- Multiplexing is not permitted and cannot be configured.
- Wire speed of 10GbE is around 1.2GB/s
- Using a single catalyst store will realize less than 50% of the performance potential of a node.
- Enabling deduplication at the media server can double the performance.
- Each node operates individually so in a full configuration there are 8 nodes.

Data Protector 7 configuration to produce multi-stream backups with HP StoreOnce Catalyst

It is essential to send multiple streams to a catalyst store to realize the performance potential of the StoreOnce B6200 system. The 'sweet' spot for an individual node is to have 6 streams per store with 8 stores in use. Aim to plan the backup specification to take advantage of this and 6 – 8 data streams (fully loaded) is ideal.

HP Data Protector can write multiple streams for data to a single catalyst store by starting multiple media agents on the backup server. The limit on the streams is set in 2 places:

1. The gateway setting. The implicit (source-side deduplication) has a default limit of 2 streams per client. This will apply to all media servers using this gateway. The advanced setting for each explicit gateway also has a setting for maximum number of streams. Figs. 21 and 22 show the advanced settings box.
2. When load balancing is selected there is also a streams limit. This overrides the gateway limit. (So if gateway stream limit is set to 6 and load balancing is set at 5 (default) then ONLY 5 streams will be possible for a job specification). Load balancing needs to be selected (default anyway) in order to modify this setting. Load balancing is not recommended to 2 different StoreOnce Catalyst stores. Therefore do not select 2 stores for a destination. The problem arises when one store runs out of space. (This is not recommended for tape either!). However if only one destination is selected there is no problem.
3. Although not a setting the source data selection will dictate how many streams are sent in a particular backup specification. If data is selected for backup from multiple mount points then each mount point will have a disk agent started. It is also possible to use the backup specifications to select multiple directory selections for backup to produce multiple streams. Fig. 23 shows 2 directory entries selected under the 'trees' tab for WinFS. Fig. 25 shows the resultant multi-stream backup in progress. Note the 2 media agents running (= 2 streams).

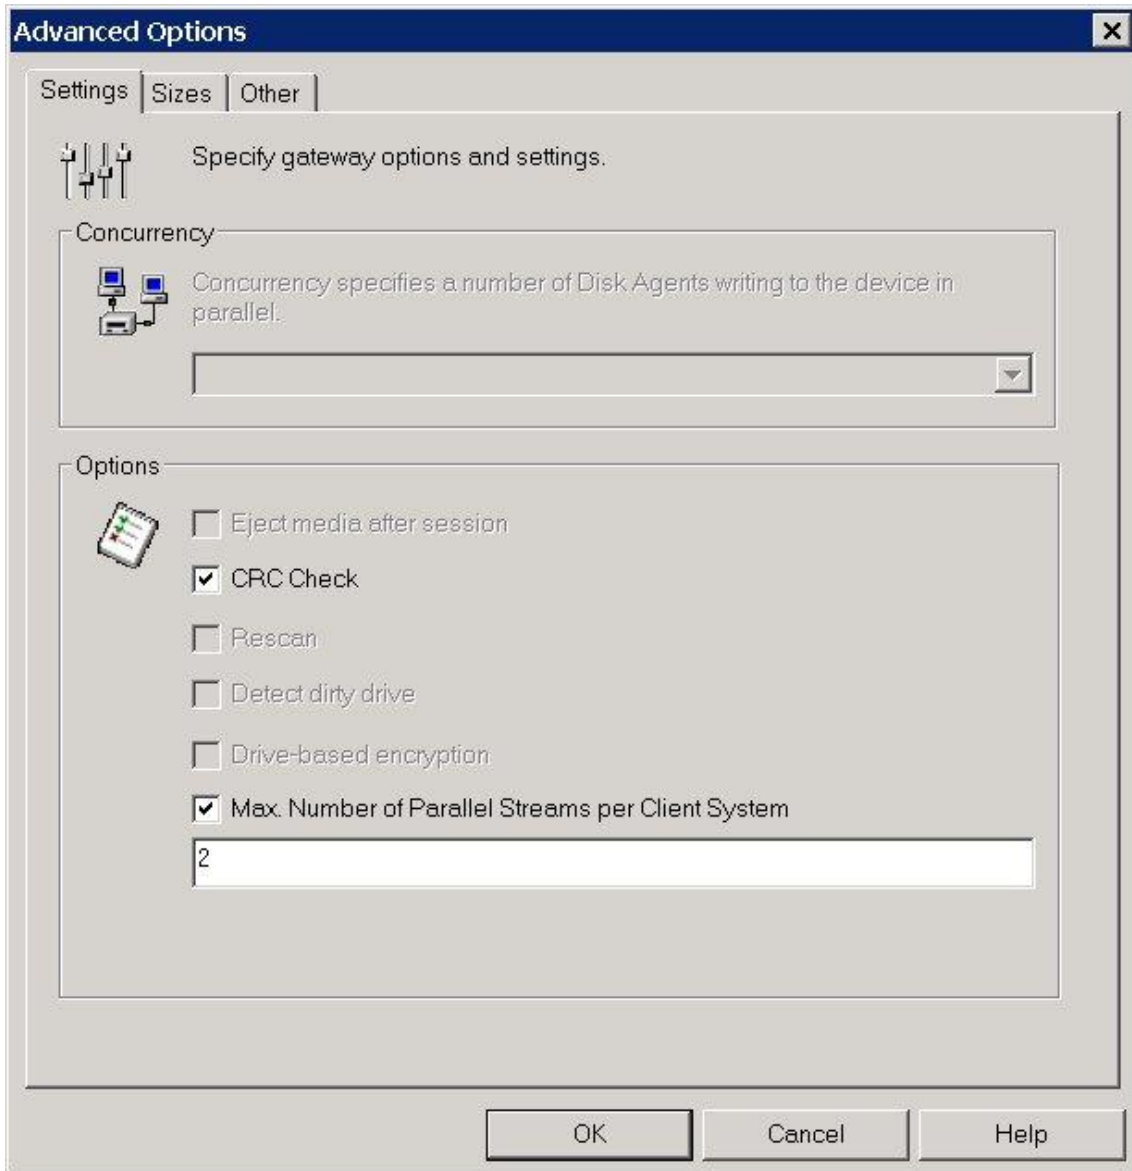


Fig. 21: Advanced settings – Source-side deduplication (default setting for max. streams)

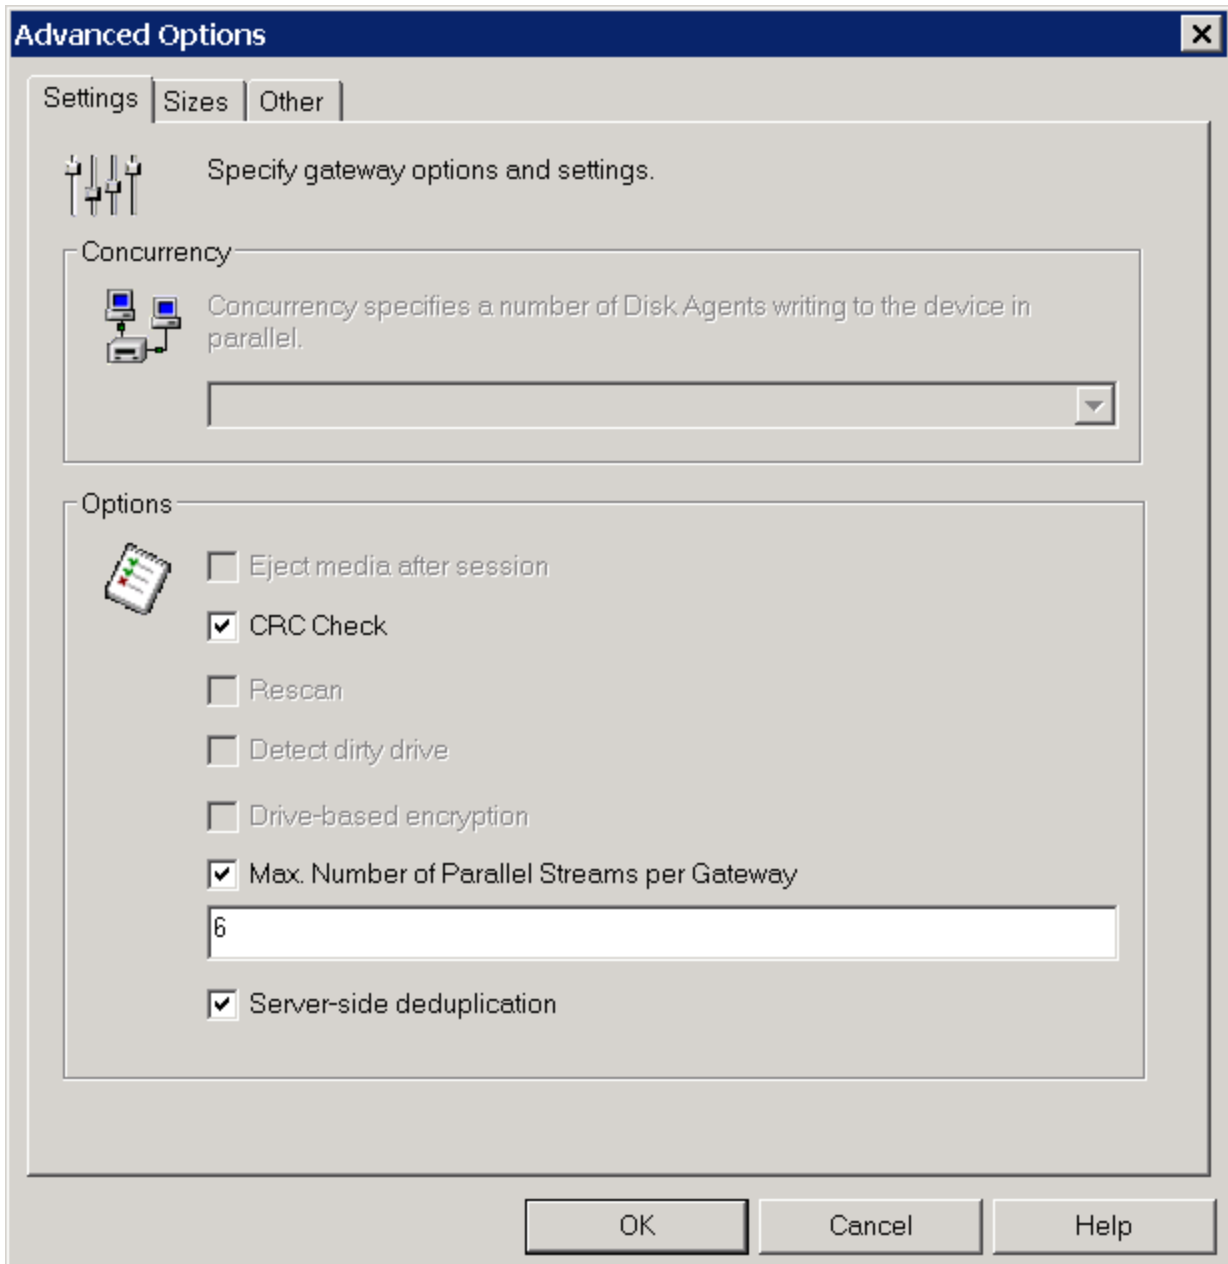


Fig. 22 Advanced settings – Server-side (explicit) gateway settings for max. streams



Fig 23. Setting the streams limit for a Backup Specification. (Default=5)

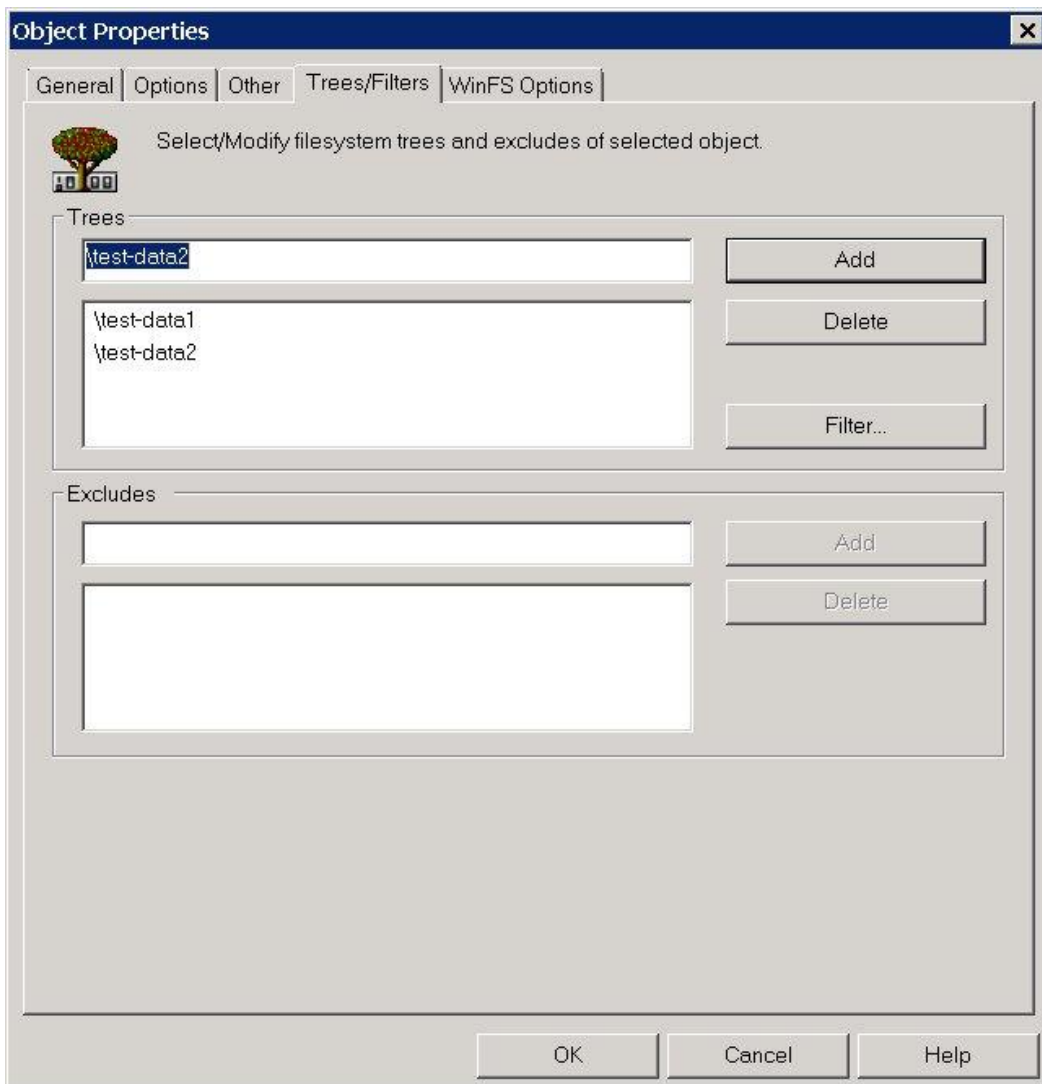


Fig. 24: Selecting 2 directories for backup as 2 streams.

(If selecting directories here DO NOT select as source items when creating a backup specification.)

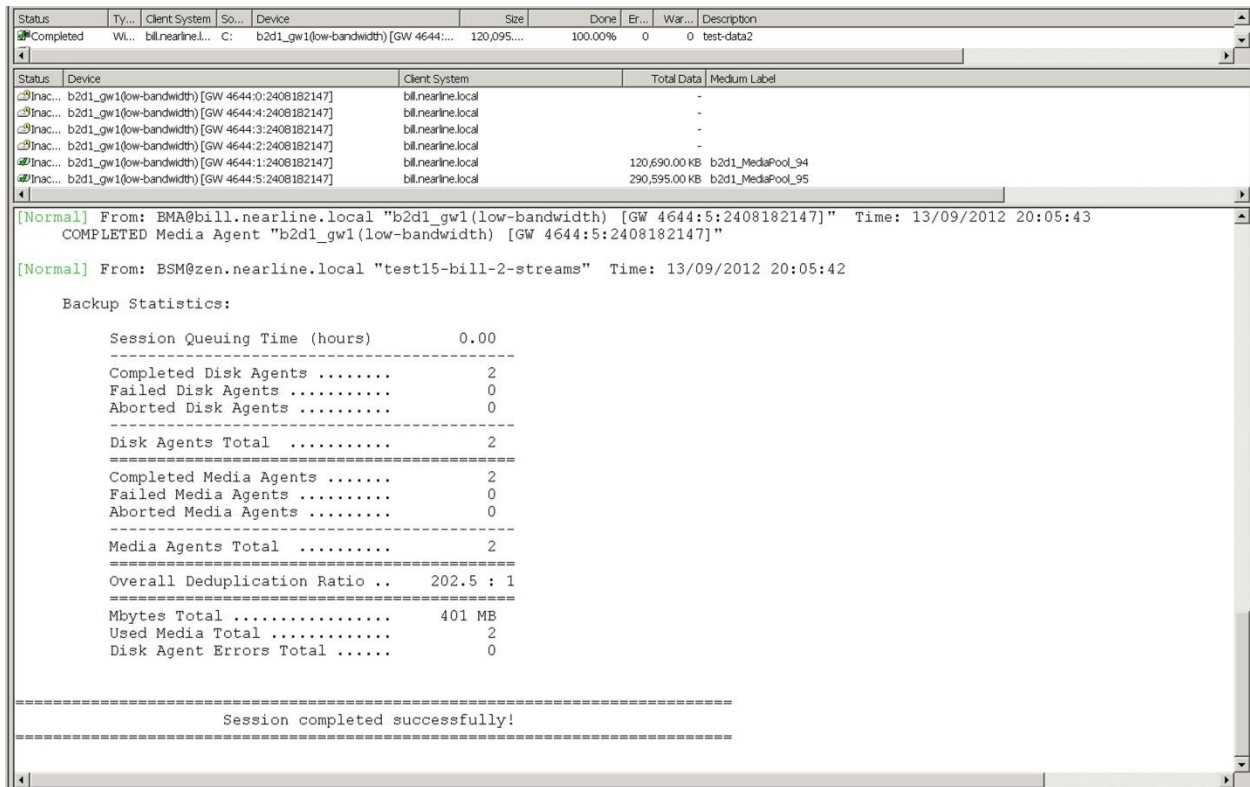


Fig. 25: backup job specification showing 2 DP7 media agents in use out of 6 available.

Reclaiming space from expired backups

HP StoreOnce Catalyst has the added advantage that expired backups can be removed as catalyst 'items' and the space made available for other users of the service set. Please not that as the data is deduplicated not all space is returned and housekeeping needs to run to actually free up the space. Hp Data Protector 7 by default returns space occupied by expired backups every 24hrs. This can be modified down to hourly by modifying the global options file. (this file is located at \ProgramData\OmniBack\Config\server\options\global). An example is shown in Fig. 26 where the delete unprotected media occurs hourly. (This frequency would be unnecessary in a real world situation but is useful for demonstration systems)

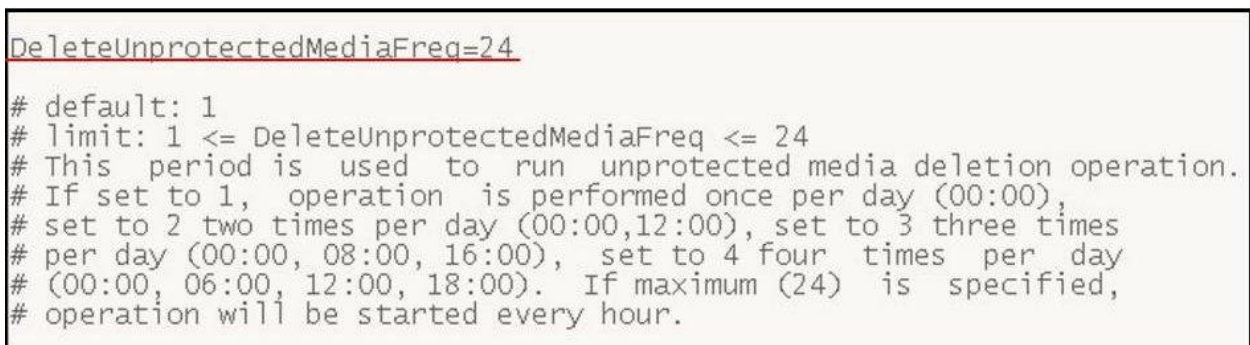


Fig. 26: Global option for delete unprotected media

As housekeeping is triggered by expiration of media best practice would be to adjust the delete unprotected media time with care.

Key Points:

- Create job specifications for server backups which use multiple streams.
- Server side or source side deduplication will result in greater throughput providing the backup server is sized correctly.
- For maximum throughput aim for 6-8 streams per store with 8 stores being used at any one time.
- Use the implicit gateway and use source-side deduplication if the source data is located on the server running the DP media agent.
- Use the explicit gateway and server-side deduplication for individual server settings and when source data may be located on other servers.
- Expired backups are deleted from the StoreOnce catalyst store at intervals specified by the global options file.
- Exporting backup media will leave 'orphaned' items in the catalyst store – avoid doing this.

Using DP7 Object Copy with HP StoreOnce Catalyst

HP StoreOnce Catalyst has the ability to move data between StoreOnce Catalyst stores without re-hydration. This means that bandwidth efficient transfers take place making it ideal for moving data quickly offsite. HP StoreOnce Catalyst differs from VTL or NAS replication in that data can be duplicated to more than one appliance for extra resilience. Additionally the HP Data Protector internal database is now aware of all StoreOnce Catalyst copies and can restore from any copy without complex scripting arrangements or imports. When required, data can also be copied to real tape for long term storage. Obviously when data is moved to tape the deduplicated data has to be re-hydrated. DP 7 object copy offers a rich selection of options. This document will only cover basic object copy functions, replicating a backup to one catalyst store then on to another store. For more detail please consult the appropriate HP Data Protector documentation.

All transfers between StoreOnce catalyst stores are bandwidth efficient (low-bandwidth). As the StoreOnce catalyst protocol can now control the StoreOnce appliance the data does not need to flow through a backup server.

Object copies can be interactive (useful for ad hoc copies), automated or scheduled. Additionally they can be set to be 'post-backup' where they are launched after the backup completes. It is not possible to make backups to multiple destinations at the same time. Copies are made sequentially from one store to another.

HP Data Protector 7 'object copy' functions.

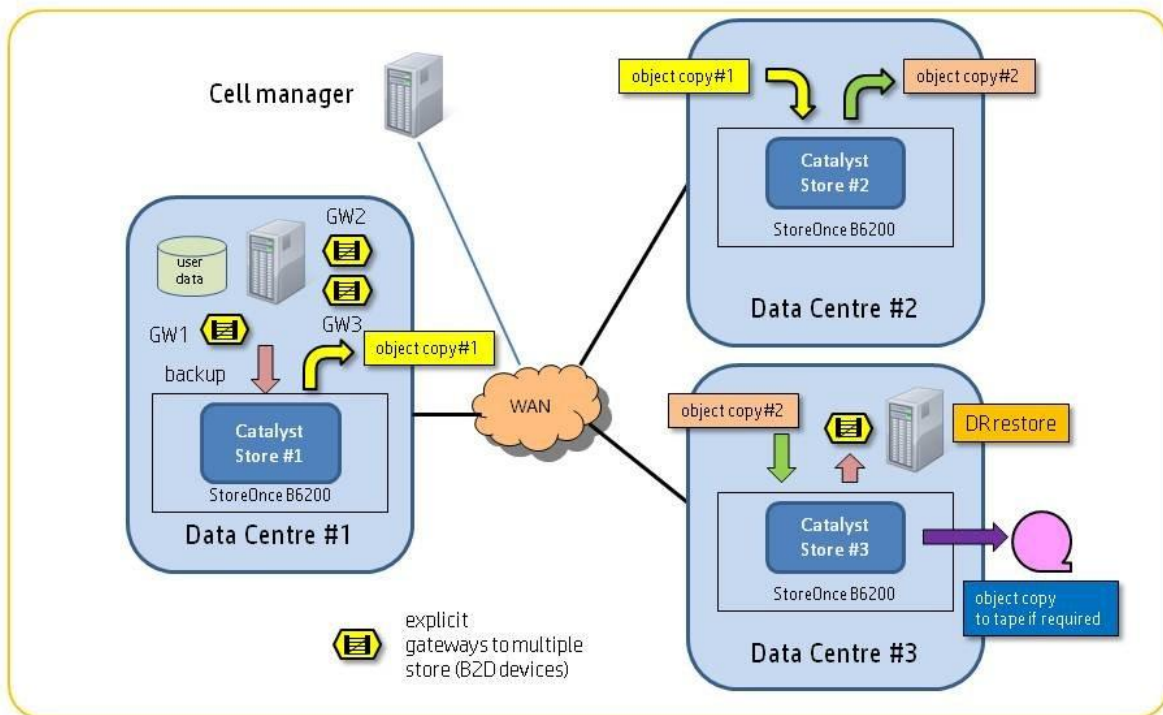


Fig. 26 HP Data Protector Object copy between StoreOnce Catalyst stores

An example of a 3 data centre setup is shown in Fig. 26. Backups of user data are performed using a server-side gateway to catalyst store #1 located in data centre #1. Post backup (or scheduled) HP DP7 object copy can move the backup offsite via the WAN to catalyst store #2 in data center #2. This is performed in a bandwidth efficient manner and after the first transfer only new data chunks will be transferred. The expiry date of the original backup can be shorter or even immediate once data is offsite. The backup can then be duplicated to catalyst store #3 in data center #3. This gives extra resilience. The object copy configuration allows post backup replication. Data Protector 7 could then move the data onto tape if required. Transfer of data is direct from StoreOnce B6200 catalyst store to catalyst store. Although it may appear the gateways are involved they only pass the commands to perform the replication. The cell manager internal database tracks the copies. It is important to backup the cell manager database after every backup session.

Key Points:

- Use Data Protector 'object copy' function to move data offsite.
- Design the WAN link using the sizer tool in order to complete the duplication in the appropriate time window. Allocate separate time slots for duplication to backup.
- Select only the server-side gateway for backups. Because of the virtual nature of the source-side (implicit) gateway it is not possible to use for duplication. Destination server-side (implicit) gateways are 'grayed' out. If you select source-side gateway for backup jobs will fail unless the object copy job is remapped to use the explicit gateways.
- Duplication can either be post backup or scheduled separately.
- Object copy to tape is referred to as a 'copy' as distinct from a catalyst copy which is 'replication'.

Creating an object copy specification – example:

Step 1 is to create a backup specification. This performs a low-bandwidth backup via an explicit gateway (server-side deduplication) to a catalyst store. (Catalyst Store#1 in Fig. 26)

Step2 – create a post backup object copy specification. Fig. 27 shows the selection of the backup specification which will be copied post execution. The next screen shown in Fig. 28 is 'copy specification'. This is used to select another object copy job after this job completes. There are no entries here as there are no other object copy jobs created yet. The next screen is not used as it is for combining virtual full/synthetic full backups. The next screen is the object filter which can be used to select objects that have less than specified copies. The next screen is the library filter. Fig .28 . This shows how libraries capable of replication are identified. You can see that the 'real' tape library is 'grayed out' as you can only move data to tape by a full copy with rehydration. It is important at this stage that if the backup is performed in a ROBO and then copied to a Data Center that devices in the ROBO are not selected.

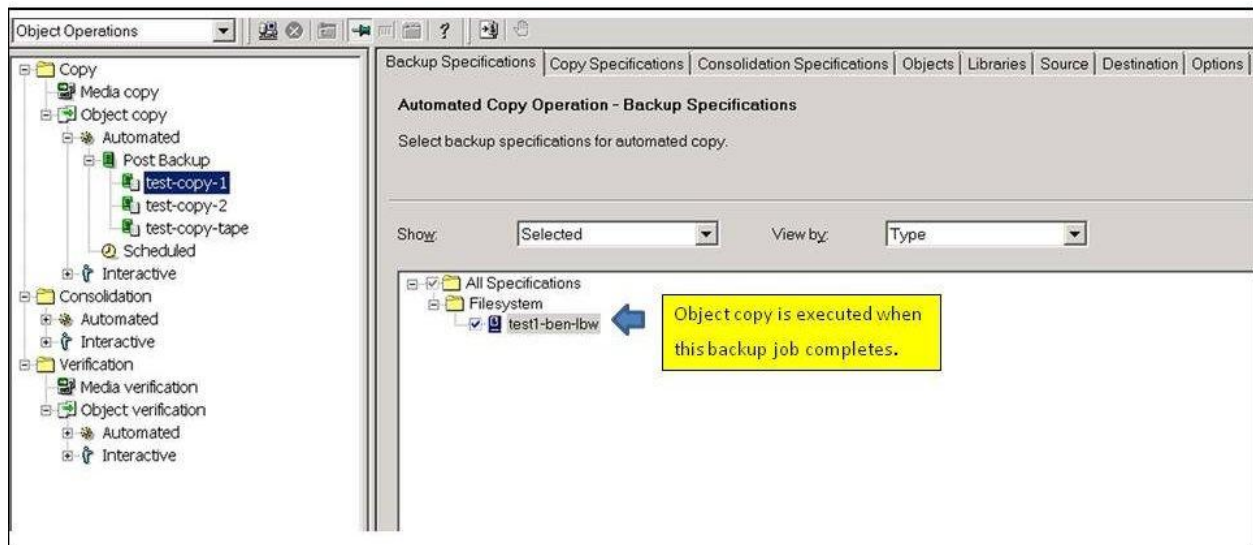


Fig. 27 HP Data Protector 7 – Object copy ‘automated post-backup’ selection

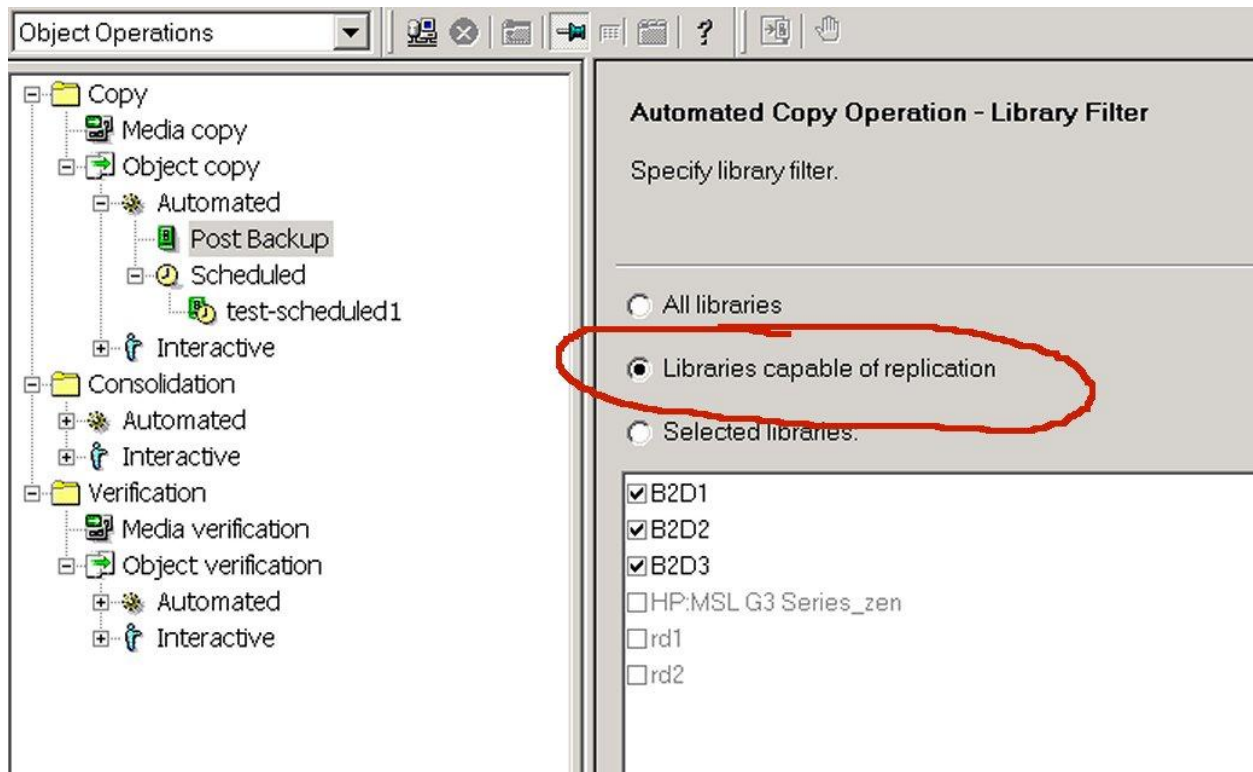
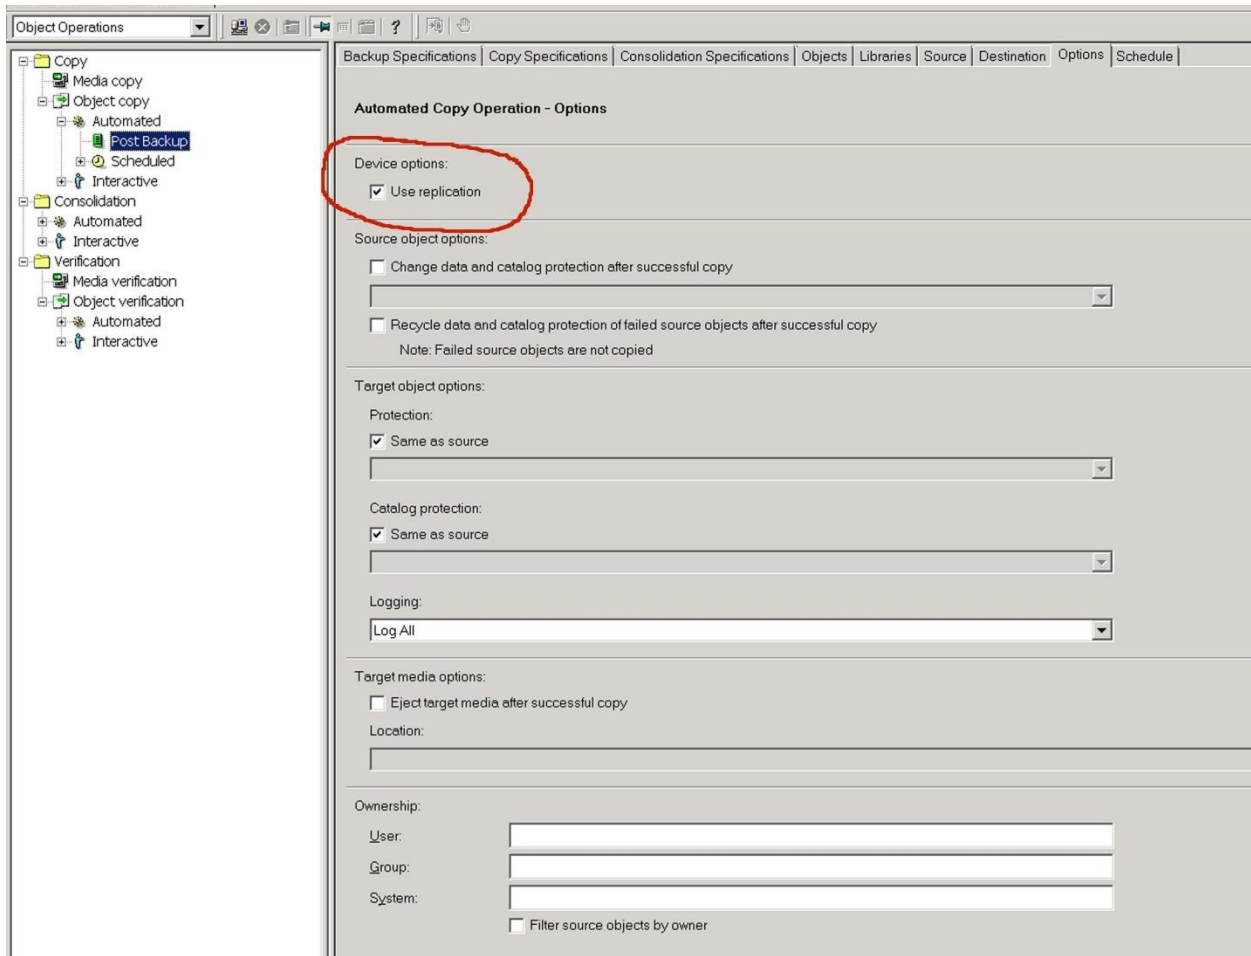


Fig. 28 Library Filter

Step 3. – Source Selection. Leave this at Automatic Device Detection. (Note: If the backup used the source side or implicit gateways then this is where the gateway and be re-mapped to use a server-side (explicit gateway).

Step 4 - Set Destination. Select the gateway to use for this function Notice all the ‘Source-side’ (implicit) gateways are not selectable.

Step 5 –Set use replication. This enables the process to move the data without rehydration. Fig. 29 shows the correct selection. This screen also allows the setting of different protection times for targets. This is useful as the original backup maybe protected for a shorter period. If you do not set replication Data Protector will copy the software which means re-hydration of data.



Fig, 29: Data Protector Object Copy – Options

It is possible to cascade object copy operations so in our example shown in Fig. 26 the data is backed up to one store (store#1) and then duplicated to store#2 and then again to store#3. This is done by creating more object copy jobs. In these jobs the backup specification is left unselected but in the copy specification the previous copy job is ticked. In this way copies are cascaded. Fig. 30 shows the selection set to follow 'test-copy-1'. The object copy specification test-copy-tape is set to follow test-copy-1 BUT replication is not selected because data is rehydrated and sent to tape.

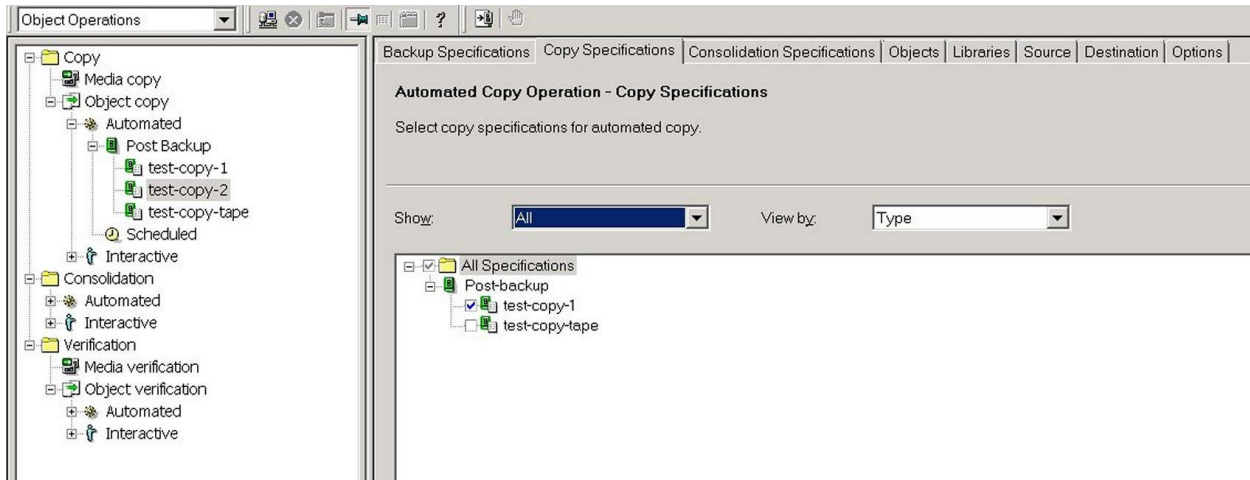


Fig. 30: Copy Specifications for test-copy-2 showing the settings for the copy to run after test-copy-1.

Fig. 31 shows the monitor screenshot for a post backup object copy which copies the backup object from the original destination of the backup (B2D1 – catalyst store#1) to a second B2D target (B2D2 – catalyst store#2) and then onwards to a third store on B2D3 (b2D3 – catalyst store#3). All performed sequentially and the data is not rehydrated. The movement of data can be seen in the B6200 management interface by selecting each store and looking at copy-out and copy-in job log. See Fig. 32.




Fig. 31: Post backup automated object copy.

Note that the post backup object copies are of the type 'replication'. This means that they have moved the data directly between the catalyst store without rehydration.




Fig. 32 StoreOnce B6200 outbound copy job log

Restoring data from a StoreOnce catalyst Backup and Replication

This section will cover restore of files from a catalyst store. Object copies are fully located by the internal database located on the cell manager. It is vital that the internal database of the cell manager is backed up after all the backups complete. In the event of a disaster it can be restored to a new server with a new cell manager. If the database is lost there is a process to 'import' from a catalyst store. This is not as straightforward as tape/virtual tape but can be done. This paper will only cover the basics of Data Protector restore process and reader are advised to use the data Protector user guide for more information. However the intention is to cover restore from multiple StoreOnce Catalyst stores.

If multiple object copies have been made of a backup session as described earlier HP Data Protector the restore process will automatically select media or it can be selected manually. If one of the backups has been lost or accidentally deleted then the one of the other copies can be used. Restore can be by object or session. Restore by session is useful for restoring a file or files from a previous backup. It is necessary to select a session and then the session content is displayed. Selecting an object by default will choose the last used media. The following steps will briefly describe how to restore a missing file called 'my-valuable-data' which has been backed up and then object copies to 2 additional catalyst stores.

1. Select 'restore' in the scope window of the DP7 management GUI screen.
2. Expand the filesystem menu and the select the filesystem for restore (see Fig.33).
3. The 'source' menu tab allows expansion of the filesystem and selection of an individual file (in this example 'valuable.txt').
4. The destination tab allows choice of restore to original location or restore elsewhere. There are also choices for overwrite or keep original etc.
5. The devices tab is used to select the appropriate gateway. This can be left at automatic which restores using the same device as the backup.
6. The media tab can be used to select a particular item form a media pool. However there is no need to do this as the 'copies' tab can be used to select a particular copy (or the original).
7. The copies tab shows the version which will be restored (Fig. 34). The 'properties' selection will allow selection of the original backup or copies. (This backup was as produced from the example described earlier with one backup and 2 copies). Select as required.
8. Preview or restore as required and the file selected will be restored.

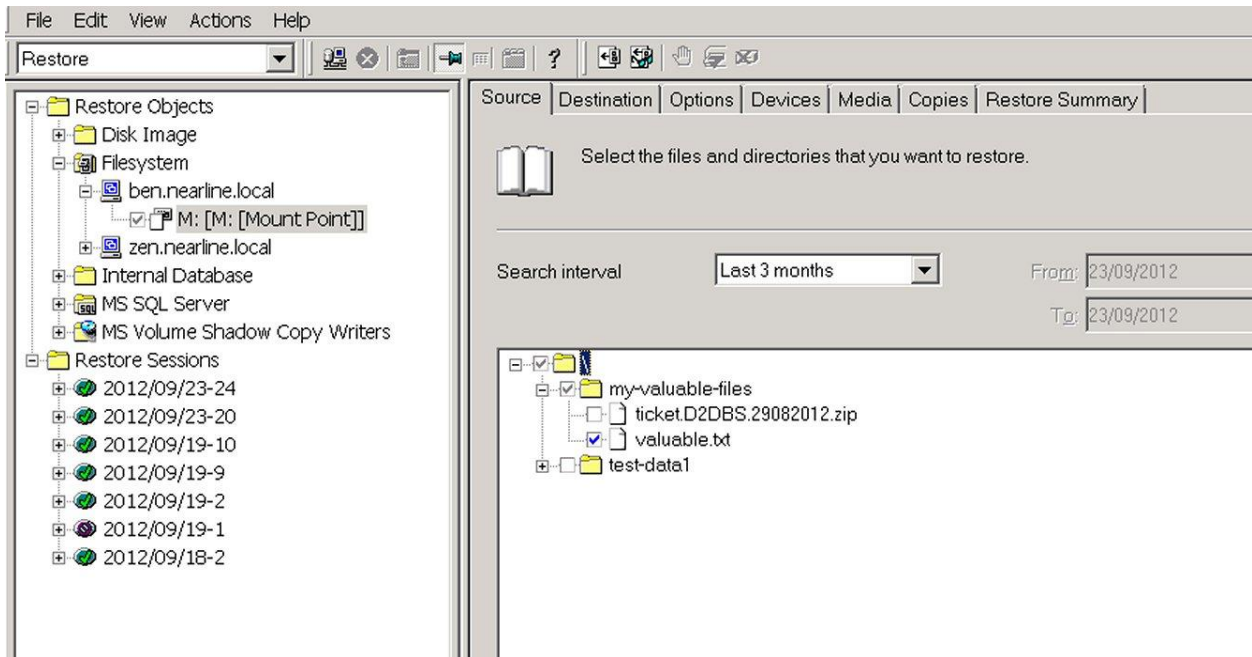


Fig. 33: Data Protector restore selection.

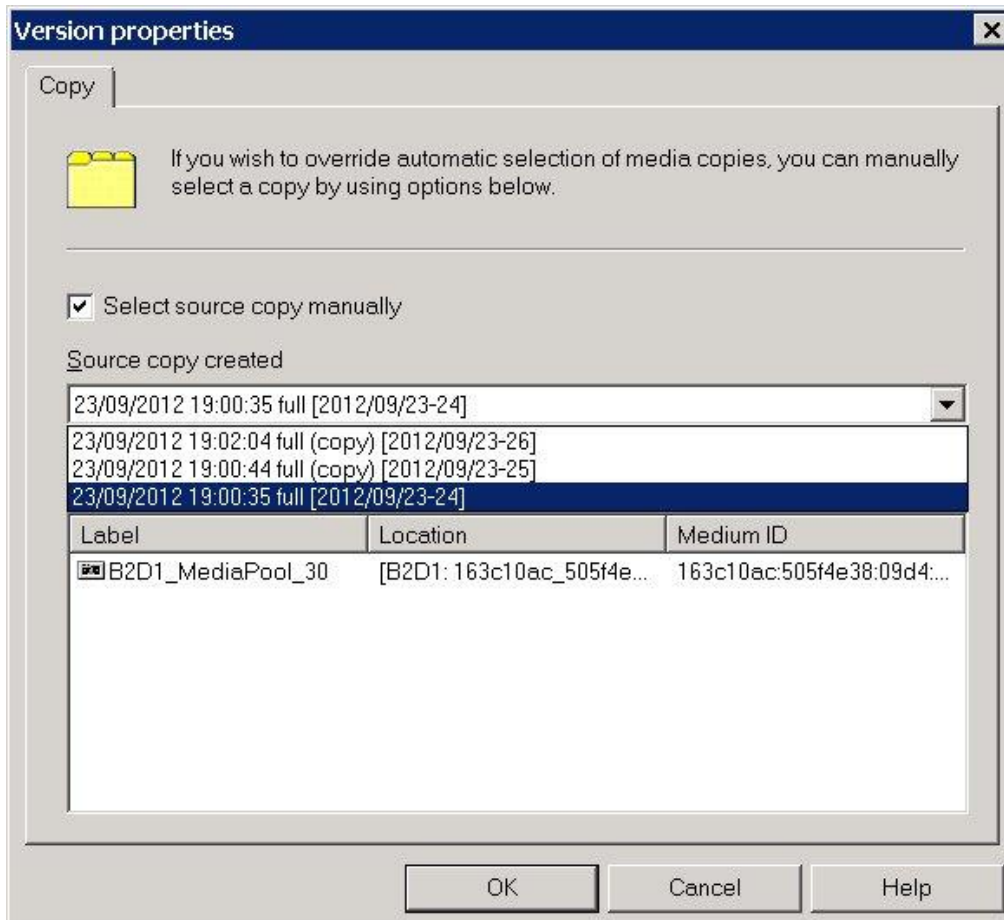


Fig. 34 Selecting copy or original for restore – note the media pool item is selected automatically.

Recovering from total loss - 'Disaster recovery' planning (DR).

Unlike virtual tape and NAS file replication it is not possible (or necessary) to connect another cell manager server at the DR site. VT and NAS when replicated are visible but are read only. The difference is that the internal database running within the cell manager is not aware of the replicated copies. However with StoreOnce catalyst the cell manager database is aware of all copies. It is really important however to keep a backup of the internal database so that in the event of a disaster you can recover to your new cell manager if the old one is lost.

In a DR situation a StoreOnce catalyst user is likely to have backup copies on multiple sites and possibly to tape as well.

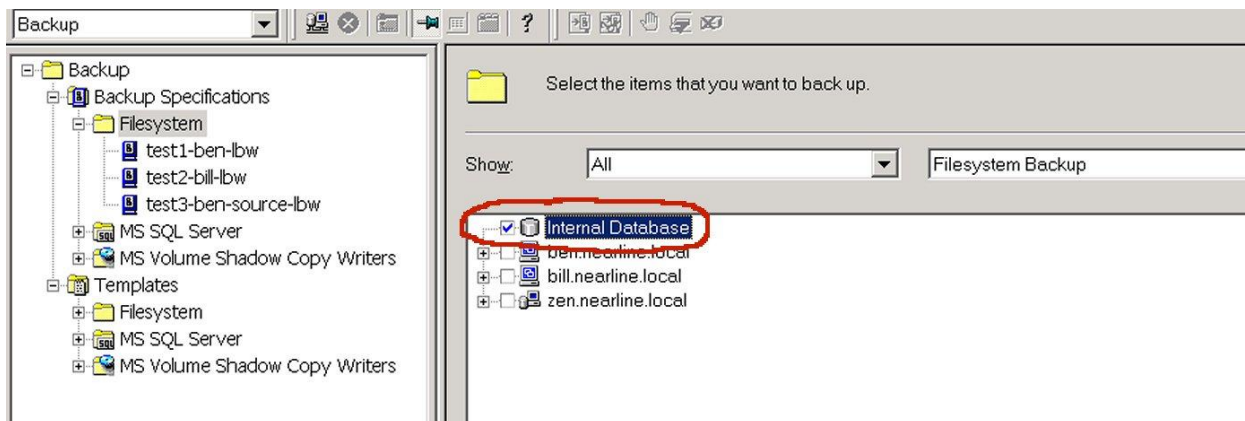


Fig. 35 Selecting IDP for backup

Best practice is to use backup the IDB (Fig.35 shows the selection in a backup specification) and then replicate offsite. It is recommended to use a separate catalyst store as it will be easier to import the IDB as there will be less objects. If the original cell manager is lost a new one can be created and then updated with the IDB. However it will be necessary to 'import' the catalyst items for the relevant store. In our example the IDB could be backed up to media in a separate pool (for easier identification) and then object copied to the 2 additional data centers. The import process for a catalyst store is as follows:

1. With a new cell manager configure the B2D device which is intact on the DR site. The original cell manager has been 'lost' in the disaster.
2. Obtain a list of StoreOnce catalyst objects in the store using the following command:
(commands in \Program Files\Omniback\bin)

```
omnib2dinfo.exe -list_objects -type OS -host << VIF of the B6200 service set>> -name <<storename>>
```

(note the storename is the name of the store on the B6200 and not the DP name)

3. Update the repository for each catalyst object.

```
omnimm -add_slots <<Store name in D2D>> <<catalyst object name>>
```

where the Store name is the name given in DP7.

Catalyst Object names are typically of the form "e74d2262_502255d2_06b8_0004".

4. Either use the GUI to import the catalog data or the command:

```
omnimm -import <<logical D2D device>> -slot SlotID
```


Networking best practice

It is very important to follow recommended best practice when using StoreOnce B6200 catalyst technology with HP Data Protector. The StoreOnce B6200 can have one of 5 different network templates. These templates specify can either be 2 discrete subnets for management and data or one subnet with management and data using the same subnet. The other key components are the Data Protector cell manager, data protector backup/application servers and DNS server (if used). Fig. 36 shows an example basic network configuration for HP Data Protector.

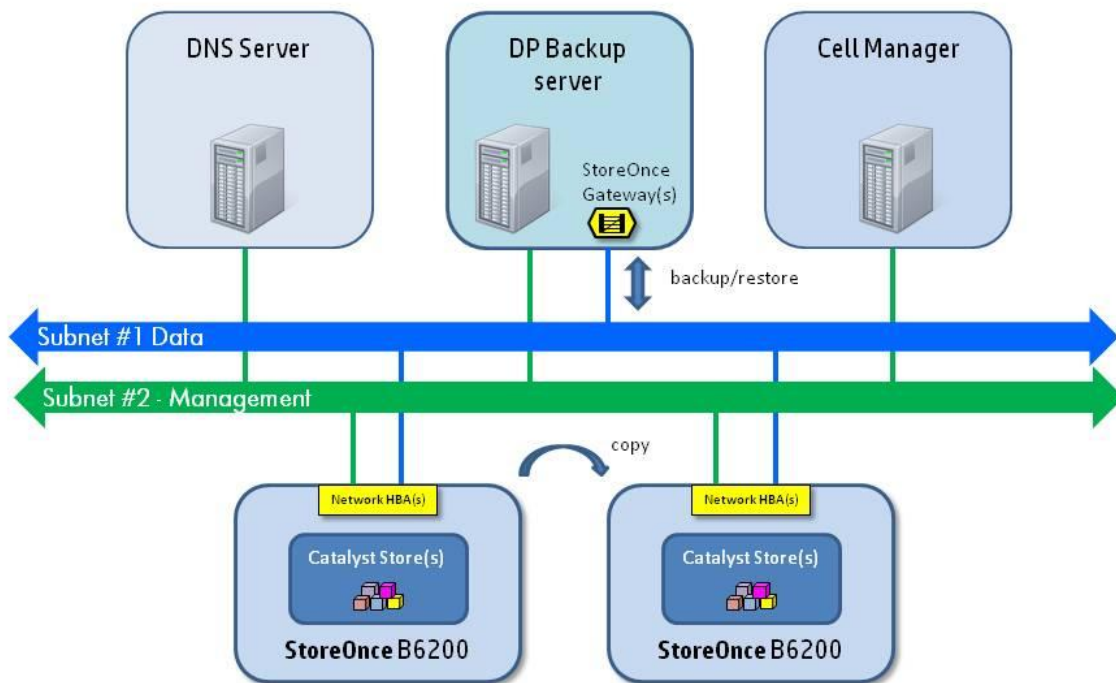


Fig. 36: StoreOnce B6200 Networking example

The StoreOnce B6200 only supports static ip addresses but can be addressed by fully qualified domain name if DNS is used and the B6200 has the ip address of the DNS server configured . Do NOT use host tables to resolve ip addresses to host names. You cannot setup a host table within the StoreOnce B6200 system. When configuring StoreOnce catalyst stores within Data Protector 7 use either fixed ip

address or FQDN if DNS is in use. Backup servers and Cell manager server then would require DNS entries.

StoreOnce Catalyst data and commands are passed between the backup server (equipped with DP media agent) and the StoreOnce B6200 node via the data subnet if configured with 2 subnets. Entries are placed in DNS for both management and data networks. For single subnet configurations there are entries for management and data addresses. The StoreOnce B6200 displays the fully qualified domain name and the ip address for each service set if DNS is used. Fig 37 shows an example of a service set virtual ip address/fqdn. The management ip address is not displayed and is set in the initial configuration. However it can be displayed with the CLI command '**net show config**'.

System Information	
Type	ProLiant DL380 G7
Node Name	hp4a9273bd8c-2
Serial Number	hp3c4a9273bd8c01
VIF IP	172.16.60.2
VIF FQDN	b6200ss1.nearline.local
Software Revision	3.2.0-1219.42

Fig. 37 StoreOnce B6200 network information

The DNS entries for a typical 4 node B6200 system operating with 2 subnets is shown in Fig. 38. Note that the management (B6200) lists an ip address on a separate subnet to the 4 node VIFs (B6200SS1 – B6200SS2). The ip addresses are all static as DHCP is not supported for the B6200 system.

Name	Type	Data	Timestamp
(same as parent folder)	Start of Authority (SOA)	[2312], reddwarf.nearline.l...	static
(same as parent folder)	WINS Lookup	[192.168.0.2]	static
(same as parent folder)	Name Server (NS)	reddwarf.nearline.local.	static
(same as parent folder)	Host (A)	192.168.0.2	19/09/2012 12:00:00
18g2	Host (A)	192.168.10.4	static
3ParSwitch	Host (A)	192.168.1.9	static
_msdcs			
_sites			
_tcp			
_udp			
Admin-Laptop3	Host (A)	192.168.200.33	11/09/2012 10:00:00
admin-PC	Host (A)	192.168.90.1	19/09/2012 14:00:00
Aldwych	Host (A)	192.168.0.239	static
Andromeda	Host (A)	192.168.0.22	static
auriga	Host (A)	192.168.0.5	19/09/2012 12:00:00
B6200	Host (A)	192.168.60.10	static
B6200DemoSwitch	Host (A)	192.168.1.6	static
B6200SS1	Host (A)	172.16.60.2	static
B6200SS2	Host (A)	172.16.60.4	static
B6200SS3	Host (A)	172.16.60.6	static
B6200SS4	Host (A)	172.16.60.8	static
beagle	Host (A)	192.168.0.230	static
BEN	Host (A)	192.168.0.238	24/08/2012 13:00:00
BILL	Host (A)	192.168.0.237	19/09/2012 13:00:00
bombe	Host (A)	192.168.0.90	static

Fig. 38: Example of DNS entries for a 4 node StoreOnce B6200 system

Each Data Protector Backup/Application server must have access to the data subnet. Data Protector clients running only disk agents and the cell manager do not require access to the data subnet of the B6200.

A single gateway can be configured. This enables access to other subnets for management OR data.

Key Points:

- Do not use host tables for hostname resolution.
- DHCP is not supported.
- StoreOnce Catalyst does not use active directory services for access control.
- B6200 does support the use of DNS.
- StoreOnce catalyst data and commands use the data subnet when 2 subnets are in use on a StoreOnce B6200.
- Only one network gateway can be configured.
- Cell manager needs a network path to backup servers and clients. Access to StoreOnce B6200 is via media agent and gateway.

Bonding network interfaces

The network interfaces used in StoreOnce B6200 are bonded. This technology sometimes referred to 'link aggregation'. Bonding combines 2 network interfaces into a single 'logical' interface with a single 'virtual' IP address (this is called a VIF in StoreOnce B6200 systems). There are different modes for bonding which can deliver additional throughput (2 x bandwidth) or high availability). Some bonding modes require special network switch support. For the HP B6200 no special switch support is required. The StoreOnce B6200 with template 2 configured uses mode 6 bonding which distributes the network traffic across the interfaces. (Adaptive load balance) All other templates use mode 1 (active backup).

Active backup provides high availability but transfers data via one network connection. If that connection fails then the 2nd connection takes over.

Key Points:

- HP StoreOnce systems do NOT require special switch settings.
- IEEE 802.3ad support is not available.
- VLAN tagging is not yet supported.

Housekeeping

Housekeeping is a process which is necessary in any deduplication storage process and is required at some time after any deletion or overwrite operations performed by the host system. Data when written is stored essentially as ‘chunks’ of unique data together with a index reference and index count. The theory being that an unchanged ‘chunk’ of data between backups is stored only once and subsequent backups just increase a counter. Using StoreOnce Catalyst and Data Protector 7 the same media item is not overwritten as a tape cartridge would be. The items are removed from the media pool and deleted from the store on expiry. There is no ‘concept’ of appending backups. A new item is created for every backup.

When a backup is overwritten or expired the housekeeping process will start providing the blackout window is not active. The process is no different to VTL or NAS share housekeeping. Housekeeping is I/O intensive so the more disk shelves installed the faster the rate. Approximate housekeeping rates for a single node are shown in table 1. They refer to how much data would have been overwritten. So if a backup is ‘expired’ then work out the time required from the volume of data. These figures are built-in to the sizer tool.

Configuration	Housekeeping rate MB/s
B6200 single node – base unit 1 disk shelf	800
B6200 single node – 2disk shelves	1019
B6200 single node – 3 disk shelves	1237
B6200 single node -4 disk shelves	1455

Table 1

It is important to periodically check that the housekeeping process is keeping up. A graphical display is available for each store. Use B6200 management GUI navigation tree to select StoreOnce Catalyst -> Housekeeping. See Fig. 39. Ensure that the ‘jobs processed’ (green) always exceed the ‘jobs received’ (red). Normally housekeeping is set to run every day outside the backup or replication window.

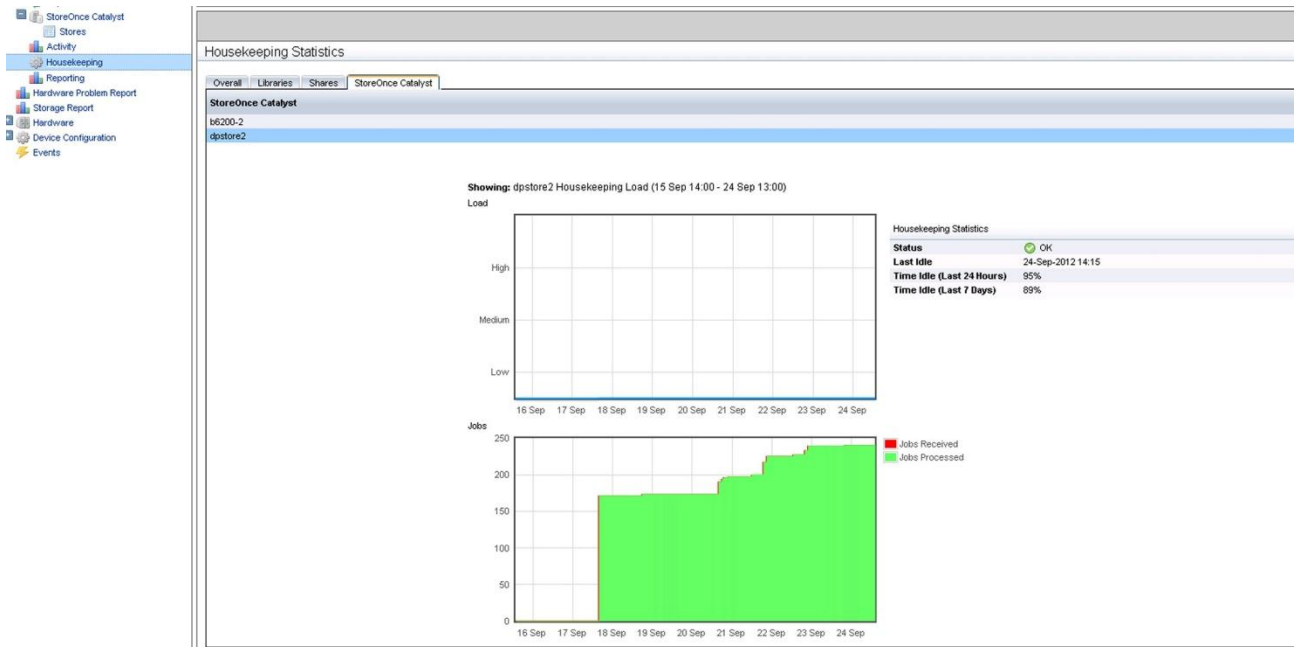


Fig. 39: Monitoring Housekeeping

Key Points:

- Adjust blackout windows so that housekeeping does not clash with backup or replication.
- Ensure that the 'delete unprotected media' will run at sometime before housekeeping. If the option is set to default unprotected media will run at 12:00hrs.
- Do NOT permanently disable housekeeping.

HP StoreOnce Catalyst Licensing

HP StoreOnce Catalyst is licensed on a per couplet basic for the B6200 system. If it is intended to move data between catalyst stores then both systems will need a license. Licenses are applied via the command line. The license product number is:

TC397AAE – HP B6200 StoreOnce Catalyst Licence for Data Protector

(Note: A separate licence is required for use with Symantec OST – TC396AAE if it is intended to use B6200 with both products.)

The cluster id# will be required when obtaining the key from HP's webware site. The cluster id# can be obtained via the management GUI or CLI.

CLI command is #system show config

This is displayed as the serial number in the top level menu. (Serial number=cluster id)

To enter a licence use the CLI command as follows:

#licence add << licence string >>

To check a licence use:

#licence show

Output should be:

License(s):

License Name: D2D Catalyst Data Protector

Key: 9CJG CQEA H9PA CHUY VRB4 HW6V Y9JL KMPL B89H MZVU DXAU 2CSM GHTG L762 YVRZ
GKZ4 KJVT D5KM EFVW TSNJ 2SXP 6TS2 JMQK 9828 UJY5 TWV5 ZWWQ Q687 RX2U G4VY 5FE6
SJ66 388L 4ZX5 XWDD XCRS ZQKL LR7M 4WBL 2N3E VQ9G RUX2 CZUH WG7Y Y2KN F8RV XYRR
HNQU T827 ANDB WVTY LTXN KSWK XUY4 NGHL E7A4 R6KH BYAB G5RB JLEF VVW4 CP6F SF9P
R7GS "IPP12DZ693996B5A0AC672F517E TC397AAE HP B6200 StoreOnce Catalyst DP E-LTU"

Description: IPP12DZ693996B5A0AC672F517E TC397AAE HP B6200 StoreOnce Catalyst DP
E-LTU

Licenses to Use (LTU): 1

Expiration Date: None

HP Data Protector licensing for use with StoreOnce Catalyst

Data Protector 7 does not require a special licence for StoreOnce Catalyst. It is only required to have the Advanced Backup to Disk licence for the required capacity. You do of course need the Data Protector Starter Pack for the appropriate operating system. The starter pack allows and unlimited number of clients to be setup at no additional cost. The Advanced backup to Disk license also allows the use of VTL. User can choose whether to license VTL as tape or disk. Remember that the DP7 media agent only runs on the Windows and Linux platforms. HP-UX will be released in due course.

Advanced Backup to disk licenses

B7038AA 1TB

B7038BA 10 TB

B7038CA 100TB

HP StoreOnce B6200 Autonomic Restart with HP DataProtector Catalyst

The StoreOnce B6200 System has the ability to 'fail over' in the rare case of node hardware failure. This process is autonomic meaning that it requires no external intervention. As the service set which provides the StoreOnce emulations and deduplication takes a period of time to start up on the good node in the couplet the Data Protector jobs will fail. However HP has produced a script which can be integrated into HP Data Protector to restart all the backup jobs automatically. The script can be customised and run as a post backup executable. A separate paper is available from HP with full instructions. (www.hp.com/go/dataprotector).

Conclusions and key learning points

HP StoreOnce Catalyst and HP Data Protector combine to give a flexible advanced high performance backup solution. The ability to deduplicate data at the backup or application server can deliver considerable performance gains and unlike competitor product will not incur additional licence charges. The StoreOnce Catalyst protocol is also much more efficient and suited to WAN or LAN than traditional NAS protocols. Backups can also be replicated to additional sites but in a bandwidth efficient manner reducing networking costs and eliminating tape transportation.