

# HP StorageWorks VLS and D2D Solutions Guide

## Design Guidelines for Virtual Tape Libraries with Deduplication and Replication

This document describes the HP StorageWorks VLS and D2D systems and their concepts, including automigration, deduplication, and replication, to help you define and implement your virtual tape library system. It includes best practices for working with specific backup applications. This document is intended for use by system administrators who are experienced with setting up and managing system backups over a SAN.



\* A G 3 0 6 - 9 6 0 2 8 \*

Part number: AG306-96028  
Seventh edition: March 2010



**Legal and notice information**

© Copyright 2005, 2010 Hewlett-Packard Development Company, L.P.

The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

**Acknowledgements**

Microsoft, Windows, Windows XP, and Windows NT are U.S. registered trademarks of Microsoft Corporation.

Oracle is a registered US trademark of Oracle Corporation, Redwood City, California.

---

# Contents

1 Introduction .....	13
2 Concepts .....	15
Disk-based Backup and Virtual Tape Libraries .....	15
Problems Addressed by Virtual Tape Libraries .....	15
Integration of Disk in Data Protection Processes .....	15
Where Virtual Tape Fits in the Big Picture .....	16
HP VLS and D2D Portfolio .....	16
Typical VLS Environments .....	17
Typical D2D Environments .....	17
What are the Alternatives? .....	18
Physical Tape .....	18
NAS .....	18
Application-based Disk Backup .....	18
Business Copy .....	19
Deduplication .....	20
Introduction .....	20
HP StorageWorks Deduplication Solutions .....	20
Deduplication Ratios .....	21
Target-based Deduplication .....	22
Tape Oversubscription .....	23
Replication .....	23
Introduction to Replication .....	23
HP StorageWorks Replication Solutions .....	24
Replication Deployment Options .....	26
Backup Application Interaction with Replication .....	28
Replication Limitations .....	29
3 Backup Solution Design Considerations .....	31
Analyze the Existing Environment .....	31
Consider How you Want to Back Up your Data to the VLS or D2D .....	31
Consider How you Want to Copy the Backup Data to an Off-site Location .....	33
Consider Speed and Ease of the Restore as well as the Backup .....	33
Single Library vs. Multiple Libraries .....	33
Single Library .....	33
Benefits of Single Library Systems .....	34
Considerations for Single Library Systems .....	34
Multiple Library .....	35
Benefits of Multiple Library Systems .....	35
Multiplexing, Multistreaming, and Multipathing .....	36
Multiplexing .....	36
Multistreaming .....	36
Multipathing .....	36
Blocksize and Transfer Size .....	37

LAN-free Backups .....	37
Retention Planning .....	37
Future Data Growth .....	38
Considerations for Copies .....	38
Copy to Physical Tape through the Backup Application .....	39
Media Server Considerations .....	40
Benefits of Copying to Physical Tape through the Backup Application .....	40
Considerations for Copying to Physical Tape through the Backup Application .....	41
Copy to Tape using VLS Automigration .....	41
Benefits of Echo Copy .....	42
Considerations of Echo Copy .....	42
Copy to Tape using D2D Tape Offload .....	43
Copy to Remote Disk Backup Device using Replication .....	43
Benefits of Replication .....	44
Considerations of Replication .....	44
Creating Archive Tapes from the Replication Target .....	44
Considerations for Restores .....	44
Restoring from Disk Backup Device .....	45
Restoring from Backup Application-created Tape Copy .....	45
Restoring from the Replication Target .....	45
Performance Bottleneck Identification .....	45
Backup SAN Design Guidelines .....	46
General SAN Design Considerations .....	46
SAN Zoning .....	46
Operating System Tape Configuration .....	47
LUN Masking and Mapping .....	49
Backup Application Basic Guidelines .....	50
HP Data Protector Application Overview .....	51
Symantec NetBackup Application Overview .....	51
IBM TSM Application Overview .....	52
EMC NetWorker Application Overview .....	53

## 4 D2D Systems ..... 55

D2D Devices .....	55
D2D Defined .....	55
D2D Technical Specifications .....	56
D2D Design Considerations .....	56
D2D Port Optimization .....	56
Multiple Backup Streams .....	57
Disable Backup Application Verify Pass .....	58
Tape Copy or Offload Performance .....	58
Tape Library Emulation .....	59
Optimizing Rotation Scheme to Reduce Housekeeping .....	59
D2D Blueprints .....	60
Single Site Cost Effective Backup Device Consolidation .....	60
Large SME Site Consolidation Requiring Fibre Channel Shared Devices .....	61
Multi-site Small Business Disaster Recovery Solution .....	63
Introduction to VMWare Terminology .....	65
Using D2D in Simple VMWare Environments with ESXi .....	66
Using D2D in VMWare Environments with ESX4 and VMWare Consolidated Backup .....	66
Using D2D in larger VMWare environments with HP Data Protector Zero Downtime Backup with Instant Recovery (ZDB + IR) .....	67
D2D Dynamic Deduplication .....	69
How it Works .....	69

Dynamic Deduplication Implementation .....	70
Restoring Data .....	71
Housekeeping .....	72
D2D Replication .....	73
How it Works .....	73
D2D Replication Implementation .....	73
Licensing .....	74
Implementing the Initialization Process .....	74
Replication Setup .....	76
Reporting .....	82
Design Considerations .....	84
Link Sizing .....	84
Telco Provisioning — Overview of Inter-Site Links .....	90
Some Idea of Telco Costs .....	92
Telco Terminology and Branding for WAN Services .....	93
D2D Replication Data Recovery Options .....	95
Restore Directly from the D2D Target Device .....	95
Reverse Replication on the D2D .....	95
Reverse Tape Initialization on the D2D .....	98
Creating Archive Tapes from the Target .....	99

## 5 Virtual Library Systems ..... 101

VLS Devices .....	101
VLS Defined .....	101
Prime Environments .....	102
VLS Benefits .....	102
VLS9000-series Configurations .....	103
VLS Technical Specifications .....	104
How it Works .....	111
VLS Scalability .....	111
VLS Automatic Performance Load Balancing .....	113
VLS Warm Failover .....	114
Implementation .....	115
Device Configuration Preparation .....	115
Capacity Licensing .....	117
Virtual Libraries/drives/cartridge Configuration .....	118
EVA Preparation .....	119
Storage Pooling .....	121
VLS Design Considerations .....	123
VLS Sizing .....	123
VLS Blueprints .....	127
Reduce Backup Window (No Deduplication) .....	127
Large Scale Backup Consolidation (No Deduplication) .....	128
Managing Data Retention and Reducing Backup Storage Costs .....	129
Full Enterprise Backup and Disaster Recovery Capability .....	130
Multi-data Center Replication to a Common Disaster Recovery Site .....	132
VMWare Backup Environment with ESXi. ....	133
VMWare Backup Environment with ESX4 and VCB-based Backups .....	134
Larger VMWare Environment with ESX4 and HP Data Protector ZDB IR Integration with VMWare Consolidated Backup .....	136
Large Enterprise Cross-site Backup with Deduplication and Inbuilt Disaster Recovery .....	138
Combining D2D and VLS in an End-to-end Solution for ROBO/Regional DCs and Main Data Center using HP Data Protector .....	139
VLS Automigration .....	140

Echo Copy Concepts .....	141
Implementation .....	141
Echo Copy Pools .....	141
Automigration Policy .....	142
Automigration Setup .....	143
Design Considerations .....	145
Automigration Use Models .....	145
Sizing the Tape Library .....	147
Restoring from Automigration Media .....	148
VLS Accelerated Deduplication .....	148
How it Works .....	149
Accelerated Deduplication Implementation .....	152
Supported Backup Applications and Data Types .....	153
Deduplication Preparation .....	153
Licensing .....	154
Migrating your Existing Backup Data .....	155
Configuration and Reporting .....	155
Handling the Device Out of Capacity Condition .....	156
Design Considerations .....	157
Performance .....	157
Capacity Sizing .....	158
Node Oversubscription .....	158
Optimum Record Sizing .....	159
Optimizing File Server Deduplication .....	159
Media Management .....	160
Client and Backup Job Naming .....	160
Disabling Deduplication on Specific Backup Policies .....	161
VLS Replication .....	161
How it Works .....	161
Replication of Incremental Backups .....	163
Multi-node Replication Scaling .....	164
VLS Replication Implementation .....	164
Licensing .....	164
Replication Preparation .....	165
Replication Setup .....	166
Implementing the Initialization Process .....	172
Reporting .....	176
Design Considerations .....	180
Advanced Replication Job Control .....	180
Link Sizing .....	182
Device Sizing .....	187
VLS Replication Data Recovery Options .....	187
Restore Directly from the VLS Target Device .....	187
Restore the VLS over the LAN/WAN .....	189
Rebuilding Source Device and Re-establishing Replication .....	190
Creating Archive Tapes from the Target .....	190
VLS Non-deduplicated Replication .....	192
Replicating a Subset of Virtual Cartridges - a Use Case .....	193
Detailed Backup Application Guidelines for VLS .....	194
HP Data Protector .....	194
Data Protector General Guidelines .....	194
Data Protector Deduplication Guidelines .....	195
Data Protector Import Example Script .....	196
Symantec NetBackup .....	197
NetBackup General Guidelines .....	197

NetBackup Deduplication Guidelines .....	198
NetBackup Import Example Script .....	198
IBM TSM .....	199
TSM General Guidelines .....	199
TSM Deduplication Guidelines .....	200
TSM Useful Queries .....	203
EMC NetWorker .....	204
NetWorker General Guidelines .....	204
NetWorker Deduplication Guidelines .....	206
<b>6 Support and Other Resources .....</b>	<b>209</b>
Related Information .....	209
Documents .....	209
Websites .....	209
Contacting HP .....	209
Before you contact HP .....	209
HP contact information .....	210
Subscription Service .....	210
Document Conventions and Symbols .....	210
<b>Glossary .....</b>	<b>213</b>
<b>Index .....</b>	<b>215</b>

---

# Figures

1	Common Backup Technologies .....	16
2	HP Virtual Tape Library Product Range .....	17
3	Basic Write-to-disk Setup .....	19
4	Unique Backup Data .....	20
5	Enterprise Deployment with Small and Large Remote and Branch Offices .....	24
6	Remote Site Data Protection Before Replication .....	25
7	Remote Site Data Protection Using Replication .....	25
8	Replication Configuration Options .....	27
9	Presenting the Replication Target to a Different Backup Application .....	28
10	Backup to the VLS in a LAN/SAN Hybrid Environment .....	32
11	Backup to VLS in a Simple Deployment (VLS9000-series with One Shared Library Shown) .....	34
12	Backup to VLS in a Simple Deployment (VLS9000-series with Four Dedicated Libraries Shown) .....	35
13	Writing to Tape in a LAN/SAN Hybrid Environment .....	40
14	Echo Copy is Managed through Automigration .....	42
15	Virtual Tape Environment .....	50
16	TSM LAN-based backups .....	52
17	TSM LAN-free backups .....	53
18	Single Site Backup Consolidation using iSCSI .....	61
19	Large SME Site Consolidation Requiring FC Shared Devices .....	62
20	Multi-site Small Business Automated Disaster Recovery Solution .....	63
21	Multi-site Small Business Automated Disaster Recovery Solution with no Physical Tape .....	64
22	VMWare Infrastructure Environment .....	65
23	Using D2D in VMWare Environment with ESX4 and VCB Backup for SAN-based Backup .....	66
24	Using D2D in a VMWare Environment with ESX4 and HP Data Protector ZDB IR .....	68
25	Hash-based Chunking .....	70
26	D2D Configuration .....	71
27	Dynamic Deduplication Ratio .....	71
28	Restoring from a Recipe File .....	72
29	Options for Initialization the Initial Data .....	74
30	Slot Mappings .....	77



31	Example Schematic .....	78
32	Example 1: 100 GB Virtual Cartridge Replication over a 2 Mb/sec Link with 100% Bandwidth Utilization .....	85
33	Example 2: 100 GB Virtual Cartridge Replication over a 2 Mb/sec Link with 25% Bandwidth Utilization .....	86
34	Example 3: 500 GB Virtual Cartridge Replication over a 2 Mbit/sec Link .....	86
35	Example 4: 500 GB Virtual Cartridge Replication over a 10 Mbit/sec Link .....	87
36	Many-to-One Example using Mixed WAN Link Speeds .....	88
37	A Basic WAN using MPLS .....	92
38	Recovery Options — Data Center Rebuild and Reverse Replication .....	95
39	Recovery Options — Physical Tape .....	99
40	Creating Archive Tapes from the Target .....	100
41	7.5 TB System with Expansion Disk Array Enclosures .....	103
42	Full VLS9000 System .....	104
43	Internal Architecture of the VLS (VLS9000 Shown) .....	112
44	Virtual Drives across Multiple Nodes .....	113
45	Storage Pool Load Balancing on the VLS .....	114
46	VLS12000 EVA Gateway Connected to an EVA .....	120
47	Single vs. Multiple Storage Pool Considerations .....	122
48	VLS9000 Partially Populated Storage Pools .....	123
49	Testing Backups to Tape vs. Backups to a VLS .....	126
50	Using VLS to Reduce the Backup Window .....	127
51	Using VLS for Large Scale Consolidation .....	128
52	Using VLS to Manage Data Retention and Reduce Backup Storage Costs .....	130
53	Using VLS with Full Enterprise Backup and Disaster Recovery Capability .....	131
54	Using VLS with Multi-data Center Replication to a Common Disaster Recovery Site .....	132
55	Using VLS with VMWare using Backup Agents .....	134
56	Using VLS with VMWare using VCB .....	135
57	Using VLS in Larger VMWare Environments with HP Data Protector ZDB IR Backup .....	137
58	Using VLS in Enterprise Cross Site Backup Scenarios .....	138
59	Using D2D and VLS in ROBO/RDC/MDC Environments .....	140
60	Linked Media .....	142
61	Automigration Policy Details .....	142
62	Automigration Media Life Cycle .....	144
63	Multiple Virtual Libraries .....	146
64	Shared Virtual Library .....	147
65	Steps of Accelerated deduplication .....	149

66 Backup Analysis .....	150
67 Physical versus Logical Data .....	150
68 Backup Sessions with Duplicate and Unique Data .....	152
69 VLS9000 and VLS12000 node oversubscription .....	158
70 How Replication Fits into the Deduplication Architecture .....	162
71 The Replication Sequence .....	163
72 Replication Scaling .....	164
73 Preparing the Network Connection for Replication in a VLS .....	166
74 VLS Replication by Numbers .....	167
75 VLS Initialization using Physical Tape .....	173
76 VLS Initialization using the WAN .....	175
77 VLS Initialization by Co-location .....	176
78 Dividing the Backup Jobs by Priority Level .....	181
79 VLS Sizing Example 1 .....	183
80 VLS Sizing Example 2 .....	184
81 VLS Sizing Example 3 .....	185
82 VLS Sizing Example 4 .....	186
83 VLS Recovery Options for a Data Center Rebuild .....	188
84 Data Recovery from the Target VLS .....	189
85 Creating Archive Tapes from the VLS .....	191
86 Site-to-site Replication .....	193

---

# Tables

1 VLS Compared to Application-based Write-to-disk .....	19
2 HP Deduplication Solutions .....	21
3 1 TB File Server Backup .....	22
4 Deduplication Ratio Impact .....	22
5 Estimated Time to Replicate Data for a 1 TB Backup Environment at 2:1 .....	26
6 D2D Technical Specifications .....	56
7 Sample Initialization and Replication Times .....	89
8 Telco Prices for a 3-year Contract Including Installation (No Existing Links in Place) .....	92
9 Prices of a Major US Telco for HP Facilities in the Houston, Texas Area .....	93
10 Prices of a Telco Provider in Colorado for HP Sites Located around 250 km (175 miles) Apart .....	93
11 Prices for Point to Point Connects from HP Facilities based in Singapore .....	93
12 Sample British Telecom Services and Speeds .....	94
13 VLS9000-series Multi-node Maximum Capacity by Configuration with 20-port Connectivity Kit and 40 TB Arrays (2:1 Compression, Non-deduplication) .....	104
14 VLS9000-series Multi-node Maximum Capacity by Configuration with 32-port Connectivity Kit and 40 TB Arrays (2:1 Compression, Non-deduplication) .....	105
15 VLS9000-series Multi-node Maximum Capacity by Configuration with Two 32-port Connectivity Kits and 40 TB Arrays (2:1 Compression, Non-deduplication) .....	105
16 VLS9000-series Backup Performance by Configuration .....	106
17 VLS6000-series Technical Specifications .....	106
18 VLS6000-series Performance by Configuration, 4 GB Models .....	107
19 VLS12000 EVA Gateway Technical Specifications .....	108
20 VLS12000 EVA Gateway Backup Throughput using Deduplication .....	108
21 VLS12000 EVA Gateway Backup Throughput without Deduplication .....	109
22 Base Capacity License by VLS Model .....	117
23 Adding Capacity Beyond the Base Configuration .....	117
24 Required Deduplication Licenses by Platform .....	155
25 VLS9000 and VLS12000 Node Oversubscription Options .....	158
26 Network Latency .....	181
27 Summary of the Calculations Involved in the Sizing Examples .....	185
28 Recommended Settings for each TSM Server's dsmserv.opt File or Storage Agent's dsmsta.opt .....	199

29 Recommended Settings for each Backup-archive Client's dsm.sys File (dsm.opt File for Windows) .....	200
30 Document conventions .....	210

---

# 1 Introduction

Welcome to virtual tape libraries. This guide describes the HP StorageWorks VLS and D2D systems and their concepts, including automigration, deduplication, and replication, to help you define and implement your virtual tape library system. It includes best practices for working with specific backup applications.

Although every user environment and every user's goals are different, there are basic considerations that can help you use the VLS or D2D effectively in your environment. The VLS and D2D are two powerful and flexible families of devices. Because they can be productively used in so many ways, there is no "best" configuration. But by asking yourself the questions and following the parameters outlined in this guide, you can define and implement a system that is best for your particular environment and applications.

Before proceeding, make sure you are familiar with the items below.

- Tape backup technologies, tape libraries, and backup software.
- SAN environments.
- Fibre Channel technology.

See the [Glossary](#) for the definition of acronyms and specific terms.



## NOTE:

This guide replaces the *HP StorageWorks Virtual Library System Solutions Guide*, the *HP StorageWorks Deduplication and replication solutions guide*, and the *HP StorageWorks Deduplication solutions guide*.

---



---

# 2 Concepts

## Disk-based Backup and Virtual Tape Libraries

### Problems Addressed by Virtual Tape Libraries

You can optimize your backup environment with VLS and D2D if you are:

- Not meeting backup windows due to slow servers.
- Not consistently streaming your tape drives.
- Dealing with restore problems caused by interleaving.
- Performing many restores (such as single file or small database restores).
- Backing up data that has a short life.
- Having issues with backup reliability.
- Using snapshot and clone technology for non-critical data (which makes the storage inappropriately expensive for the nature of the data).
- Looking to deemphasize tape in your environment. Bear in mind that removable media remains valuable in its own right and for particular purposes such as site protection and protection from malicious attack (e.g., viruses and hackers), data distribution, data copy, archive, and regulatory compliance.
- Improving media management. You can keep incremental backups on virtual tape and send full backups straight to tape.

### Integration of Disk in Data Protection Processes

Globalization, 24x7 environments, and consolidation are driving more rigorous data protection requirements. To address these requirements, disk is frequently introduced into the backup process.

In disk solutions, data is backed up from an application server (disk) over a dedicated SAN to a disk-based system and from there to a traditional tape library. This provides enhanced solutions for slow servers, single-file restores, and perishable data.

One of the particular benefits of the VLS and D2D is that they make a disk array look to your backup server like a tape library. Implementation requires no new software and no significant redesign of your backup processes. On the VLS300 and VLS12000 Gateways, because they are attached to an EVA, the existing Fibre Channel infrastructure and management framework is used and there is no new management server required.

---

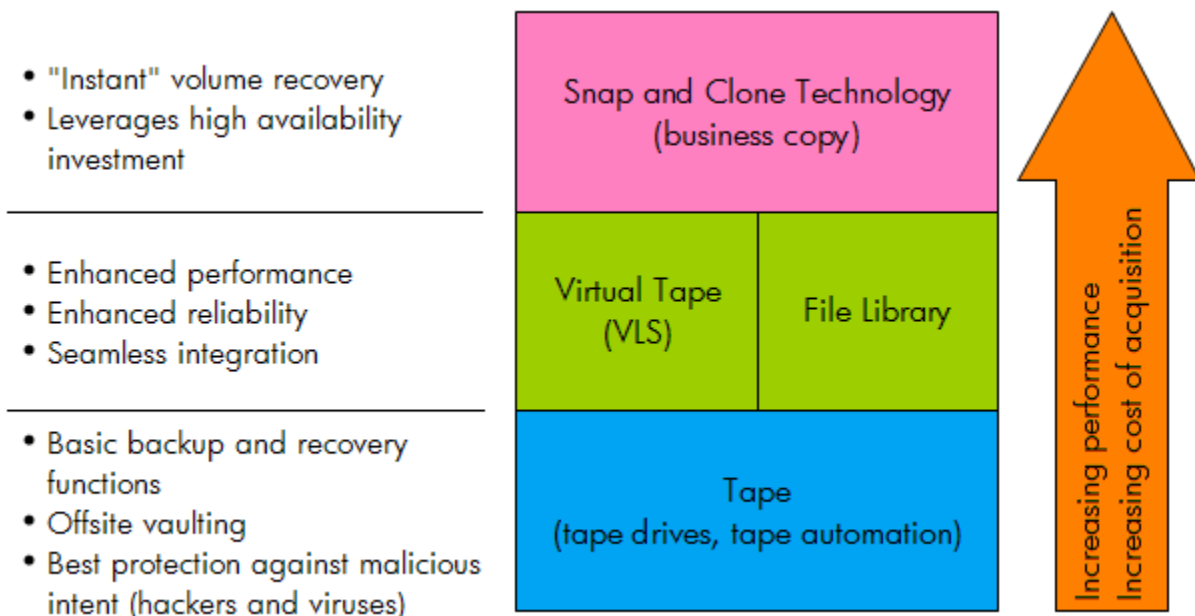
 **NOTE:**

Tape holds its value for ease of vaulting, economical long-term retention, and immutability (with WORM). It is the last step in your data's storage cycle.

---

## Where Virtual Tape Fits in the Big Picture

Virtual libraries are not necessarily the only piece of your backup plans, but they can be an integral piece of a successful solution. [Figure 1](#) illustrates the common backup technologies and their relative benefits and costs.



**Figure 1 Common Backup Technologies**

See [What are the Alternatives?](#) for more discussion of the other potential players in your backup environment.

## HP VLS and D2D Portfolio

HP offers a wide range of disk-based backup products to help organizations meet their data protection challenges. Moving the front line of data protection from tape to disk reduces administrative overhead; daily backups are entirely automated and involve no tape handling to provide better backup reliability and less worry.

The entry level D2D100 series Backup System meets the needs of small businesses as a low-cost solution that does not incorporate deduplication technology. The D2D2500 is well suited for remote and branch offices and small IT environments, while the more powerful D2D4000 and D2D4100 are designed for medium-sized companies and small data centers. The D2D2500, D2D4000 and D2D4100 products include HP Dynamic deduplication, which provides low cost and flexibility to meet the needs of smaller IT environments.

HP Virtual Library Systems are known for their easy integration, simple management, performance and capacity scalability, and fast restores. The VLS6000, VLS9000 and VLS12000 EVA Gateway are designed for medium to large-scale enterprises. They feature Accelerated deduplication, available by license, to deliver the best backup performance and scalability for high availability data center environments.



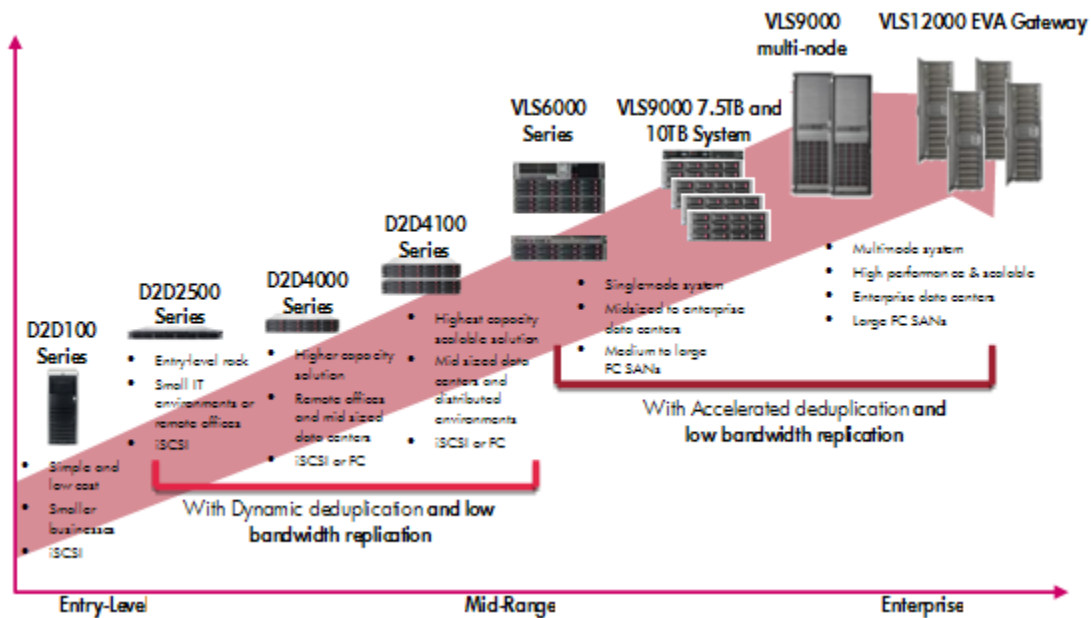


Figure 2 HP Virtual Tape Library Product Range

## Typical VLS Environments

In a typical enterprise backup environment, there are multiple application servers backing up data to a shared tape library on the SAN. Each application server contains a remote backup agent that sends the data from the application server over the SAN fabric to a tape drive in the tape library. However, because backup over the SAN is single-threaded (a single host is backing up to a single tape drive), the speed of any single backup can be limited. This is particularly true when the environment has high-speed tape drives such as Ultrium 2 or Ultrium 3. The hosts simply cannot keep the drives streaming at capacity.



### NOTE:

HP Ultrium drives will adjust the tape speed to match the data stream to prevent “back-hitching.” However, the tape drive is still not operating at optimal performance and cannot share bandwidth with another backup job.

Enterprise data centers with slow SAN hosts in the environment may be unable to utilize the full performance of high-speed tape drives. Also, shared tape libraries on the SAN can be difficult to configure both in the hardware and in the data protection software.

## Typical D2D Environments

In a typical entry-level or mid-range backup environment, the backup application is performing LAN backups to a dedicated (non-shared) backup target such as a tape library connected to the single backup server. Multiple instances of the backup application will generally each require their own dedicated backup target. These environments may also be remote branch offices, each with their own local backup application.

As with the VLS, the backup speed of a single host backing up to a single tape drive is normally limited by the host (which cannot stream high-speed tape drives such as LTO), so currently tape backups use multiplexing to interleave multiple hosts’ backups together into a single tape drive impacting

restore performance. The addition of a D2D device to these environments allows de-multiplexing of the backups so that restore performance is improved, the deduplication allows for a longer retention time on disk without needing significantly higher disk capacities, and the deduplication-enabled replication allows cost-effective off-site copying of the backups for disaster protection.

## What are the Alternatives?

Alternatives to virtual tape solutions include:

- [Physical Tape](#)
- [NAS](#) (network attached storage)
- [Application-based Disk Backup](#) (disk to disk, backup to disk, disk to disk to tape)
- [Business Copy](#) (snapshot and clone solutions)

### Physical Tape

Tape is the foundation for data protection and should be a part of most data protection solutions (except those with highly perishable data). Consider a direct-to-tape scheme if:

- You are doing large image backups (i.e., databases), *or*
- Your servers can stream the tape drives.

*and*

- You do not need fast single file restore, *or*
- Your current backup window is not strained.

### NAS

An alternative to a virtual library is a NAS device acting as a backup target (via NFS or CIFS network file system protocols). However, this protocol has significant performance and scaling limitations; writing backups over TCP/IP and NFS/CIFS to the NAS target uses much more CPU on the backup infrastructure compared to Fibre Channel SAN. In addition, a NAS mount point does not scale to the size of an enterprise virtual tape library. For example, a VLS can present a single virtual library target containing multiple petabytes of tape capacity with all backup jobs configured to use the one common shared high-performance high-capacity VLS backup device.

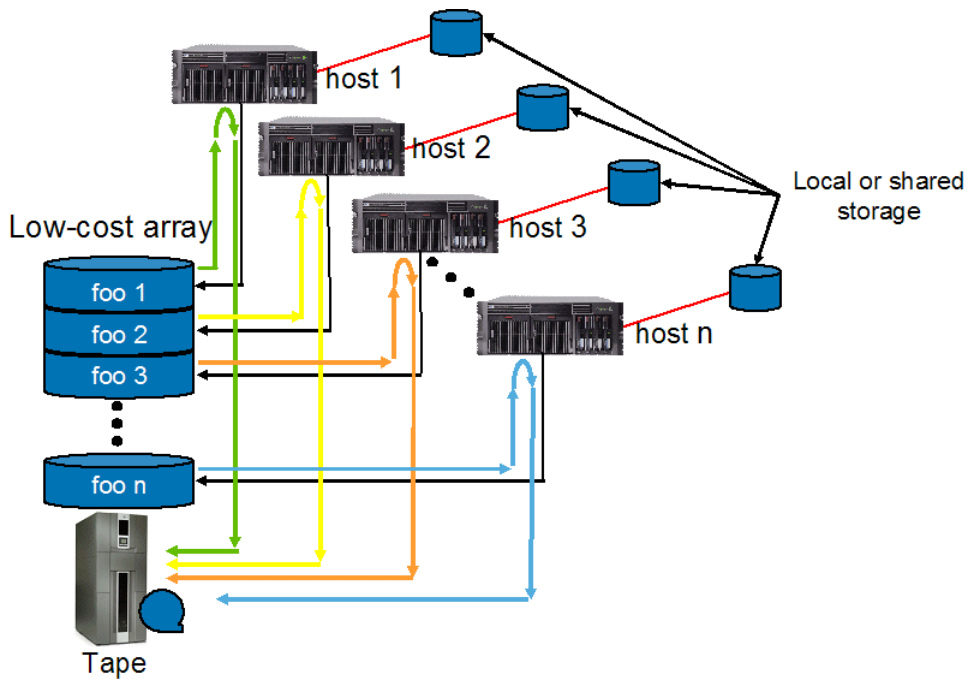
Consider a NAS target if you:

- Do not have high performance requirements.
- Do not want to run SAN backups.
- Do not need the backup target to significantly scale capacity or performance.
- Want to run Data Protector “virtual full backups.”

### Application-based Disk Backup

Utilizing the file library functionality of backup applications is good for small or isolated jobs. When a large-scale implementation is required, virtual tape offers a more easily managed, higher performing solution. Consider a file library system if:

- The application is in a LAN or LAN/SAN hybrid configuration.
- Fewer than four servers write data to secondary disk storage.
- You can redeploy existing arrays as secondary disk storage.
- Your environment is static.



**Figure 3 Basic Write-to-disk Setup**

**Table 1 VLS Compared to Application-based Write-to-disk**

	<b>Virtual tape devices</b>	<b>Write-to-disk</b>
Setup and management complexity	Sets up just like a physical tape library.	Requires configuration of RAID groups, LUNs, volumes, and file systems.
Data compression	Software or hardware enabled (software compression generally decreases performance).	No device-side data compression available.
Performance	Hardware devices are tuned for sequential read and write operations.	Performance dependent on target array or server.
Cost	More expensive acquisition cost. Backup software licenses as if physical library or per TB. Storage efficiency gained through compression. Lower management overhead.	Free or licensed per TB in most backup applications. Higher management overhead.

## Business Copy

Using a business-copy solution (array snapshots/clones) generally involves a much higher cost than a virtual library system. You might, however, implement such a solution if:


- Virtually instant recovery is critical.
- You need to leverage a high-availability investment.
- You are doing image recovery rather than file recovery.
- You need a zero downtime solution.

# Deduplication

## Introduction

In recent years, the amount of data that companies produce has been steadily increasing. To comply with government regulations, or simply for disaster recovery and archival purposes, companies must retain more and more data. Consequently, the costs associated with data storage – labor, power, cooling, floor space, transportation of physical media – have all risen. Virtual tape libraries have become a cornerstone in modern data protection strategy due to their many benefits; chief among these is cost. The list of virtual tape benefits also includes seamless integration into existing backup solutions, improved SAN backup performance, and faster single file restores than those performed with physical tape.

Deduplication, one of the most significant storage enhancements in recent years, promises to reshape future data protection and disaster recovery solutions. This technology is ideal for virtual tape libraries. Deduplication technology references blocks of data that have been previously stored, and *only stores new backup data that is unique*. Data that is not unique is replaced with a pointer to the location of the original data. Because there is often a great deal of duplicate data present from one backup session to the next, disk space is consumed by similar or identical iterations of data. Deduplication greatly improves storage efficiency by only storing an instance of data once, while still allowing backup streams to be restored as if they had been retained in their entirety. See [Figure 4](#).



Item	Description
1	Data from the first backup stream is stored to disk.
2	Duplicate data (in blue) as well as unique data (in red) in a second backup stream is identified.
3	Duplicate data in the second backup stream is eliminated.
4	Unique data in the second backup stream is stored to disk.

**Figure 4 Unique Backup Data**

## HP StorageWorks Deduplication Solutions

HP offers two deduplication technologies: *HP Accelerated deduplication*, a licensed feature available with HP StorageWorks Virtual Library Systems (VLS), and *HP Dynamic deduplication*, an integrated feature with HP StorageWorks D2D Backup System. Both HP deduplication solutions offer the following benefits:

- Longer retention of data.
- Faster, less expensive recoveries and improved service levels.
- Fewer resources consumed, reducing operational costs.

- Completely transparent to host.
- No data is lost – backup streams can be fully restored.
- Block or chunk level deduplication, providing greater reduction of data.
- Even greater reduction of data when combined with traditional data compression.

HP Accelerated deduplication and HP Dynamic deduplication are designed to meet different needs, as shown in [Table 2](#).

**Table 2 HP Deduplication Solutions**

HP Accelerated deduplication	HP Dynamic deduplication
<ul style="list-style-type: none"> <li>• Intended for enterprise users.</li> <li>• Uses <i>object-level differencing</i> technology.</li> <li>• Fastest possible backup performance.</li> <li>• Fastest restores.</li> <li>• Most scalable solution in terms of performance and capacity.</li> <li>• Potentially higher deduplication ratios.</li> </ul>	<ul style="list-style-type: none"> <li>• Intended for mid-sized enterprise and remote office users.</li> <li>• Uses <i>hash-based chunking</i> technology.</li> <li>• Integrated deduplication.</li> <li>• Lower cost and a smaller RAM footprint.</li> <li>• Backup application and data type independence for maximum flexibility.</li> </ul>

See [VLS Accelerated Deduplication](#) and [D2D Dynamic Deduplication](#) for more details on HP deduplication technologies.

## Deduplication Ratios

The storage capacity saved by deduplication is typically expressed as a ratio, where the sum of all pre-deduplicated backup data is compared with the actual amount of storage the deduplicated data requires. For example, a 10:1 ratio means that ten times more data is being stored than the actual physical space it would require.

The most significant factors affecting the deduplication ratio are:

- How long the data is retained.
- How much the data changes between backups.

[Table 3](#) provides an example of storage savings achieved with deduplication. However, many factors influence how much storage is saved in your specific environment. Based on the retention policies shown below, six months of data without deduplication requires 12.75 TB of disk space. With deduplication, six months of data requires less than 1.25 TB of storage.

Retention policy:

- 1 week, 5 daily incremental backups
- 6 months, 25 weekly full backups

Data parameters:

- Data compression rate = 2:1
- Daily change rate = 1% (10% of data in 10% of files)

**Table 3 1 TB File Server Backup**

	Data stored normally	Data stored with deduplication
1st daily full backup	500 GB	500 GB
1st daily incremental backup	50 GB	5 GB
2nd daily incremental backup	50 GB	5 GB
3rd daily incremental backup	50 GB	5 GB
4th daily incremental backup	50 GB	5 GB
5th daily incremental backup	50 GB	5 GB
2nd weekly full backup	500 GB	25 GB
3rd weekly full backup	500 GB	25 GB
...		
25th weekly full backup	500 GB	25 GB
<b>Total</b>	<b>12,750 GB</b>	<b>1,125 GB</b>
<b>Approximately 11:1 reduction in data stored</b>		

Table 4 is an example that may not reflect the savings that all environments achieve using deduplication. As shown, deduplication ratios depend on the backup policy and on the percentage of change between backups.

**Table 4 Deduplication Ratio Impact**

Daily change rate	Backup policy					
	Daily full and weekly full			Daily incremental (10%) and weekly full		
	4 months*	6 months	1 year	4 months*	6 months	1 year
0.5%	15:1	19:1	25:1	12:1	16:1	23:1
1.0%	12:1	13:1	16:1	10:1	11:1	15:1
2.0%	8:1	9:1	9:1	7:1	7:1	9:1
*4 months = 5 daily + 17 weekly backups						

See [Performance](#) for additional information on optimizing your deduplication performance.

## Target-based Deduplication

VLS and D2D deduplication is target-based; the process is running transparently inside the hardware. This means that when the data is read (by copying to physical tape, restoring a backup, etc.), the device rebuilds the data. The data that is read is identical to the data that was originally written (like tape drive compression); there are no pointers in the read data.

## Tape Oversubscription

Deduplication requires more virtual tape capacity than physical disk; this is sometimes called tape oversubscription. The purpose of deduplication is to reduce the amount of disk required to store multiple generations of backups. Be sure to create enough virtual tape capacity to contain your entire retention policy, and the amount of physical disk will be much less capacity due to deduplication.

For example, if you are backing up 50 TB per week and retaining four weeks, you need to create enough virtual tape capacity (after compression) to store 200 TB of backups. If you have 2:1 compression, you must create 100 TB of virtual tape capacity to hold the four weeks of backup data. Given deduplication across the four weeks of backup versions, the amount of physical disk required for this 100 TB of virtual tape would be significantly less.

---

### NOTE:

Do not create too much virtual tape capacity or your backup application may be set to prefer to use blank tapes instead of recycling older tapes. You would likely run out of disk space because the older backups are not being recycled/overwritten and thus the disk space used by these old backups is not freed up. As in the example above, you should create enough virtual tape capacity to hold backups for your entire retention policy but no more.

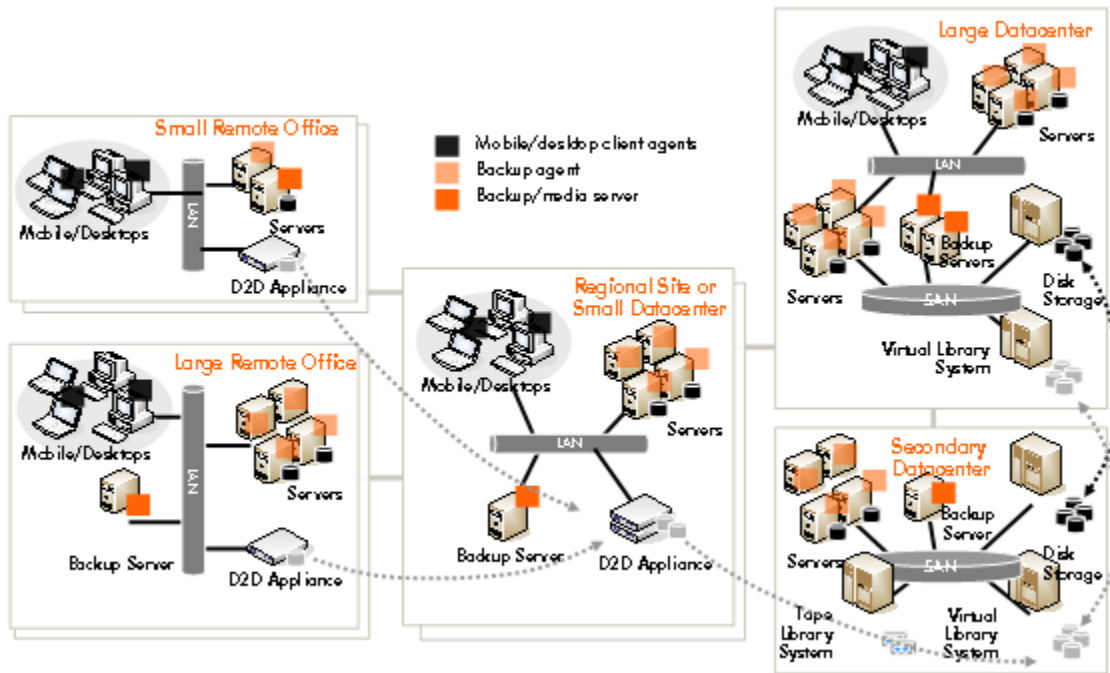
---

## Replication

### Introduction to Replication

Deduplication can automate the off-site process and enable disaster recovery by providing site to site deduplication-enabled replication at a lower cost. Because deduplication knows what data has changed at a block or byte level, replication becomes more intelligent and transfers only the changed data instead of the complete data set. This saves time and replication bandwidth, and is one of the most attractive features that deduplication offers. Replication enables better disaster tolerance with higher reliability but without the operational costs associated with transporting data off-site on physical tape.

You can take control of your data at its furthest outposts and bring it to the data center in a cost-effective way. Using replication, you can protect data anywhere.



**Figure 5 Enterprise Deployment with Small and Large Remote and Branch Offices**

Replication provides end-to-end management of backup data from the small remote office to the regional site and finally into the primary data center, all controlled from the primary data center, while providing local access to backup data as well. Note that replication is within device families (VLS to VLS, D2D to D2D).

## HP StorageWorks Replication Solutions

Most companies now recognize the importance of a robust backup and restore data protection strategy, although only enterprise level users tend to invest in site disaster recovery. In most cases, data protection is in the form of daily off-siting of physical tapes. However, even the off-siting of physical tapes has its down sides—a high level of manual intervention, tape tracking requirements, etc. The physical transfer of tapes off-site is not very automated.

In addition, one of the pain points for many companies large and small is protecting data in remote offices. Untrained IT staff manage a daily backup process involving changing of physical tapes, and the process is prone to human error.

HP replication, available on its VLS and D2D systems, now offers the solution to both these problems. You can replicate local backup data (virtual cartridges) between sites in a reliable, automated manner at a fraction of the costs previously required when using high bandwidth links or in some cases physical tape off-siting.

Consider the “Before” and “After” scenarios detailed below.



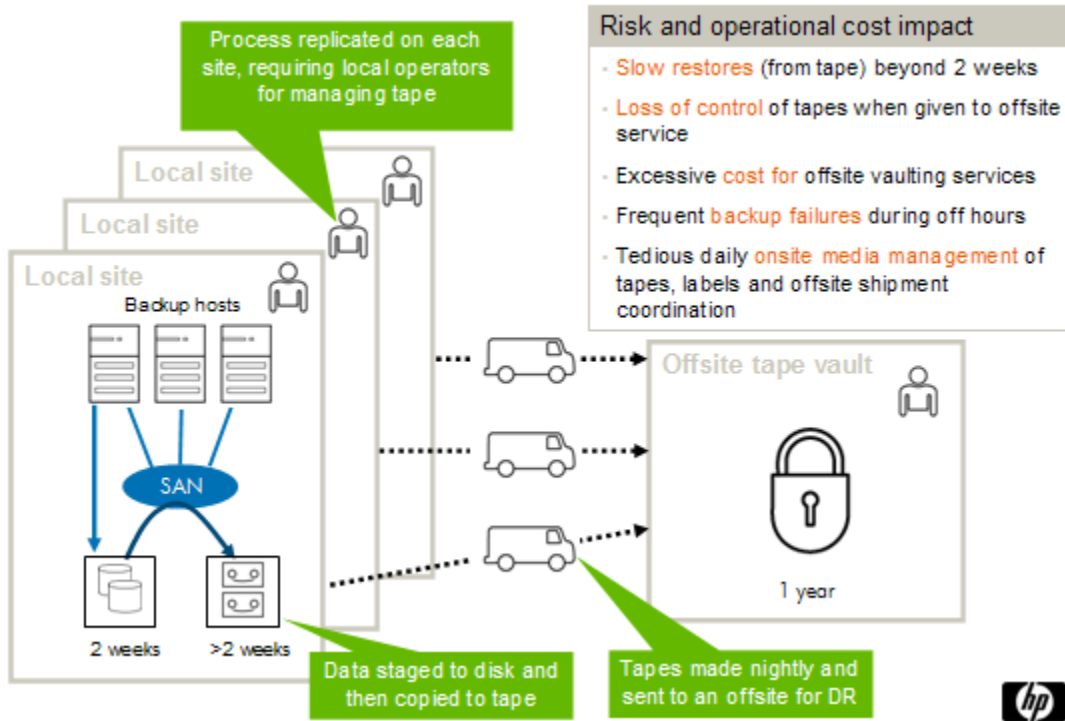


Figure 6 Remote Site Data Protection Before Replication

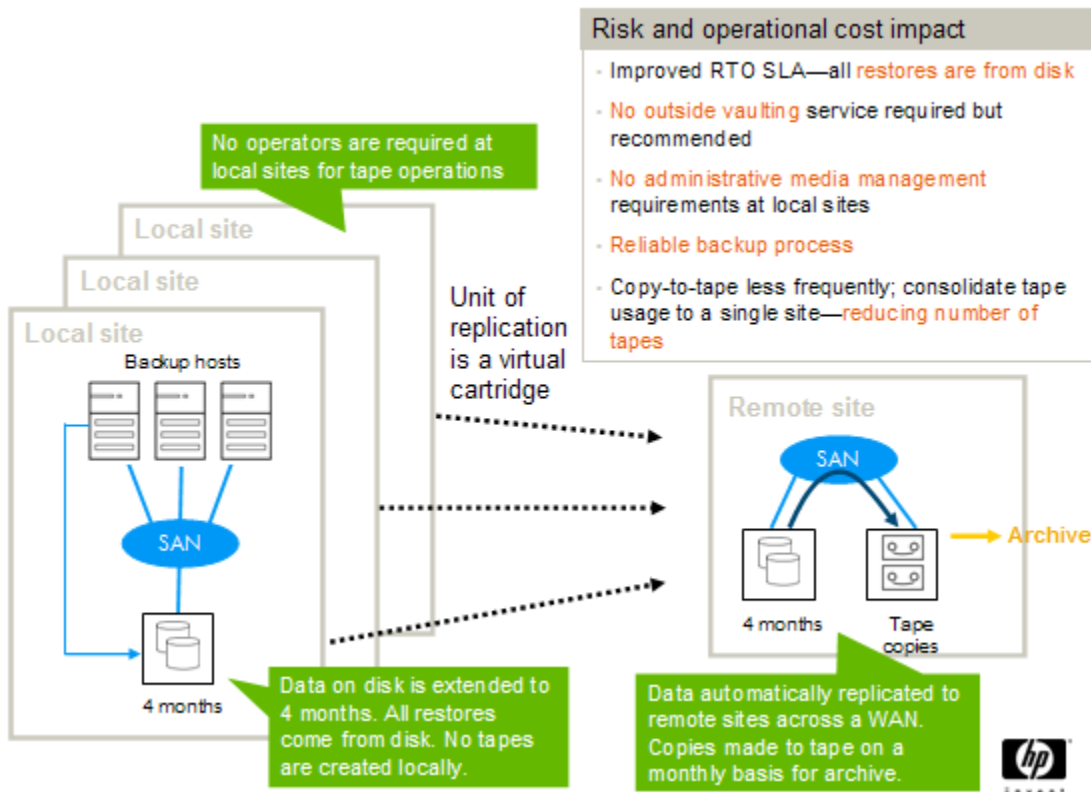


Figure 7 Remote Site Data Protection Using Replication

Deduplication is the key technology enabler for replication on HP VLS and D2D systems. (VLS systems use HP Accelerated deduplication, and D2D systems use Dynamic deduplication.) The same technology

that allows duplicate data to be detected and stored only once on the HP VLS or D2D system also allows only the unique data to replicate between sites. Because the volume of data being replicated between sites is much less than if the full data set was replicated, you can use lower bandwidth links at correspondingly lower price points. In addition, backup at remote offices can be automated to a local virtual tape library and then replicated back to a regional data center or primary data center allowing end-to-end management from the data center of all data in the remote offices.

This transformation is shown in [Table 5](#) which compares the amount of data to transfer both with and without deduplication. The amount of data to back up in this example is 1 TB.

**Table 5 Estimated Time to Replicate Data for a 1 TB Backup Environment at 2:1**

		Link Type		
	Data Sent	T1	T3	OC12
Link Rate (66% efficient)		1.5 Mb/s	44.7 Mb/s	622.1 Mb/s
Backup Type	Without deduplication			
Incremental	50 GB	4.5 days	3.8 hours	16 minutes
Full	500 GB	45.4 days	1.6 days	2.7 hours
Change Rate	With deduplication			
0.5%	13.1 GB	29 hours	59 minutes	4.3 minutes
1.0%	16.3 GB	35 hours	73 minutes	5.3 minutes
2.0%	22.5 GB	49 hours	102 minutes	7.3 minutes



**NOTE:**

T1/T3 and OC12 are old terms with respect to WAN link terminology. Many link providers use their own names (e.g., IP Clear, Etherflow). This document distinguishes them by their speed using 2 Mbits/sec, 50 Mbits/sec, etc.

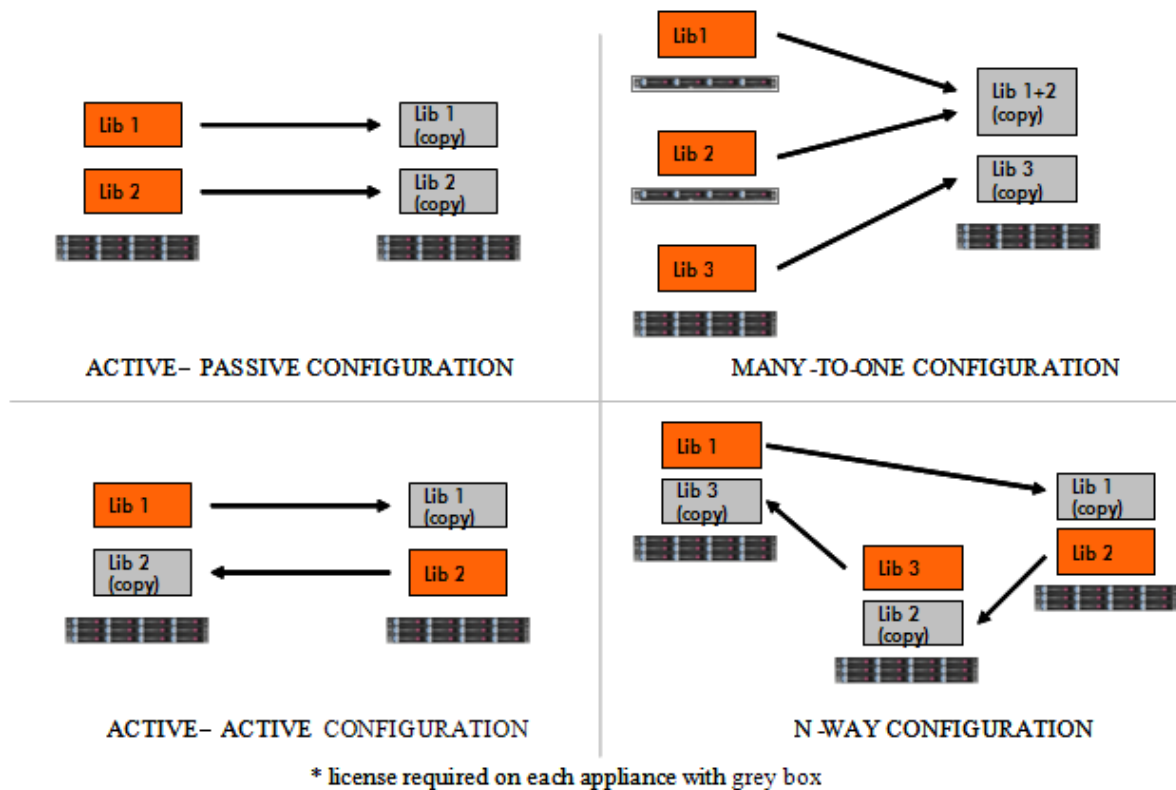
One consideration with replication is that you must “initialize” the Virtual Tape Libraries with data prior to starting the replication. This ensures that the source and target devices are both synchronized with the relevant reference data to allow them to interpret the changed data (deltas) that comes across the WAN link during replication.

## Replication Deployment Options

You can deploy the HP VLS and D2D systems for replication in many ways depending on your requirements. You should understand the terminology associated with deduplication and replication. The key terminology for replication deployment:

- **Source:** A series of slots/cartridges in a virtual library that act as the source data to replicate. This is the original copy of the backup data, written to and managed by the source site’s backup application.
- **Target or LAN/WAN destination:** A series of corresponding slots in another virtual library on another site in another location which receives data from the source library. This is the secondary (disaster recovery) copy of the backup data, managed by the replication system.

For both HP VLS and D2D systems, the unit of replication is a virtual cartridge and the replication link is TCP/IP (one GbE connection per node on the VLS system, and one to two GbE connections on the D2D system). Figure 8 shows how you can configure the system to replicate all of the cartridges or just a subset of cartridges from the source virtual library to the target virtual library.



**Figure 8 Replication Configuration Options**

- **Active-Passive:** The best deployment for a single site disaster recovery protection. The active source device receives the local backup data and then replicates it to a passive target device at the disaster recovery site dedicated to receiving replication data.
- **Many-to-one:** The best deployment for several smaller sites consolidating their replication to a central disaster recovery site. Each source device in the smaller sites replicate to a central target device which you configure to have multiple sources replicate to a common virtual library on the central device (each source replicates to its own subset of the cartridges in the virtual library). Alternatively, each source can have its own dedicated virtual library. Up to four remote VLS sites can copy to a single HP VLS at the central site at launch, and this will be increased over time. Up to 16 remote D2D sites can copy to a single D2D4000 at the central site, and up to 24 remote D2D sites can copy to a D2D4100.
- **Active-Active and N-Way:** The best deployment for sharing your VLS or D2D system hardware for both receiving backups and receiving replication data (so each device is both a source and a target as shown in the above diagram). Active-active is one way to implement cross-replication between sites, but you can use two active-passive deployments to achieve the same result. Choosing between either active-active or 2x active-passive deployments for cross-replication depends on which provides the lowest cost. Active-active is only recommended if the backup traffic on each device is only using up to half of the device's maximum performance and capacity, because you need additional performance and capacity for the replication target operations.

For example, if you have two VLS9000 sites that each requires 2-nodes/2-arrays for just their backup performance/capacity and 2-nodes/2-arrays for their replication target performance/capacity, then you have the following choices for cross-replication deployment:

- Active-Active: Each site requires a 4-node/4-array VLS9000 (with deduplication) shared between backups and replication target, one rack, and four replication LTUs.
- 2x Active-Passive: Each site requires a 2-node/2-array VLS9000 (with deduplication) for backup and a separate 2-node/2-array VLS9000 (with deduplication) for replication target, two racks, and two replication LTUs.

In this example, it costs less to use active-active because it adds two replication LTUs but saves the hardware/power/footprint cost of a second rack and the cost of a second VLS connectivity kit. However, if your backups required more than half of the maximum device performance (for example, more than two nodes out of a maximum configuration of four nodes), you may have to deploy two devices per site. In this case, it would be cheaper licensing (and better future device scalability) to use 2x active-passive deployment.

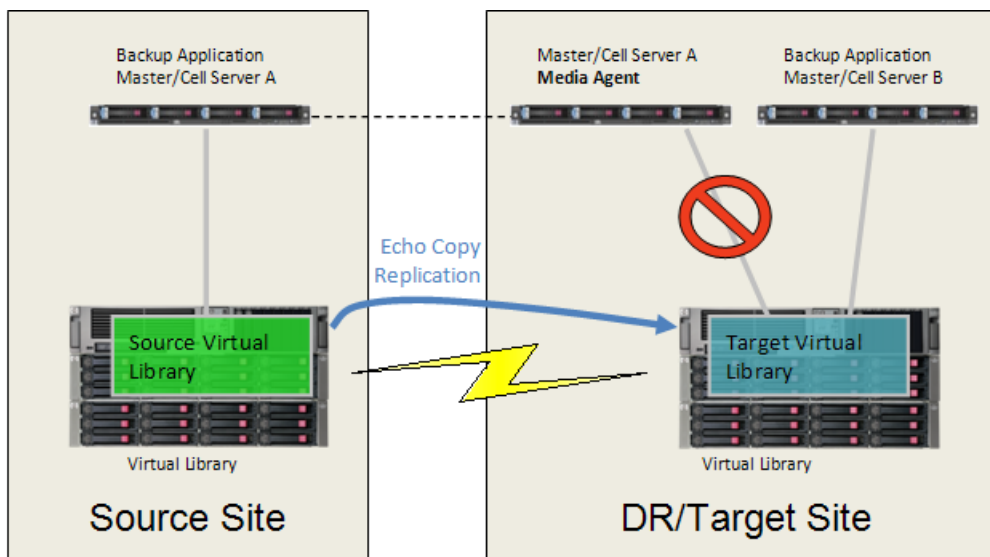
 **NOTE:**

Multi-hop replication (replicating a cartridge from device A to device B, and then replicating the replicated cartridge from device B to device C) is not yet supported.

## Backup Application Interaction with Replication

The replication in both the VLS and D2D systems is mirroring the source cartridge to its matching target cartridge so both cartridges have the same barcode, the same tape contents, etc. Backup applications currently cannot handle seeing two copies of the same cartridge at the same time (because to the backup application, the cartridge is a single entity in the media database). Given this limitation, you must hide the target virtual library from the source device's backup application:

- For VLS systems, the replication target is a subset or an entire virtual library that is presented on front-end Fibre Channel ports, so if the source backup application is running media agents on the target site you either need to use SAN zoning or the device's LUN mapping feature to hide this replication target virtual library from the source device's backup application.
- For D2D systems, this is currently automatic because the replication target is hidden from all external host access (until it is converted into a non-replicating library in the event of a disaster recovery).



**Figure 9 Presenting the Replication Target to a Different Backup Application**

On a VLS (by default) or a D2D (if you enable the read-only mode on the target library), you can still present the replication target to a different backup application instance (i.e., a separate backup application master/cell server on the target site with its own media database), which you can use to “import” replicated cartridges into its media database and then perform restores or copy to physical tape, etc. See “[Creating Archive Tapes from the Target](#)” on page 190 (VLS) or “[Creating Archive Tapes from the Target](#)” on page 99 (D2D) for an example on automating this.

---

 **NOTE:**

With HP Data Protector, if you have a cell server in each site that can share library devices across sites through a MoM/CMMDB, you still need to ensure that each cell server only sees its local virtual library (i.e., the source cell server must not be configured to see the target virtual library and vice-versa).

---

## Replication Limitations

VLS and D2D replication may not work in every environment. Understand the possible limitations:

- Do not confuse Virtual Tape Library replication with “high availability/continuous access” which is a type of full bandwidth replication used on Disk Array technology whereby primary application data can be accessed within hours of a disaster at one site from the other site. Virtual tape replication is *not* high availability; it is a means of automating the offsite of data resulting in better disaster recovery coverage.
- System data rate change. The higher the data change rate, the more data requires replicating. Systems with very high change rates and slower links may not be able to replicate all the data off-site within 24 hours.
- High latency links. For very large distance replications with many routers involved, the latency of the WAN link may lead to high inefficiency of the link where throughput is limited and replications cannot be performed in a timely manner.
- Current link speed is too slow or the implementation of replication on the existing link will cause unacceptable delays in application response times. Using the HP StorageWorks sizer tool and some of your inputs, you can evaluate if you will need to increase an existing link speed to be able to benefit from replication. See <http://www.hp.com/go/storageworks/sizer>.
- Some additional financial investment will be required as increased bandwidth links, hardware additions, and/or deduplication and replication licenses, but in general the increased robustness of the data protection process should pay for itself within 2–3 years.
- On the VLS, the HP Accelerated deduplication relies on understanding the metadata format of the incoming data stream. It does not currently support all data formats and backup API’s. In the case where an HP VLS cannot deduplicate the data type, the data is sent untouched to the VLS. This data is replicated as “whole cartridge;” the entire tape contents are replicated and not the delta’s or unique data objects. If a high percentage of your data cannot deduplicate, the volume of data to replicate will be very large. If you do not have very large volumes of data to replicate, you should consider using HP whole cartridge replication. This works essentially in the same way as replication using echo copy pools; it requires no tape transfer or initialization and no deduplication or replication licenses. However, all data is transferred between sites and this means the WAN links will have to considerably higher performance at an associated higher cost.



---

# 3 Backup Solution Design Considerations

This section uses use models to explore many of the concepts you must consider when designing your system. Use models are organizational schemes that provide a basic organizational framework that you can use to delineate your environment and visualize how to implement the VLS or D2D for your best results. You might also think of the use models as decision trees in which each branch junction is a decision point that comes with value judgments and trade-offs. The decisions you make lead to an implementation plan.


The choices you make are economic as well as functional. Different implementations can have different direct costs such as investment in licensing or equipment purchase/usage, and indirect costs such as investment in maintenance or configuration. You need to balance time, money, and use policies against one another. As with the major categories (tape, VT/D2D, clone), there is no one size fits all solution. The optimum solution will be a balance of the models.

## Analyze the Existing Environment

Carefully consider your environment in general, and your topology. Think about what you have, how you use it, and most importantly, what you want to optimize or improve. Determine your goals for your environment.

## Consider How you Want to Back Up your Data to the VLS or D2D

- Identify the data that you want to backup to the VLS or D2D. For example:
  - Data that needs a better service level than going directly to tape but does not need the highest service level (split mirror, for example).
  - Slow servers in your current backup topology. Supported topologies include SAN or LAN/SAN hybrid:



12001

**Figure 10 Backup to the VLS in a LAN/SAN Hybrid Environment**

This includes servers that contain lots of small files that are likely to slow down backup performance on the host (i.e., those found on Windows file servers and web servers), blade servers, etc. Slower servers on the LAN typically interleave the backup through the media server and go out to one tape drive. These are good candidates for VLS and D2D because you can configure multiple virtual tape drives, disable multiplexing, and then backup each slow server in parallel (with multiplexing disabled the restores will run faster). See [Multiplexing, Multistreaming, and Multipathing](#). You can check the true performance of your servers in a variety of ways, including using the tools found at in the Software Downloads section of the Developer & Solution Partner Program web page: <http://h21007.www2.hp.com/dev/>.

- Data that you want to electronically off-site using deduplication—enabled replication (and thus must be backed up to a local VLS or D2D which can then replicate the data automatically to another VLS or D2D in another site).
- Archive data.
- Look for aggregate bandwidth bottlenecks between hosts and backup devices. One common bottleneck is the LAN bandwidth (for LAN backups), so identify which of the larger application



servers could perform LAN-free backup (see [LAN-free Backups](#)) and thus move their backup traffic from the LAN to the SAN.

- List your backup applications. In a heterogeneous environment, you can have multiple applications writing to different virtual libraries on the same VLS or D2D (which is highly flexible with regard to the libraries that can be configured). This is similar to library partitioning but more flexible and easier to setup and maintain. To reduce backup administration, always use the fewest number of virtual libraries. For example, use one virtual library for each backup application. (See [Single Library vs. Multiple Libraries.](#))
- How long do you want to retain the data on the VLS or D2D. (See [Retention Planning.](#))
- How fast is your data growing. (See [Future Data Growth.](#)) This is critical in deciding which disk backup technology best fits because you should choose a technology that will continue to scale in performance and capacity so that it can provide a single backup target even after several years of data growth.

## Consider How you Want to Copy the Backup Data to an Off-site Location

Think about copying the backup data to an off-site location for site disaster protection or long-term archival. See [Considerations for Copies](#) for more details.

## Consider Speed and Ease of the Restore as well as the Backup


If you are still within your backup retention period, you can restore the data directly from the disk backup device. If you are outside of your retention period, you can restore from your long-term retention tapes if you have created them. If your local disk backup device has been destroyed (i.e., site disaster), you can recover from your off-site copies (could be tape or could be another replicated disk backup device). See [Considerations for Restores.](#)

## Single Library vs. Multiple Libraries

The VLS and D2D configurations are highly flexible with regard to the number and types of libraries. There are two basic backup use models (single library and multiple library). The choice of which primarily depends on the architecture of the backup applications. To reduce backup administration, use the fewest number of virtual libraries. For example, use one virtual library per backup application.

### Single Library

The single-library use model is a many-to-one configuration – all hosts see one communal library. This is the default configuration.



12000

**Figure 11 Backup to VLS in a Simple Deployment (VLS9000-series with One Shared Library Shown)**

### Benefits of Single Library Systems

- This use model is easy to manage.
- It is easy to copy through the backup application because you already have shared devices and they all see and are seen by a copy engine or server.
- Some configuration tasks are easier; you don't, for example, have to worry about assigning specific jobs to specific libraries.
- You don't have to worry about performance tuning. If, for example, you have enough bandwidth for ten backup streams, you configure ten drives and let the device load balance automatically. The backup application decides how to allocate; you don't have to specify.
- Multiple hosts are mixed together in one media set. When you copy to tape, you're using less media.
- With the VLS9000 or VLS Gateway, you can consolidate your storage system to use only one large VLS (greater than 500 usable TB) instead of multiple smaller appliances.

If you use Data Protector, licensing is relatively economical — media servers do not require individual licenses, object copy doesn't require additional licensing, and virtual tape is licensed per TB. Setup is largely automated and does not require manual intervention.

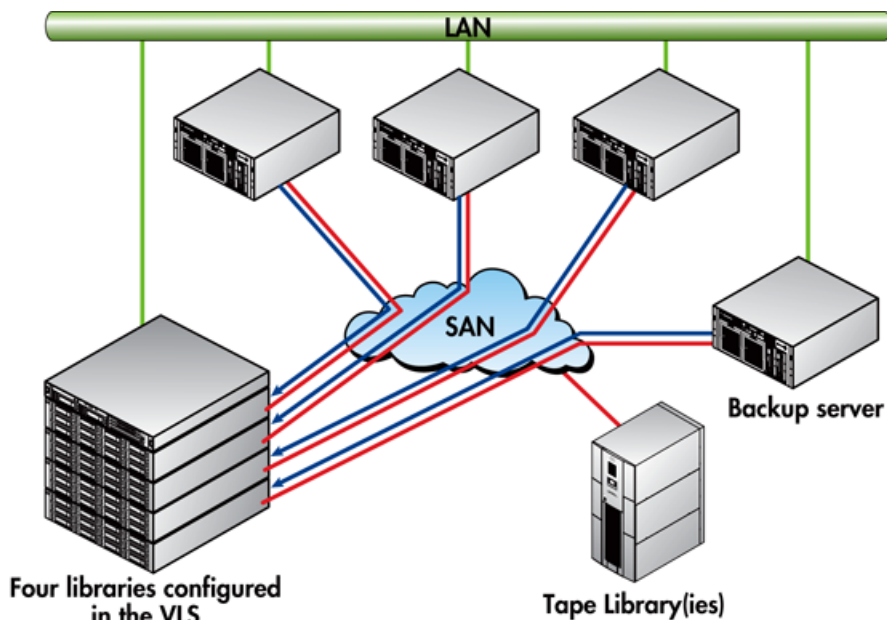
### Considerations for Single Library Systems

- Not all backup applications are well-suited to this setup. This depends (at least in part) on the application's licensing schemes. If you have to purchase shared drive licenses, this model may be more expensive than a dedicated setup.
- Setup can be challenging in some applications. You may have to configure paths to the tape drives from all the backup hosts. The backup application may do some of this automatically, but this, in itself, can be time consuming. And if you change the configuration, you may have to repeat some management steps and do some re-configuration.

- Multiple hosts are mixed together in one media set. Although you may be using less physical media, this may not meet your media tracking and pooling needs.

## Multiple Library

In this model, multiple hosts map to discrete libraries in a one-to-one or some-to-one configuration.



11201

**Figure 12 Backup to VLS in a Simple Deployment (VLS9000-series with Four Dedicated Libraries Shown)**

## Benefits of Multiple Library Systems

The assignment of drives and media in this configuration is more individualized. This model is helpful when you need to control the organization and grouping of media. Particular backup jobs go to particular media (or pools) and no other hosts are writing to that media. Note that backup jobs can be assigned to particular media in the shared model, also, but it's easier to group physical media in the dedicated model because you can group media by host.

You might choose a restricted use model if:

- Your backup application won't handle SAN sharing or if your backup application's licensing is prohibitively expensive in a shared environment. You may be able to spend less on licenses because you can pick the hosts that you want to be visible to particular LUNs
- You have multiple backup applications. Configure one virtual library per backup application
- You have multiple SAN fabrics. Configure one virtual library per SAN
- Your backup application does not handle SAN sharing
- You have licensing concerns with your backup application. You can adjust library slots and drives to optimize license fees

If you choose to copy to tape through the backup application, there are two basic options:

- Shared copy library – There are two scenarios: 1) You have a dedicated copy host that sees the copy library and all of the dedicated virtual libraries. There is SAN sharing and you may have to

do LUN mapping and masking 2) All hosts write to the library as a shared device. This involves copies going on during prime backup or operational time on the hosts.

- Dedicated copy library for each virtual library -- This may be expensive with multiple virtual libraries.

## Multiplexing, Multistreaming, and Multipathing

This section explains the concepts of [Multiplexing](#), [Multistreaming](#) and [Multipathing](#), and notes which of these technologies are useful for disk based backup devices and which are not recommended.

### Multiplexing

This technology (sometimes called interleaving) is where multiple backup streams (from multiple servers) are concurrently mixed together into a single tape drive (and thus a single tape cartridge). This is commonly used for physical tape drives such as LTO where the speed of a single backup stream is too slow to maximize the tape drive performance and you must run multiple concurrent streams together into the drive to achieve full drive performance. However, this also means that restore performance is severely affected because to restore a single backup stream (i.e., restore one server) the backup application still has to read *all* of the intermixed data on the tape cartridge and then throw away the data from the other backups. If you have five servers' backups multiplexed together on one tape cartridge, you only get 20% restore performance when you restore one server.

Multiplexing is not recommended. When you have a disk backup device, you should disable multiplexing on all backups going to the device because the main advantage of a disk backup device is the ability to run many more concurrent backup streams than an equivalent physical library with limited numbers of tape drives. For example, on the VLS you can create 32 virtual tape drives on every Fibre Channel port and allow 32 non-multiplexed backup streams to be written concurrently to each Fibre Channel port. This improves backup performance by allowing more backups to run in parallel and improves restore performance because the backups are non-multiplexed. Disabling multiplexing also improves deduplication efficiency and performance.

### Multistreaming

This technology is where multiple objects in a backup job can be written concurrently to multiple tape drives (each object going to a different drive). For example, if you have a file server with C:, D:, and E: volumes, with multistreaming this could be concurrently written to three separate tape drives (C: going to drive1, D: going to drive2, E: going to drive3). This increases the backup performance of that server. Another example is a database backup where you may have multiple tablespaces or files within the same database and these objects can be written in parallel to the tape drives, or some database backup agents such as Oracle RMAN allow multiple streams even from one object (e.g., one database tablespace).

Multistreaming is recommended for disk based backup, particularly because many virtual tape drives can be created thus allowing more concurrent streams to run from multistream backup jobs.

### Multipathing

This technology uses two Fibre Channel paths from the backup server to the virtual tape library, allowing for higher availability; for example, if you have dual SAN fabrics and wish to continue operations even if one entire fabric fails. The VLS and D2D support the option of presenting any virtual library device (changer or tape drive) on two Fibre Channel ports.

The main use of this is to present the virtual library changer device over two Fibre Channel ports because the backup application must always be able to access the changer to perform any backup/restore operations. Many enterprise backup applications (such as HP Data Protector, Symantec

NetBackup, etc.) support dual paths to the library change device and will automatically switch over to the alternate path if the primary path fails.

Multipathing virtual tape drives is not recommended. Many enterprise backup environments (operating systems or backup application) do not support dual path to tape drives because they see this as two separate drives. Given the fact that a virtual library can create many virtual tape drives over multiple different Fibre Channel ports, if a SAN fabric fails (and you have a dual-path changer configured) the only impact of this is that half the virtual tape drives cannot be accessed but the remaining half of the drives in the library can still be used and can still access every virtual cartridge in the library. This also implies that when a SAN fabric fails this halves the available Fibre Channel performance so if full backup performance is required even after a SAN fabric failure then the VLS needs to be configured with double the number of Fibre Channel ports. Instead of configuring dual path virtual tape drives, it is better to configure single-path virtual tapes across multiple Fibre Channel ports that are connected to both SAN fabrics.

---

 **NOTE:**

The D2D can configure its two LAN ports in high availability port mode (see [D2D Ethernet Ports](#)), but this is transparent to the backup application because it only sees one "bonded" path to the virtual library.

---

## Blocksize and Transfer Size

As with physical tape, larger tape block sizes and host transfer sizes are of benefit; they reduce the amount of overhead of headers added by the backup application and the transport interface.

- D2D: HP recommends a minimum of 64 KB blocksize and suggests up to 1 MB.
- VLS: HP recommends a blocksize of 256 KB (the maximum supported).

## LAN-free Backups

All enterprise backup applications provide the ability to run "LAN-free" backups. This is where application servers connected to the SAN run the backup media agent and can perform backups directly over the SAN to the tape library. This removes this backup traffic from the LAN and speeds up backups (by removing the LAN performance bottleneck). One of the disadvantages of LAN-free backup to a physical library is that it cannot be multiplexed because each backup stream is written directly from the data source to the tape drive, and so it could generally only be used with the largest application servers which can supply high speed backup streams.

One of the key advantages of a SAN virtual library is the ability to increase the amount of LAN-free backup used. With a virtual library you do not have to provide maximum performance per LAN-free backup stream as needed on physical tape drives. You can create many virtual drives and thus enable LAN-free backup on any SAN-connected appliance servers (regardless of their performance) that are large enough to bother with LAN-free backup. This improves the backup performance of the LAN-free backups and also improves performance of the remaining LAN backups because the LAN-free backup traffic has been removed from the LAN.

## Retention Planning

Retention planning and sizing go hand in hand. How long do you need to keep data on disk? How many full backups do you want to keep on disk? How many incremental backups? How do you want

to optimize retention times of the VLS or D2D? Retention policies help you recycle virtual media. Bear the following considerations in mind as you plan retention policies:

- If you are not using deduplication, you can retain the data on the disk backup device for a shorter period such as 1-2 weeks (because more than 90% of restores generally occur within the first week's retention of backups) and then use tape copies to retain data for longer periods.
- If you are using deduplication, you can retain data on your disk backup device with the same level of retention times as you would have had on tape. This provides a more granular set of recovery points with a greater likelihood that a file that you need to recover will be available for longer and in many more versions. To use deduplication-enabled replication you must have at least two full backups retained on the disk backup device.
- Once the retention period expires, the virtual media is automatically recycled.
- You should set the tape expiration dates (that is, when the tape is marked as worn out) high because virtual media does not wear out.
- Backup-job retention time is for virtual media.
- Copy-job retention time is for physical media.
- When copying through the backup application, the virtual and physical pieces of media are tracked separately and the retention times should be considered and set individually.

## Future Data Growth

How fast is your data growing? This is critical in deciding which disk backup technology best fits; choose a technology that will continue to scale in performance and capacity so that it still can provide a single backup target even after several years of data growth.

If you choose a disk backup technology that starts close to its maximum performance/capacity, then as data grows you must add another device (and thus another backup target) and so on, which means significant increases in backup administration over time due to having to manually balance all backup jobs across multiple backup targets. This manual balancing is even more difficult when you have multiple targets that are each separate deduplication domains, because then when you switch a backup job from one target to another it “resets” the deduplication for that backup making capacity planning very complex.

## Considerations for Copies

Generally the disk backup device is located at the same site as the data being backed up, which means you must create copies of the backup data that are stored in a different remote site to protect against a site disaster. There are also sometimes long-term (e.g., multi-year) retention requirements for some backup data that may require copying the backups to physical tape. (There may be legal requirements to use tape, plus a tape on a shelf uses zero power compared to storing the long-term retention backup data on spinning disks). When designing your solution, consider the following options available for copying backups. You can implement all of these without affecting the application servers and with the minimum of impact on data center backup processes.

- Copy your backups from your virtual library to physical tape and then ship these physical tapes to off-site storage. There are different methods of creating the copy on physical tape:
  - Use the backup application to copy data from the virtual library to the physical library (preferred method).
  - Use the automigration functionality within the VLS (which turns the VLS into a disk-cache of the physical library) or the tape offload functionality within the D2D.
- If you are using deduplication, another option is to use the deduplication—enabled replication technology which cost-effectively copies the backups to another remote device (e.g., VLS replicates

to another VLS, D2D replicates to another D2D) and removes the need to move tapes between sites. See [VLS Replication](#).


## Copy to Physical Tape through the Backup Application

You can migrate data from virtual media in the VLS or D2D device to physical tape media using the various tape copy/clone mechanisms that exist in enterprise backup applications. Using the data protection software to copy from virtual to physical media has the following advantages compared to any transparent migration to tape:

- The virtual and physical media are separately tracked in the backup application's database, so you can have different retention times for virtual and physical media and the two pieces of media can be in different locations. This provides flexibility when restoring data; you can restore from tape or (if you have configured your system appropriately) you can restore from the virtual library because it can retain virtual media online after the physical media is sent off-site.
- The tape copy process is an integrated part of the backup application.
- Improved media management and utilization. You can, for example, specify which backups are copied to physical tape, so fewer tapes are used (the difference can be as much as half).
- The tape library can be used for other tape backup processes. The tape library is a common resource on the SAN because it is not hidden behind the VLS or D2D.

After installing the VLS or D2D device and configuring/redirecting the backups to the device, you will create tape copy/clone jobs in the backup applications to perform the migration to physical tape. Most backup applications include several copy/clone options, ranging from creating the tape copy at the same time as the backup (mirroring), scheduling a copy of all backup media to run when all the backups are expected to finish (scheduled copy), or creating the tape copy immediately after each backup completes (triggered copy). Scheduled copy is the recommended option so the copy jobs run after the backup window; otherwise, having copies and backups running at the same time results in significant performance loss.

When copying virtual cartridges to physical cartridges through the backup application, the data is read from the virtual library into one of the available media servers which then writes the data back out again to the physical library. The physical tape library is independent of the virtual tape library and is entirely under the control of the backup application. This means that the virtual library type and virtual drive type and virtual cartridge size do not need to match the physical library in any way; the backup applications all perform "object copy" where they effectively backup the backup and therefore append and span copy jobs onto the physical tapes just like they would with backup jobs.



12003

**Figure 13 Writing to Tape in a LAN/SAN Hybrid Environment**

## Media Server Considerations

To support the background copy of backup data from the VLS or D2D to a tape library, use one or more of your existing backup media servers (running the backup application media agents) to copy the data directly from the device media onto the tape media with data passing over the SAN from the VLS or D2D to the device server and then back again to the tape library. If you are running the copies after the backup window is complete, all of the media servers that were performing LAN backups are now idle and can be used for copies. Do not use media servers that are performing LAN-free backup because these will be active application servers and could affect application performance if used for copy operations. Ensure that the media servers used for tape copies are plugged into the same SAN switch as the virtual library and physical library so the copy traffic does not use any SAN bandwidth. In this configuration, the tape copy does not impact the running application servers, so it can be run in the background after the backups are completed.

## Benefits of Copying to Physical Tape through the Backup Application

Using a dedicated media server to migrate data stored on a VLS or D2D to a physical library allows you to:



- Use the functionality of a backup application.
- Copy only the specific files you need onto physical tape.
- Not waste tape storage as physical tapes can be fully filled.
- Monitor and track copy jobs from the backup application.
- Use any tape library that is supported by the backup application.

## Considerations for Copying to Physical Tape through the Backup Application

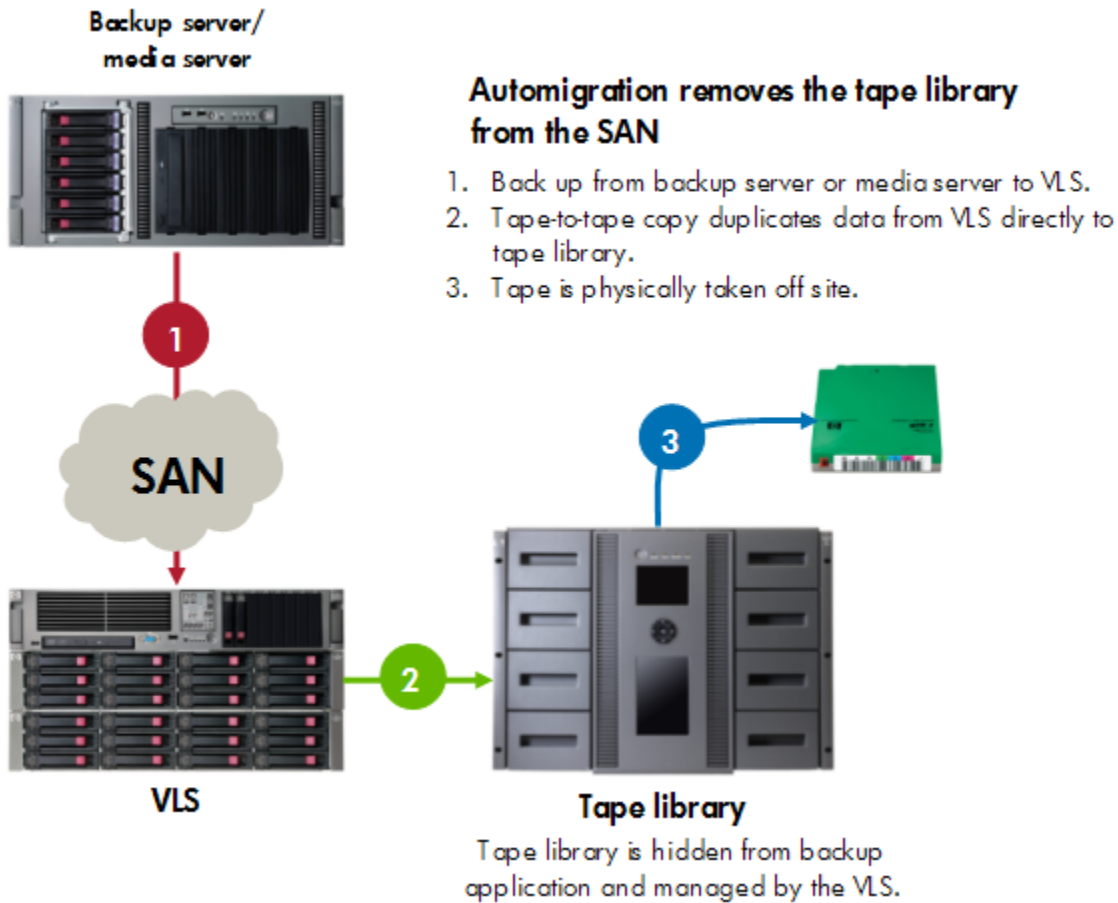
There are potential drawbacks to using a backup application instead of caching to copy data on the VLS or D2D to physical tape. Consider the following when deciding which data migration method to use:

- You must purchase additional backup application licenses for the physical library.
- You must configure a backup application with a scheduled copy job to perform the migration.
- If you have the media servers on different SAN switches from the libraries, copy bandwidth is added to the SAN during data migration.

## Copy to Tape using VLS Automigration

You can migrate data from virtual media in the VLS device to physical tape media using automigration/tape-offload to perform transparent tape migration. The automigration echo copy feature allows the VLS to act as a tape copy engine that transfers data from virtual cartridges on disk to a physical tape library connected to the VLS device. See [VLS Automigration](#).

Echo copy essentially acts as a transparent disk cache to the physical library so echo copy jobs can be performed at times other than the peak backup window. Once automigration is set up, echo copy operations are automatic and seamless to the user. Echo copy is managed through the automigration software on the VLS instead of through the backup application.



**Figure 14 Echo Copy is Managed through Automigration**

### Benefits of Echo Copy

- The destination library is not visible to the backup application, so it does not need licensing.
- There is no need to license/configure copy jobs in the backup application.
- Copy traffic is no longer routed through the SAN.
- Maximizes performance of high-speed tape drives by freeing bandwidth during the normal backup window.
- All destination tapes created by echo copy are in the native tape format, so they can be restored directly from any tape drive or library that is visible to the backup application.

### Considerations of Echo Copy

- The destination library can only be used for copy operations.
- The copy is a full tape copy, rather than an incremental change copy, so it can be an inefficient use of media if you are using non-appending media pools in your backup jobs.
- The backup application will not be aware of any copy failures.
- Echo copy can only use the tape libraries supported by the VLS firmware (for example, HP MSL, EML, ESL-E libraries).
- Because echo copy automatically links the virtual and physical tapes (to maintain media management with enterprise backup applications), human error is a factor to keep in mind. With linked media management, any mistakes with the destination library's media will also affect the virtual cartridges. For example, if new tapes are not loaded into the destination library the new, matching

virtual tapes will not be created. Subsequent backups will fail because the virtual tapes present are protected.

---

 **NOTE:**

Automigration echo copy is not suitable for use with deduplication because:

- You cannot use echo copy to create archive tapes from the replication target device because these must have a different barcode, retention time, cartridge size, and contents. They must be created by another instance of the backup application. (See [Creating Archive Tapes from the Replication Target](#).)
  - Echo copy acts as a disk-cache to the physical tapes, so when physical tapes are ejected from the library (for example, to be sent off-site) this also ejects the matching virtual cartridges into the Firesafe; they are disabled in the deduplication system.
  - Echo copy only copies whole cartridges and because the size of deduplicated virtual cartridges is generally 50-200 GB, a large amount of tape capacity would be wasted if LTO3 or LTO4 physical tapes were used.
- 

## Copy to Tape using D2D Tape Offload

You can offload cartridges to physical tape using a tape drive or library that is physically attached to the D2D system via SAS or pSCSI and is not visible to the backup application. This solution offers the following benefits:

- Tape offload can be conducted during normal working hours without affecting network performance because no data is sent over the network during the offload.
- Cartridges can be offloaded at the maximum read performance for a virtual library which makes this process relatively fast.
- Offloaded cartridges are in backup application format, so you can use it to directly restore data using a tape drive attached to a media server if necessary.

However, you must take into account a number of shortfalls of this configuration:

- The backup application cannot track the physical cartridges because it played no part in their creation. A cartridge copied by the D2D is only valid for as long as the virtual cartridge remains current; if the virtual cartridge is modified (overwritten or appended), the physical cartridge content no longer has a valid entry in the backup application database. However, you could still use it for disaster recovery if you lost the backup application database. Creating long rotation schemes where cartridges are overwritten infrequently works well with this model.
- Even with the scheduling features provided in the D2D GUI, offloads are hard to accurately schedule to coincide with completion of the original backup.

Do not use physical tape offload as a way to free up space on the D2D by removing the virtual cartridge after an offload. Doing this will not save space due to the deduplication effect and will result in backup application database entries becoming invalid. In addition, HP does not recommend using physical media to extend the retention period of a backup (i.e., keeping a physical cartridge beyond the point where its associated virtual cartridge has been overwritten) because the backup application database will become inconsistent with the backup on tape.

## Copy to Remote Disk Backup Device using Replication

Deduplication can automate the off-site process and enable disaster recovery by providing site to site deduplication-enabled replication at a lower cost. Because deduplication knows what data has changed at a block or byte level, replication becomes more intelligent and transfers only the changed

data instead of the complete data set. This saves time and replication bandwidth, and is one of the most attractive features that deduplication offers. Replication enables better disaster tolerance without the operational costs associated with transporting data off-site on physical tape. See [Introduction to Replication](#).

The replication system on both VLS and D2D creates and maintains a cartridge mirror between the source and target devices. Once replication is set up, operations are automatic and seamless to the user. Replication is managed through the firmware on the VLS or D2D, not the backup application.

## Benefits of Replication

- Improves the overall reliability in offsite data.
- Completely automates offsite data.
- Improves the resilience of remote offices and regional data centers while maintaining local data recovery capabilities.
- Reduces the overall costs of offsite data compared to physical tape offsite when you consider all manual costs.
- The replication target library is not visible to the backup application and does not need licensing.
- There is no need to license/configure copy jobs in the backup application.

## Considerations of Replication

- The backup application will not be aware of any copy failures.
- You cannot have different retention times between the source and target media because cartridges are mirrored. However, with deduplication reducing the amount of disk space required to store tape retention times, you do not generally need different retention times between source and target.
- You cannot present both the source and target media to the same backup application at the same time. (See [Backup Application Interaction with Replication](#).)

## Creating Archive Tapes from the Replication Target

Using replication you cannot present both the source and target media to the same backup application at the same time. (See [Backup Application Interaction with Replication](#).) You cannot use automigration (which mirrors whole cartridges) to create the archive tapes because they will have a different barcode, retention time, contents, and potentially a different size from the replication target media. To create archive tapes from the replication target media, you must use a second instance of the backup application to perform this. (See [Creating Archive Tapes from the Target](#).)

The backup application must have the ability to “import” tapes created on one backup application domain into another domain (e.g., importing a cartridge from a Data Protector cell server to another cell server). This can be done manually, or it can also be done automatically via a script. On the VLS the script can be driven by an “ISV Email” email report. See [Creating Archive Tapes from the Target](#) for details and example scripts.

## Considerations for Restores

There are two restore use models: from tape and from the VLS or D2D. In either case, copies are in native tape format and so can be restored directly. Retention policies drive copy and restore use models.

## Restoring from Disk Backup Device

Restore from the VLS or D2D is performed in the same way that restore from tape is done. Simply direct your backup application to do the restore.

---

### NOTE:

A restore from one tape can be performed at the same time as a backup to a different tape.

---

## Restoring from Backup Application-created Tape Copy

The format of physical tapes used in the VLS and D2D environments is the same as the format in environments not using a virtual tape library. To restore from tape, place the tape in a library or drive that the host has access to and restore.

- Time to retrieve media is a factor in recovery time. Depending on the location of your vault and the media management within the vaulting process, retrieval time can be a key part of recovery time.
- Recovery from tape can be slower than recovery from disk.

## Restoring from the Replication Target

Consider a use case after a site disaster where the local VLS or D2D device is unavailable and the restore must be performed from the replication target device instead. There are three main methods of performing disaster recovery restore from a replication target:

- Restore directly from target device.  
The replicated cartridges in the target device are all native format (so they can be restored directly by a backup application). After presenting the replication target to a replacement backup application and restoring its media database, you can restore your servers from the VLS or D2D device. See [Restore Directly from the VLS Target Device](#) and [Restore Directly from the D2D Target Device](#).
- Restore over LAN/WAN.  
This option is where some or all of the source device is rebuilt by restoring the replicated cartridges over the LAN/WAN back from the target to the source. In the VLS, this is a wholesale (non-deduplicated) restore so can only restore a subset of the source device such as the last backup set. (See [Restore the VLS over the LAN/WAN](#).) In the D2D, this is a deduplicated reverse replication so can rebuild the entire source device. (See [Reverse Replication on the D2D](#).)
- Reverse tape initialization.  
This option is where the replicated cartridges on the target device can be exported to physical tape which can then be imported back into a new source device to rebuild it. This is currently only supported on D2D devices. (See [Reverse Tape Initialization on the D2D](#).)

## Performance Bottleneck Identification

In many cases, backup and restore performance using the VLS or D2D is limited by external factors. For example, performance is affected by the speed at which data can be transferred to and from the source disk system (the system being backed up), or by the performance of the Ethernet or Fibre Channel SAN link from the source to the VLS or D2D. To locate bottlenecks in the system, HP provides some performance tools which are part of the Library and Tape Tools package available at <http://www.hp.com/support/tapetools>.

The HP Library and Tape Tools:

- Dev Perf – This provides a simple test which will write data directly from system memory to a cartridge in a library on the VLS or D2D system. If this tool is run on a server being used as the backup media server, it can provide the maximum data throughput rate for a single backup or restore process. This isolates any backup application or source disk system from the environment and helps identify whether the bottleneck is the VLS or D2D system or the data transport link (Ethernet or Fibre Channel SAN) to the VLS or D2D system. For example, if this test reports 30-40 MB/s on a D2D, there is not a problem with the D2D or transport link. This test should report 100-150 MB on a VLS. Before starting the test, HP recommends creating a new virtual library to avoid overwriting any data on the current libraries. Use the `frontpanel` function on the Library and Tape Tools to manually "move" a cartridge into the tape drive.
- Sys Perf – This tool provides two tests which are conducted on the source disk system to perform backup and restore performance tests. These tests either read (for backup) or write (for restore) from and to the system disks to calculate how fast data can be transferred from disk and therefore whether this is a bottleneck. These tests should be run on the media server which backs up to the VLS or D2D. In order to test how fast any client servers can transfer data, the same performance tests can be used by mounting a directory from any of the client servers to the media server then running the test from the media server against these mounted directories; this will show how quickly data can be transferred from the client server disk right through to the media server.

## Backup SAN Design Guidelines

The design of your SAN environment will affect the performance, efficiency, and reliability of your backup and recovery scheme. Good SAN design is conducive to good VLS and D2D performance. Inefficient SAN design can degrade the performance and efficiency of all members of the SAN.

## General SAN Design Considerations

See the reference materials at [www.hp.com/go/ebs](http://www.hp.com/go/ebs) for information about:

- SAN design and configuration
- Synchronizing equipment and firmware (heterogeneous SAN support)
- Working across operating systems (SAN in heterogeneous environments)

## SAN Zoning

Due to complexities in multi-hosting tape devices on SANs, it helps to make use of zoning tools to help keep the backup/restore environment simple and less susceptible to the effects of changing or problematic SANs. Zoning provides a way for servers, disk arrays, and tape controllers to only see what hosts and targets they need to see and use. See the *Enterprise Backup Solutions Design Guide* available at <http://www.hp.com/go/ebs> for details.

The benefits of zoning include:

- The potential to greatly reduce target and LUN shifting.
- Limiting unnecessary discoveries on the FC interfaces.
- Reducing stress on backup devices by polling agents.
- Reducing the time it takes to debug and resolve anomalies in the backup/restore environment.
- Reducing the potential for conflict with untested third-party products.

Zoning may not always be required for configurations that are already small or simple. Typically the bigger the SAN, the more zoning is needed. HP recommends the following for determining how and when to use zoning.

- Small fabric (16 ports or less) may not need zoning. If no zoning is used, make sure that the virtual tape external Fibre Channel connections resides in the lowest ports of the switch.
- Small to medium fabric (16 - 128 ports) can use host-centric zoning. Host-centric zoning is implemented by creating a specific zone for each server or host, and adding only those storage elements to be utilized by that host. Host-centric zoning prevents a server from detecting any other devices on the SAN or including other servers, and it simplifies the device discovery process.
- Disk and tape on the same pair of HBAs is supported along with the coexistence of array multipath software (no multipath to tape, but coexistence of the multipath software and tape devices).
- Large fabric (128 ports or more) can use host-centric zoning and split disk and tape targets. Splitting disk and tape targets from being in the same zone together will help to keep the tape controllers free from discovering disk controllers which it does not need to see. For optimal performance, where practical, dedicate HBAs for disk and tape.

---

 **NOTE:**

Overlapping zones are supported.

---

## Operating System Tape Configuration

Enterprise backup applications generally support media agents (which communicate with tape and virtual tape devices) on various operating systems to provide a heterogeneous enterprise backup environment. The VLS and D2D products support many of these operating system variations. (The *EBS Compatibility Matrix* available at <http://www.hp.com/go/ebs> details which operating systems with each backup application are supported by VLS and D2D). The *Enterprise Backup Solutions Design Guide* available at the same web location has detailed instructions and best practices for configuring the tape drivers for the following operating systems:

- HP-UX
- Microsoft® Windows
- Tru64 UNIX®
- Linux
- NetWare
- Sun Solaris
- IBM AIX

When adding a VLS or D2D virtual library to a server running a media agent, you must:

- Ensure that the minimum operating system patches are installed.
- Configure the Fibre Channel HBAs in the operating system which generally means configuring persistent binding so that the virtual library device LUNs are always presenting on the same SCSI path regardless of SAN configuration changes.
- Detect the virtual tape drives and create the tape drive device paths (generally the virtual tape robot is discovered and configured by the backup application itself rather than the operating system). Some operating systems limit the number of tape LUNs that can be discovered to eight LUNs on each device port (remember there are multiple Fibre Channel ports on the backup devices), but there are workarounds in some cases:
  - Windows Large LUN support.  
Enable when using more than 8 LUNs per port. Information is available at <http://support.microsoft.com/kb/310072/>.
  - HP-UX Large LUN support (on v11.31 or higher).

For an HP-UX 11.31 server using persistent DSFs, the maximum number of LUNs per bus on the server can be increased by entering `# scsimgr set_attr -a max_lunid=32`.

- Disable unnecessary polling/diagnostic applications that can interfere with backups or impact performance.
  - Windows: Disable Windows Removable Storage Manager (RSM).
    1. Disconnect the Windows node from the SAN (unplug all Fibre Channel cables).
    2. Delete all files and subfolders under the `.. \system32\NtmsData` folder (location of the `system32` folder varies with different Windows versions).
    3. Enable and start the Removable Storage service in the Microsoft Computer Management applet.
    4. Access the Removable Storage (`My Computer/Manage/Storage/Removable Storage`) in the Microsoft Computer Management applet.
    5. Verify that there are no tape or library devices listed.
    6. Stop and disable the Removable Storage service in the Microsoft Computer Management applet.
    7. Reconnect the Windows node to the SAN (plug all Fibre Channel cables back in).
    8. Reboot.
  - Windows using the HP LTO tape driver: Disable Windows Removable Storage Manager (RSM).
    1. Install the 1.0.4.0 or later driver.
    2. Complete one of these steps:
      - a. The driver package contains a `DisableAutoRun.reg` file that can be used to modify the system registry. Log into the system as a user with Administrative privileges and double-click the `DisableAutoRun.reg` file. The system will prompt you to confirm. Enter `Y` or `Yes` to modify the appropriate registry entry and disable polling.
      - b. If the driver package does not include the `DisableAutoRun.reg` file, manually edit the system registry using `RegEdit`. Log into the system as a user with Administrative privileges, run `RegEdit`, and navigate to the following registry key:  
`HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\hplto`  
Edit the `AutoRun` value found in this key. A value of 0 (zero) indicates that polling is disabled; a value of 1 indicates that polling is enabled.
    3. After completing steps 1 and 2, reboot the affected system.

---

❗ **IMPORTANT:**

Adding or removing tape drives from the system may cause an older driver `inf` file to be re-read, which in turn can re-enable RSM polling. If tape drives are added or removed, check the registry for proper configuration and, if necessary, repeat step 2.

---

- Windows: Disable Windows Test Unit Ready polling.  
Information is available at <http://support.microsoft.com/kb/842411>.
- HPUX: Ensure the EMS is not polling tape drives by either disabling it or by copying the archived `dm_stape.cfg` file to the `/var/stm/config/tools/monitor` folder, and the polling interval can be set to 0 to disable polling.



- HPUX: prevent utilities like `mt` from rewinding a tape that is in use by the backup application by turning on the HPUX kernel tunable parameter `st_san_safe` (which disables tape device special files that are rewind-on-close)


## LUN Masking and Mapping

D2D devices with the Fibre Channel option do not require LUN masking/mapping because every virtual device LUN (library or drive) has its own unique Fibre Channel WWPN (using what is called N-port virtualization). This means you use SAN zoning to present each virtual device's WWPN to the backup hosts (but you must use a SAN switch that supports this and it makes sharing virtual devices across multiple media servers more difficult).

On VLS devices, by default all hosts on the SAN can access all the virtual devices (libraries and tape drives) of the VLS and this mode is used to easily share virtual drives across multiple media servers. (Enterprise backup applications support the ability to share tape drives across multiple servers).

- The backup applications take care of drive locking to ensure only one media server is using one shared drive at a time.
- VLS software manages the LUNs assigned to the VLS shared virtual devices by default. The VLS changes device LUN assignments as needed to make sure that each Fibre Channel host port has a LUN0, that there are no duplicate LUNs, and that there are no gaps in the LUN numbering on a Fibre Channel host port (these are operating system requirements).
- This is the easiest scheme for the backup administrator because they can simply configure all their backup jobs across their media servers to use the shared library with shared drives and the backup application takes care of deciding which drive is used for which backup job and automatically scheduling all the jobs onto the available drives.
- One disadvantage of this scheme is that if any one media server goes wrong (for example, mixes up its tape drives paths), then it can interfere with the backups running from the other media servers to the shared drives and can be very difficult to diagnose when there are many media servers.

The VLS has an option to enable LUN masking in the device which is used to assign dedicated virtual libraries/drives on a VLS. LUN masking lets you restrict which hosts you want to be visible to particular virtual devices so that if a media server has any problem it only affects the tape drives dedicated to that media server; it will not affect any backups on other media servers. Because LUN masking prevents a host from seeing selected virtual devices on the Fibre Channel host ports, a LUN-masked host will not see the virtual devices to which it does not have access and to which it cannot write. This requires the administrator to manually allocate specific virtual drives in the virtual library to specific media servers to meet their performance requirements.



**Figure 15 Virtual Tape Environment**

In the VLS firmware version 1.x/2.x, the LUN mapping mode is specific to each host. It lets you manually assign new LUN numbers to the virtual devices visible to the host, and hosts that are not LUN masked continue to use the default (shared device) LUN numbering. In the firmware version 3.x, the LUN mapping mode is a global setting across the entire device, and when LUN mapping is enabled all hosts will only see the virtual devices presented to those hosts (so by default any new host will see no virtual devices). Also in v3.x, the LUN numbering is automatically generated for any virtual devices presented to a host.

## Backup Application Basic Guidelines

The following are basic configuration guidelines specific to integrating an HP virtual library with enterprise backup applications:

- Improve performance by using larger tape blocks such as a 256 KB block size
- Improve performance by disabling multiplexing/interleaving, because multiplexing will dramatically reduce performance for restores and will also reduce deduplication performance. Instead of multiplexing, create more virtual tape drives and use multistreaming (so that backup jobs run multiple streams concurrently to multiple tape drives)
- Disable client backup compression because this will defeat target-based deduplication/compression in the backup device

For more information on backup application configuration optimizations needed for VLS deduplication, see [Detailed Backup Application Guidelines for VLS](#). Additional guidelines are available in the HP StorageWorks backup application implementation guides at <http://www.hp.com/go/ebs> detailing

- HP Data Protector
- IBM Tivoli Storage Manager
- EMC NetWorker
- Symantec Backup Exec

- Symantec NetBackup

## HP Data Protector Application Overview

The Data Protector cell is a network environment that includes a Cell Manager and client systems that run agents. Client systems are imported into a cell and belong to a single Cell Manager. Multiple cells may exist, each with their own Cell Manager. This environment may be managed by a single Manager of Managers or “MoM.”

- **Cell Manager:** the main system in the cell. The Cell Manager is the “traffic cop” that controls the activities and contains the Internal Database (IDB) within the Data Protector cell. It is not necessary to administer the backup and restore activities directly from the Cell Manager itself, because any client within the cell (as supported) can connect to the Cell Manager over the network and be used to administrate the activities of the cell.
- **Disk Agent:** install the Disk Agent on client systems you want to back up. The Disk Agent reads or writes data from a disk on the system and sends or receives data from the Media Agent. The Disk Agent is also installed on the Cell Manager, allowing you to back up data on the Cell Manager, the Data Protector configuration, and the IDB.
- **Media Agent:** (For servers that have direct access to tape drives.) During a backup session, the Media Agent receives data from the Disk Agent and sends it to the tape device (which can be directly attached or allocated over a SAN). During a restore session, the Media Agent locates data on the backup medium and sends it to the Disk Agent. The Disk Agent then writes the data to the disk. The Media Agent also manages the robotics control of a library. If a system has both the client agent and the media agent installed then it can perform LAN-free backups.

## Symantec NetBackup Application Overview

Symantec NetBackup Enterprise Server is based on a client/server architecture, and comprises several distinct components:

- **Master server:** The master server (one per storage domain) manages data protection operations. The NetBackup master server is responsible for containing the backup configurations and policies, running the scheduler that initiates automated backups, maintaining catalogs that track the location and contents of all backups, communicating with media servers to initiate backup and restore processes, and providing both a command line interface and a graphical user interface to administer NetBackup.
- **Enterprise Media Manager server:** NetBackup Enterprise Server 6.0 introduced the Enterprise Media Manager (EMM) server, which centrally manages Media Manager data that had previously been distributed across multiple media servers. You can only have one EMM server per storage domain. The master server can also be designated as the EMM server. The NetBackup EMM server is responsible for managing a consolidated media and device database for the NetBackup storage domain and maintaining run-time status information, NDMP credentials, and a managed server list.
- **Media server:** A media server can be any standalone server running NetBackup server software that receives requests for backup and restore operations from the master server. (A master server can also be configured as a media server.) A storage domain can contain multiple media servers. Media servers communicate with the master server and EMM server to initiate backup and restore operations on its attached storage devices, communicate with NetBackup clients during a backup or restore operation, monitor the status of storage devices, and provide robotic control.
- **Client:** NetBackup clients are servers that have NetBackup client software installed. A client performs backups through a media server (over the LAN if the media server is running a separate server or directly as LAN-free backup if the media server is in the client server).
- **Storage Unit:** A NetBackup storage unit is a storage device attached to a NetBackup server. A storage unit consists of the media server and tape devices where NetBackup stores files and data.

If a storage unit contains two drives and one is busy, NetBackup can use the other drive without administrator intervention. To send backups to a storage device, the administrator must define storage units using the Device Configuration Wizard. For virtual library devices such as VLS and D2D, the type of storage unit used is the "Media Manager storage unit." (Even though VLS and D2D are using disk internally, they are seen by NetBackup as a tape library.)

## IBM TSM Application Overview

Figure 16 shows an overview of TSM (Tivoli Storage Manager) LAN-based backups that incorporate virtual libraries. The steps are described below.

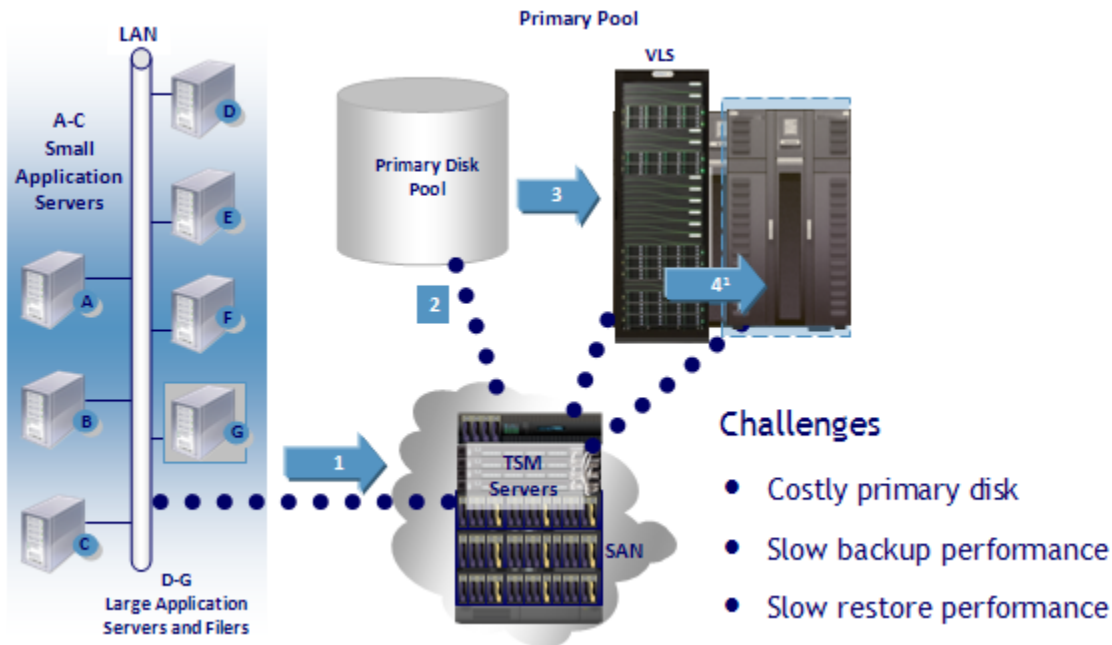
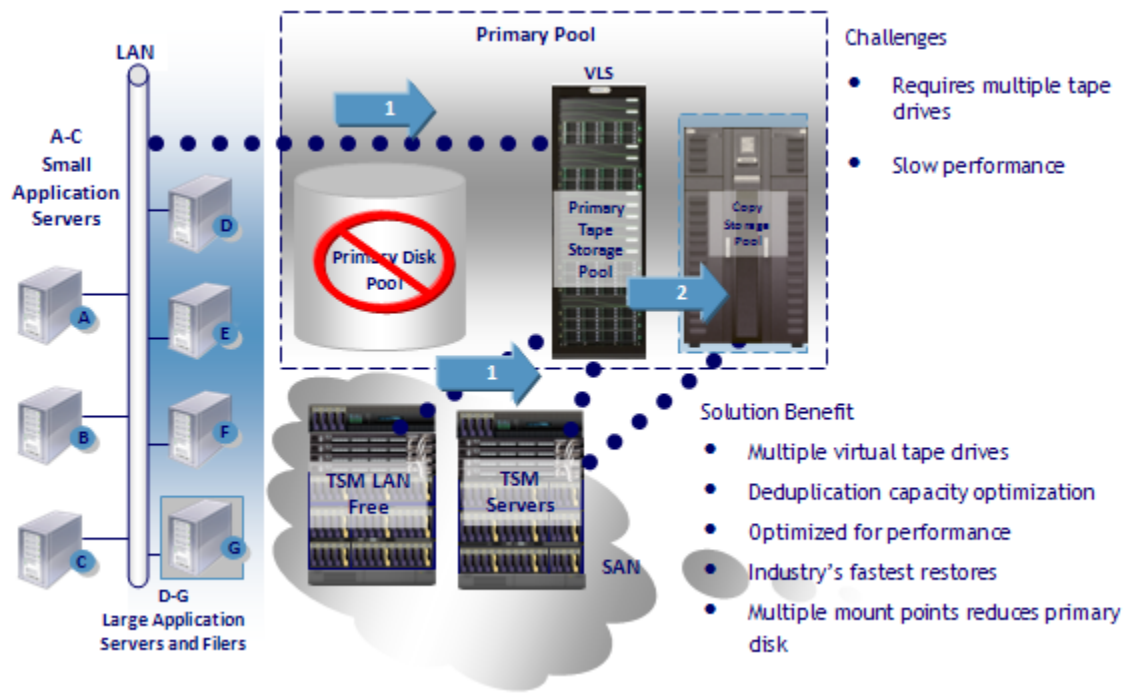


Figure 16 TSM LAN-based backups

1. Data on application servers (A-G) is backed up by TSM server via the SAN to the TSM disk pool. A typical TSM primary disk pool is sized for a minimum of one nightly backup.
2. As the primary disk pool fills, TSM moves data from it via the SAN to the primary tape pool. Unscheduled Migrations Undersized disk storage pools; a common mistake is to define disk pools that do not have the capacity to hold one night's worth of backup data. Doing so results in data migration to tape starting when the disk pool has reached its high migration threshold, rather than starting as scheduled using the "migrate" function. Unscheduled migrations can interfere with backups causing them to fail or run beyond the allowable window.
3. Data is copied (migrated) from the primary disk pool to the primary tape storage pool on a tape library. This migration is an I/O and CPU-intensive operation, which is typically limited by the speed of tape drives. If a migration occurs during a backup, the backup performance slows dramatically.
4. Data is copied from the primary tape storage pool to a copy storage pool, making a set of tapes to send off-site for disaster protection. By replacing the primary tape storage pool with a virtual library, you can dramatically improve backup and restore times, reduce the need for primary disk, and enable more use of advanced TSM functions such as LAN-free backups.

Figure 17 shows an overview of TSM LAN-free backups that incorporate virtual libraries to improve backup performance and reduce LAN load. The steps are described below.



**Figure 17 TSM LAN-free backups**

1. Data on large application servers (D-G) is backed up by TSM server via the SAN directly to the tape storage pool.
2. Data is copied from the primary tape storage pool to a copy storage pool, making a set of tapes send off-site for disaster protection.

## EMC NetWorker Application Overview

EMC NetWorker is based on a client/server architecture and comprises several distinct components:

- Console server: all EMC NetWorker servers and clients are managed from the NetWorker Console server. The Console server also provides reporting and monitoring capabilities for the servers and clients.
- Web browser: the Console server is accessed through a graphical interface that can be run from any computer that has a web browser. Multiple users can access the Console server concurrently from different browser sessions. A computer that hosts the web browser can also be a NetWorker client.
- Datazone: a datazone is a single NetWorker server and its client computers. Datazones can be added as backup requirements increase.
- NetWorker server: NetWorker servers provide services to back up and recover data for the NetWorker client computers in a datazone.
- NetWorker storage node: a NetWorker storage node can be used to improve performance by offloading much of the data movement from the NetWorker server involved in a backup or recovery operation.
- NetWorker client: a NetWorker client computer is any computer whose data must be backed up. The NetWorker Console server, NetWorker servers, and NetWorker storage nodes are also NetWorker clients.



---

# 4 D2D Systems

## D2D Devices

### D2D Defined




The entry level D2D100 series Backup System meets the needs of small businesses as a low-cost solution that does not incorporate deduplication technology. The D2D2500 is well suited for remote and branch offices and small IT environments, while the more powerful D2D4000 and D2D4100 are designed for medium-sized companies and small data centers. The D2D2500, D2D4000 and D2D4100 products include HP Dynamic deduplication, which provides low cost and flexibility to meet the needs of smaller IT environments.

All HP D2D products are RAID disk-based backup devices using serial ATA drives. All D2D devices include iSCSI virtual library interfaces over LAN, and the D2D4000/4100 have the option of adding a Fibre Channel interface to broaden connectivity. The D2D emulates a variety of physical tape libraries including the tape drives and cartridges inside the libraries. D2D accommodates mixed IT-platform and backup-application environments. For more insight into mixed-IT solutions and environment analysis, visit the HP StorageWorks Enterprise Backup Solutions web site: <http://www.hp.com/go/ebs>.

The virtual tape solution uses disk storage to simulate a tape library. That is, a virtual library simulates libraries with tape drives and slots as if they were physical devices. Backup servers see the virtual libraries as physical libraries. Because you can create many more virtual libraries and drives than you have physical tape drives, many more SAN-based backups can run concurrently from the application servers, reducing the aggregate backup window. After the backups are complete, the tape-offload feature or the data protection software can migrate backup data from the virtual media to physical tape for off-site disaster protection or long-term archival. This migration to physical tape can happen outside of the backup window at whatever time you determine will work optimally with your backup environment. The deduplication-enabled replication feature also provides the ability to create a second off-site copy in another D2D via a TCP/IP link.

## D2D Technical Specifications

**Table 6 D2D Technical Specifications**

D2D2500 Backup Systems	D2D4000 Backup Systems	D2D4112 Backup Systems
		
For remote or branch offices and smaller data centers	For remote or branch offices and mid-size data centers	For mid-size data centers and distributed environments
1U rack-mount	2U rack-mount	2U rack-mount plus 2U upgrade
Up to 2, 3, or 4 TB raw capacity (1.5, 2.25, or 3 TB usable) depending on model	4.5 TB and 9 TB raw capacity (3 TB and 7.5 TB usable)	Scalable from 12 to 24 TB raw capacity
Speeds of up to 75 MB/sec	Speeds of > 90 MB/sec	Speeds of > 150 MB/sec
2 x iSCSI interfaces	2x iSCSI or 2x FC interfaces	2x iSCSI and 2x FC interfaces
Protects up to 6 servers	Protects up to 16 servers	Protects up to 24 servers
Hardware RAID 5	Hardware RAID 6	Hardware RAID 6
Dynamic deduplication	Dynamic deduplication	Dynamic deduplication
Deduplication—enabled replication	Deduplication—enabled replication	Deduplication—enabled replication
Up to 6 source appliances per target	Up to 16 source appliances per target	Up to 24 source appliances per target

## D2D Design Considerations

### D2D Port Optimization

The D2Ds have two 1Gbit Ethernet ports, and some models also have two 4 Gbit Fibre Channel interfaces. Correct configuration of these interfaces is important for optimal data transfer.

### D2D Ethernet Ports

The Ethernet ports are used for both data transfer over the iSCSI protocol and management access to the web GUI. The two ports have the following configuration modes:

- Single port – Port 1 must be used for all data and management traffic. Use this mode only if no other ports are available on the switch network or if the appliance is used to transfer data over Fibre Channel ports only.
- Dual port – Both ports are used, but must be in separate subnets. Both ports can access the web GUI but virtual libraries are split across the two ports. Use this mode if:



- Servers to back up are split across two physical networks which need independent access to the D2D. You must assign the virtual libraries to whichever port is on the network of the device to back up.
- Separate data and management LANs are used, i.e., each server has a port for business network traffic and another for data backup. One port on the D2D can be used solely for GUI access with the other used for data transfer.
- High Availability port – Both ports are used but are bonded to appear as a single port. This is the recommended configuration, there is no special switch configuration required other than to ensure that both Ethernet ports from the D2D are connected to the same switch. This mode sets up “bonded” network ports where both network ports are connected to the same physical switch and behave as one network port. This mode provides some level of load balancing across the ports and also port failover.

## Backup Server Networking

When considering backup performance, consider the whole network. Any server acting as a backup server should be configured where possible with multiple network ports which are bonded in order to provide a fast connection to the LAN. Client servers (those that backup via a backup server) may be connected with only a single port if backups are to be aggregated through the backup server.

Ensure that no sub 1Gbit network components are in the backup path because this will significantly restrict backup performance.

## Fibre Channel SAN

Two Fibre Channel ports are provided at 4Gbit. These should both be connected to a SAN switch which supports 4Gbit Fibre Channel.

Virtual library devices are assigned to an individual interface; therefore, balancing virtual devices across both interfaces will ensure that one link is not saturated while the other is idle.

Fibre Channel devices should be zoned on the switch to only be accessible from a single backup server device to ensure that other SAN events (such as the addition and removal of other Fibre Channel devices) do not cause unnecessary traffic to be sent to devices. It also ensures that SAN polling applications cannot reduce the performance of individual devices.

## Multiple Backup Streams

The HP D2D Backup System performs best with multiple backup streams sent to it simultaneously. For example, a D2D4009i will back up data at approx 40 MB/s for a single stream; multiple streams can deliver an aggregate performance in excess of 80 MB/s.

There are two ways that the D2D can be used to achieve multiple stream backups:

- Several virtual library devices each with a single virtual tape drive receiving one backup stream each.  
Use this configuration if you are backing up multiple servers, each of which is a backup application “media server,” which can directly connect to a virtual library target. Also use this configuration if you are backing up one or more servers but wish to keep the data backed up from each logically separate for easier management.
- A single virtual library with multiple virtual tape drives (up to four on HP D2D4000 models) receiving backup jobs to each tape drive simultaneously.  
Use this configuration if you are backing up a single server (or multiple servers using a single “media server”), and you wish to achieve some deduplication of backup data across all of the

servers. In this case, also ensure that bonded/trunk network ports are configured for both the media server and the D2D device.

With either configuration there is a requirement to correctly configure the backup application to target the devices correctly; this will usually involve creating multiple backup jobs to back up specific data to selected virtual libraries.

Configuring backups to run in parallel rather than serially also improves replication and tape offload operations because these run after backup jobs have completed. Having backups finish at roughly similar times will ensure that the replication and tape offload jobs are not in contention with backup jobs.

However, backup application vendors often base their licensing model on the number of tape devices configured, so you may need to balance between performance and cost. Some applications (e.g., HP DataProtector) offer capacity based licensing (per TB) for virtual tape appliances; this licensing model may be more cost effective because it allows you to configure any number of libraries or tape devices.

## Disable Backup Application Verify Pass

Most backup applications will default to performing a verify operation after a backup job. While this offers a very good way to ensure that data is backed up successfully, it will also heavily impact the performance of the whole backup job. Performing a verify operation will more than double the overall backup time due to the fact that restore performance (required for verify) is slower for deduplication enabled devices.

Disabling verify for selected backup jobs can be done relatively safely because D2D Backup Systems perform CRC (Cyclic Redundancy Check) checking for every backed up chunk to ensure that no errors are introduced by the D2D Backup System. However, if data is incorrectly written by the backup application but is not corrupt, this will not be detected by the D2D. For this reason, verifying some backup jobs is still recommended.

## Tape Copy or Offload Performance

Copies of cartridges to physical media using the HP D2D Backup System can be achieved by using either a directly attached tape drive or library where the offload is controlled by a manual or scheduled operation from the D2D Web Management Interface, or by using a backup application on the media server to copy data from a cartridge on the HP D2D Backup System to a tape device or library connected to the media server. In either case, take the deduplication process into consideration when scheduling the offload.

The D2D allows a tape offload to be automatically scheduled to start immediately after a backup to the virtual tape completes; most backup applications provide a similar configuration option if performing a backup application copy. Take care when using these options because they can result in slower than expected offload performance due to some level of housekeeping running; also, if replication is configured the two operations will clash and result in a delay to replication.

You should delay any tape offload for several hours after the backup completes if the backup is doing a significant amount of data overwrite. Otherwise, the housekeeping process will have a detrimental effect on the tape offload performance. As a general rule, for every 100 GB of data on the cartridge being overwritten, allow 30 minutes to pass before starting an offload operation. Because it depends on the deduplication ratio and amount of unique chunks contained in that backup, it is hard to determine how many housekeeping operations will be generated (and therefore how long it will take to run) by an overwrite backup.

In addition, consider other operations running on the D2D. For example, if multiple backups finish at different times, each may result in some housekeeping work. Add all of the data backup amounts together to get an idea of the time that housekeeping will take.

## Tape Library Emulation

### Emulation Types

The HP D2D Backup Systems can emulate several types of HP tape library devices, and the maximum number of drives and cartridge slots is defined by the type of library configured. However, performance is not related to library emulation except in the ability to configure multiple drives per library and thus enable multiple simultaneous backup streams.

### Cartridge Sizing

The size of a virtual cartridge has no impact on its performance. HP recommends that cartridges are created to match the amount of data being backed up. For example, if a full backup is 500 GB, you should select 800 GB, the next larger configurable cartridge. If backups span multiple cartridges, you risk a performance impact because housekeeping operations will start on the first backup cartridge as soon as the backup application spans to the next cartridge.

### Number of Libraries per Appliance

The D2D appliance supports creating multiple virtual library devices. If large amounts of data are backing up from multiple hosts or for multiple disk LUNs on a single host, you should separate these across several libraries (and consequently into multiple backup jobs). Each library has a separate deduplication “store” associated with it, and reducing the amount of data and complexity of each store will improve its performance.

Creating a number of smaller deduplication stores rather than one large store which receives data from multiple backup hosts can have an impact on the overall effectiveness of deduplication. However, generally the cross server deduplication effect is quite low unless you are storing a lot of common data. If you store a large amount of common data on two servers, HP recommends that you back these up to the same virtual library.

## Optimizing Rotation Scheme to Reduce Housekeeping

Reducing the frequency of full backup cartridge reduces the amount of deduplication housekeeping overhead required. You can optimize it using the following:

- Longer retention policy: Deduplication has little penalty for using a large number of virtual cartridges in a rotation scheme and therefore a long retention policy for cartridges (most data will be the same between backups and therefore deduplicated), which reduces the number of overwrites.
- Full vs. incremental/differential backups: The requirement for full or incremental backups is based on how often offsite copies of virtual cartridges are required and on speed of data recovery. If regular physical media copies are required, the best approach is that these are full backups on a single cartridge. Speed of data recovery is less of a concern with a virtual library than with physical media. For example, if a server fails and needs to be fully recovered from backup, the recovery will require the last full backup plus every incremental backup since (or the last differential backup). Finding and loading multiple physical cartridges is time-consuming. However, with virtual tape there is no need to find all of the pieces of media. Because the data is stored on disk, the time to restore is lower due the ability to randomly seek more quickly within a backup and load a second cartridge instantly.

- Overwrite vs. append: Overwriting and appending to cartridges is also where virtual tape has an advantage. With physical media you may want to append multiple backup jobs to a single cartridge in order to reduce media costs; the downside is that cartridges cannot be overwritten until the retention policy for the last backup on that cartridge has expired. With virtual tape, you can configure a large number of cartridges for “free” and you can configure their sizes appropriately to the amount of data stored in a specific backup. Therefore, appended backups provide no benefit. (This does not apply to VLS deduplication where appending is required to fill up tapes and perform space reclamation on them.)

Considering the above factors, the following provides an example of a good rotation scheme. This user requires weekly full backups sent offsite and recovery point objectives of every day in the last week, every week in the last month, every month in the last year, and every year in the last five years:

- Four daily backup cartridges, Monday to Thursday, incremental backup, overwritten every week.
- Four weekly backup cartridges, on Fridays, full backup, overwritten every fifth week.
- 12 monthly backup cartridges, last Friday of the month, overwritten every 13th month.
- Five yearly backup cartridges, last day of the year, overwritten every five years.

In the steady state daily backups will be small, and while they will always overwrite the last week the amount of data overwritten will be small. Weekly full backups will always overwrite, but housekeeping has plenty of time to run over the following day or weekend; the same is true for monthly and yearly backups. The user can also offload a full backup to physical tape every week, month, and year after the full backup runs for offsite storage.

## D2D Blueprints

The following blueprints of D2D virtual tape libraries with deduplication and replication start from specifying company requirements, then defining the HP blueprint for the solution, and finally defining any solution caveats or ISV dependencies associated with the solution. This will help you make informed decisions and allow you to quickly assess areas of concern and possible implementations.

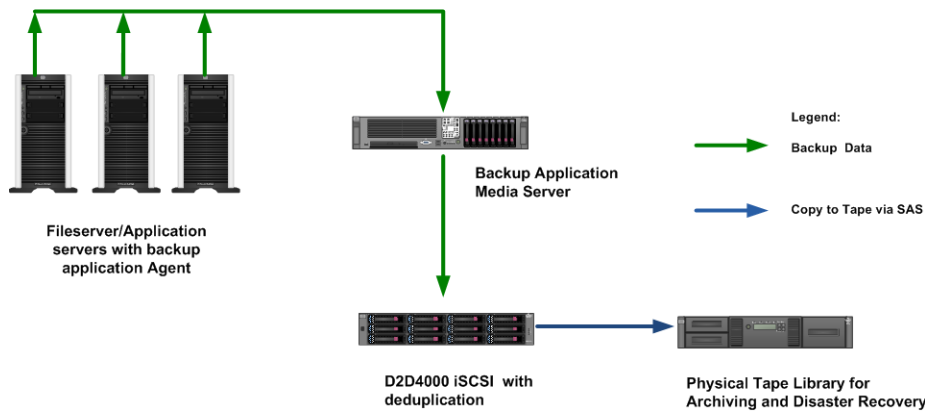
### Single Site Cost Effective Backup Device Consolidation

**Company requirements:** small site with fewer than 6 servers, fewer than 5 TB, no SAN, wants a single device for backup that can grow with the business.

**Solution:** HP D2D Backup System (iSCSI) with Tape Offload, typically used with HP Data Protector Express or Symantec Backup Exec

**Caveats:** The smaller D2D units shown here are a fixed size and not directly expandable in capacity. The initial sizing is very important.

**ISVs:** HP Data Protector Express or Symantec Backup Exec are typical backup applications in this segment. See <http://www.hp.com/go/connect> and <http://www.hp.com/go/ebs>.



**Figure 18 Single Site Backup Consolidation using iSCSI**

**Device configuration highlights:**

- Configure iSCSI ports in high availability mode to give added redundancy.
- Configure a single library for best deduplication ratio or multiple libraries for best throughput.
- Consider configuring multiple virtual libraries for different data types (office files, MExchange databases) across multiple host servers to get the best possible performance and deduplication ratio.

**Backup application highlights:**

- D2D devices have the highest throughput with multiple concurrent backup streams. Configure multiple jobs to run in parallel (at least 4) to obtain best throughput.
- No additional licensing required for physical tape library because the backup application is connected directly to the D2D unit (SAS or SCSI HBA required).

**Recovery options:** Depending on retention periods set, it is possible to recover all data directly from D2D device over extended periods (3-6 months). If backup data on D2D is regularly copied to physical tape and off-sited, there is a Disaster Recovery solution available also.

**Solution advantages:**

- Simple to implement
- Cost effective consolidation (iSCSI is free)

**Solution trade-offs:** Tape Offload functionality needs monitoring as a separate process.

When to use direct attach tape and when to use Backup ISV copy functionality?

If cost is a major driver, use direct tape attached to the D2D unit (because this does NOT attract extra licencing costs with the backup software), and then use D2D tape offload capability to transfer data from D2D to physical tape. However, this requires closer daily management by checking D2D through the GUI. This copy is termed a media copy. If D2D tape has 50 GB of data, then 50 GB of data will be transferred to physical tape and cannot be appended to.

Use Copy from D2D to Physical Tape using Backup Application when a single point of management is required and where better use of physical media is required. Object copies via the ISV software can merge several virtual cartridges onto a single physical cartridge saving physical media costs but incurring additional licencing costs.

**Large SME Site Consolidation Requiring Fibre Channel Shared Devices**

**Company requirements:** fewer than 20 TB, up to 24 servers, growth proof, 12 hour backup window, reduce dependency on physical tape.

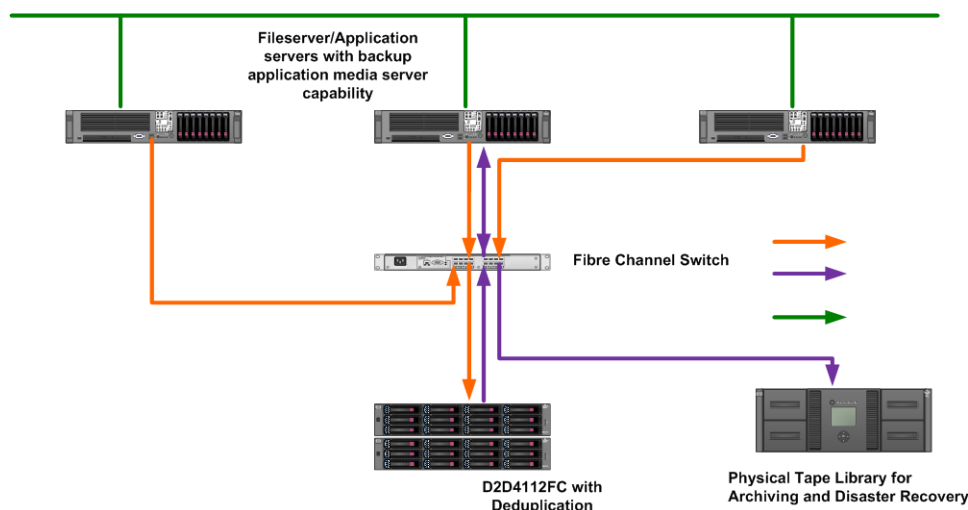
**Solution:** HP D2D 4112 “expandable” storage, good performance, backup application Object Copy to small physical tape library for archiving and disaster recovery.

**Caveats:** Preferred offload to physical tape is via backup application “object copy” functionality for maximum physical tape usage efficiency.

**ISVs:** Typically HP Data Protector, Symantec Netbackup, Tivoli Storage Manager, EMC Networker, Commvault, Backup Exec, etc. For a full list see <http://www.hp.com/go/connect> and <http://www.hp.com/go/ebs>.

**Device configuration highlights:** Consider configuring multiple virtual libraries for different data types (office files, MExchange databases) across multiple host servers to get the best possible performance and deduplication ratio.

**Backup application configuration highlights:** Run multiple backups in parallel for best throughput.



**Figure 19 Large SME Site Consolidation Requiring FC Shared Devices**

**Recovery options:** Depending on retention periods set, it is possible to recover all data directly from D2D device over extended periods (3-6 months). If backup data on D2D is regularly copied to physical tape and off-sited, there is a Disaster Recovery solution available also.

**Solution advantages:**

- Effective consolidation for SME
- Allows backup window to be reduced
- Scalable using D2D4112 – up to 24 virtual devices can be created
- Single management point with Backup Application for D2D & physical tape
- Flexible - some backups to physical tape, some to D2D if required
- Disaster recovery solution included via Physical tape offsiteing

**Solution trade-offs:** Using backup software to do D2D to tape copies attracts extra licencing costs.

**Other information:**

In this example we are using the backup example to copy data from D2D to physical tape because the environment is significantly larger and a single point of management is required for all backup and copy operations. Physical media usage is also a priority. These requirements are best served by using object copy functionality within the backup application itself. This will attract higher licencing costs, however.

## Multi-site Small Business Disaster Recovery Solution

**Company requirements:** remove physical tape and any human intervention for backups at remote sites – completely automate backup and disaster recovery. Fewer than 1 TB at each remote site and up to 24 remote sites.

**Solution:** Multi-remote site D2D iSCSI devices replicating into a central consolidated D2D at a disaster recovery site.

**Caveats:** Some scripting is required for automated replicated data catalog imports at the target site. (HP Data Protector). Other ISVs require manual import. Sizing replication links and replication window is key; use the HP StorageWorks Backup Sizer.

**ISVs:** HP Data Protector Express, HP Data Protector, Backup Exec, Netbackup, Networker, TSM, Commvault Galaxy, BakBone NetVault – D2D4000 and D2D2500, Syncsort Backup Express – D2D4000 and D2D2500.

**Device configuration highlights:** Remote sites typically will have only one virtual library configured. The central site can have up to four slot mappings per virtual library. 6 x 4 gives support for up to 24 remote sites.

Consider configuring multiple virtual libraries for different data types (office files, MExchange databases) across multiple host servers to get best possible performance and deduplication ratio.

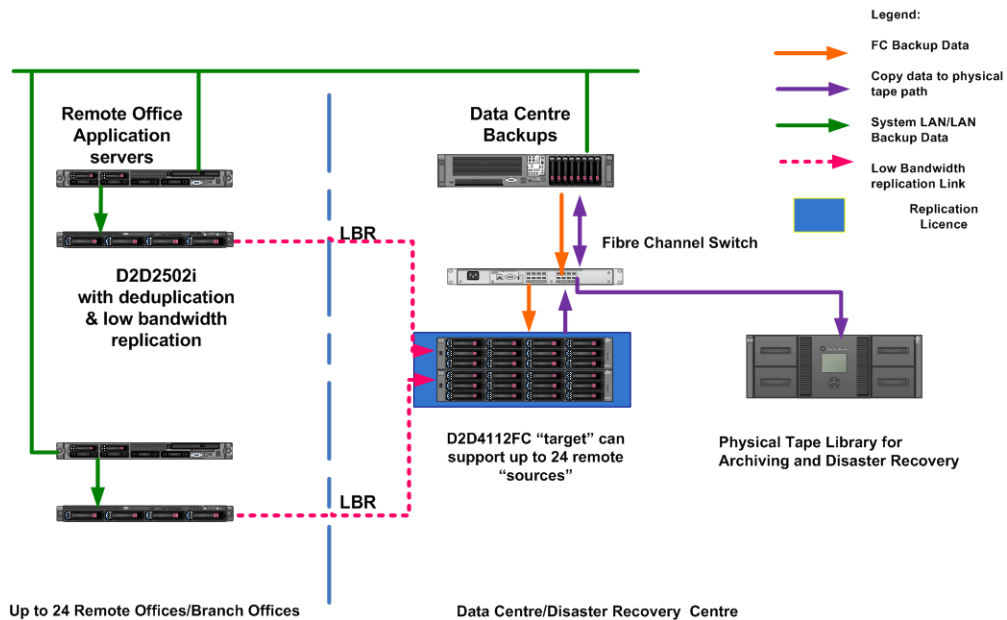
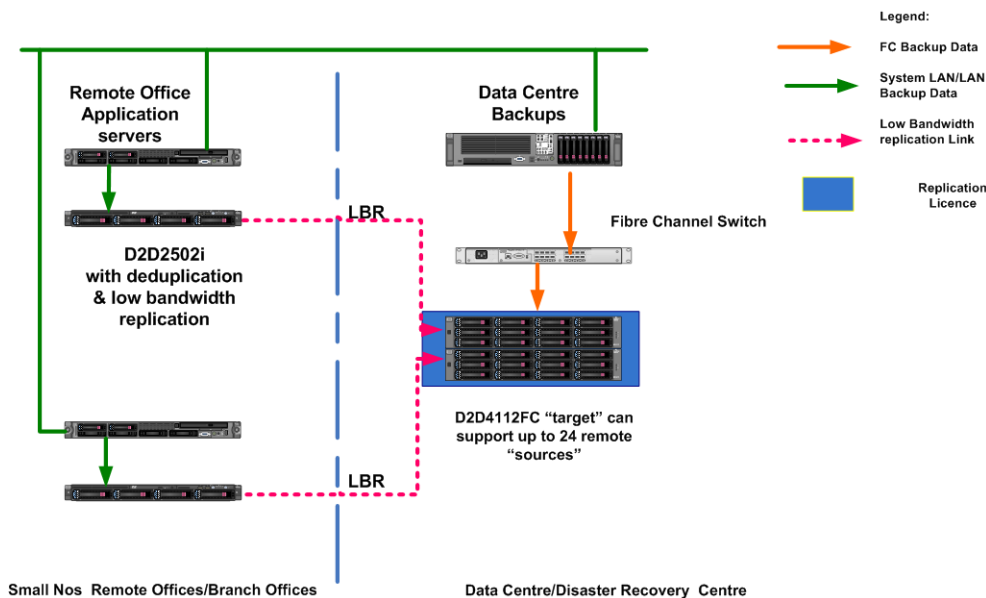


Figure 20 Multi-site Small Business Automated Disaster Recovery Solution



**Figure 21 Multi-site Small Business Automated Disaster Recovery Solution with no Physical Tape**

For a smaller business where there may be limited budgets or no compliance requirements dictating data to be archived or stored on tape for several years, the solution shown without physical tape above is perfectly adequate. The D2D4112 provides RAID 6 data protection plus hot spare disk capability so data integrity is guaranteed.

**Backup application configuration highlights:** All replication is controlled via the D2D units. Remote offices and the central site are configured as normal with the backup application. Multiple parallel backup streams will always give better performance.

Object copy functionality at the Data centre copies both local data centre backups and replicated data from the remote offices onto physical tape.

**Recovery options:** Remote sites can recover data from the local D2D in a matter of minutes. In the event of a major disaster at the remote site there are several recover options:

- Rebuild server data from replicated data held on the D2D at the data centre.
- Rebuild remote office and “reverse replicate” data from the data centre back to the remote office over the WAN link (not recommended for high volumes of data such as hundreds of GB).
- Take physical tape copy at the data centre and temporarily connect the tape device to the server at the remote site and recover data from physical tape.

The main data centre can be recovered either from D2D at the main data centre or from copies of data made to physical tape.

**Solution advantages:**

- Removes all human intervention from the backup at remote sites where there may be no IT expertise.
- Cost effective replication links as low as 2 Mbit/sec because D2D only replicates differences in data.
- D2D at the data centre also acts as both a source for data centre backups and a target for remote office replications offering significant benefits of consolidation.
- Physical tape for archiving and a disaster recovery solution for the data centre.

**Solution trade-offs:**

Before replication can start, a one-time “seeding process” has to be followed to ensure remote sites and data centre D2D have the same data reference points. Seeding can be done via WAN link, via

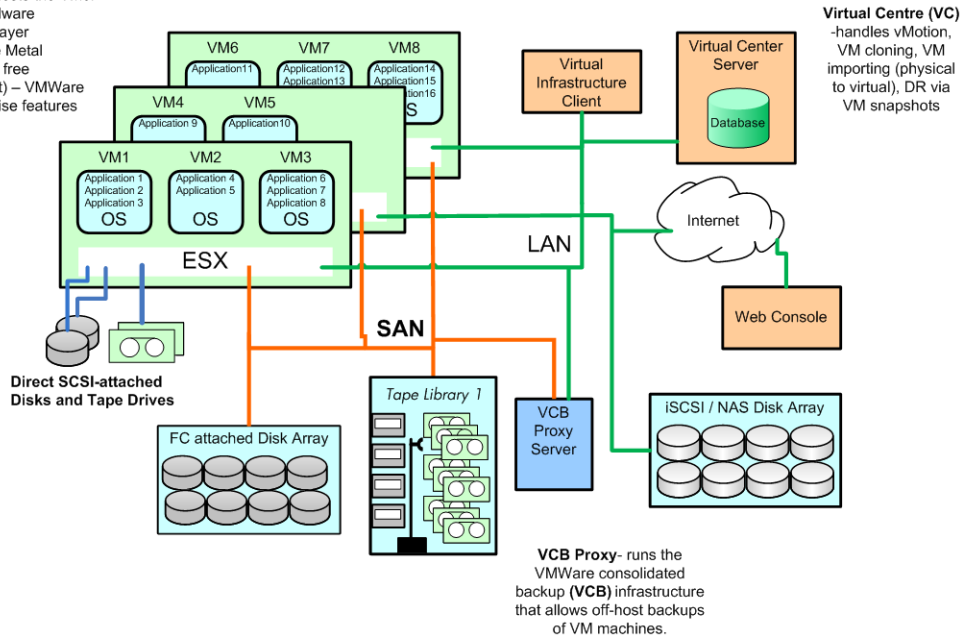


co-location of units and then splitting of units, or by using physical tape to transfer data from remote office to seed unit in data centre. This process only has to be done once.

**Other information:** WAN link sizing and data centre D2D sizing is critical. Size for existing link or new link using the HP StorageWorks Backup Sizer.

## Introduction to VMWare Terminology

**ESX** - is the Customised Linux OS that runs on the server that hosts the VMs. Acts as hardware abstraction layer  
**ESXi** = Bare Metal Hypervisor - free  
**ESX4** (64 bit) – VMWare with Enterprise features



**Figure 22 VMWare Infrastructure Environment**

The above diagram shows the different components that make up a virtualized server environment and its connections to external primary storage and data protection devices.

ESX acts as a hardware extraction layer on one side providing virtual resources to the VM machines whilst on the other side interfacing to the physical hardware.

In this next set of solutions we will talk about possible methods of providing data protection to this virtualized environment:

- Base Support – Agents inside Virtual Machines
- Virtual Machine snapshots – ESX Console
- Virtual Machine snapshots – VCB
- Zero Downtime Backup and Instant Recovery

Each technique has its pros and cons as listed below. It is important to point out that “snapshot” backups (sometimes called image or point-in-time backups) cannot always provide single file restore capability (for example with Linux the whole image has to be restored). In addition, snapshots are not best suited to complex online backups of applications such as Oracle, SQL, and Exchange because they do not quiesce the application in the appropriate way before taking the snapshot. Snapshots are best suited to being aligned with a VMWare machine OS where the whole environment can be captured in a single image.

## Using D2D in Simple VMWare Environments with ESXi

This configuration is currently being certified with VMWare. See future supported configurations at the end of this document.

## Using D2D in VMWare Environments with ESX4 and VMWare Consolidated Backup

**Company requirements:** Mid-size VMWare environment, wants VMWare based snapshot backup with individual windows file based recovery using VMWare consolidated backup (VCB) in ESX4. This is the backup solution promoted by VMWare themselves. The main idea behind VCB is to offload the workload for backup from the ESX server onto a separate proxy server giving potentially better performance.

**Solution:** HP D2D with VTL FC interface and separate VCB proxy server.

**Caveats:** This is mainly a snapshot based backup technology where complete images have to be restored. File level recovery from a file system backup is only possible with Windows . The proxy server has to have sufficient storage to store the snapshots if image based backups are being used because the images are copied from the primary storage to the proxy server for these types of backup.

**ISVs:** VCB in ESX4 takes snapshots of the VMWare machines which are then presented to the proxy server for backup so most backup applications will work when they are installed on the Proxy server. HP has tested Data Protector and Symantec Netbackup to great depth in this area.

**Device configuration highlights:** Consider creating one virtual library per virtual machine (best throughput) or one large library for all machines to share (best deduplication ratio). Consider configuring multiple virtual libraries for different data types (office files, MExchange Databases) across multiple host servers to get best possible performance and deduplication ratio.

**Backup application configuration highlights:** Backup several streams concurrently for best performance. Backup is a combination of VCB doing the snapshots and backup applications writing the snapshots to a backup device. VCB is hardware independent hence can accommodate snapshots with all array types.

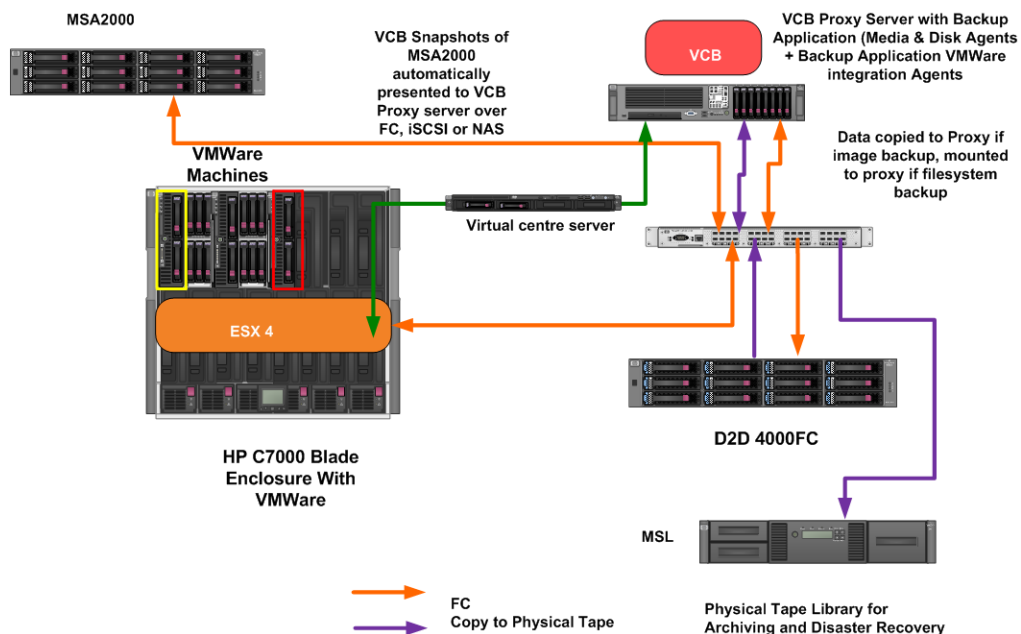


Figure 23 Using D2D in VMWare Environment with ESX4 and VCB Backup for SAN-based Backup

**Recovery options:** Recovery is a two stage recovery: 1) the image is restored to the VCB proxy or the ESX server, and 2) the image is restored as a VM using VMware Converter if restored to the VCB proxy, or vcbRestore if restored to the ESX server.

Restart the VM and you are back in business as of the last snap. This is the fastest way to do a complete disaster recovery. For example HP DP on Proxy can restore directly from D2D/Tape onto VM machines.

For Windows OS, you can also mount the filesystem backup to VCB proxy and restore single files or restore the complete system.

**Solution advantages:**

- VMWare approved method of backing up, takes the load off the ESX console for backup and restore and passes it to the proxy server
- Allows all HP supported Fibre Channel backup devices to be connected to Proxy server
- VCB is disk array independent snapshot technology.
- Automation of backup and restore of VM snapshots.
- Disaster recovery of ESX environments
- Easy transfer to physical tape if required, if backup application used.
- With HP Data Protector and most other Backup software the following approaches are possible:
  - Online and offline snapshots of VMs.
  - File level backup and recovery of data in the VMs
  - VCB Image Backup: Online and offline snapshots, Full Backup only.
  - VCB File Backup:
    - Full Backup
    - Incremental and Differential Backups based on per-file modification time — incremental contains changes since previous backup, differential contains changes since latest full backup.

**Solution trade-offs:**

- Full snapshot based recovery – single file recovery only on Windows OS file system backups.
- Distinctly separate proxy server required at additional cost – requires additional storage for copy of snapshots.
- Backup application integrations such as Oracle, SQL, Exchange, etc. cannot be used with VCB.

**Other Information:**

After the snapshot has been copied to proxy and backed up, it is deleted from the primary storage array where it was first taken.

Further enhancements are possible by providing multi-site disaster recovery by using low bandwidth replication on HP D2D devices.

## Using D2D in larger VMWare environments with HP Data Protector Zero Downtime Backup with Instant Recovery (ZDB + IR)

**Company requirements:** Large VMWare environment, high throughput required to meet backup window (Fibre Channel). High SLA s dictate an “Instant Recovery.” High proportion of online backups required.

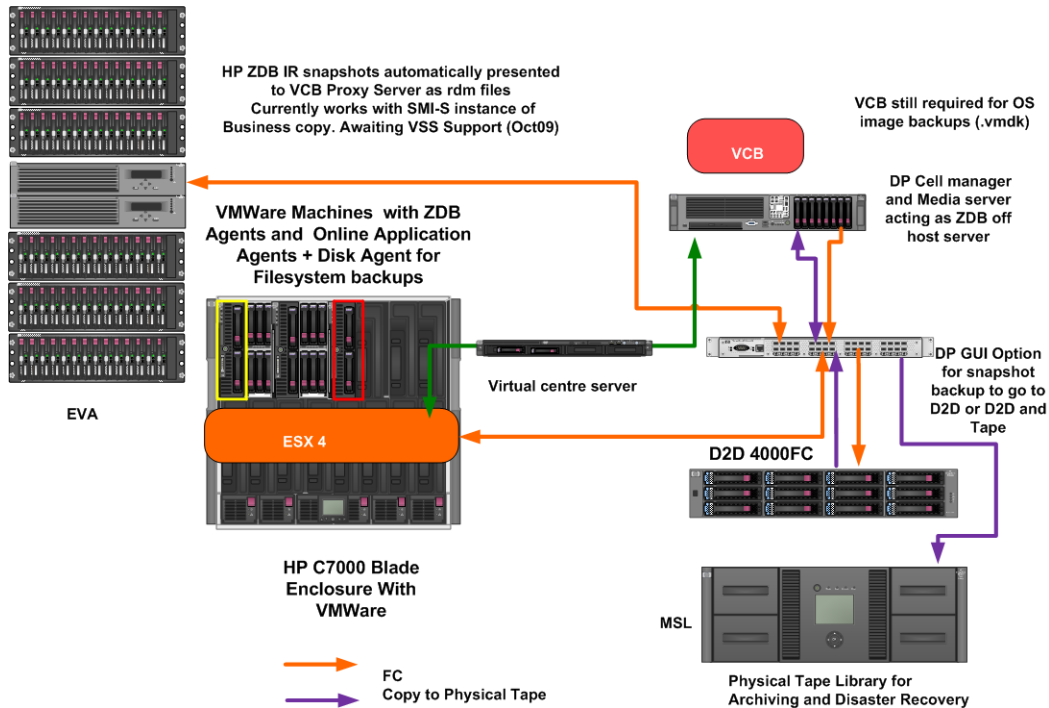
**Solution:** HP D2D with VTL interface and separate VCB proxy server with Data Protector VCB ZDB IR. Note that the VCB backup and the ZDB backup are completely independent of each other.

**Caveats:** Limited to arrays that support DP ZDB IR – XP, EVA and LHN (future).

**ISVs:** HP Data Protector with Zero Downtime Backup Option and Symantec Netbackup with Snapshot Client are both deep dive tested by HP in this mode of operation.

**Device configuration highlights:** Consider configuring multiple virtual libraries for different image types.

**Backup application configuration highlights:** Configuring Zero Downtime Backup within Data Protector and Symantec Netbackup requires some additional scripting work. See the HP EBS SAN Design Guide at <http://www.hp.com/go/ebs>.



**Figure 24 Using D2D in a VMWare Environment with ESX4 and HP Data Protector ZDB IR**

**Recovery options:**

- Two step restore process that provides full recovery of VM and data:
  - Recover the VM image from VCB backup.
  - Required only in case the whole VM is lost or restore to a new ESX server.
- Recover Application Data:
  - For replication based disc backups (IR), perform Instant Recovery.
  - For disc/tape backups, restore directly to the VM.
  - Recovery performed automatically by Data Protector as specified by user.

**Solution advantages:**

- VCB snapshot based backup is not really “application aware” whereas HP Data Protector ZDB quiesces the application prior to the hardware based snapshot and creates a crash consistent snapshot. It is much more suitable for live databases and email servers than VCB based backups.
- Remove the backup load from VM to a backup server with Data Protector Zero Downtime Backup.
- Reduce application recovery times from hours to minutes with Data Protector Instant Recovery.
- Easy copy to physical tape if required.

**Solution trade-offs:**

- Limited to EVA & XP & LHN (future)
- Snapshots could take up large amount of EVA and XP space depending on frequency of snapshots at extra cost.
- Note that to do the ZDB backup of VM volumes it is a requirement that the VM volumes use the VMware raw device mappings (RDM) feature.

**Other information:**

A simple two step backup for complete protection:

- VM Backup
  - Data Protector Integration with VMware VCB provides the first layer of protection.
  - Provides snapshots that can be used for restoring VM.
  - Needed only when VM configuration changes.
- Application Backup
  - Provides the second layer that enables protection of application data.
  - Can be performed using ZDB/IR feature of Data Protector by using HP Business Copy for EVA, XP (LHN to follow).

Further enhancements possible by providing multi-site disaster recovery by using low bandwidth replication on HP D2D devices.

## D2D Dynamic Deduplication

HP Dynamic deduplication technology is designed around compatibility and cost for users with smaller IT environments. It offers the following features and benefits:

- Uses hash-based chunking technology, providing deduplication technology at a lower cost.
- Operates independently of backup format, avoiding issues with backup application support.
- Automatic and manageable.

Configuring Dynamic deduplication is simple. There is no need to tell the deduplication software which backup application you are using or where file markers are. Dynamic deduplication works with any backup application and with any data type.

See <http://www.hp.com/go/ebs> for more information on backup applications and data types recommended for use with Dynamic deduplication.

### How it Works

Dynamic deduplication uses hash-based chunking technology, which analyzes incoming backup data in “chunks” that average to 4K in size. The hashing algorithm generates a unique hash value that identifies each chunk, and points to its location in the deduplication store.

---

 **NOTE:**


Using small sections of about 4K in size allows the Dynamic deduplication software to locate a greater number of instances of duplicated data. With smaller chunks of data, more hash values are generated and compared against each other. If larger chunks of data are analyzed, and a pattern does not repeat within of the boundaries of the chunk, the duplication is not caught.

---

Hash values are stored in an index that is referenced when subsequent backups are performed. When data generates a hash value that already exists in the index, the data is not stored a second time. Rather, an entry with the hash value is simply added to the “recipe file” for that backup session.

Over the course of many backups, numerous instances generating the same hash values occur. However, the actual data is only stored once, so less storage is consumed by duplicate backup data.

Figure 25 illustrates the process that backup data undergoes with Dynamic deduplication. The numbered list that follows corresponds with the image.



Item	Description
1	The backup stream is analyzed in 4K chunks that generate unique hash values. These hash values are placed in an index in memory.
2	When a subsequent backup contains a hash value that is already in the index, a second instance of the data is not stored.
3	When new hash values are generated, they are added to the index and the data is written to the deduplication store.

Figure 25 Hash-based Chunking

## Dynamic Deduplication Implementation

Dynamic deduplication is enabled per library on the D2D. See Figure 26.

When you configure the library, it defaults to deduplication enabled. If you disable it, you cannot selectively apply deduplication to any data on the library device. Compression is also disabled if deduplication is disabled.

HP D2D Backup System

Home Configuration Status Tape Attach Settings

ssdcl405 : IPV4 15.38.74.17 : IPV6 fe80::21c:c4ff:fe5f:513a/64 Logout Help

Devices Cartridges Network Email Alerts

Devices

Name	Cartridges	Connection	Device Type	Port
Library 1	48	Connected	MSL G3 Series (4x48) /LTO 3	iSCSI Port

Create library

Library 1

Libraries World Wide Names iSCSI Authentication

Library Name	Library 1		
Emulation Type	MSL G3 Series (4x48)	Drive emulation	LTO 3
Total No Of Cartridge Slots	48	Number of drives	1
Deduplication	<input checked="" type="checkbox"/>		
Port	iSCSI Port		
iSCSI Initiator Name	iqn.1991-05.com.microsoft:ssdcl347		

Delete Edit

**Figure 26 D2D Configuration**

The current deduplication ratio for the D2D device is calculated and displayed in the GUI. This figure updates automatically. See [Figure 27](#).

HP D2D Backup System

Home Configuration Status Tape Attach Settings

ssdcl405 : IPV4 15.38.74.17 : IPV6 fe80::21c:c4ff:fe5f:513a/64 Logout Help

Disk RAID iSCSI Log

Disk Space Allocated

Total Disk Space Allocated	
Used 466 GB	Max 7451 GB
6 %	
Deduplication Ratio (1.9 : 1)	

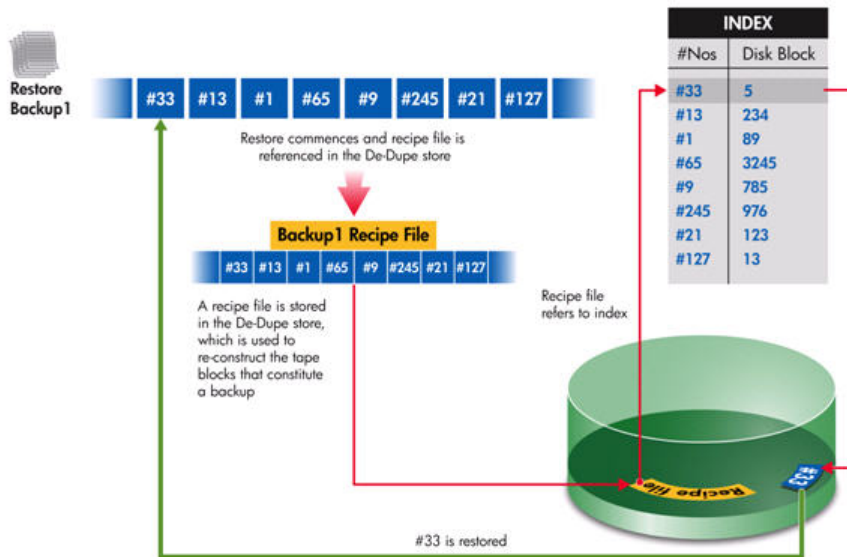
**Figure 27 Dynamic Deduplication Ratio**

The D2D device fully emulates a tape library, including its drive-type and supported compression. Dynamic deduplication can be used in combination with data compression for even greater storage savings.

You can directly connect a tape drive to the D2D device for copying and exporting data to physical media. Data that has been deduplicated can be exported to physical tape without being reassembled first, and without using the backup application.

## Restoring Data

On receiving a restore command from the backup system, the D2D device selects the correct recipe file and begins reading it, starting with the first hash. This process is illustrated in [Figure 28](#) and the ordered list that follows.



**Figure 28 Restoring from a Recipe File**

1. The first hash in the recipe file is located in the index, which provides the location on disk of the original chunk of data.
2. The original chunk is found and returned to the restore stream.
3. The D2D moves on to the second hash, and repeats this process for all subsequent hash values in the recipe file until the entire file is returned to the restore stream.

## Housekeeping

If data is deleted from the D2D system (e.g., a virtual cartridge is overwritten or erased), any unique chunks only applicable to the deleted data will be marked for removal and non-unique chunks will be de-referenced. The process of removing chunks of data is not an inline operation because this would significantly impact performance. Instead, the process of “housekeeping” runs as a background operation on a per cartridge basis and runs as soon as the cartridge is unloaded and returned to its storage slot.

While the housekeeping process can run as soon as a virtual cartridge returns to its slot, this could cause a high level of disk access and processing overhead affecting other operations such as further backups, restores, tape offload jobs, or replication. In order to avoid this, the housekeeping process checks for available resources before running; if other operations are in progress the housekeeping will delay to prevent impacting performance of other operations. The delay is not binary (i.e., on or off), so even if backup jobs are in process some low level of housekeeping will still occur which may have a slight impact on backup performance.

Housekeeping is important in maximizing the deduplication efficiency of the appliance. Therefore, you must ensure that it has enough time to complete. Running backup, restore, tape offload, and replication operations with no break will result in housekeeping never completing. As a general rule, housekeeping needs 30 minutes per day for every 100 GB of data overwritten on a virtual cartridge.

For example, if on a daily basis the backup application overwrites two cartridges in different virtual libraries with 400 GB of data on each cartridge, the system will need two hours of inactivity over the course of the next 24 hours to run housekeeping in order to de-reference data and reclaim any free space.



Configuring backup rotation schemes correctly is very important to ensure the maximum efficiency of the product by reducing the amount of housekeeping required. See [Optimizing Rotation Scheme to Reduce Housekeeping](#).

## D2D Replication

The D2D replication technology (which leverages Dynamic deduplication) offers:

- **Ease-of-Use.** Unlike some competitors, D2D has a web based GUI to configure replication targets, implement policies, and provide comprehensive reporting. Wizards for replication and recovery make the configuration easy to accomplish in only five steps.
- **Tape Offload functionality.** Only HP offers easy migration of data from virtual tape to physical tape either for off-site/archiving of data for compliance reasons, or for use in tape initialization during replication setup, and in the disaster recovery process.
- **HP D2D can run up to eight replications concurrently,** allowing optimal use of the bandwidth available so your replication completes as soon as possible.
- **The D2D implementation offers link throttling** so you can determine what percent of the available link speed is to be utilized for replication. This prevents D2D replication swamping the inter-site link and maintains acceptable application performance across sites whilst replication is also taking place.
- **With many-to-one and N-way replication,** the D2D4100 devices can support up to 24:1 fan in from D2D2500s located at remote sites, allowing you to make replication very cost effective by consolidating up to 24 remote sites into a single replication target. Replication can also be arranged to go Site 1 to Site 2 to Site 3 to Site 1 in N-way fashion.
- **Minimizes the amount of data to replicate.** One of the features of HP D2D replication is that if the backup software is configured to append to cartridges rather than overwrite them, HP D2D only transmits manifests to match from the append point, which saves further bandwidth compared to matching the manifest for the full cartridge each time (i.e., only transmit the manifest for the new 'session' appended to the tape).
- **Price point.** HP D2D2500/4000/4100 Backup Systems support deduplication as standard and are better value than products from other leading competitors. Add to that competitive license costs for replication and the fact that you can replicate up to 16 D2D2500s into a single D2D4000 but only have to purchase a single replication license for the D2D4000. The final result is a very compelling cost structure for deduplication and replication.

## How it Works

HP D2D Backup Systems with HP Dynamic deduplication use a technology called hash-based chunking to save space. Every unique 4K of data within the D2D deduplication store has its own unique hash code. After the two D2Ds have been "initialized" with reference data (see [Implementing the Initialization Process](#)) then only new data created at the source site needs to replicate to the target site. This efficiency in understanding precisely which data needs to replicate can result in bandwidth savings in excess of 95% compared to having to transmit the full contents of a cartridge from the source site. In addition, there is some overhead of control data that also needs to pass across the replication link. This is known as manifest data, a final component of any hash codes that are not present on the remote site and may also need to be transferred. Typically the "overhead components" are less than 2% of the total virtual cartridge size to replicate.

## D2D Replication Implementation

The following must be performed to implement replication on D2D systems.

## Licensing

On HP D2D 2500/4000/4100 Backup Systems, Dynamic deduplication is included as standard. Replication is a licensable feature at extra cost. Only one license is required for the replication target.

## Implementing the Initialization Process

Before D2D replication can take place the reference data (the first full backup cycle performed with deduplication enabled) needs to be sent to the target. This initial reference copy of the data is large and could take a significant amount of time to transfer (weeks or months) if the WAN/LAN replication link was sized around the daily deduplicated data to be transferred. Once the target is initialized with the reference data the daily deduplication-enabled replication can commence over the WAN/LAN link. The following diagram shows the initialization options:

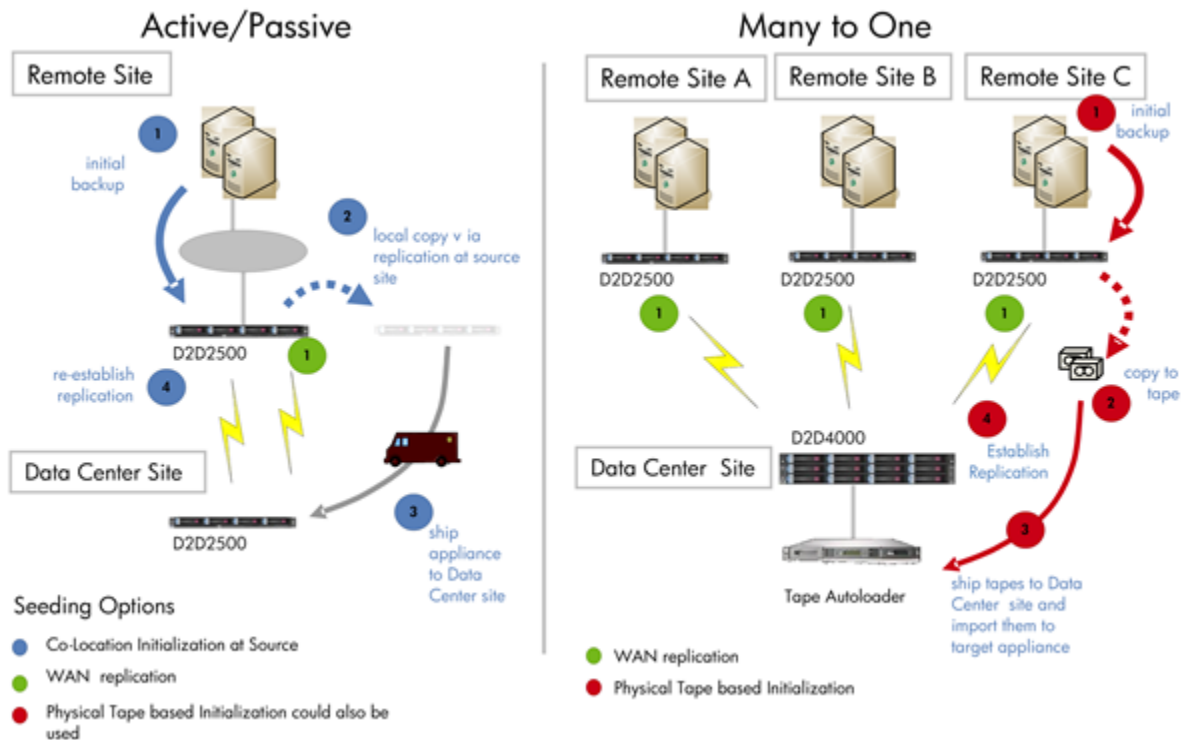


Figure 29 Options for Initialization the Initial Data

- Active/Passive initialization options:
  - Option 1: Co-locate the units and use a local high bandwidth (GbE) link, then separate units and ship one to the target site.
  - Option 2: If the link bandwidth is high enough and you are willing to wait several days then use the WAN replication link to initialize the target device.
  - Option 3: Use physical tape initialization. Physical tape library access is required at both sites. Export tapes from the source, import tapes into the target. During the import, the deduplication store is automatically created. Data is exported in backup software format.
- Many-to-one initialization options:

Co-location is probably not an option here because once remote sites are replicating to a single target device you do not want to disturb it to ship to a remote site to initialize it.

  - Option 1: If the link bandwidth is high enough and you are willing to wait several days then use the WAN replication link to initialize the target device.

- Option 2: Use physical tape initialization. Physical tape library access is required at both sites. Export tapes from the source, import tapes into the target. During the import, the deduplication store is automatically created. Data is exported in backup software format.

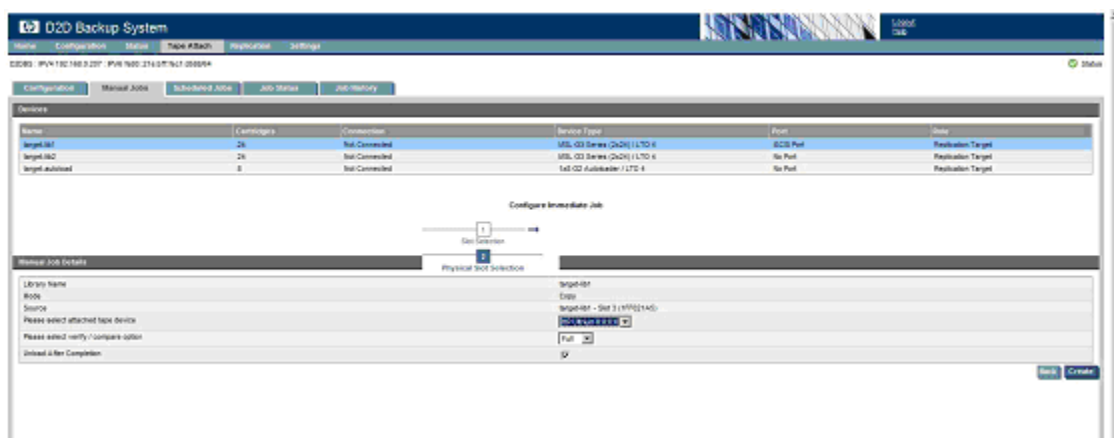
## Initialization — Co-location and WAN Replication

See [Configuring Replication](#) for the process of configuring replication on HP D2D devices. In this case the two D2D units can be directly connected through a GbE in a single location. The first replication to a target involving a full local copy will take some significant time depending on size of cartridge being replicated because the full set of hash codes and deduplication store have to replicate to the target.

## Physical Tape Initialization — Export and Import

With tape initialization:

- Attach a tape drive or MSL Tape library to the source D2D unit using the Tape Attach tab in the GUI; perform a scan to discover the device.
- Perform your first full backup to the source D2D device.
- Using the Manual jobs tab under tape attach select the slots on your D2D device that require copying and perform a virtual tape to physical tape copy.
- The virtual barcode on the source device is copied into the cartridge memory of the LTO physical media and retrieved during the import process, so even the virtual barcode is transferred correctly.
- Select “Create.”

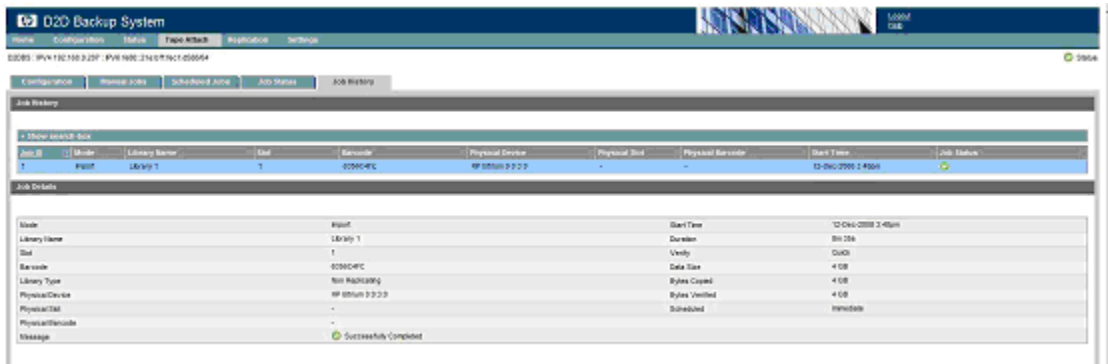


Once the tapes have been created and shipped to the remote site where the target D2D is located, the reverse process can take place using “import.”

### NOTE:

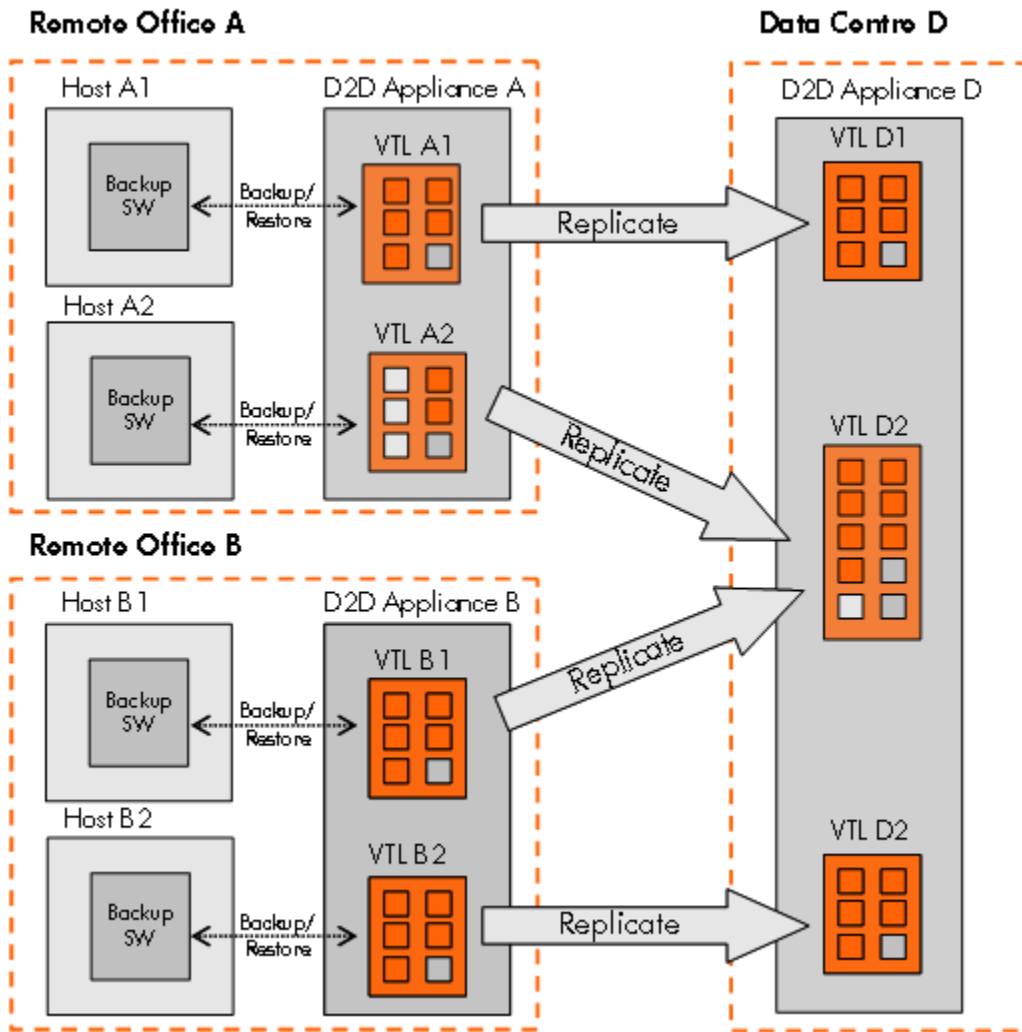
The import for the tape initialization process *must* take place before the replication slot mappings are established. (See [Configuring Replication](#).)

As part of the import process the hash codes and deduplication store are automatically created on the target device. The import is forced into the virtual mail slot of the virtual library on the target device and can then be moved into a usable slot.



## Replication Setup

Before you set up the D2D for replication it is important to understand the concept of slot mapping, because this is how replication is administered. It is also important to understand your cartridge rotation scheme because the replication is all predicated on knowing that weekly backups go to Tape X in Slot Y, that daily incrementals go to Tape A in Slot B and so on. Replication is triggered by a tape being ejected from a virtual tape drive in the D2D. HP D2D replication ensures that minimum bandwidth is used for incremental backups being replicated from the same tape, by only transferring the incremental manifests associated with each "session" of incremental backups appended to the same tape.



**Figure 30 Slot Mappings**

- A “slot mapping collection” is any number of slots within a source library that are configured to replicate to a target library.
- There is a maximum of *one* slot mapping collection per source library (you can add slots to mapping after initial creation).
- There is a maximum of *four* slot mapping collections per target library (one for the D2D2500).

The preceding example shows the flexibility of HP D2D replication technology.

In remote office A, daily “fulls” are taken each day so all six tapes in VTL A1 are required to replicate to an identical target device at data center D because this is considered critical data. All six slots must be mapped, but it is only when data changes on a cartridge in VTL A1 that replication will take place because replication is only “triggered” on cartridge eject. A similar situation applies to VTL B2 in remote office B.

However, for host A2 only the weekly full and monthly full require replicating. Rather than create a target replication library at remote site D for just two cartridges, you can combine this with the five cartridges from source library B1. This is because there is a limit of four slot mappings on the target device. So A2 and B1 have to share a slot mapping on VTL D2. Notice how at data center D, target replication libraries can be shared among sources, allowing for consolidation.

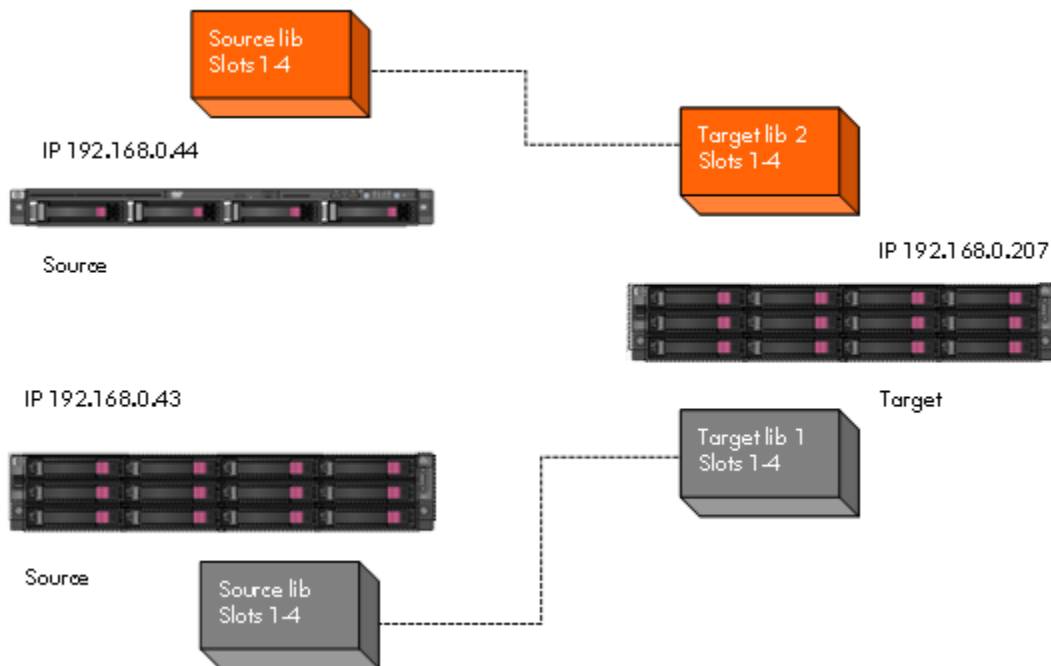
## Configuring Replication

The following steps will configure D2D replication.

If using tape initialization, the initialization process has to happen before the replication is configured. Prior to configuring replication, you must enter the replication license on the target device.

Use the diagram below when reviewing the steps to configuring replication. This example uses a D2D2500 source library IP 192.168.0.44 and a D2D4000 Fibre Channel target library at 192.168.0.207. The target already has another source mapped into it from D2D4000 at 192.168.0.43.

All the network parameters must be set correctly before starting the replication configuration. The IP addresses are set as part of the initial D2D installation using either a "Discovery CD" or by connecting a monitor and keyboard to the rear of the D2D unit. You must set your DNS and gateway address. DHCP is permitted but the network administrator should ensure that there is a permanent lease assigned.



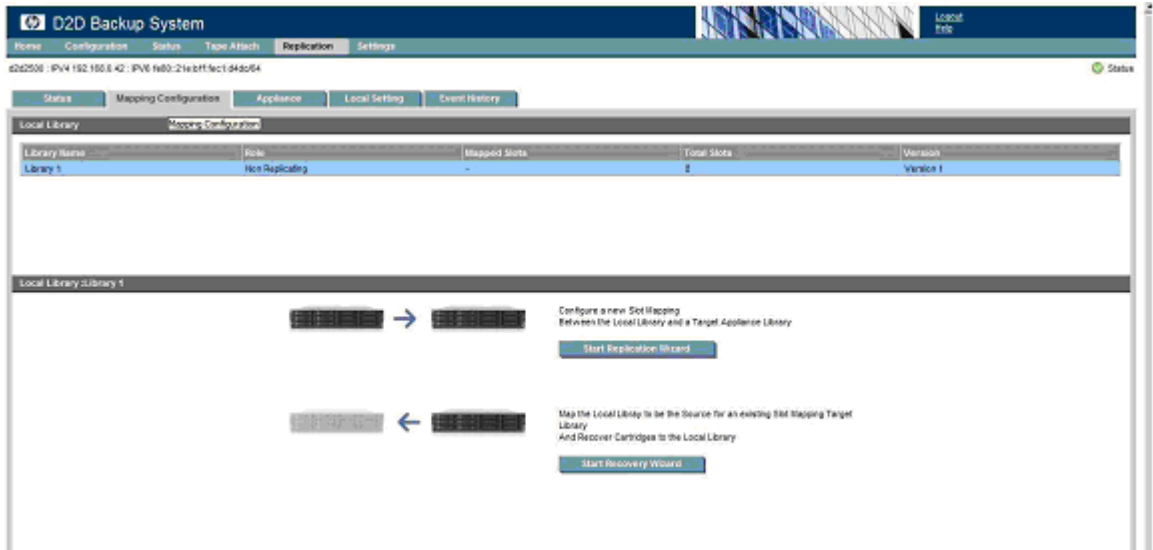
**Figure 31 Example Schematic**

1. Start the replication wizard.

- a. Select the Replication tab in D2D GUI and then the Mapping Configuration tab. The wizard can then be initiated. All the configuration talks place on the source library. The only actions necessary on the target library are to configure it with an IP address. Note also that this is the starting point for the recovery wizard. Library 1 is the new source library in the D2D2500.

Notice how the source library is initially “non-replicating” until it becomes part of a replication pair.

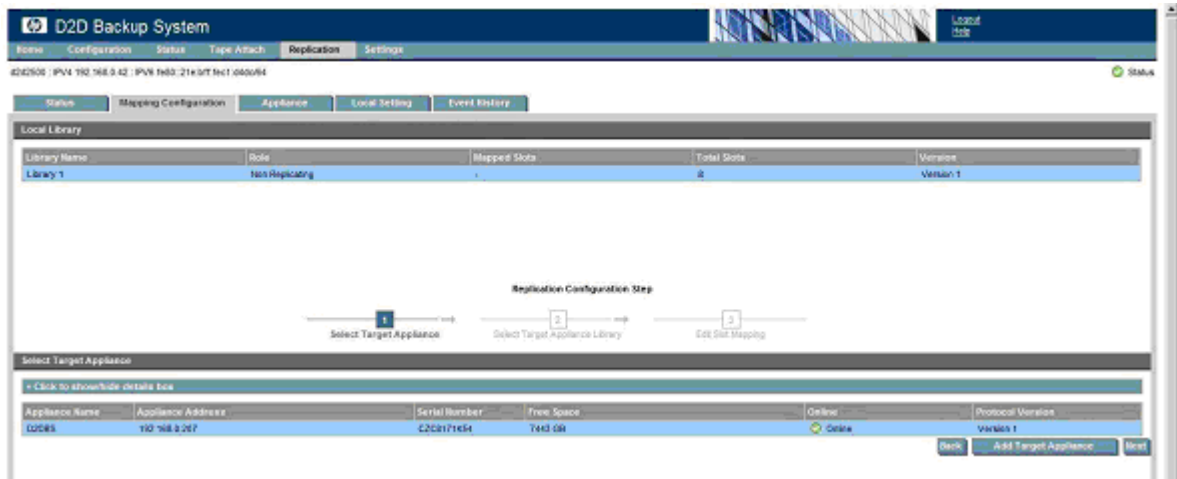
- b. Select **Start Replication Wizard**.



2. Add a target device to source device.

- a. Using the D2D replication wizard you can set the IP address of the target library device. You must know the target library IP address. In this case, 192.168.0.207.
- b. Input the target appliance IP address.
- c. Select **Add Target Appliance**.

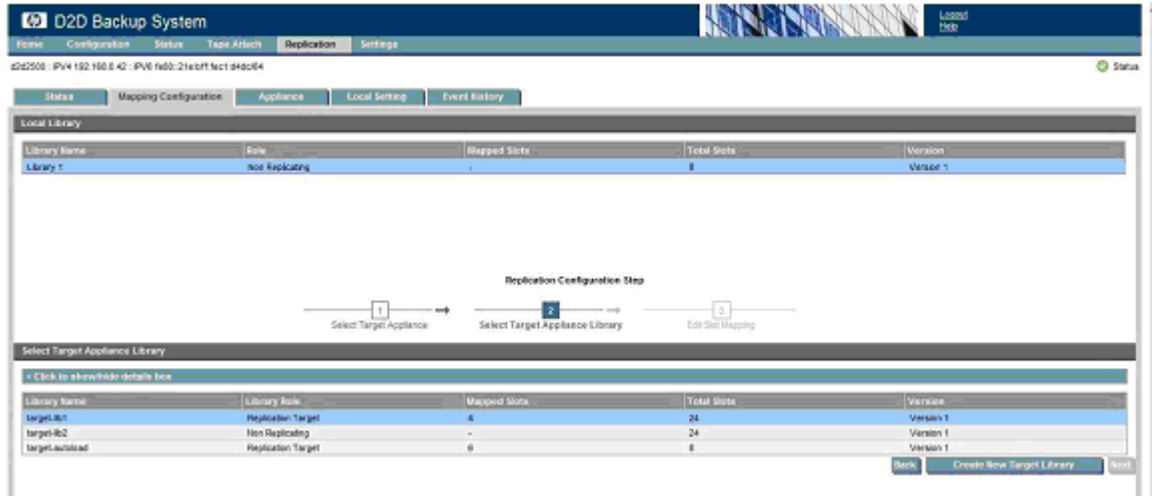
The target library appears to the source and is online.



3. Assign target slots.

a. Select **Next**.

Use part of an existing library on the target or create a new library on the target directly from the source for the purpose of replication. In this case map four slots from the local (source) library onto target Lib 2 on the D2D4000 which has already been created. Alternatively you can create a brand new target library by selecting **Create New Target Library**.



If a new target library is created from the source the bar codes match identically source and target. If the target library exists already the bar codes do not match.

Note how target-lib2 is non-replicating until the slot mappings have been configured.

b. Select target-lib2 as the target replication library to use.

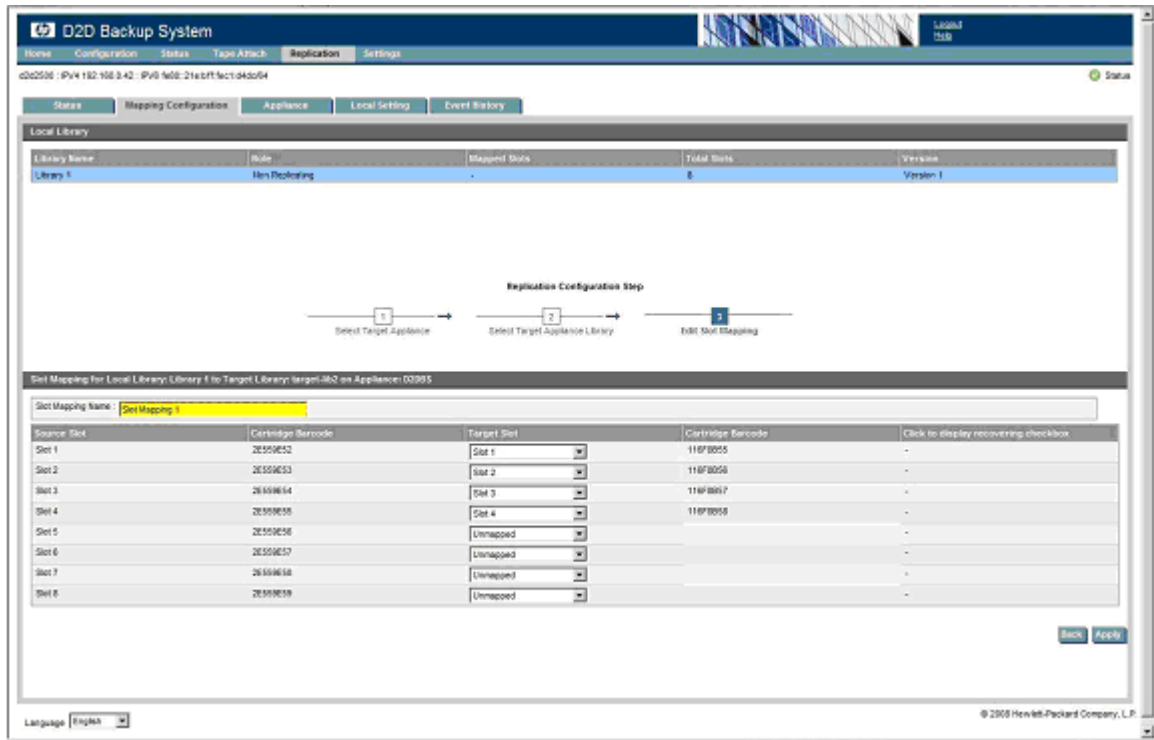
c. Select **Next**.



#### 4. Configure slot mappings.

Set up the slot mappings from the source library slots to the target-lib2 slots and give the slot mappings a name.

- a. Provide a name for the slot mapping.
- b. Select **Apply**.



Notice how the virtual bar codes at both the source and the target are displayed.

In this case there is a 1:1 mapping from the source onto the target replication library using the drop down lists.

## 5. Configure throttling.

- a. Configure the necessary Blackout windows and Bandwidth throttling settings unique to your requirements. This example does not allow replication between 08:00 and 18:00 during the week.
- b. Select **Finish**.

The screenshot shows the 'Replication' settings page for a D2D Backup System. The 'Appliance Name' is 'D2D500'. The 'Configure Blackout Windows' section has a table with the following data:

Day	Apply Time Restriction to Jobs	Do Not Allow to Start Between	Do Not Allow to Start Between
Sunday	<input type="checkbox"/>	00:00	00:00
Monday	<input checked="" type="checkbox"/>	08:00	18:00
Tuesday	<input checked="" type="checkbox"/>	08:00	18:00
Wednesday	<input checked="" type="checkbox"/>	08:00	18:00
Thursday	<input checked="" type="checkbox"/>	08:00	18:00
Friday	<input checked="" type="checkbox"/>	08:00	18:00
Saturday	<input type="checkbox"/>	00:00	00:00

The 'Bandwidth Limiting' section shows 'Bandwidth Limiting Enabled' as  (unchecked), 'Bandwidth Limiting (Kbps)' as 0, 'Maximum WAN Usage (%)' as 100, and 'WAN Speed' as 1Mbps.

The 'Configure Replication Ports' section shows 'Command Protocol Port Number' as 4833 and 'Data Protocol Port Number' as 4834.

Note in this example the bandwidth is not limited so there is 100% utilization of the available link.

## Reporting

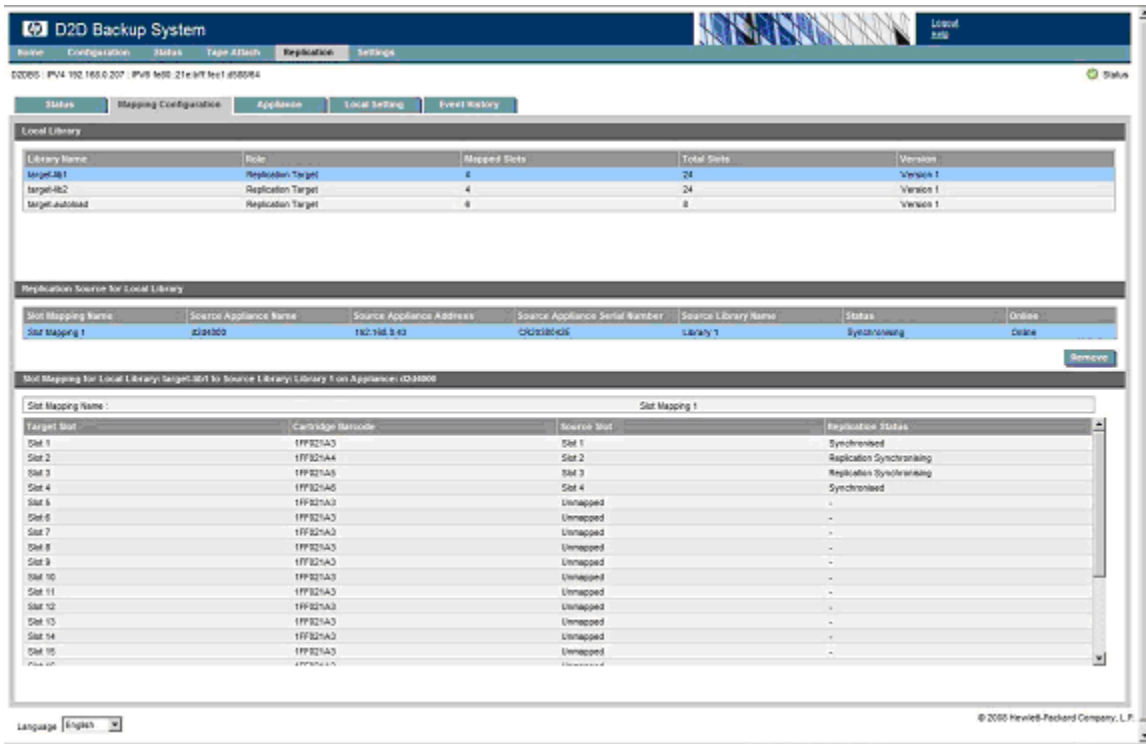
It is possible to view the replication status from both the source and target D2D devices.

### Target Device Reporting

On the target device you can view the target replication library to see the “synchronization status” of each virtual cartridge. For example, the screenshot below shows a D2D4000 target replication library with three replication target libraries set up (Target-lib1, Target-lib2, Target-autoload). The top part of the screen shows how many of the slots in the target libraries are already mapped to different sources.

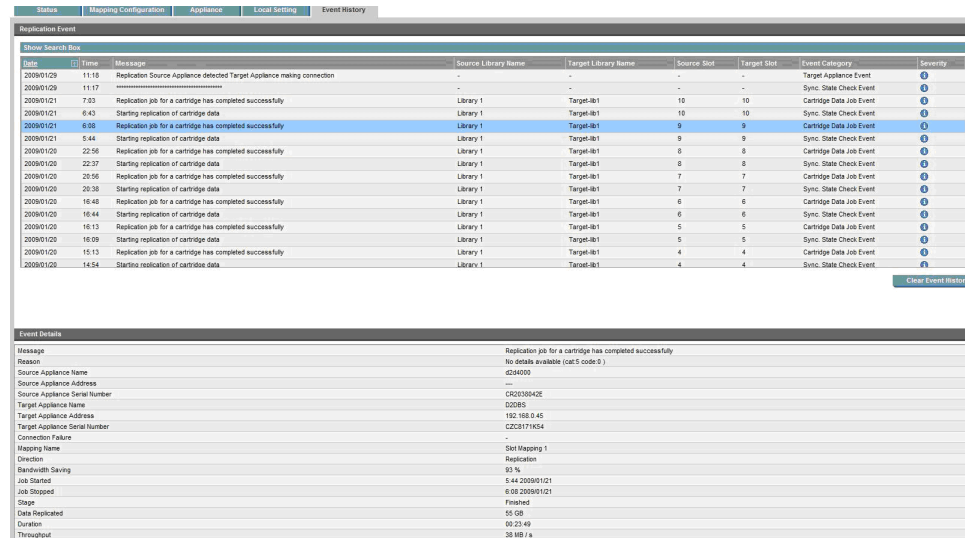
Imagine that you just performed a backup to the four cartridges in slots 1-4 on source library 192.168.0.43. As the cartridges are ejected from the drive, replication starts and reports “synchronized” when the replication of a particular cartridge is complete.

By clicking on target lib 1, then on slot mapping 1 (for target lib 1) you can see the status of the four cartridges configured for replication; two are already synchronized and two are in the processing of replicating (replication synchronizing).



## Source Device Reporting

From a source perspective, it is important to know if your cartridges have replicated successfully. From each source library device on the replication tab there is an “Event History” tab which displays all the replication activity concerning that source in a time stamped Event History log. For the source library configured in the example above, you can see the stages and times as each of the four mapped cartridges is replicated. For the activity selected, the details are shown in the bottom of the screen shot. Using this Event History you can determine how long a particular cartridge took to replicate, average replication speed, and so on.



## Design Considerations

Because every enterprise backup environment is different, there is no single answer for optimizing replication. This section discusses replication design considerations to help you optimize replication in your environment.

### Link Sizing

Knowing the rate of change of data on your systems and hence being able to size the replication link efficiently is the key to a successful implementation. Size the link too small and you cannot off-site all your data in a reasonable time. Size the link too big and you are paying for bandwidth you may not be able to use. The most likely implementation scenario is either to utilize unused bandwidth in an existing link or to provision a bandwidth increase on an existing link to accommodate virtual tape replication. For most users somewhere between a 2 Mb/sec and 50 Mb/sec link should be sufficient for the amounts of typically data transferred using HP D2D replication. This section focuses on understanding how to size the replication link bandwidth required.

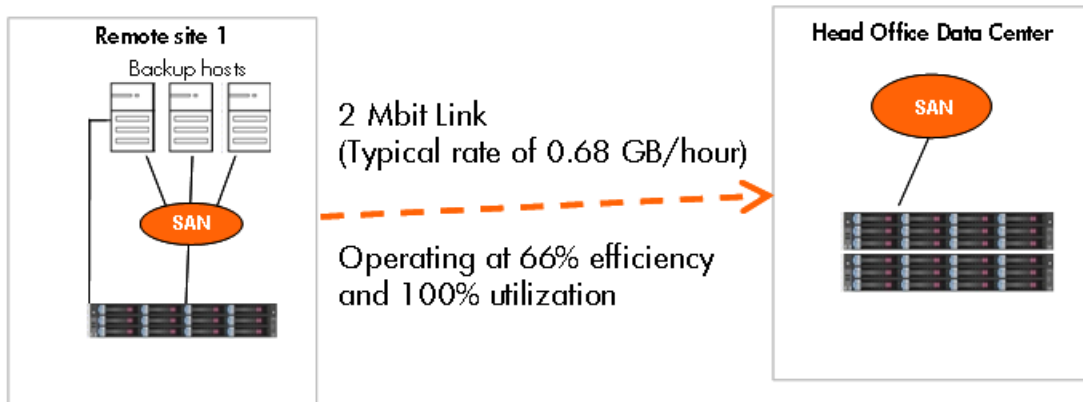
The size of the replication link will depend on the usage model for replication.

- Backup and replication take place in a fixed window (up to 12 hours) preferably when the WAN link is not heavily used by applications (i.e., outside business hours). The replication process has almost the full bandwidth of the WAN link available to it. Within the HP D2D Backup System configuration for replication you can set a “blackout window” during which time replication is NOT allowed to operate.
- The business runs 24x7 and there are no fixed windows in which to do replication. In this example the replication can be limited and the HP D2D devices can be configured to only use a percentage of the available bandwidth so as not to adversely affect application performance. The replication will take longer using WAN link throttling.

HP also provides a Backup Sizing Tool to help you size any backup solution including deduplication and replication: <http://www.hp.com/go/storageworks/sizer>.

Sizing the bandwidth required is best illustrated by some examples:

- 100 GB virtual cartridge replication over a 2 Mb/sec link with 100% bandwidth utilization, data off-sited in a 12 hour fixed window.
- 100 GB virtual cartridge replication over a 2 Mb/sec link with 25% bandwidth utilization, data off-sited over 24 hours.
- 500 GB virtual cartridge replication over a 2 Mbit/sec link, data off-sited in a 12 hour fixed window.
- 500 GB virtual cartridge replication over a 10 Mbit/sec link with 20% bandwidth utilization, data off-sited over 24 hours.
- Many-to-one example with mixed WAN speeds.



**Figure 32 Example 1: 100 GB Virtual Cartridge Replication over a 2 Mb/sec Link with 100% Bandwidth Utilization**

Calculating data to replicate:

- 100 GB daily backup to replicate within a 12-hour fixed window (between 20:00 and 08:00).
- 1% daily change rate = 0.66 GB compressed
- Overhead (manifest, imperfect match, etc.) is 1.5% = 1.5 GB
- **Total data to transmit is 2.16 GB**

Result:

- Initial 100 GB backup complete in 30 minutes.
- **Replication time for 2.16 GB is 3.46 hours over this link at 100% bandwidth utilization.**

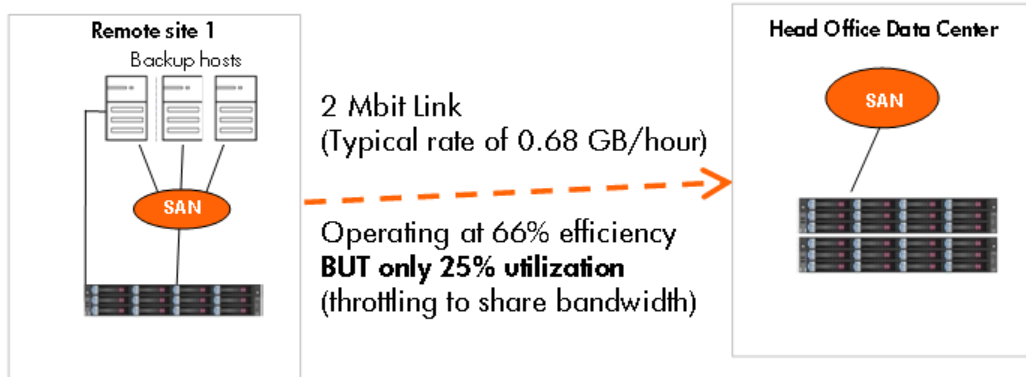
The above example shows a 100 GB tape to replicate daily between sites with a Symmetric Digital Subscriber link (SDSL) between the two sites.

There are several factors involved in sizing the link:

- Full backup size in this example is 100 GB.
- Daily rate of change of data at primary site, in this example 1%.
- Replication required in fixed window or in shared bandwidth mode using bandwidth limiting (throttling).
- New data being replicated is compressed, communication overhead data (manifest and imperfect hash code matches) are already compressed.
- Communication overheads typically add up to about 1.5% of the full backup capacity.
- WAN links are not 100% efficient because of TCP/IP overheads contention and so on. This example assumes 66% link efficiency (typical).

The net result is that the system can replicate a total of 2.16 GB in 3.64 hours with full bandwidth available at 66% efficiency.

To replicate in shared bandwidth mode, perhaps because this business is a 24x7 operation, then set the D2D device up to use no more than 25% of the bandwidth available. In this case the replication will take 14.4 hours which is acceptable to most businesses.



**Figure 33 Example 2: 100 GB Virtual Cartridge Replication over a 2 Mb/sec Link with 25% Bandwidth Utilization**

Calculating data to replicate:

- The same data calculations as the example in [Figure 32](#).
- **The total data to transmit is 2.16 GB.**

Results:

- Initial 100 GB backup complete in 30 minutes.
- **Replication time for 2.16 GB is 14.4 hours over this link at 25% bandwidth utilization.**


The time required for replication is derived from the equation:

$$\frac{\text{New replication data (x comp ratio) + manifest (0.5\% full backup size) + New hash codes (1\% of full backup size)}}{\text{Rated WAN link speed} \times \text{Efficiency} \times \text{Throttling Factor}}$$

**NOTE:**

To set this in context, in a high bandwidth model you would have to replicate the whole 100 GB between sites on a 2 Mb/sec link which would take a long time. Instead, using deduplication and replication, you only have to replicate 2.16 GB of data to achieve the same result.

Now consider a larger backup requirement where 500 GB of backup requires replicating each day.



**Figure 34 Example 3: 500 GB Virtual Cartridge Replication over a 2 Mbit/sec Link**

Calculating data to replicate:

- 500 GB daily backup to replicate within a 12-hour fixed window (between 20:00 and 08:00).
- 1% daily change rate = 3.3 GB compressed
- Overhead (manifest, imperfect match, etc.) is 1.5% = 7.5 GB
- **Total data to transmit is 10.8 GB**

Result:

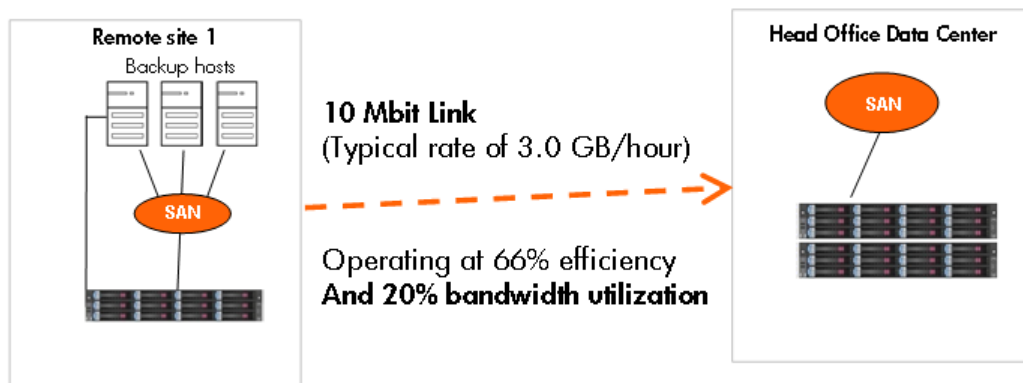
- Initial 500 GB backup complete in 2.75 hours.
- **Replication time for 10.8 GB is 18.1 hours over this link at 100% bandwidth utilization – but this does not meet the service level agreement.**

In this example, start off with the an 2 Mbit SDSL link. You soon realize that the 10.8 GB of data to replicate even using 100% of the available bandwidth is 18.1 hours which is outside most users' fixed window of available time. It would take  $4 \times 18.1$  hours using 25% throttling of the 2 Mb/sec link. Clearly a link speed upgrade is required to meet either of your criteria.

Upgrading to a 10 Mbit/sec link resolves the problem.

The replication time now within a fixed window is  $18.1/5 = 3.62$  hours.

The replication time now with shared bandwidth mode at 20% = 18.1 hours (within the required 24 hours required).



**Figure 35 Example 4: 500 GB Virtual Cartridge Replication over a 10 Mbit/sec Link**

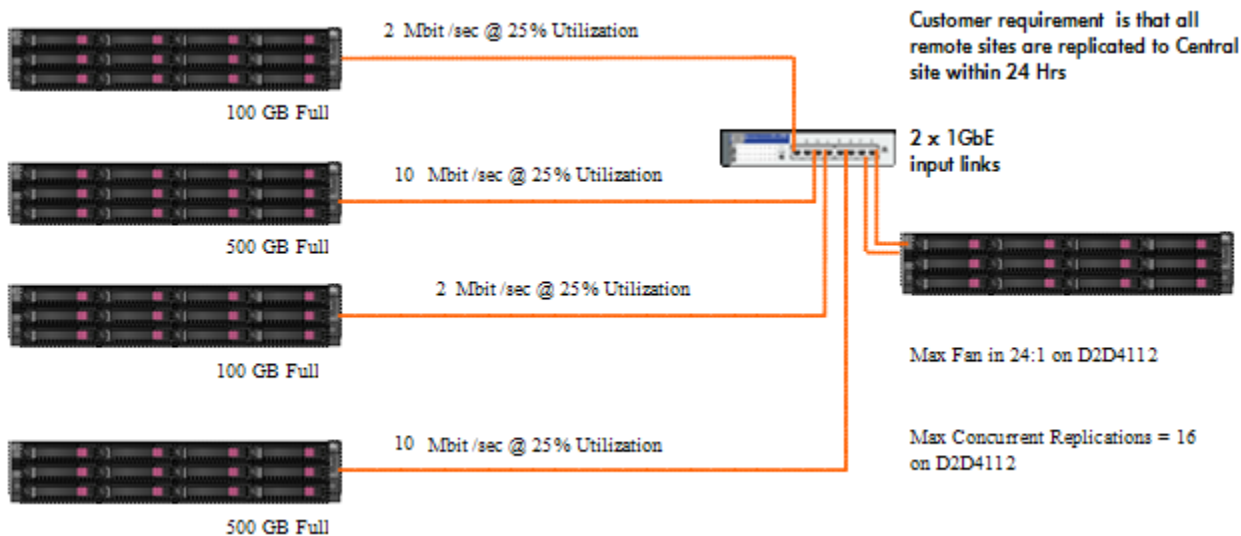
Calculating data to replicate:

- The same data calculations as the example in [Figure 34](#).
- **The total data to transmit is 10.8 GB.**

Results:

- Initial 500 GB backup complete in 2.75 hours.
- **Replication time for 10.8 GB is 18.1 hours over this link at 25% bandwidth utilization.**

Finally, look at a many-to-one example which is a combination of the above scenarios.



**Figure 36 Many-to-One Example using Mixed WAN Link Speeds**

Here are four remote sites replicating 100 GB and 500 GB backups daily into a central D2D4000.

Note also that the central D2D4000 has 2 x 1 GbE connections from the central router for link aggregation.

Follow these guidelines for the best approach to a many-to-one scenario:

- Size each of the individual links separately.
- If possible use the “blackout window” settings to stagger the replications throughout the day to ensure there are not greater than eight replications active at any one time.
- Be careful with bandwidth limiting values. There is a minimum of 512 KB/sec and the throttling value selected should not cause impaired application response times.
- If possible, provision some extra bandwidth above what is required for daily replication to allow for “catch-up” mode if replications are paused/stopped due to a link failure.
- Wherever possible do a proof of concept before implementing into production.



Table 7 shows the typical ongoing replication times depending on link speeds and how long the initialization process will take depending on if initialization is done via co-location or the WAN link. Use this table as a reference guide.

**Table 7 Sample Initialization and Replication Times**

Initial backup size in GB	Link speed in Mb/sec	Backup time at 25 MB/sec in hours	Initial. time via WAN in hours (66% efficiency uncontented), compression at 1.5:1, deduplication at 1.5:1, local backups in progress	Initial. time via co-location in hours, replication on GbE	Ongoing replication time in hours (1% Data change rate) – uncontented link (100% utilization)	Limiting factor in replication	Ongoing replication time in hours (1% data change rate) – 25% utilization link	Limiting factor in replication
50	2	0.6	37.1	0.5	1.5	WAN link	6.1	WAN link
100		1.1	74.2	1.0	3.1		12.3	
500		5.6	370.9*	5.0	15.3*		61.3*	
1000		11.1	741.8*	9.9	30.6*		122.6*	
100	10	1.1	14.8	1.0	0.6	WAN link	2.5	WAN link
500		5.6	74.2	5.0	3.1		12.3	
1000		11.1	148.4*	9.9	6.1		24.5	
100	15	1.1	9.9	1.0	0.4	WAN link/D2D appliance	1.6	WAN link
500		5.6	49.5	5.0	2.1		8.2	
1000		11.1	98.9*	9.9	4.3		16.3	
100	30	1.1	4.9	1.0	0.4	D2D appliance	0.8	WAN link
500		5.6	24.7	5.0	2.1		4.1	
1000		11.1	49.5	9.9	4.3		8.2	
100	50	1.1	3.0	1.0	0.4	D2D appliance	0.5	WAN link/D2D appliance

Initial backup size in GB	Link speed in Mb/sec	Backup time at 25 MB/sec in hours	Initial. time via WAN in hours (66% efficiency uncontented), compression at 1.5:1, deduplication at 1.5:1, local backups in progress	Initial. time via co-location in hours, replication on GbE	Ongoing replication time in hours (1% Data change rate) – uncontented link (100% utilization)	Limiting factor in replication	Ongoing replication time in hours (1% data change rate) – 25% utilization link	Limiting factor in replication
500		5.6	14.8	5.0	2.1		2.5	
1000		11.1	29.7	9.9	4.3		4.9	
100	1000	1.1	1.0	1.0	0.4	D2D appliance	0.4	D2D appliance
500		5.6	5.0	5.0	2.1		2.1	
1000		11.1	9.9	9.9	4.3		4.3	
* = Unacceptable usage model								

Depending on the size of the full backup and the link speeds deployed, D2D can deliver the replication in either a 12 hour window (at 100% link utilization) or 24 hour window (with 25% link utilization). Areas marked with an asterisk are outside of these stated usage models. A 1000 Mbit/sec link represents the link speed used when co-locating the devices for initialization.

## Telco Provisioning — Overview of Inter-Site Links

This section provides an overview of how Telcos provide WAN links. With HP D2D device replication the typical parameters required to be specified to the Telco are:

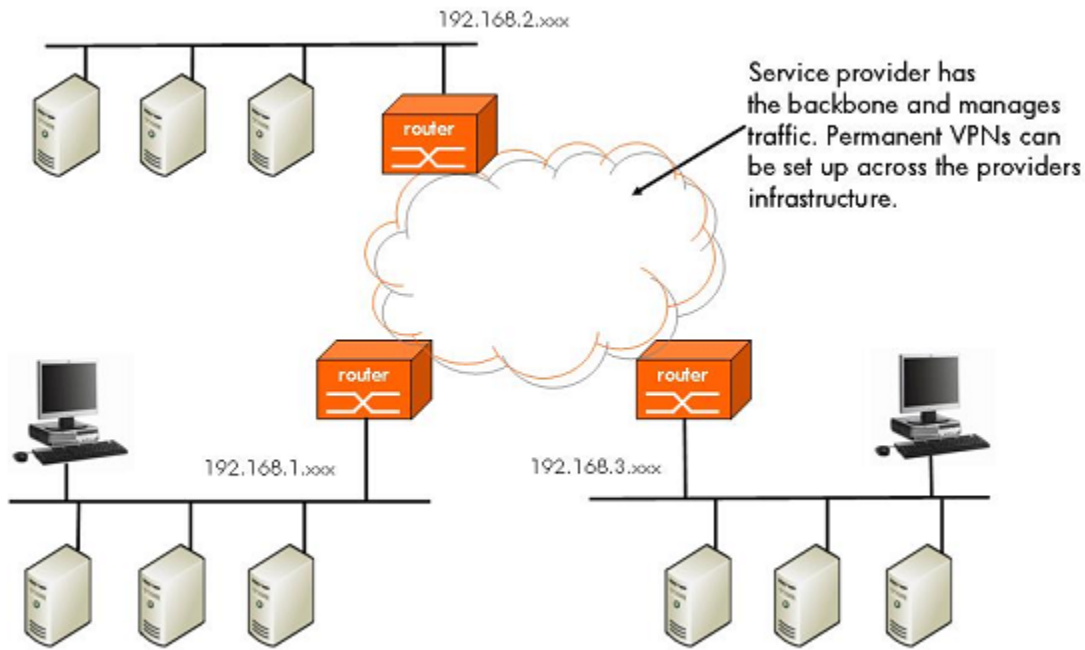
- Physical location with respect to a Telco Point of Presence. It is often said that it is the last 100 m of a WAN link where all the cost is incurred. A Point of Presence (PoP) is where the Telco has a connection point to their Core infrastructure. Generally this is housed in the same building as their telephone exchanges. Sometimes laying in the copper or fibre from a Telco's point of presence to your building is where the majority of the installation costs are incurred. The physical location of your office or disaster recovery site will definitely have an effect of the link price. Needless to say if your premises already have a link in place then generally increasing the bandwidth of an existing link is relatively easy. Most Telcos require a business to be within 25 km of a point of presence for Ethernet connections.
- Speed required—Most D2D replication implementations will not require links above 50 Mbits/sec. Many will be able to function adequately on speeds as low as 2 Mbit/sec links.

- Resilience—dependent on business criticality of replication. Even with the most basic package (single line unprotected) most Telcos will guarantee 99.9% availability. Above this you can opt for master and standby paths (only one active at a time) that link to a single NTE (Network Termination Equipment) at your premises. If one path fails YOU are responsible for moving your equipment over to the standby path. Finally you can have 2 x NTE at your site with Master and Standby paths and fail-over can be automated.
- Distance—inside a single country for example UK, there is generally no additional charge for distance (unless very high bandwidth requirements > 1 Gb/sec) but in large countries where a LATA (Local Access and transport and area) boundary is crossed then additional costs (to use different Telcos lines in different states) to establish a point to point connection will be incurred. Inter country or Intercontinental distances require a more consultative engagement with your Telco Provider.
- Contention/Uncontended—Contended links are those shared with other users and bandwidth cannot always be guaranteed. Even some business services have a certain amount of contention. For example, BT Etherflow Standard Class of Service offers 80% bandwidth guaranteed with 20% contented. Whereas BT Ethernet Premium Class of Service allows 100% uncontended bandwidth. HP recommend an 80% uncontended or above link for D2D replication.
- What is not acceptable is to try and run an HP D2D Replication solution on a Home broadband type of link where contention rates of 20:1 or 50:1 can be present. Upgrade to business broadband Class of Service. This is a way of prioritizing traffic on the Telco backbone. Voice over IP for example requires a high priority to maintain coherence. Replication is not considered business critical in this respect and so typically default class of service is acceptable. Above default class there will be a premium class (Assured Forwarding) where data queues in an MPLS environment the data are marked with an identifier to allow it to have some level of priority within the Telco's infrastructure. Finally data such as VOIP is given a strict priority queue identifier so it always takes precedence; this is sometimes known as expedited forwarding. Class of Service may also have an effect on the Latency (see below) of your environment.
- Security. Another reason for choosing an uncontended service is that your replication data is not mixed with other users. Today's Ethernet connections typically take the form of a VPN (Virtual private network) so an element of securing exists within the transfer protocol already. A further enhancement if security is a real issue should be to consider implementing IPSEC (Internet Protocol security). HP D2D will support this in future releases. The addition of IPSEC security however may reduce overall replication throughput because of the increased overhead involved in encryption.
- Latencies. These are delays introduced in sending data over a distance or delays involved in routers, etc., used to direct the data paths. A round trip delay (RTD) in milliseconds is a measure used to indicate latency.

The longer the latency the longer the replication can take, because the HP D2D will only put so much information on the line before waiting for a response from the target device. HP has tested D2D with latencies up to 280 ms with no issues.

A Typical Telco example, British Telecom provides a Core Network performance SLA on Round Trip Delay, this may vary with class of service. For example:

- Expedite Forwarding 20 ms (RTD)
  - Assured Forwarding 23 ms (RTD)
  - Default 30 ms (RTD)
- Packet Delivery: This is a measure of the packet loss due to noise or transients, packets are retransmitted which again may slow replication down if the link is particularly bad, but most Telcos offer an SLA on packet delivery of above 99.8%.
  - Most Telcos using the latest Networking technology such as MPLS deliver WAN links via a "cloud" of VPN connections as if their core infrastructure was a massive router.



**Figure 37 A Basic WAN using MPLS**

### Some Idea of Telco Costs

It is very difficult to generalize on typical link costs because of the variables listed above, but as a practical example a major UK Telco operator quoted costs for links at various speeds for HP Sites in Bristol, UK and Warrington, UK, for a distance of around 250 km. This assumes they are within the 25 km of a point of presence and no additional construction is required to lay cables/Fibre.

HP D2D will typically require link speeds from 2 to 50 Mb/sec depending on how much data needs to replicate.

**Table 8 Telco Prices for a 3-year Contract Including Installation (No Existing Links in Place)**

Bandwidth Mbit/sec	Installation Cost (UK £)	Rental per year (UK £)	Average per year (over 3 years) – (UK £)
2	7663	13004	15558
10	7532	15732	18242
100	11787	24426	28355
200	23172	45560	53284
500	23172	55755	63479
1000	23172	72781	80505

**Table 9 Prices of a Major US Telco for HP Facilities in the Houston, Texas Area**

Bandwidth Mbit/sec <sup>1</sup>	Installation Cost (US \$)	Rental per year (US \$)	Average per year (over 3 years) – (US \$)
2	2000	11400	12066
10	2000	18000	18666
100	3000	36000	37000
1000	4000	72000	73333

<sup>1</sup>This company did not offer 200 and 500 Mb/s lines.

**Table 10 Prices of a Telco Provider in Colorado for HP Sites Located around 250 km (175 miles) Apart**

Bandwidth Mbit/sec	Installation Cost (US \$)	Rental per year (US \$) <sup>1</sup>	Average per year (over 3 years) – (US \$)
5	600	8160	8360
10	600	9180	9380
100	600	15300	15500
200	1200	21300	21700
500	1200	39324	39724
1000	1200	69360	69760

<sup>1</sup>These prices are all within the same local access and transport area.

**Table 11 Prices for Point to Point Connects from HP Facilities based in Singapore**

Bandwidth Mbit/sec	Installation Cost (S \$)	Rental per year (S \$)	Average per year (over 3 years) – (S \$)
5	4000	26400	27733
10	4000	43200	44533
100	10000	96000	99333
200	10000	144000	21700
500	10000	216000	219333
1000	10000	336000	339333

There is clearly a wide variety of costs. Be sure to have a thorough discussion with your Telco provider to fully investigate all the options available.

## Telco Terminology and Branding for WAN Services

Below are a sample of the services provided by BT (British Telecom) in the UK for inter-site WAN replication links:

- BTnet Premium (Internet access)

This is a flexible bandwidth solution with speeds between 64 Kbps and 1 Gbps. The router is provided by BT and is designed primarily for internet access. The connection is typically to a BT point of presence POP and optional features such as failover and load balancing are available. The connection to the BT POP would be copper wire for the lower speeds (SDSL type connection) (64 Kbps, 2 Mbps) and optical fibre at the higher speeds.

A VPN would need to be set up to connect to another sites subnet. This would be necessary for a D2D replication connection. Contention may be an issue with this type of service.

- BT Etherflow

This service is a high-performance WAN services delivered as an ethernet connection but with the end user retaining end to end control over the IP architecture and the ability to support non-IP applications such as AppleTalk, IPX, DECNet and SNA.

Speeds are available between 1 Mbps and 1 Gbps. Building blocks are available to create VPNs and separate traffic classes. (Traffic could be separated for different groups within a company or for IP telephony.)

This service is optical fibre based. This service could be used for D2D replication but would typically provide other services on site already (connections to servers located at a central site).

This technology also runs on BT 21st century Network—MPLS.

- BT IP Clear

This service provides network integration using the Telcos MPLS infrastructure and is suitable for the larger corporate network. It provides a full range of access and is delivered to user locations by leased line from 128 Kbps up to 1 Gbps through optical fibre. All routing and VPNs are handled by the MPLS core and typically a user would be provided with an edge router. Voice circuits (IP telephony) can be fully integrated into this service.

- BT Ethernet extension services

This service provides high speed dedicated links using permanent optical fibre connections. LAN speed from 10 Mbps to 1 Gbps is supported. Alternative protocols can be used such as ATM and fibre channel.

Using this service the LAN could span several sites with no VPNs.

There is a distance limit of 25 km between sites (can be expensive).

**Table 12 Sample British Telcom Services and Speeds**

Service	Available speeds	Suitable for D2D low bandwidth replication
BTnet Premium	64 kbps to 1 Gbps	Yes (2 Mbps upwards) need VPN and fixed IP. Contention?
BT Etherflow	1 Mbps to 1 Gbps	Yes—2 Mbps upwards
BT IP Clear	128 Kbps to 1 Gbps	Yes—2 Mbps upwards
BT Ethernet Extension	10 Mbps to 1 Gbps	Yes but limited distance and likely to be expensive

When to use BT Etherflow and when to use BT IP Clear?

- BT Etherflow is the right choice for users who want an Ethernet (Layer 2) hand-off and seek to retain end to end ownership of their IP Architecture. It best suits users with a small numbers of high bandwidth sites.
- IP Clear is the right choice for users who want an IP (Layer 3) hand off and wish to share their IP architecture with BT. It offers any to any connectivity and is best suited to users with a large number of low bandwidth sites.

## D2D Replication Data Recovery Options

Replication on D2D enables easier recovery of data in the event of a major site disaster. The data is not instantly available but has to be recovered through a standard restore process using a backup application. These scenarios look at what happens in the event of a total remote site disaster and lose the server and the D2D appliance that was protecting it.

In the event of a total disaster at a remote site where the D2D Backup System and the servers it is protecting are destroyed or damaged, the replicated data at the disaster recovery (target) site can be accessed by means of a backup server on the remote site and transferred either to an application or onto physical tape for distribution. Alternatively, when the damaged site is repaired the Virtual Tape Library at the disaster recovery site can be used to “reverse replicate” to the Virtual Tape Library at the previously damaged site.

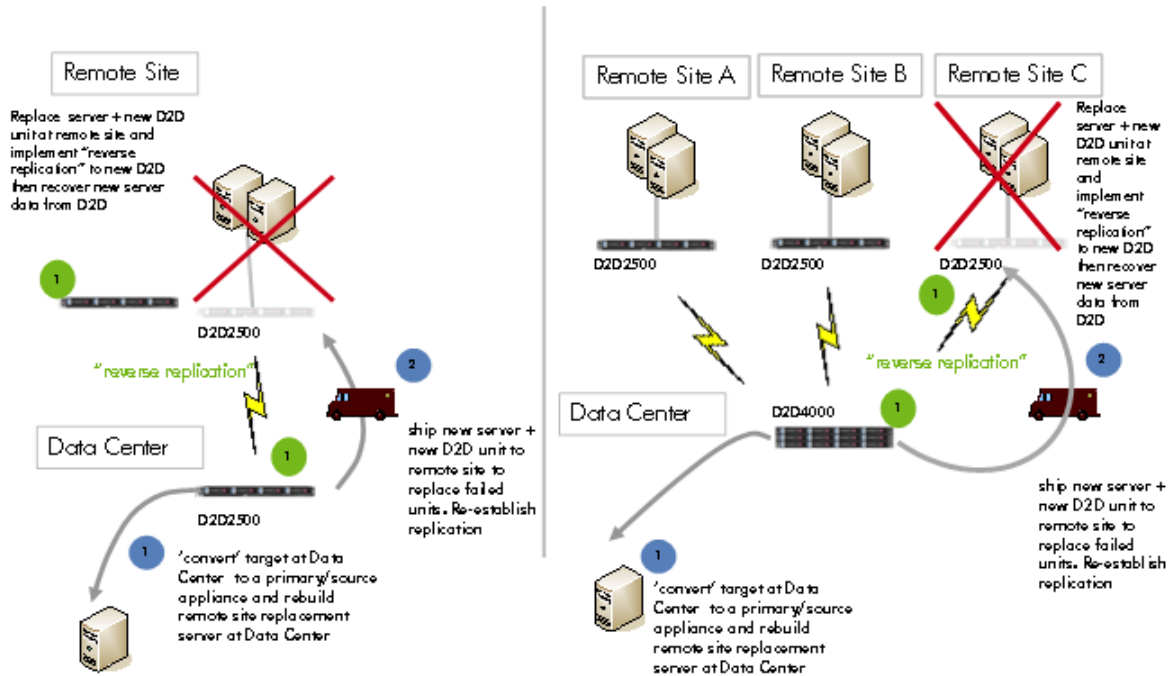


Figure 38 Recovery Options — Data Center Rebuild and Reverse Replication

### Restore Directly from the D2D Target Device

The first method of recovery uses the fact that there are probably more trained IT personnel at the data center and equipment such as the spare servers may well be kept there. The HP D2D system with replication can act as a source device, a target device or a non-replicating library (a library that can be connected to a system and data recovered directly from the device). You can convert the target library into a non-replicating library (by deleting the relevant replication slot mappings), construct a new server at the data center, hook up the D2D device, and recover the unique server data. Once complete, the new server and a new D2D can be shipped to the remote site and replication reestablished.

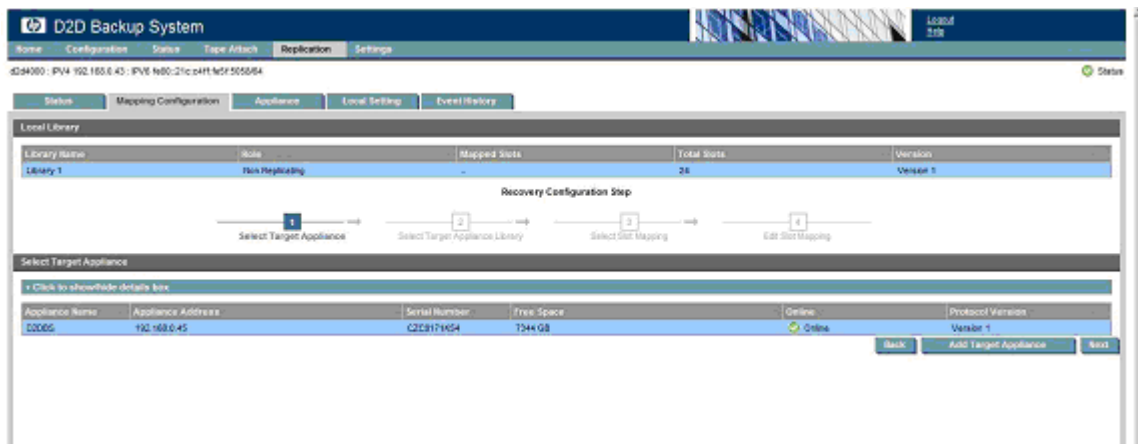
### Reverse Replication on the D2D

A second method of recovering is to perform a function called reverse replication. The data is copied back from the target to the new source appliance at the remote site. When this is complete, the remote site can be rebuilt and the server recovered from the data now available on the remote site. The effect

of reverse replication is the same as if initializing the device at the remote site and so depends on the link speed and amount of critical data to be reverse replicated. Once the server is recovered and able to perform backups again, configure replication back to the target device in the data center all over again. This option is easy to implement even in many-to-one replication scenarios.

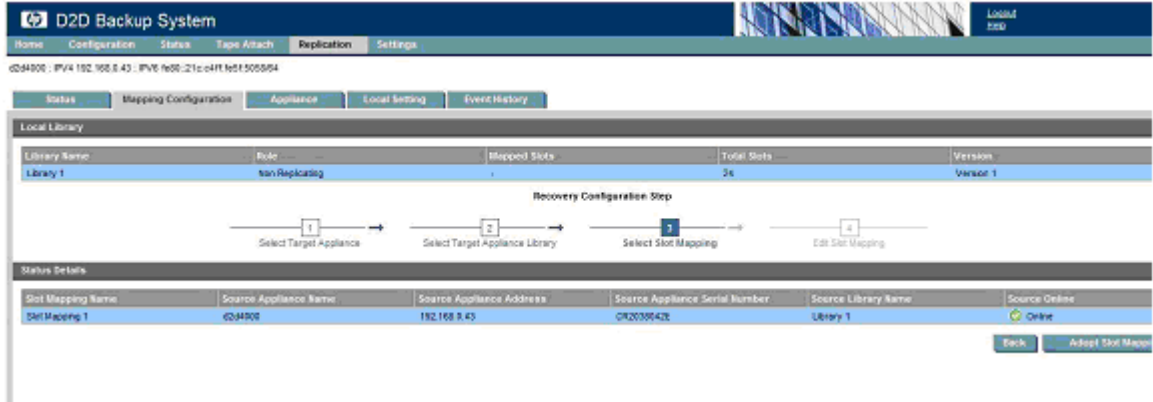
Reverse replication is a means whereby the source and target libraries maintain their identity, but where at the point of rebuilding the source using a special “recover first” mode to force the source to go to the target and reverse replicate data from the target back to the source. The basic process takes place entirely from the remote site where the new D2D appliance has been installed:

1. With a new replacement D2D at the remote site set the appropriate network settings.
2. Create an identical virtual library (or libraries if more than one was originally configured).
3. At the source select the Replication tab from the GUI. Then select **Mapping Configuration**. Begin as if configuring a replication target, but start the recovery wizard instead of the replication wizard.
4. Now create a new target appliance. Select **Add Target Appliance**. The IP address of the target appliance (which still exists) is entered and when found the target appliance is shown. Then select **Next**. If there are multiple target libraries within the target appliance select the relevant one that was configured prior to the disaster at remote site C.

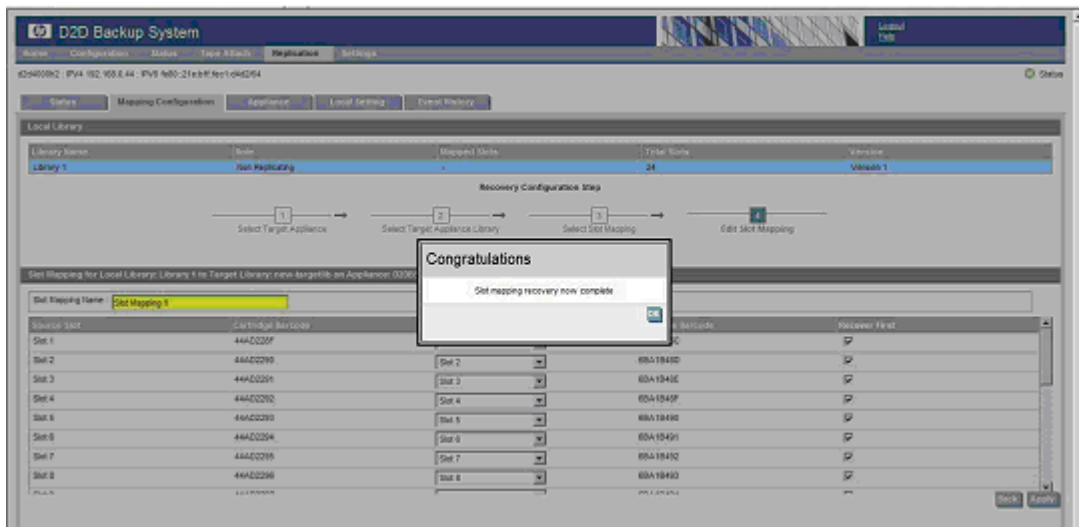




- The recovery wizard will then provide slot mappings. Select the required tapes which were last replicated prior to the disaster at remote site C. Select **Adopt Slot Mapping** and then **Next**. Choose the target cartridges you wish to be reverse replicated and select **Apply**.



The system displays a message indicating that reverse replication has been established.



The source will now commence synchronizing with the target and target virtual tapes will be copied to the source complete with their original source barcodes, to aid easy recovery of the data at the source using the backup application. The recover first box is automatically enabled. Reverse replication starts and can be tracked by clicking on the slot mapping tag.

The screenshot shows the HP D2D Backup System web interface. The top navigation bar includes 'Home', 'Configuration', 'Status', 'Tape Attach', 'Replication', and 'Settings'. The main content area is titled 'Local Library' and contains a table with the following data:

Library Name	Role	Mapped Slots	Total Slots	Version
Library 1	Replication Source	10	24	Version 1

Below this table is a section for 'Slot Mapping for Local Library' with a table showing the following data:

Slot Mapping Name	Target Appliance Name	Target Appliance Address	Target Appliance Serial Number	Target Library Name	Status	Online
Slot Mapping 1	C0305	192.168.0.45	C2CB171C14	none.targetlib	Pending Synchronization	Online

At the bottom, there is a detailed table for 'Slot Mapping for Local Library: Library 1 to Target Library: none.targetlib on Appliance: C0305':

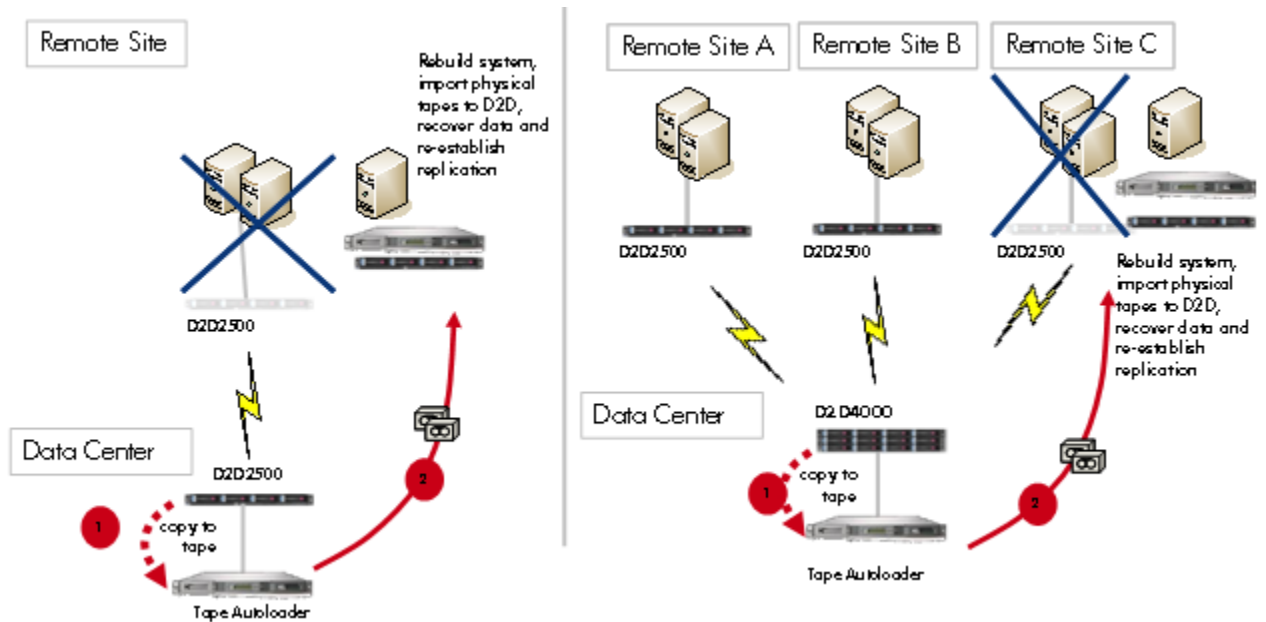
Source Slot	Cartridge Barcode	Target Slot	Replication Status
Slot 1	44A0228P	Slot 1	Replication Synchronizing
Slot 2	44A0228O	Slot 2	Replication Synchronizing
Slot 3	44A0228I	Slot 3	Replication Synchronizing
Slot 4	44A0228J	Slot 4	Replication Synchronizing
Slot 5	44A0228K	Slot 5	Pending Replication Synchronization
Slot 6	44A0228L	Slot 6	Pending Replication Synchronization
Slot 7	44A0228M	Slot 7	Pending Replication Synchronization
Slot 8	44A0228N	Slot 8	Pending Replication Synchronization
Slot 9	44A0228P	Slot 9	Pending Replication Synchronization
Slot 10	44A0228Q	Slot 10	Pending Replication Synchronization
Slot 11	44A0228R	Unmapped	-
Slot 12	44A0228A	Unmapped	-
Slot 13	44A0228B	Unmapped	-
Slot 14	44A0228C	Unmapped	-
Slot 15	44A0228D	Unmapped	-

With reverse replication the deduplication store at the source has to be repopulated again from the beginning so reverse replication will take longer than “normal replication” after initialization. After reverse replication, the normal source to target relationship is reestablished automatically. Further replication will take place from source to target in the normal manner.

It is also possible to reverse the replication even if the target D2D system has lost the mapping. This is the case if a source library is accidentally deleted. Recovery is similar to the above except that only the replication wizard will be available. The recover first mapping tick boxes have to be manually ‘ticked’ for each tape which is required to be reverse replicated.

## Reverse Tape Initialization on the D2D

A third method is also available. At the data center, copy the tape contents required at remote site C directly onto physical tape using the tape offload functionality within HP D2D Backup Systems. Typically these tapes would be the last fully synchronized full backup and any incrementals required. These tapes can then be transported to remote site C along with the tape library they were prepared on (if necessary). Importing the tapes back onto a brand new replacement D2D2500 at remote site C effectively initializes the deduplication store at remote site C and the required data can be recovered from the replacement D2D device. The data transferred on the physical tapes is in 100% backup application format. When it is imported into the replacement D2D at remote site C, the import process of reading the physical tape regenerates all the necessary hash codes and the deduplication store at remote site C. The data can even be recovered directly from the physical tapes at remote site C if required; this is exactly the same process as tape initialization.



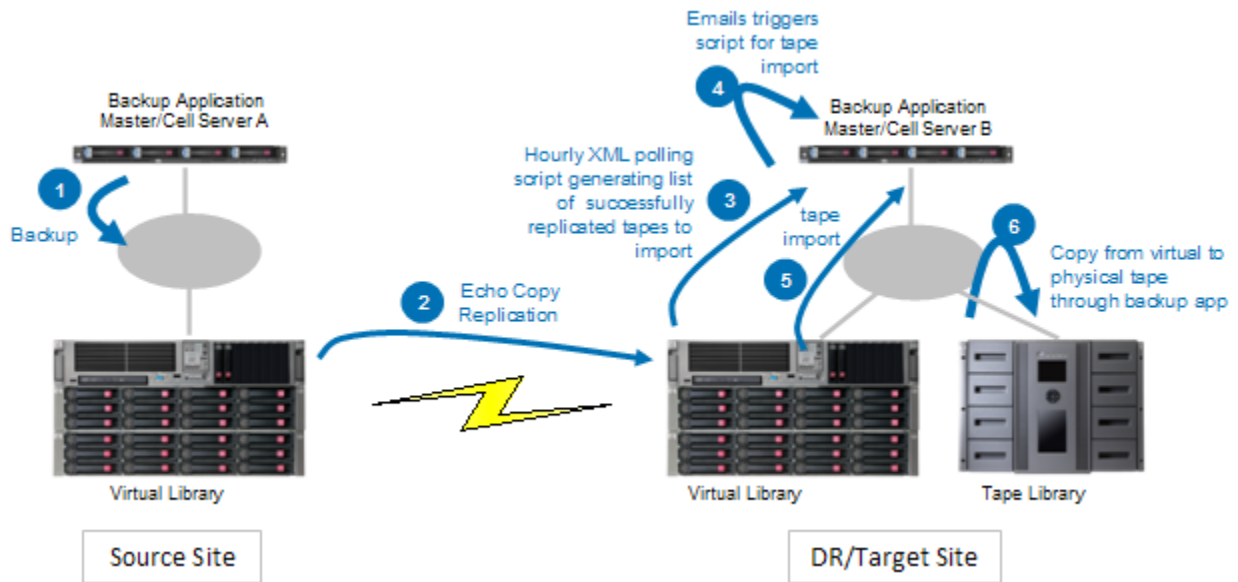
**Figure 39 Recovery Options — Physical Tape**

## Creating Archive Tapes from the Target

Even with replication removing the need for physical tape offsite, there are still many users who wish to use physical tape for archive or test recoveries, etc. An archive tape will have the following differences from the original backup:

- A different retention time from the original backup (for example, you may keep backups and offsite replicas for three months but keep tape archive for several years).
- Different cartridge contents because it will be a different size (optimum virtual cartridge size for deduplication is 50-100 GB but physical tapes such as LTO-4 are 800 GB).
- A different barcode and be tracked separately in the backup application.

These characteristics mean that the backup application must control the creation of the archive tape and must use an object-copy based system (cannot use a whole-tape copy system). This is easy if the physical library is on the source site because you can use the normal backup application object-copy mechanisms to simply copy from the source virtual library to the physical library via the backup application media servers. However, if the physical library is at the target site and you want to perform the tape archive from the replicated virtual cartridges in the target device, this requires additional steps. This is because the backup application at the target device that sees the replicated virtual cartridges cannot be the same backup application that performed the backups at the source site (see [Backup Application Interaction with Replication](#)). The target site backup application does not have any information about the replicated cartridges in its media database because it did not back up to those tapes.



**Figure 40** Creating Archive Tapes from the Target

In order to have the target backup application copy from the replicated cartridges to a physical tape library, it must be “taught” what is on the replicated cartridges. This can be done automatically as shown in [Figure 40](#) by setting the target virtual library to “read-only” mode (to allow the Backup Application B to access it) and then polling the new XML report (`https://<D2Daddress>/ReplicationStatus.xml`) on the D2D with a script to generate a listing of which virtual cartridges have been successfully replicated since the last poll. This cartridge list can then be fed into a script that automatically triggers tape import jobs in the backup application (which read the new cartridge data and import this content into the media database). The target backup application can then restore from the imported virtual cartridges or copy them to physical tape, etc. For backup application import script examples see [Detailed Backup Application Guidelines for VLS](#).

---

# 5 Virtual Library Systems

## VLS Devices

### VLS Defined

The HP StorageWorks VLS12000 Gateway, the VLS6000-series, and the VLS9000-series are RAID disk-based SAN backup devices. The Gateway platform and the VLS6000 and VLS9000-series devices emulate physical tape libraries, allowing you to perform disk-to-virtual tape (disk-to-disk) backups using your existing backup applications. The VLS12000 is FATA or Fibre Channel disk-based; the VLS6000-series and VLS9000-series are serial ATA disk-based. A multi-node VLS behaves as if it were a single device, making it easier to manage and integrate into the overall backup strategy.

The VLS emulates a variety of physical tape libraries, including the tape drives and cartridges inside the libraries. You determine the number and types of tape libraries a VLS emulates, and the number and type of tape drives and cartridges included in each tape library to meet the needs of your environment. You configure the size of the virtual cartridges in your VLS, which provides even more flexibility. With no deduplication, each node on a VLS12000, VLS6000, or VLS9000 can emulate up to 16 tape libraries and 128 tape drives. With deduplication enabled, each node on a VLS12000, VLS6000, or VLS9000 can emulate up to 96 virtual devices (0-16 tape libraries plus tape drives). The maximum number of virtual cartridges that can be configured on a VLS device depends on the firmware version and the VLS model:

- A VLS running firmware version 1.x/2.x can create up to 8192 virtual tapes, with each array LUN adding up to a maximum of 1024 tapes capability. For example, on a VLS9000 with one 30 or 40 TB array (which contains one LUN per array enclosure and four enclosures) you can create up to 4096 tapes; if you have two or more arrays you can create 8192 tapes. On a VLS6227 (which contains one array enclosure with 750 GB disks) you can create up to 3072 tapes (with three LUNs per 750 GB array enclosure) in RAID 5 mode and up to 4096 tapes (four LUNs per 750 GB array enclosure) in RAID 6 mode.
- A VLS running firmware version 3.x can create up to 32000 virtual tapes, with each array LUN adding up to a maximum of 8000 tapes capability. For example, on a VLS9000 with one or more 30 or 40 TB arrays (each full array contains one LUN per array enclosure and four enclosures) you can create up to 32000 tapes, but if you have a VLS9000 with a 10 TB array (one LUN) you can create up to 8000 tapes. On a VLS6227 (which contains one array enclosure with 750 GB disks) you can create up to 24000 tapes (with three LUNs per 750 GB array enclosure) in RAID 5 mode and up to 32000 tapes (four LUNs per 750 GB array enclosure) in RAID 6 mode.

VLS accommodates mixed IT-platform and backup-application environments. For more insight into mixed-IT solutions and environment analysis, visit the HP StorageWorks Enterprise Backup Solutions web site: <http://www.hp.com/go/ebs>.

All of your servers and backup applications access the virtual media simultaneously. You can specify which servers are allowed to access each virtual library and tape drive you configure. You can change the default LUNs assigned to the virtual library and tape drives for each host, as needed, to accommodate different operating system requirements and restrictions.

The virtual tape solution uses disk storage to simulate a tape library. That is, a virtual library simulates libraries with tape drives and slots as if they were physical devices. Backup servers see the virtual

libraries as physical libraries on the SAN. Because you can create many more virtual libraries and drives than you have physical tape drives, many more SAN-based backups can run concurrently from the application servers, reducing the aggregate backup window. After the backups are complete, the automigration feature or the data protection software can migrate backup data from the virtual media to physical tape, for off-site disaster protection or long-term archival. This migration to physical tape can happen outside of the backup window at whatever time you determine will work optimally with your backup environment. The deduplication-enabled replication feature also provides the ability to create a second off-site copy in another VLS via a TCP/IP link.

## Prime Environments

VLS is particularly suited to environments in which:

- The backup targets are in a SAN or LAN and are backed up to a SAN device via a media server.
- The VLS12000 Gateway has eight or more servers; the VLS6000-series and the VLS9000-series have four or more servers.
- The servers are slower than the tape drives.
- Backup performance is a necessity because backup windows are currently strained.
- Many restores are expected for the data set being backed up.
- Enhanced reliability is desired.
- Low impact, seamless integration is important because the virtual library looks to the backup application like a physical library and retooling of the backup jobs is not necessary.
- There is rapid growth, so quick and easy provisioning of data protection services is frequently needed.
- There is a need or wish to reduce the amount of physical media being consumed.
- There are blade servers installed in your environment.

## VLS Benefits

- Integrates seamlessly into existing backup applications and processes.
- Easy to manage.
- Accelerates backup performance in complex SAN environments while improving overall reliability. Because it can receive many data streams at the same time, the VLS maximizes the number of backup jobs occurring simultaneously, reducing backup window duration. (Note that this applies to slower SAN hosts. Faster SAN hosts bypass this issue.) Migration to tape then happens outside the backup window.
- Fast single-file restore and recovery - recent backups kept on disk for restore and recovery, fast single-file restore without the need for media loading or positioning.
- Adaptable to environmental changes. As storage requirements grow, the VLS can match your environment by scaling both capacity and performance.
- Utilizes LUN masking and mapping to protect backup and recovery jobs from disruptions caused by SAN events.
- Optimizes bandwidth use of the backup environment.
- Removes tape errors from the critical backup window.
- Less expensive than buying more fiber channel arrays.
- Can require fewer tape drives than other solutions and get better utilization of the tape drives/libraries you have.

The benefits of specific VLS models:


- Because the VLS12000 is a Gateway to an EVA, it can take advantage of Instant Support Enterprise Edition – if a data pool fails, the “phone-home” service activates.

- With VLS12000 Gateway, VLS6600, and VLS9000, data compression is provided by a hardware compression card that does not impede performance.
- The VLS12000 Gateway and the VLS9000 provide high reliability, dual paths between the node and the storage pools, and redundant power supplies and cooling.
- The VLS9000-series is scalable enough to be a single or multi-node device — yet with a multi-node configuration, it retains the manageability of a single device. VLS9000 also offers the ability to easily expand the device's performance and capacity after the initial purchase.

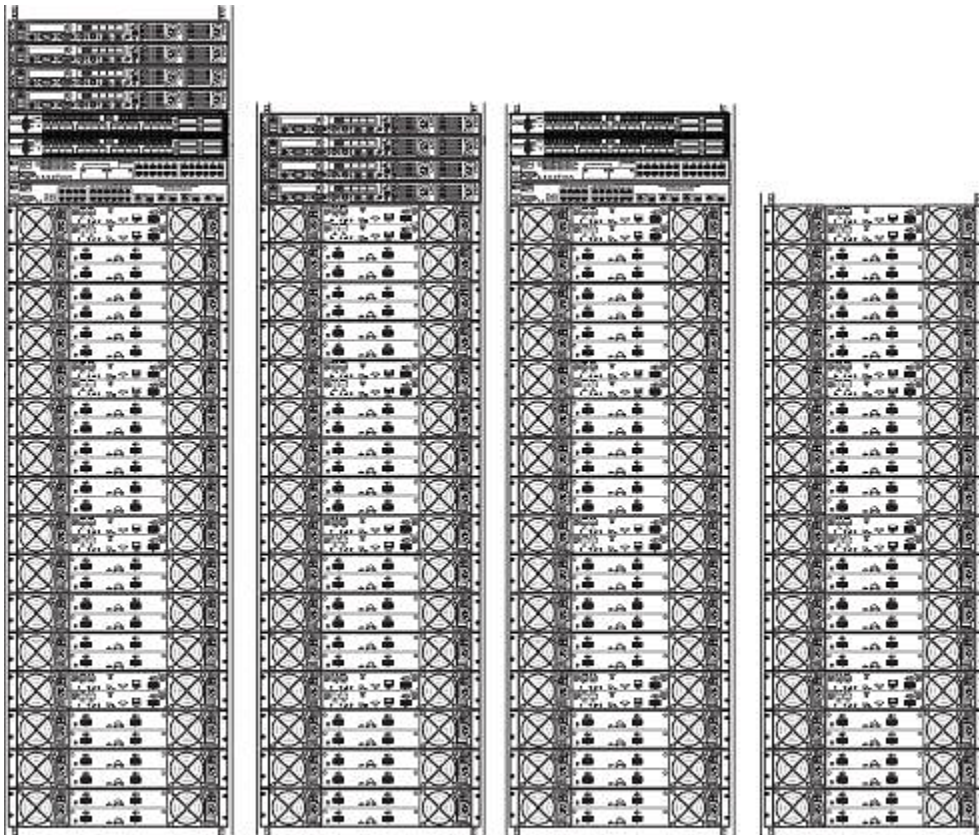
## VLS9000-series Configurations

At the smaller end, the VLS9000 7.5 TB and 10 TB systems consist of one node, one Ethernet switch, one base disk array enclosure, and up to three expansion disk array enclosures. See [Figure 41](#). A typical, larger VLS9000-series system consists of at least one VLS9000 node, at least one VLS9000 array (one base disk array enclosure and three expansion disk array enclosures), and one VLS9000 connectivity kit (two Ethernet switches for internal inter-node connections and two Fibre Channel switches for disk storage connections).

If you have a 20-port connectivity kit, you can install a total of ten devices, consisting of a combination of arrays and nodes. If you have a 32-port connectivity kit, you can install a total of 16 devices. By interconnecting two 32-port connectivity kits with an interswitch link kit, you can further expand the scalability of VLS9000-series. [Figure 42](#) shows a full VLS9000-system with two 32-port connectivity kits. This system fills four racks and includes eight nodes, 16 arrays, four Fibre Channel switches, and four Ethernet switches.



**Figure 41 7.5 TB System with Expansion Disk Array Enclosures**



11874

Figure 42 Full VLS9000 System

## VLS Technical Specifications

Table 13 VLS9000-series Multi-node Maximum Capacity by Configuration with 20-port Connectivity Kit and 40 TB Arrays (2:1 Compression, Non-deduplication)

	1 node	2 nodes	3 nodes	4 nodes	5 nodes
1 array	80 TB	n/a	n/a	n/a	n/a
2 arrays	160 TB	160 TB	n/a	n/a	n/a
3 arrays	n/a	240 TB	240 TB	n/a	n/a
4 arrays	n/a	320 TB	320 TB	320 TB	n/a
5 arrays	n/a	n/a	400 TB	400 TB	400 TB
6 arrays	n/a	n/a	480 TB	480 TB	n/a
7 arrays	n/a	n/a	n/a	n/a	n/a



	1 node	2 nodes	3 nodes	4 nodes	5 nodes
8 arrays	n/a	n/a	n/a	n/a	n/a

**Table 14 VLS9000-series Multi-node Maximum Capacity by Configuration with 32-port Connectivity Kit and 40 TB Arrays (2:1 Compression, Non-deduplication)**

	1 node	2 nodes	3 nodes	4 nodes	5 nodes	6 nodes	7 nodes	8 nodes
1 array	80 TB	n/a	n/a	n/a	n/a	n/a	n/a	n/a
2 arrays	160 TB	160 TB	n/a	n/a	n/a	n/a	n/a	n/a
3 arrays	n/a	240 TB	240 TB	n/a	n/a	n/a	n/a	n/a
4 arrays	n/a	320 TB	320 TB	320 TB	n/a	n/a	n/a	n/a
5 arrays	n/a	n/a	400 TB	400 TB	400 TB	n/a	n/a	n/a
6 arrays	n/a	n/a	480 TB	480 TB	480 TB	480 TB	n/a	n/a
7 arrays	n/a	n/a	n/a	560 TB	560 TB	560 TB	560 TB	n/a
8 arrays	n/a	n/a	n/a	640 TB	640 TB	640 TB	640 TB	640 TB
9 arrays	n/a	n/a	n/a	n/a	720 TB	720 TB	720 TB	n/a
10 arrays	n/a	n/a	n/a	n/a	800 TB	800 TB	n/a	n/a
11 arrays	n/a	n/a	n/a	n/a	n/a	880 TB	n/a	n/a
12 arrays	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
13 arrays	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
14 arrays	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
15 arrays	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
16 arrays	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a

**Table 15 VLS9000-series Multi-node Maximum Capacity by Configuration with Two 32-port Connectivity Kits and 40 TB Arrays (2:1 Compression, Non-deduplication)**

	1 node	2 nodes	3 nodes	4 nodes	5 nodes	6 nodes	7 nodes	8 nodes
1 array	80 TB	n/a	n/a	n/a	n/a	n/a	n/a	n/a
2 arrays	160 TB	160 TB	n/a	n/a	n/a	n/a	n/a	n/a
3 arrays	n/a	240 TB	240 TB	n/a	n/a	n/a	n/a	n/a
4 arrays	n/a	320 TB	320 TB	320 TB	n/a	n/a	n/a	n/a
5 arrays	n/a	n/a	400 TB	400 TB	400 TB	n/a	n/a	n/a
6 arrays	n/a	n/a	480 TB	480 TB	480 TB	480 TB	n/a	n/a
7 arrays	n/a	n/a	n/a	560 TB	560 TB	560 TB	560 TB	n/a

	1 node	2 nodes	3 nodes	4 nodes	5 nodes	6 nodes	7 nodes	8 nodes
8 arrays	n/a	n/a	n/a	640 TB	640 TB	640 TB	640 TB	640 TB
9 arrays	n/a	n/a	n/a	n/a	720 TB	720 TB	720 TB	720 TB
10 arrays	n/a	n/a	n/a	n/a	800 TB	800 TB	800 TB	800 TB
11 arrays	n/a	n/a	n/a	n/a	n/a	880 TB	880 TB	880 TB
12 arrays	n/a	n/a	n/a	n/a	n/a	960 TB	960 TB	960 TB
13 arrays	n/a	n/a	n/a	n/a	n/a	n/a	1040 TB	1040 TB
14 arrays	n/a	n/a	n/a	n/a	n/a	n/a	1120 TB	1120 TB
15 arrays	n/a	n/a	n/a	n/a	n/a	n/a	n/a	1200 TB
16 arrays	n/a	n/a	n/a	n/a	n/a	n/a	n/a	1280 TB

**Table 16 VLS9000-series Backup Performance by Configuration**

	with uncompressible data	with 2:1 compressible data
1 node	540 Mb/s	600 Mb/s
2 nodes	1080 Mb/s	1200 Mb/s
3 nodes	1620 Mb/s	1800 Mb/s
4 nodes	2160 Mb/s	2400 Mb/s
5 nodes	2700 Mb/s	3000 Mb/s
6 nodes	3240 Mb/s	3600 Mb/s
7 nodes	3780 Mb/s	4200 Mb/s
8 nodes	4320 Mb/s	4800 Mb/s
Notes:		
<ul style="list-style-type: none"> <li>• Maximum theoretical performance values.</li> <li>• Assumes at least one capacity array per node.</li> <li>• 1MB = 1,000,000 bytes.</li> </ul>		

**Table 17 VLS6000-series Technical Specifications**

	VLS6218	VLS6227	VLS6636	VLS6653
Usable capacity (standard configuration, non-compressed)	4.4 TB	6.6 TB	8.8 TB	13.2 TB
Usable capacity (2:1 compression)	8.8 TB	13.2 TB	17.6 TB	26.4 TB
Maximum usable capacity (standard configuration, non-compressed)	28.4 TB	30.6 TB	56.8 TB	61.2 TB

	VLS6218	VLS6227	VLS6636	VLS6653
Maximum usable capacity (2:1 compression)	56.8 TB	61.2 TB	113.6 TB	122.4 TB

**Table 18 VLS6000-series Performance by Configuration, 4 GB Models**

	VLS6200	VLS6600	
	with compression disabled	with uncompressible data	with 2:1 compressible data
1 disk enclosure	0.4 TB/hr (100 MB/s)	n/a	n/a
2 disk enclosures	0.7 TB/hr (200 MB/s)	0.7 TB/hr (200 MB/s)	1.4 TB/hr (400 MB/s)
3 disk enclosures	1.1 TB/hr (300 MB/s)	1.1 TB/hr (300 MB/s)	2.2 TB/hr (600 MB/s)
4 disk enclosures	1.4 TB/hr (400 MB/s)	1.4 TB/hr (400 MB/s)	2.2 TB/hr (600 MB/s)
5 to 8 disk enclosures	n/a	2.0 TB/hr (560 MB/s)	2.2 TB/hr (600 MB/s)
Notes: <ul style="list-style-type: none"> <li>• Maximum theoretical performance values. Actual performance is dependent on environment. Software compression on VLS6200 impacts performance by up to 50%.</li> <li>• 1MB = 1,000,000 bytes.</li> </ul>			

**Table 19 VLS12000 EVA Gateway Technical Specifications**

	VLS12000 EVA Gateway
Usable capacity for base (2–node) configuration	50 TB
Maximum usable capacity	512 TB and 1 PB (with hardware compression)
Number of virtual library nodes	2 to 8
Number of virtual tape libraries per node	1 to 16
Number of virtual tape drives per node	1 to 128

**Table 20 VLS12000 EVA Gateway Backup Throughput using Deduplication**

		Number of VLS12000 EVA Gateway Nodes with uncompressible data or deduplication enabled						
		2	3	4	5	6	7	8
<b>Number of EVA3000 or EVA5000 Arrays</b> (assumes 56 drives per array controller)	1	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	2	200 MB/s	n/a	n/a	n/a	n/a	n/a	n/a
	3	300 MB/s	300 MB/s	n/a	n/a	n/a	n/a	n/a
	4	400 MB/s	400 MB/s	400 MB/s	n/a	n/a	n/a	n/a
	5	500 MB/s	500 MB/s	500 MB/s	500 MB/s	n/a	n/a	n/a
	6	600 MB/s	600 MB/s	600 MB/s	600 MB/s	600 MB/s	n/a	n/a
	7	700 MB/s	700 MB/s	700 MB/s	700 MB/s	700 MB/s	700 MB/s	n/a
	8	800 MB/s	800 MB/s	800 MB/s	800 MB/s	800 MB/s	800 MB/s	800 MB/s
<b>Number of EVA4000 or EVA4100 Arrays</b> (assumes 56 drives per array controller)	1	200 MB/s	200 MB/s	200 MB/s	200 MB/s	200 MB/s	200 MB/s	200 MB/s
	2	400 MB/s	400 MB/s	400 MB/s	400 MB/s	400 MB/s	400 MB/s	400 MB/s
	3	600 MB/s	600 MB/s	600 MB/s	600 MB/s	600 MB/s	600 MB/s	600 MB/s
	4	800 MB/s	800 MB/s	800 MB/s	800 MB/s	800 MB/s	800 MB/s	800 MB/s
	5	1000 MB/s	1000 MB/s	1000 MB/s	1000 MB/s	1000 MB/s	1000 MB/s	1000 MB/s
	6	1080 MB/s	1200 MB/s	1200 MB/s	1200 MB/s	1200 MB/s	1200 MB/s	1200 MB/s

		Number of VLS12000 EVA Gateway Nodes with uncompressible data or deduplication enabled						
<b>Number of EVA6000 or EVA6100 or EVA4400 Arrays</b> (assumes 112 drives per array controller for EVA6x00 and 96 drives per array for EVA4400)	1	400 MB/s	400 MB/s	400 MB/s	400 MB/s	400 MB/s	400 MB/s	400 MB/s
	2	800 MB/s	800 MB/s	800 MB/s	800 MB/s	800 MB/s	800 MB/s	800 MB/s
	3	1080 MB/s	1200 MB/s	1200 MB/s	1200 MB/s	1200 MB/s	1200 MB/s	1200 MB/s
	4	1080 MB/s	1600 MB/s	1600 MB/s	1600 MB/s	1600 MB/s	1600 MB/s	1600 MB/s
	5	1080 MB/s	1620 MB/s	2000 MB/s	2000 MB/s	2000 MB/s	2000 MB/s	2000 MB/s
	6	1080 MB/s	1620 MB/s	2160 MB/s	2400 MB/s	2400 MB/s	2400 MB/s	2400 MB/s
<b>Number of EVA8000 or EVA8100 Arrays</b> (assumes 112 drives per array controller)	1	400 MB/s	400 MB/s	400 MB/s	400 MB/s	400 MB/s	400 MB/s	400 MB/s
	2	800 MB/s	800 MB/s	800 MB/s	800 MB/s	800 MB/s	800 MB/s	800 MB/s
	3	1080 MB/s	1200 MB/s	1200 MB/s	1200 MB/s	1200 MB/s	1200 MB/s	1200 MB/s
	4	1080 MB/s	1600 MB/s	1600 MB/s	1600 MB/s	1600 MB/s	1600 MB/s	1600 MB/s
	5	1080 MB/s	1620 MB/s	2000 MB/s	2000 MB/s	2000 MB/s	2000 MB/s	2000 MB/s
	6	1080 MB/s	1620 MB/s	2160 MB/s	2400 MB/s	2400 MB/s	2400 MB/s	2400 MB/s

**Table 21 VLS12000 EVA Gateway Backup Throughput without Deduplication**

		Number of VLS12000 EVA Gateway Nodes with 2:1 compressible data or deduplication disabled						
		2	3	4	5	6	7	8
<b>Number of EVA3000 or EVA5000 Arrays</b> (assumes 56 drives per array controller)	1	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	2	400 MB/s	n/a	n/a	n/a	n/a	n/a	n/a
	3	600 MB/s	600 MB/s	n/a	n/a	n/a	n/a	n/a
	4	800 MB/s	800 MB/s	800 MB/s	n/a	n/a	n/a	n/a
	5	1000 MB/s	1000 MB/s	1000 MB/s	1000 MB/s	n/a	n/a	n/a

		Number of VLS12000 EVA Gateway Nodes with 2:1 compressible data or deduplication disabled						
	6	1200 MB/s	1200 MB/s	1200 MB/s	1200 MB/s	1200 MB/s	n/a	n/a
	7	1200 MB/s	1400 MB/s	1400 MB/s	1400 MB/s	1400 MB/s	1400 MB/s	n/a
	8	1200 MB/s	1600 MB/s	1600 MB/s	1600 MB/s	1600 MB/s	1600 MB/s	1600 MB/s
<b>Number of EVA4000 or EVA4100 Arrays</b> (assumes 56 drives per array controller)	1	400 MB/s	400 MB/s	400 MB/s	400 MB/s	400 MB/s	400 MB/s	400 MB/s
	2	800 MB/s	800 MB/s	800 MB/s	800 MB/s	800 MB/s	800 MB/s	800 MB/s
	3	1200 MB/s	1200 MB/s	1200 MB/s	1200 MB/s	1200 MB/s	1200 MB/s	1200 MB/s
	4	1200 MB/s	1600 MB/s	1600 MB/s	1600 MB/s	1600 MB/s	1600 MB/s	1600 MB/s
	5	1200 MB/s	1800 MB/s	2000 MB/s	2000 MB/s	2000 MB/s	2000 MB/s	2000 MB/s
	6	1200 MB/s	1800 MB/s	2400 MB/s	2400 MB/s	2400 MB/s	2400 MB/s	2400 MB/s
<b>Number of EVA6000 or EVA6100 or EVA4400 Arrays</b> (assumes 112 drives per array controller for EVA6x00 and 96 drives per array for EVA4400)	1	800 MB/s	800 MB/s	800 MB/s	800 MB/s	800 MB/s	800 MB/s	800 MB/s
	2	1200 MB/s	1600 MB/s	1600 MB/s	1600 MB/s	1600 MB/s	1600 MB/s	1600 MB/s
	3	1200 MB/s	1800 MB/s	2400 MB/s	2400 MB/s	2400 MB/s	2400 MB/s	2400 MB/s
	4	1200 MB/s	1800 MB/s	2400 MB/s	3000 MB/s	3200 MB/s	3200 MB/s	3200 MB/s
	5	1200 MB/s	1800 MB/s	2400 MB/s	3000 MB/s	3600 MB/s	4000 MB/s	4000 MB/s
	6	1200 MB/s	1800 MB/s	2400 MB/s	3000 MB/s	3600 MB/s	4200 MB/s	4800 MB/s
<b>Number of EVA8000 or EVA8100 Arrays</b> (assumes 112 drives per array controller)	1	800 MB/s	800 MB/s	800 MB/s	800 MB/s	800 MB/s	800 MB/s	800 MB/s
	2	1200 MB/s	1600 MB/s	1600 MB/s	1600 MB/s	1600 MB/s	1600 MB/s	1600 MB/s
	3	1200 MB/s	1800 MB/s	2400 MB/s	2400 MB/s	2400 MB/s	2400 MB/s	2400 MB/s
	4	1200 MB/s	1800 MB/s	2400 MB/s	3000 MB/s	3200 MB/s	3200 MB/s	3200 MB/s


		Number of VLS12000 EVA Gateway Nodes with 2:1 compressible data or deduplication disabled						
	5	1200 MB/s	1800 MB/s	2400 MB/s	3000 MB/s	3600 MB/s	4000 MB/s	4000 MB/s
	6	1200 MB/s	1800 MB/s	2400 MB/s	3000 MB/s	3600 MB/s	4200 MB/s	4800 MB/s

## How it Works

### VLS Scalability

The VLS9000 and VLS12000 models have a multi-node “grid-like” architecture that provides the ability to scale both performance and capacity of the virtual library while still presenting it as a single backup target to the backup application. This means that all the backup jobs can be configured to use the same shared backup target which dramatically improves the amount of day-to-day administration required to run the backups, which is one of the main advantages of a consolidating to an enterprise-class scalable backup device. This is significantly better than having to divide up the backup solution into multiple smaller lower performance backup targets because this would require manual administration of which backups use which target, and as backups grow this also means switching backups from one target to another as each target starts to run out of performance or capacity. (This is similar to using one enterprise tape library such as an ESL-E compared to using multiple autoloaders.)

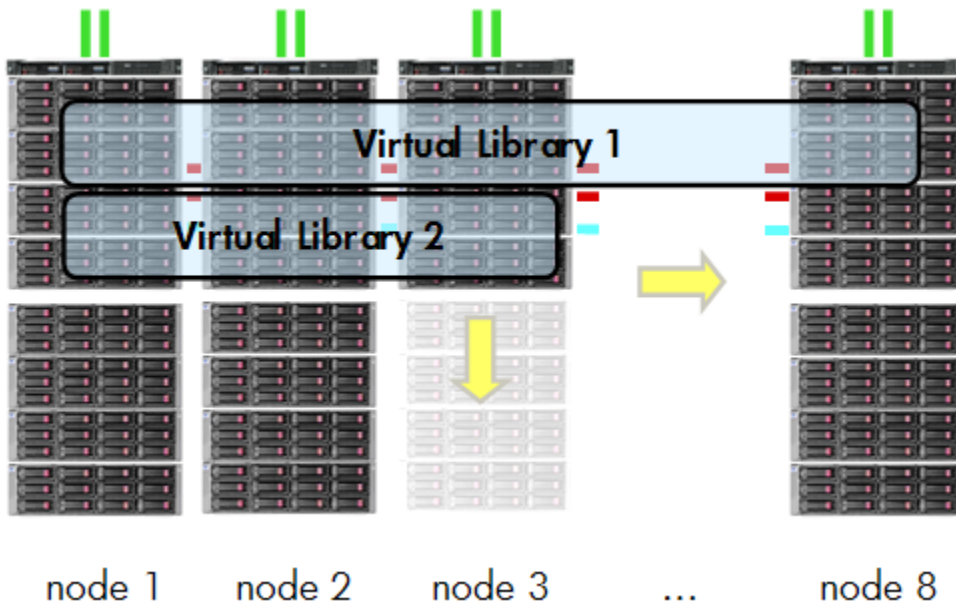
The VLS nodes present the virtual library with its virtual tape drives and the virtual cartridges are stored on the back-end array LUNs. The internal architecture of the VLS has every VLS node connected (via dual paths) to every back-end array LUN as shown in [Figure 43](#):



**Figure 43 Internal Architecture of the VLS (VLS9000 Shown)**

This allows any cartridge to be loaded into any virtual tape drive on any node, so the virtual drives in a library can be on different nodes from the virtual library robot. Therefore, the virtual drives for a library can be configured across multiple nodes as shown in [Figure 44](#), and as the device is scaled up (by adding more nodes and arrays) the virtual library can be expanded with it by adding more drives and more cartridges to the existing virtual library. (This requires the robot to be initially created with additional spare empty drive bays and empty slots.)






**Figure 44 Virtual Drives across Multiple Nodes**

### VLS Automatic Performance Load Balancing

The VLS always “thin provisions” its virtual cartridges and dynamically load balances all the incoming new backup data across the available back-end arrays to ensure that there are never any performance hotspots. When any virtual cartridge is written (in any virtual tape drive), it dynamically assigns which array LUNs will store the cartridge data based on which array LUNs are least active. When you have multiple virtual cartridges written concurrently, the load is always evenly balanced across the back-end arrays and will always provide maximum performance regardless of which virtual cartridges are used for backup.



**Figure 45 Storage Pool Load Balancing on the VLS**

As [Figure 45](#) shows, the VLS storage pool:

- Provides clustered virtual cartridge file system across one or more arrays.
- Automatically load-balances incoming backup data across all available array LUNs (in 32MB extents).
- Prevents performance and capacity hot-spots.

If a virtual cartridge is overwritten by the backup application (e.g., because it reached the end of the retention period and was reused for new backups), all of its existing blocks are freed up and then new writes are dynamically assigned array LUNs as usual. When a new array is added to the VLS, the load balancing system will automatically start using this new storage on subsequent virtual cartridge writes (i.e., backups performed after the array was added will automatically load balance across all storage including the new array). Any existing data that was present before the new array was added will be rebalanced as virtual cartridges are recycled/expired and then overwritten by the backup application.

## VLS Warm Failover

A VLS feature (in firmware version 3.2 or higher) provides an automatic mechanism for faster recovery from node0 failures, which is the ability to automatically restore the previous virtual library configuration and licenses after a complete node0 replacement. This warm failover feature works as follows:

- The VLS will automatically save, within 1 hour, the current configuration and licenses to a hidden virtual cartridge stored on the back-end disk arrays.
- In the event of a node0 disaster (node failed and was replaced with a new one, lose of node0 operating system, Quick Restore of new operating system, etc) the node0 recovery process is:
  - Replace/repair node0 and set it to “master” ID through console/serial as usual.
  - From that point on, the rest of the device recovery process is fully automatic. The device will automatically restore its previous configuration and previous licenses (which will generate an extra automatic reboot).
  - The resulting repaired device will have the identical configuration as before including serial numbers, front-end FC WWPNs, virtual device configuration, licenses, etc. When using the deduplication-enabled replication feature, you must still manually re-enter the SLAVE node TCP/IP addresses because these are currently not part of the save/restore configuration system and thus not part of the warm failover).

This process makes it much quicker and easier to replace a failed node0. One warm failover scenario: purchase a spare unconfigured node (so it must be ordered separately and *unracked*). Pre-install and pre-wire this spare node into the VLS but leave it *unconfigured* and not powered on. In the event of a node0 failure, manually unplug the external Fibre Channel and LAN connections from the failed node0 and move them to the replacement node0 and then power it on and configure the replacement node as the master. The automatic configuration restore takes care of the rest.

---

 **NOTE:**

Manually transferring the external Fibre Channel connections from the original node0 to the replacement node0 guarantees that there will be no backup SAN IOP problems because it will effectively be the same node0 with the same WWPNs plugged into the same customer Fibre Channel switch ports. If you have the replacement node0 pre-wired with its own external Fibre Channel and LAN connections (so you do not have to transfer the Fibre Channel connections from original to replacement node0) this can cause IOP problems because this will not be identical to the original node0. It will be using different Fibre Channel switch ports, and some operating systems that use switch-based SAN addressing (e.g., HPUX older than v11.31) will see this as a brand new device with new device paths; backups configured to the old paths will no longer work.

---

## Implementation

### Device Configuration Preparation

Before installing and configuring the VLS device, review the following steps needed to scope the configuration requirements for the VLS and the backup application following the best practices.

If deduplication and replication will be configured on the VLS, review the [Deduplication Preparation](#) and [Replication Preparation](#) for more information.

### Basic VLS device configuration

The following steps highlight the considerations for designing the required virtual libraries/drives/cartridges and standard device configuration.

- Identify the number of required virtual libraries. This could be based on the number of separate SAN fabrics or separate backup application instances/types. See [Single Library vs. Multiple Libraries](#).

- Select a storage pooling strategy for the VLS9000 and VLS12000. Identify which pools are used for which virtual libraries. This could be based on separating cartridges across virtual libraries, or decreasing the impact of array failure. See [Storage Pooling](#). If you are planning to use deduplication, note that this currently supports only one storage pool per VLS.
- For each virtual library, decide if dual SAN fabric design is needed. If so, this requires dual port virtual library robots and will affect the virtual drive assignment strategy when the system is in LUN masking mode. See [Multiplexing, Multistreaming, and Multipathing](#).
- Select the virtual library/drive types. The HP ESL-E library type is generally preferred due to its higher emulation limits, but the HP VLS library type (with the same emulation limits as ESL-E) is recommended for Netbackup because this is officially certified.
- Select whether to use shared drives or LUN masked drives. This is based on the number of media servers that will be using each virtual library and on O/S limitations (such as HPUX versions prior to 11.31 not supporting more than eight library/tape LUNs per target port). See [Backup SAN Design Guidelines](#).
- Select how many virtual tape drives are needed per library. See [Virtual Libraries/drives/cartridge Configuration](#).
  - If you use LUN masking, determine how many virtual drives need to be assigned to each media server based on the approximate number of concurrent backups needed for each media server, and develop a strategy for assigning virtual drives to each media server.
  - If you use shared drives, determine the number of drives based on the total number of concurrent backup streams required. (HPUX versions prior to 11.31 cannot see more than eight library/tape LUNs per target port.)
- Determine the cartridge sizing. See [Virtual Libraries/drives/cartridge Configuration](#).
  - Generally, HP recommends the default 100 GB cartridge size.
  - If you plan to enable deduplication on this device, assess the average amount of backup data per virtual tape drive and decide if the default 100 GB cartridge size should be decreased or increased. See [Deduplication Preparation](#).
- Determine if each virtual library is used as a replication source library; if so, decide if cartridges need to be created first in this source library or in the target library. See [Replication Preparation](#).
- Gather virtual cartridge barcode template requirements per library.
- Assess the required number of slots in each virtual library. Methods of calculation:
  - Use the total logical amount of tape capacity per library needed to store the retention policy divided by compression divided by the virtual cartridge sizes.
  - Check the current total physical tape usage. If you are not changing the backup retention scheme when moving to VLS, the amount of virtual tape will be approximately the same as the current total amount of physical tape (already compressed) divided by the virtual cartridge sizes.
  - If this virtual library will be a replication target, it must contain at least as many slots as the total number of source slots that are replicating to this target. See [Replication Preparation](#).
- Gather the required basic device management settings:
  - TCP/IP settings for device management (node0). You only need TCP/IP settings for the other nodes if you will be using replication (every node requires a TCP/IP address for replication).
  - SNMP or email notification settings.
  - NTP server settings.

## Basic Backup Application Configuration

The following steps highlight the considerations for designing the changes needed in the backup application to optimize for use with a VLS:

- Set the tape block size to 256 KB for all VLS tape drives.
- Disable multiplexing/interleaving on all backup jobs to the VLS.
- Identify clients on the SAN that are currently performing LAN-based backup but could be converted to LAN-free backup to improve performance. See [LAN-free Backups](#).
- Determine if there are performance bottlenecks in the current backup (e.g., LAN or SAN ISL bottlenecks) that you must resolve to allow full performance to the VLS.
- Based on your offsite copy requirements, assess the configuration changes needed to implement offsite copy. See [Considerations for Copies](#).
  - Copying through the backup application means creating copy jobs, ensuring there are good copy paths between the virtual library and the physical library on the SAN (e.g., plug into same switch), etc.
  - Copying with automigration means the system needs no changes to the backup application.
  - Copying with replication (deduplication-enabled) means the system needs no changes to the local backup application, but if there is a separate backup application at the target device this may require setting up an import script.

## Capacity Licensing

The VLS devices are licensed by their usable disk capacity (in terms of array LUNs). Each VLS has an embedded “base capacity” that varies depending on the VLS model:

**Table 22 Base Capacity License by VLS Model**

VLS model	Base capacity license
VLS6218, 500 GB drives	Covers the one MSA20 disk array enclosure that comes bundled.
VLS6218, 750 GB drives	Covers the one MSA20 disk array enclosure that comes bundled.
VLS6636, 500 GB drives	Covers the two MSA20s disk array enclosures that come bundled.
VLS6653, 750 GB drives	Covers the two MSA20 disk array enclosures that come bundled.
VLS9000-series	Each node adds a base license for one full VLS9000 array (48 drives, 30 TB or 40 TB). The entry-level VLS9000 7.5 and 10 TB systems do not need additional licenses to add capacity kits up to the full array configuration.
VLS12000-series	Each node adds a base license for 10 EVA LUNs. The base 2-node VLS12000 kit includes a 5-LUN upgrade license giving a total license of 25 LUNs of up to 2 TB each.

If capacity is added to the VLS beyond the above base license level, you must install capacity upgrade licenses (which you can install via the VLS GUI). There is a different upgrade license for each VLS range:

**Table 23 Adding Capacity Beyond the Base Configuration**

VLS model	Capacity upgrade license
VLS6000-series	Enables an additional MSA20 array with 250/500 GB drives or half an MSA20 array with 750/1000 GB drives per license (so you need two licenses to add an MSA20 array with 750/1000 GB drives).

VLS model	Capacity upgrade license
VLS9000-series	Enables an additional 30 or 40 TB VLS9000 array per license.
VLS12000-series	Enables an additional 2 TB (1 LUN) per license.

## Virtual Libraries/drives/cartridge Configuration

The VLS provides a very flexible virtual library configuration, allowing you to create many more virtual tapes drives than a physical library could contain (and therefore run many more concurrent backup jobs), and to create virtual cartridges of any required size. On the multi-node devices such as VLS9000 and VLS12000, you also have the key ability to create a virtual library across the VLS nodes (see [VLS Scalability](#)) thus providing a single backup target with the entire device performance and capacity if required along with the ability to increase the virtual library performance by adding more nodes at a later date.

When designing the virtual library configuration keep in mind the following considerations:

- The virtual library emulation type can be whatever is certified on the backup application. For example, Symantec NetBackup certification requires an emulation that is not the same as a physical library so the “HP VLS” library type is used for this. For other backup applications any library emulation can be used but generally the “HP ESL-E” is the most flexible emulation because it can scale to 10000 slots and 1024 drives (the same as the HP VLS emulation). Remember that the virtual library does not in any way have to match any existing physical library because it is a brand new library on the SAN. It will have very different geometry because it will have many more drives and slots than the physical library and will also have different SAN WWPNs and serial numbers (i.e., it is not a replacement for the physical library).
- You always need to have multiple concurrent backups going to the VLS to achieve maximum throughput. Therefore, analyze how many concurrent backup jobs will be needed to saturate the VLS performance and create at least that many virtual tape drives in the virtual library. Do not create more drives than you need because there is a minor performance impact (e.g., 5-10%) if you use the maximum number of drives per node compared to the optimum performance at 4-8 drives per Fibre Channel port.
- Load balance your virtual devices across the ports. For example, if you have a four-port configuration and you configure a 24-drive library, you should spread the drives evenly across the ports. Ensure that your backup application also balances the drive distribution.
- With VLS9000 and VLS12000, you can upgrade capacity as needed after the initial setup. To prepare for future expansion, the user should configure additional drive bays and slots on the virtual robot to avoid having to delete and recreate the virtual library when expanding capacity.
- If you initially deploy without deduplication (with a maximum virtual drives per node of 128) but you intend to enable deduplication in the future, limit the maximum number of virtual tape drives per node to the deduplication limit which is 96 virtual devices (robots and libraries) per node. This will then make the migration to deduplication easier.
- If possible, configure eight or fewer LUNs (virtual devices, including the robot) per port. Some operating systems will not see more than this. If you need more than eight devices on a port, you must use LUN mapping and masking.
- HP recommends that you do not delete individual tape drives from inside the virtual library. Deleting drives in the virtual library introduces gaps in the LUN numbering for shared (unmapped) virtual tape drives, and many operating systems fail to scan the virtual tape drives that follow a gap. In multi-node configurations, it can change drive element addresses. In both cases, when the device reboots it will reorder the LUNs/elements of the tape drives and may require reconfiguration of the media server tape paths.

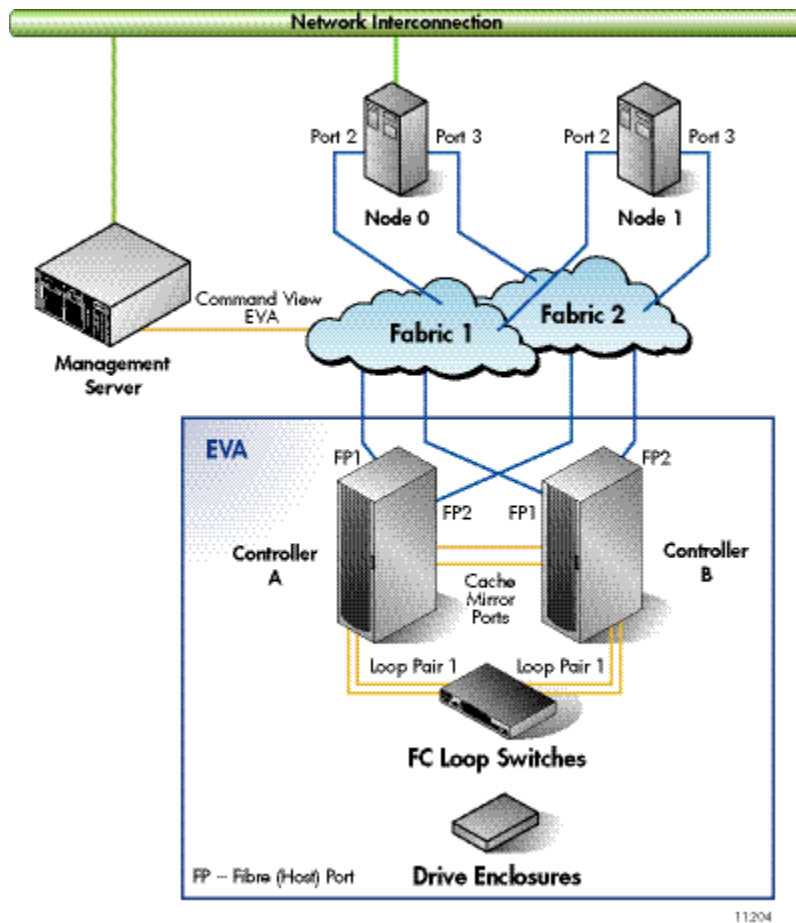
If you use LUN mapping and masking to assign specific virtual drives to specific backup media servers, make sure that the LUN numbering stays in sequence with no gaps. Instead of deleting tape drives, you can disable them. (You can present them to a new host in the future if required.)

When creating virtual cartridges inside the virtual library, it is good practice to keep the virtual cartridge size smaller than the physical tape. Generally it is always better to create smaller cartridges than larger ones because it frees up the cartridge quicker for use by other processes such as restore, tape copy, or deduplication processing. Cartridge sizes of 100 GB (the default size) are common and this size can be used regardless of the tape emulation type because the cartridge is virtual (i.e., you can create a 100 GB cartridge even if you selected LTO4 drive emulation). If copying to physical tape via the backup application, the virtual cartridge size does not have any to match the physical tape in any way (type, size, library type, drive type) because the backup applications all perform “object copy” where they effectively backup the backup and therefore append and span copy jobs onto the physical tapes just like they would with backup jobs. If copying to physical tape via automigration, this will require the virtual cartridges to contain the same amount of data as the physical tape and so automigration will automatically create the correct size virtual cartridges to match the physical tapes.

## EVA Preparation

Before you connect the VLS12000 Gateway to the external EVA consider these concepts and perform the related tasks:

- The VLS12000 attaches to a properly functioning EVA running firmware v. 5.0 or higher. Ensure that your EVA meets these requirements before connecting to a VLS12000.
- Before installing the VLS12000, ensure that the external Fibre Channel switch/fabric has either two external Fibre Channel switches/fabrics or two zones. This is required so that there are two data pathways from the VLS to the EVA. Dual pathing provides path balancing and transparent path failover on the VLS12000.



**Figure 46 VLS12000 EVA Gateway Connected to an EVA**

- Ensure that the zoning configuration is complete and that storage ports 2 and 3 on each VLS node connect to different switches/fabrics or zones. The EVA controllers must also be connected to both switches/fabrics or zones.
- VLS nodes can be configured through the EVA to have dual Fibre Channel paths to the back-end disk arrays. With dual paths, if one path fails the node can still access each LUN through the remaining path(s). To get the best performance from the EVA, you must have balanced performance across the array controllers. To address these requirements:
  - VLS software will automatically recognize the multiple paths to the disk array LUNs. If it does not, revisit your configuration and make changes as necessary so that the software does recognize all paths upon rebooting the system.
  - Path failure is reported through the array information presented in the graphical user interface.

**NOTE:**

Path failure can reduce throughput for backup and restore operations.

- Balance data traffic across both controllers of the array (called A and B in Command View VLS). To do this, ensure that the preferred path for half of the VRaids in the array is set to **Path A-Failover only**, and the preferred path for the other half of the VRaids in the array is set to **Path B-Failover only**.



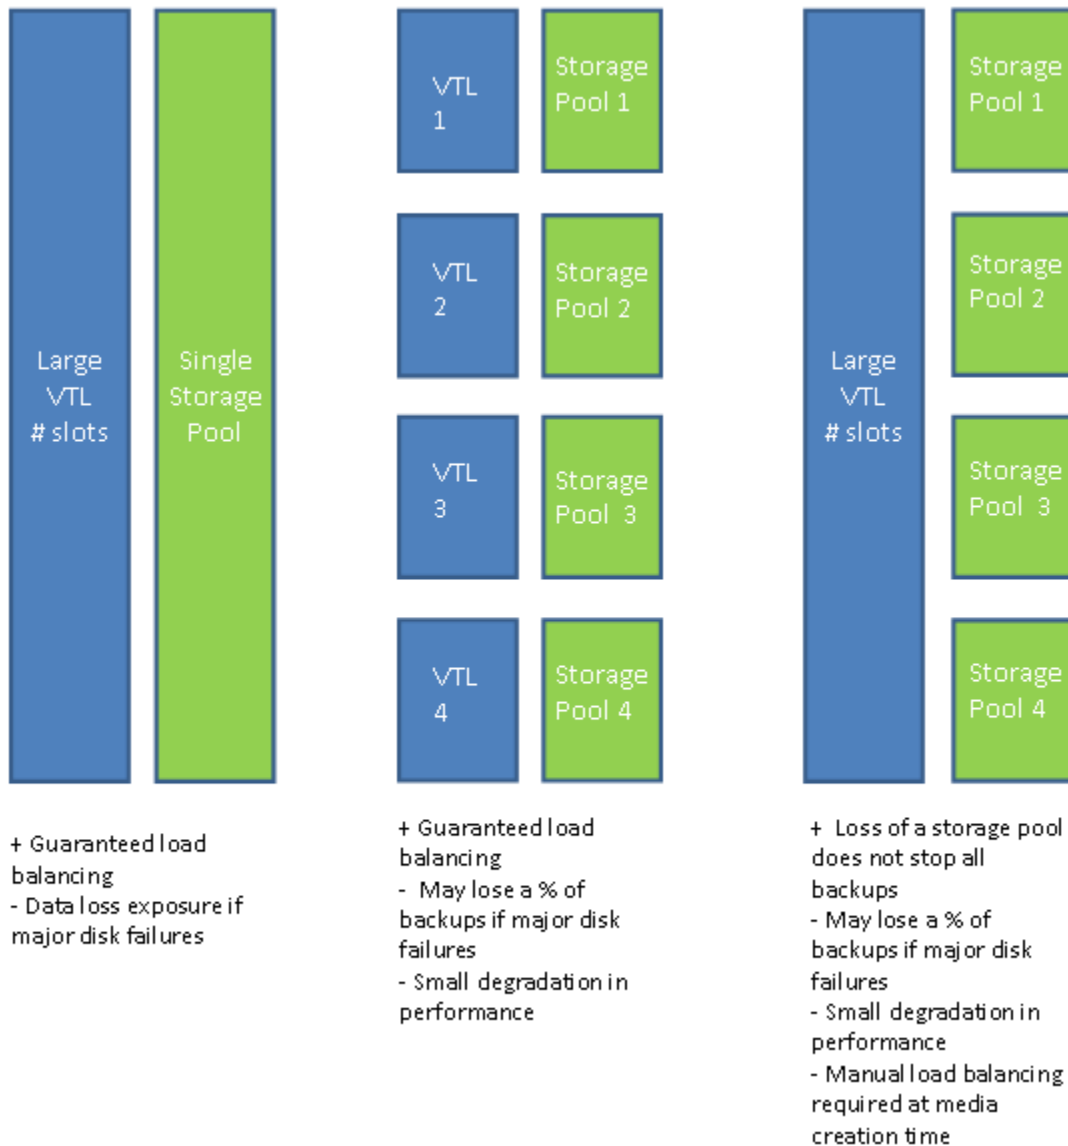
- The EVA LUNs must be fully initialized and ready for use before attaching the VLS to the EVA. It can take several hours to initialize an entire EVA; therefore, HP advises you to create the LUNs in advance of the installation.
- Ensure that all VRaid LUNs required for the VLS have been created on the EVA such that the LUNs are RAID 6 (preferred) or RAID 5, are not configured as read-only, and are roughly the same size.
- The maximum supported EVA LUN size is 2047GB. Do not create LUNs for the VLS larger than this even if the EVA model supports it.
- On the EVA 8000, there are 4 Fibre Channel ports per controller which allows up to eight paths to each LUN; however, configuring more than four paths per LUN will limit the maximum capacity of the device without adding performance benefits. For this reason, HP recommends use of only four paths per LUN.
- It is possible to share an EVA between the VLS12000 and other applications. If the EVA is shared, VLS traffic may overwhelm the array and significantly reduce performance across all applications sharing the array. Alternately, the other applications may slow down the performance of the VLS. If this happens, manual load-balancing between applications is required (that is, you will have to schedule the VLS traffic only when all other applications are inactive). For this reason, HP recommends using a dedicated EVA for the VLS12000.

## Storage Pooling

On the VLS12000 and VLS9000, physical disk capacity can be separated into storage pools. Virtual tapes can then be assigned to a specific storage pool. Having multiple storage pools allows you to:

- Minimize the impact if multiple disks fail at the same time — what is considered a catastrophic array failure. If a catastrophic array failure destroys one or more LUNs on an array, only the storage pool containing that array loses its data; all others retain their data.
- Keep specific backup data physically separate from other backup data. For example, businesses can keep data from one customer separate from that of another customer, or data from different organizations can be kept separate.

Despite the benefits of creating multiple storage pools, maintenance is easier on a system with only one storage pool because the VLS balances performance automatically for a single pool. Load balancing the backup jobs across the pools with the VLS12000 and VLS9000 must be done manually. See [Figure 47](#).



**Figure 47 Single vs. Multiple Storage Pool Considerations**


### VLS9000 Storage Pooling

With VLS9000, the user establishes a storage pool policy which defines how the arrays are pooled. VLS9000 then automatically creates storage pools based on the policy. Storage pools are defined in terms of whole arrays, which consist of one base disk array and three expansion disk arrays. The user can set the storage pool width across the arrays in the system as follows:

- All (default) — One storage pool across all arrays
- 1 — One storage pool across each array
- 2 — One storage pool across every two arrays.
- 4 — One storage pool across every four arrays

Partially populated storage pools may result if you select storage pool policy 2 or 4. For example, if the storage pool policy is 2 and there are three arrays in the system, the VLS divides the three arrays into two storage pools – one storage pool spans across two arrays, and the other storage pool spans

across one array. See [Figure 48](#). If a fourth array is added, the VLS automatically adds it to the storage pool with only one array.



**Figure 48 VLS9000 Partially Populated Storage Pools**

The flexibility of the VLS9000 storage pooling effectively allows the user to partition the disk arrays to fit the user's environment. For example, if a company has four SANs, each with its own backup application, the user can configure the storage pools so that there are two arrays per SAN. The performance of each storage pool is independent and does not impact the performance of the other pools.

### VLS Gateway Storage Pooling

When creating and assigning storage pools, use the following information to ensure even load balancing:

- Storage pools must be manually configured onto the array LUNs.
- Limit the storage pools to 256 LUNs across the VLS device.
- Keep arrays with the same performance characteristics within separate pools (i.e., place one array type in one pool and another array type in another pool).
- In configurations with multiple arrays, choose to balance performance so that either the pool comprises LUNs from every array, or each array is allocated to a separate pool.
- When pooling across multiple LUNs, try to keep the same number of LUNs in each array in the pool.
- Try to keep approximately the same capacity on each LUN within a storage pool. If LUNs have different capacities, there may be performance hot-spots when smaller LUNs are filled.
- When presenting LUNs, present only the disks that you wish to use with the VLS12000 because existing data will be destroyed during storage pool configuration.

## VLS Design Considerations

### VLS Sizing

Sizing is the art of deciding how much capacity you need and how to best use it. Understanding capacity requirements requires you to understand the raw capacity of the VLS (in its various models), the size and nature of your backups, and the retention periods associated with your data.

### Outlining the Steps

Here are some basic exercises that can help you understand your capacity needs:

- Identify slow SAN hosts
- Compare the capacity required to hold the data that represents what is backed up from them

- If you use 'x%' of the VLS to better fit into your backup window, what will that gain you in your backup window and is that the best use of your VLS?
- You should do some backup tests to understand exactly how much the backups are sped up through the introduction of the VLS into the environment before you can make this estimate reasonably accurately.
- Compare the capacity required to hold the data that is most often restored to the VLS capacity. If you use 'x%' of the VLS to speed up your restores, is that the best use of the VLS? How much time in restores will this save? What does the ROI look like for this?
- Compare the capacity required to hold data that is backed up but has the shortest retention periods.
  - If you use 'x%' of the VLS to reduce the amount of physical tape you create, is that the best use of your VLS?
  - To answer this, you should capture the cost of writing, tracking, storing (local and off-site), and retrieving physical tapes.

## Considerations

As you approach the sizing/capacity questions for VLS, see [VLS Technical Specifications](#), and bear in mind:

- When tuning performance, remember that the throughput of the VLS device is only one factor in performance. The throughput may be very fast, but there are a myriad of other environmental concerns that may gate throughput. The host/server, itself, may be slow, for example, or block size may be an issue. Remember to take a holistic view of your environment.
- Hardware data compression is available with many VLS products, and it does not impede data throughput. However, for VLS nodes that have been migrated to a VLS device with hardware compression, the user must overwrite virtual tapes for the hardware compression to work. Using data compression (2:1 ratio) will increase the backup capacity of the virtual tapes. For VLS models without hardware compression, it can also reduce data throughput by up to 50% on a VLS6000-series (in maximum configuration) compared to its native performance. For this reason, the user must decide between maximum performance (compression disabled) or maximum capacity (compression enabled).
- Excessive polling by other devices on the SAN (Windows 2003, for example) can have an extremely detrimental effect on performance.
- Above the entry-level VLS configuration, you need to have multiple concurrent backups going to the VLS to achieve maximum throughput. Host speed multiplied by the number of concurrent host backups equals the aggregate throughput requirement.
- Consider how many hosts you need to have running to keep the VLS optimized but not overloaded. That is, don't try to back up more hosts than you have bandwidth for. If the throughput requirements exceed the bandwidth, you may want to stage backups or increase the VLS capacity.

## EVA Performance Sizing (VLS12000 Gateway)

To get the performance you need from the VLS12000 Gateway, you must also keep in mind the EVA performance considerations:

- EVA tuning must be sized to the desired performance and must take into account the performance of the arrays plus the compressibility of the data and whether deduplication is enabled or not (deduplication requires double the array performance).
- An array at its maximum configuration will give full performance.
- The number of disks on the array can impact performance.
- To make the most of the system, there must be enough disk shelves in the EVA (12/14 disks per shelf) to reach the required performance.

- Use RAID 5 configuration on the EVA in order to get faster sequential access and better capacity efficiency.

## Sizing/implementation Examples

There are no standard, generic rules for sizing your environment but there are general parameters that you can use to help you understand your capacity needs.

### Example 1 VLS Sizing for Performance and Capacity

Assume that you have 12 servers being backed up across a SAN. Each is pushing data at 15 MB/sec. You are writing to four LTO2 tape drives, which have a maximum write speed of about 60 MB/sec with 2:1 compressible data. You are not fitting within your backup window.

Naturally, you want to use your servers, backup window, SAN, and tape drives/library as efficiently as possible. If the servers are only backing up at 20 MB/sec, you are not streaming the LTO2 drives and the backups are taking four times as long to complete, because you are only using the tape drives at one quarter of their performance. You do not want to send the backups over the LAN because it will cause LAN bottlenecks and you do not want to buy more tape drives.

Introduce VLS into your environment. How do you size the VLS in terms of performance and capacity? The average throughput speed of the VLS 6510 base configuration is approximately 225 MB/sec, and the aggregate throughput required by your backups is 180 MB/sec. This fits. But you have additional options.

Now that you have determined the minimum configuration of the VLS (by making sure that the performance of the VLS accommodates your aggregate throughput), consider how much capacity you need to have and how you reach that. Of course this requires not only that you understand your base storage needs (what is the size of the full and incremental backups) but that you factor in retention policies.

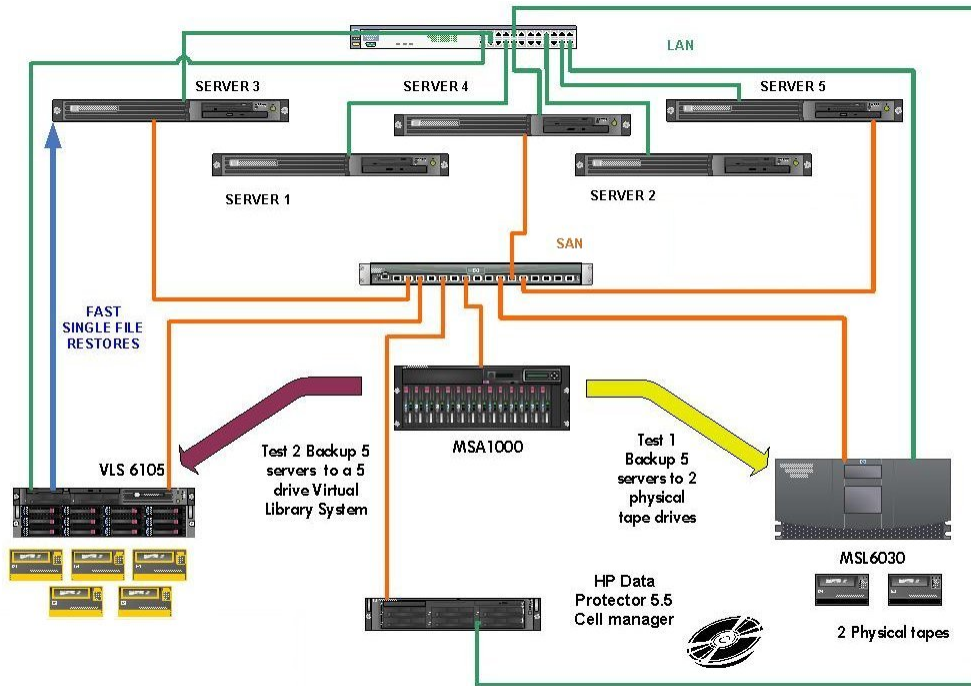
For example, if each host is backing up 400 GB in a full backup and 40 GB for each incremental backup, and if the retention policy for the virtual cartridges is to store one week of backups, and a “week’s worth” is one full and five incremental backups, then each host will require storage for 600 GB of data. This would therefore require 7,200 GB of backup capacity in the VLS. There are two methods for achieving this. You can enable compression when you create virtual tape drives, so that the capacity of the base VLS 6510 would be sufficient, but this would not allow for future performance expansion (e.g., adding more hosts), since compressed performance is limited to 190 MB/sec even if you add capacity upgrade. The alternative, if you want to allow for future performance growth, is to leave the data uncompressed and add a capacity upgrade to the VLS 6510 to increase the capacity to 7,500 GB (this increases the uncompressed performance as well).

Finally, you have to configure copy jobs. If you use the four-drive library for the copies, the VLS will be able to stream the drives. Despite the fact that the copy jobs are going to fewer drives than the backup, the copy time will be approximately equal to the backup time because you’re using the drives efficiently.

### Example 2 Testing Backup to Tape vs. Backup to VLS

In this example:

- One virtual library (VLS 6105) has been configured with five drives
- Servers 1 and 2 have LAN-attached data
- Servers 3, 4, and 5 have shared storage resident on the MSA 1000
- The backup application is Data Protector



**Figure 49 Testing Backups to Tape vs. Backups to a VLS**

If you run two tests, in both cases backing up approximately 13 GB of data from the five servers, the following results:

- Backing up to physical tape takes 6 minutes 41 seconds
- Backing up the same data to the VLS takes 3 minutes 34 seconds

### Example 3: VLS12000 EVA Sizing Without Deduplication

When sizing the EVA array requirements for a VLS12000 without deduplication enabled, the key factor to take into account is the compressibility of the data since the hardware compression built into the VLS12000 nodes will reduce the data rate going to the EVA arrays by the compression factor. For example, a VLS12000 with 5 nodes has a 3000 MB/s backup performance capability. With 2:1 data compressibility this would only require 1500 MB/sec of EVA array performance. Therefore, you would need at least three EVA 4400/6xxx/8xxx arrays (each array with 8 shelves full of disks) to meet that requirement.

### Example 4: VLS12000 EVA Sizing With Deduplication

When sizing the EVA array requirements for a VLS12000 with deduplication enabled, first calculate the base array performance required without deduplication (see above) and then for deduplication you double this performance to get the required array performance. This is because the VLS deduplication is post-process and therefore will add additional disk I/O requirements on top of the backup I/O. For example, a VLS12000 with 3 nodes and deduplication has a 1500 MB/s backup performance capability. With 2:1 data compressibility this would only require 750MB/sec of base EVA array performance plus another 750MB/sec for the deduplication, giving a total of 1500MB/sec array performance. Therefore, you would need at least six EVA 4400/6xxx/8xxx arrays (each array with 8 shelves full of disks) to meet that requirement.

## VLS Blueprints

The following blueprints of VLS virtual tape libraries with deduplication and replication start from specifying company requirements, then defining the HP blueprint for the solution, and finally defining any solution caveats or ISV dependencies associated with the solution. This will help you make informed decisions and allow you to quickly assess areas of concern and possible implementations.

### Reduce Backup Window (No Deduplication)

**Company requirements:** Company has Fibre Channel SAN, needs to reduce backup window, already using tape-based backup application, wants low risk easy integration of disk-based backup.

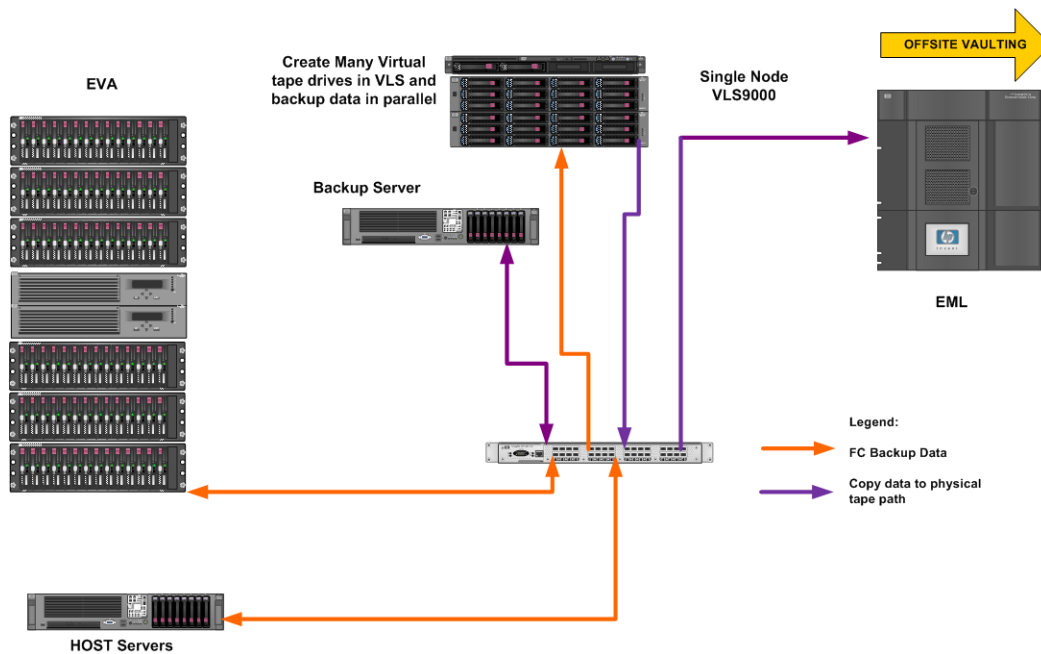
**Solution:** Entry Level VLS9000 with ISV-based copy to tape. No deduplication.

**Caveats:** None.

**ISVs:** Data Protector, Netbackup, Networker, TSM, Commvault, etc. See <http://www.hp.com/go/ebs> for a full list.

**Device configuration highlights:** Highly scalable in capacity and performance and numbers of libraries and drives that can be emulated. Virtual media can be configured to any size.

**Backup application configuration highlights:** The whole premise of VLS integration is that it is seamless when integrating into backup applications, so no changes are required other than configuring object copies for transfer to physical tape.



**Figure 50 Using VLS to Reduce the Backup Window**

**Recovery options:** Typically direct from VLS, single file restores will be particularly fast; set retention periods accordingly. Older copies of data can be recovered directly from physical tape if required.

**Solution advantages:**

- Very scalable up to 1280 TB, 4800 MB/sec with 8 nodes and 16 arrays
- Simple to integrate into existing backup environment
- Up to 16 libraries and 64 drives per node can be emulated

- Cost effective
- Easy disaster recovery solution if physical tape copy is regularly off-sited

**Solution trade-offs:** Take care to size the solution correctly. VLS9000 entry level scales to 40 TB before you need to purchase additional connectivity kits.

**Other information:** HP strongly recommends using the backup application to perform the copy from the VLS to physical tape. This provides a single management console for all activities and better utilization of physical media through object copy technology.

## Large Scale Backup Consolidation (No Deduplication)

**Company requirements:** Up to 500 TB backup storage required in a single unit, high throughput required to meet backup windows.

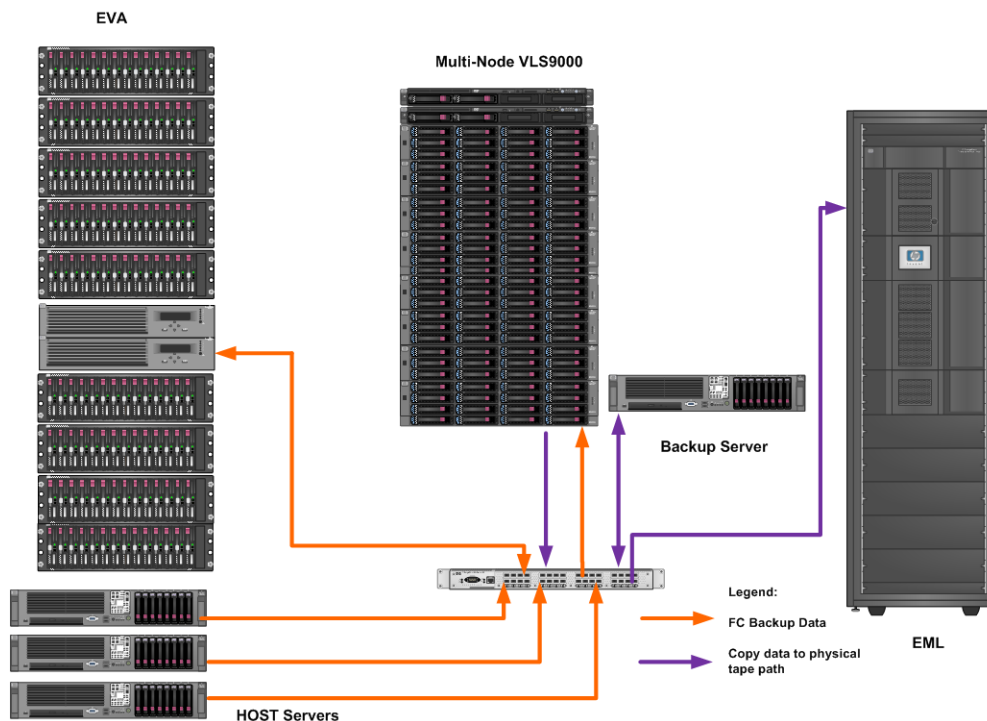
**Solution:** Multi-node VLS9000 using Object copy to physical tape.

**Caveats:** Very large systems (greater than 300 TB) might be better split into two smaller systems for increased availability. With a single large VLS, typically one storage pool is configured which means if any full shelf in the VLS is lost (highly unlikely) then all virtual media is lost. But configuring multiple storage pools to overcome this means increased manageability overhead.

**ISVs:** Data Protector, Netbackup, Networker, TSM, Commvault. See a full list at <http://www.hp.com/go/ebs>.

**Device configuration highlights:** Create as many virtual devices as you need and dedicate devices to certain SAN hosts by using the LUN mapping functionality.

**Backup application configuration highlights:** For best throughput, schedule as many backups in parallel as you can. Use object copy functionality to schedule copy to physical tape.



**Figure 51 Using VLS for Large Scale Consolidation**

**Recovery options:** Typically direct from VLS, single file restores will be particularly fast; set retention periods accordingly. Older copies of data can be recovered directly from physical tape if required.



**Solution advantages:**

- Massive scalability in terms of capacity and performance, “hot add” capability for adding capacity and nodes
- High availability features in nodes and back-end disk arrays
- Flexible emulations in terms of number of libraries, drives, slots, and media sizes
- Easy disaster recovery solution if physical tape copy is regularly off-sited

**Solution trade-offs:** To guarantee scalable performance, each array module of 30 TB or 40 TB is a single component and cannot be split. Sizing capacity and throughput correctly is important to the solution being priced competitively.

**Other information:** HP strongly recommends using the backup application to perform the copy from VLS to physical tape. Single management console for all activities and better utilization of physical media through object copy technology.

## Managing Data Retention and Reducing Backup Storage Costs

**Company requirements:** To improve RPO objectives, reduce expenditure on ever increasing backup storage, better manage data growth for backup, provide more data online, reduced dependency on tape.

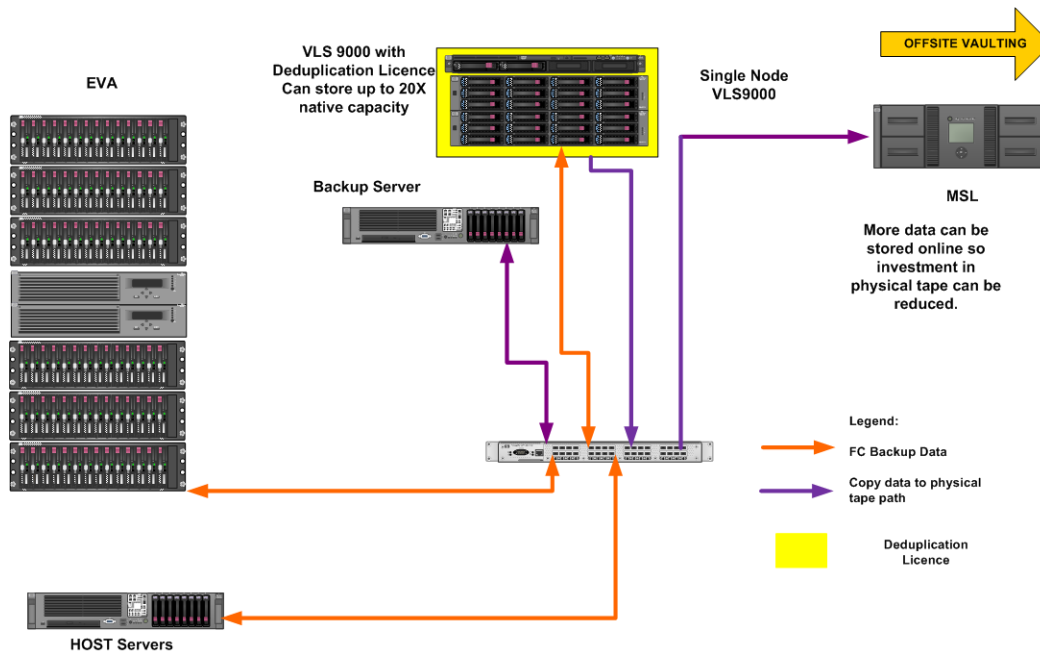
**Solution:** VLS9000 with deduplication (licensable feature).

**Caveats:** Needs specific data type support with ISV software. Best practices include using small cartridges (fewer than 300 GB). Some scripting changes will be required for stream naming conventions with RMAN Oracle backups. Can take several days for space to be fully reclaimed.

**ISVs:** Data Protector, Netbackup, TSM – with Unix files, Windows files, SQL, Exchange, Oracle and NDMP. See <http://www.hp.com/go/ebs>.

**Device configuration highlights:** Size using the HP StorageWorks Backup Sizer; the smaller the virtual media the faster deduplication and space reclamation will happen.

**Backup application configuration highlights:** The only change required is that RMAN scripts need to be changed to ensure consistent stream naming. Needs strict control over backup policy naming.



**Figure 52 Using VLS to Manage Data Retention and Reduce Backup Storage Costs**

**Recovery options:** All recovery from VLS with deduplication is fast, especially the last full backup (typically 90% of restores are from the last backup). Copies from the VLS to physical tape is done by expanding deduplicated data so full content is copied to physical tape. Therefore, data can be directly recovered from physical tape.

**Solution advantages:**

- Deduplication is done post process so it does not slow down the backup.
- Deduplication can be enabled/disabled by data type or policy name.
- Fast restores of deduplicated data.
- Scalable up to 320 TB currently.

**Solution trade-offs:** Deduplication has to be certified with backup applications *and* data types.

**Other information:** HP strongly recommends using the backup application to perform the copy from VLS to physical tape. Single management console for all activities and better utilization of physical media through object copy technology.

**Full Enterprise Backup and Disaster Recovery Capability**

**Company requirements:** Full disaster recovery capability for an Enterprise Data center to replicate data automatically to another site in a cost effective manner with reduced fixed costs (intersite link costs).


**Solution:** VLS with active-to-passive or active-to-active configuration using low bandwidth replication. Deduplication and target replication are licensable features at extra cost. Copy to physical tape at disaster recovery site for archiving and further disaster recovery options.

**Caveats:** Allow 24 hours for replication. Scripting may be required on the target site for trickle background import of replicated cartridges to ISV software backup catalog.

**ISVs:** HP Data Protector, Netbackup, and TSM only at this time. (Deduplication is the key enabler for deduplication-enabled replication.)

**Device configuration highlights:** Setting up replication is done by creating echo copy pools. Identical cartridges are kept on source and target devices and report “up to date” when they are fully synchronized.

**Backup application configuration highlights:** Two separate master servers are required with one at each site.



**Figure 53 Using VLS with Full Enterprise Backup and Disaster Recovery Capability**

**Recovery options:** Primary site data can be recovered directly from the VLS at primary site. If there is a total disaster at the primary site and all servers and VLS are lost, data can be recovered from the disaster recovery site in several ways:

- At the disaster recovery site using a separate backup application, master server data can be recovered to the new servers.
- A replacement VLS and servers can be installed at the primary site and data can be reverse replicated from the disaster recovery site over the low bandwidth link, although this is not recommended for high volumes of data (TBs) because of the time to replicate data wholesale over a relatively small link speed.
- Data can be copied to physical tape at the disaster recovery site then transported to the primary site and imported to new servers and arrays using physical tape at the primary site. Tape could also be copied to the VLS and then recovered from there.

**Solution advantages:**

- Automated disaster recovery over cost effective low bandwidth links
- Flexible recovery options
- Scalable in terms of capacity and performance
- Flexible replication implementations: active-to-passive, active-to-active, and many-to-one (up to four)
- Only target nodes need to be licensed for replication

**Solution trade-offs:**

- Deduplication has to be certified with backup applications and data types. Data Protector, Net-backup, and TSM currently supported.

- A one-time initialization process is required the first time replication runs so the source and target can synchronize. This can be done by co-location, replication over low bandwidth link, or by using physical tape transport – export and import.

**Other information:** HP strongly recommends using the backup application to perform the copy from VLS to physical tape. Single management console for all activities and better utilization of physical media through object copy technology.

## Multi-data Center Replication to a Common Disaster Recovery Site

**Company requirements:** Full disaster recovery capability for a number of data centers to replicate data automatically to a single disaster recovery site in a cost effective manner by consolidating storage at the disaster recovery site and with reduced fixed costs (link costs). Copy to physical tape at disaster recovery site for archiving and further disaster recovery options.

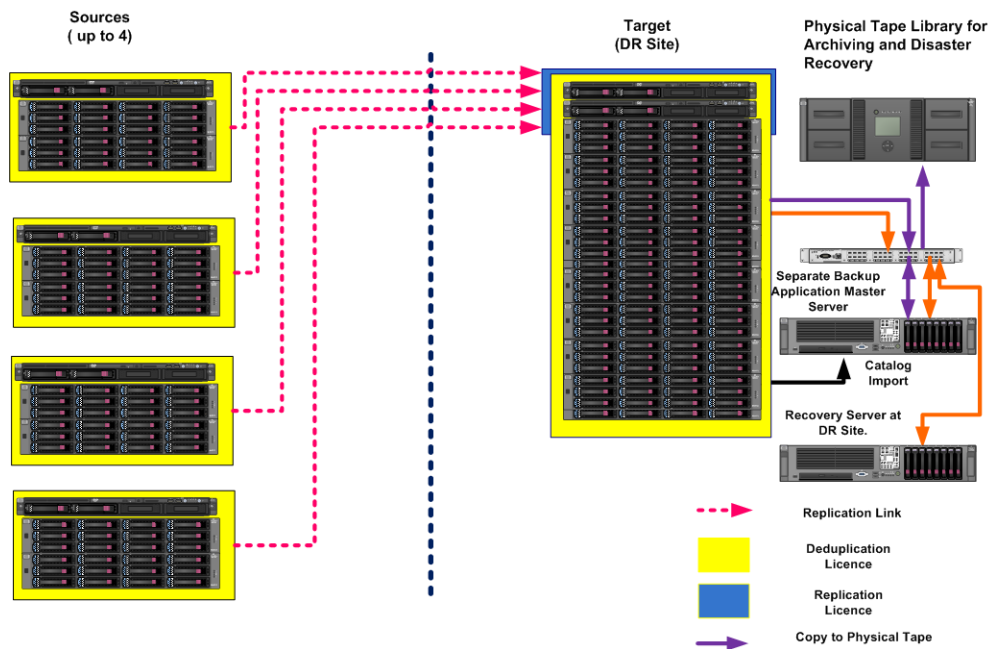
**Solution:** VLS with deduplication and many-to-one replication.

**Caveats:** Allow 24 hours for replication. Scripting may be required on target site for trickle import of replicated cartridges to ISV software backup catalog.

**ISVs:** HP Data Protector, Netbackup, and TSM only at this time. (Deduplication is the key enabler for deduplication enabled replication.)

**Device configuration highlights:** Setting up replication is done by creating echo copy pools. Identical cartridges are kept on source and target devices and report “up to date” when they are fully synchronized.

**Backup application configuration highlights:** Two separate master servers are required with one at each site.



**Figure 54 Using VLS with Multi-data Center Replication to a Common Disaster Recovery Site**

**Recovery options:** Primary site data can be recovered directly from VLS at the primary site. If there is a total disaster at one of the primary sites and all servers and VLS are lost, data can be recovered from the disaster recovery site in several ways:

- At the disaster recovery site using a separate backup application master server, data can be recovered to new servers.

- A replacement VLS and servers can be installed at the primary site and data can be reverse replicated from the disaster recovery site over the low bandwidth link, although this is not recommended for high volumes of data (TBs) because of the time to replicate data wholesale over a relatively small link speed.
- Data can be copied to physical tape at the disaster recovery site then transported to the primary site and imported to new servers and arrays using physical tape at the primary site. Tape could also be copied to the VLS and then recovered from there.

**Solution advantages:**

- Automated disaster recovery over cost effective low bandwidth links
- Flexible recovery options
- Scalable in terms of capacity and performance
- Flexible replication implementations: active-to-passive, active-to-active, and many-to-one (up to four)
- Very effective disaster recovery consolidation in many-to-one implementation
- Only target nodes need to be licenced for replication

**Solution trade-offs:**

- Deduplication has to be certified with backup applications *and* data types. Data Protector, Net-backup, and TSM currently supported.
- A one-time initialization process is required the first time replication runs so the source and target can synchronize. For many-to-one implementations this is typically by WAN link or by using Tape Transport.

**Other information:** HP strongly recommends using the backup application to perform the copy from VLS to physical tape. Single management console for all activities and better utilization of physical media through object copy technology.

For many-to-one implementations especially, sizing the links, capacity required, and replication windows required is quite complex. Please use the HP StorageWorks Backup Sizer.

## VMWare Backup Environment with ESXi.

**Company requirements:** Entry-level VMWare environment. The company requires a single device for backup consolidation across all VMWare machines and other SAN hosts in this infrastructure. Wants deduplication to provide room for growth and longer retention online.

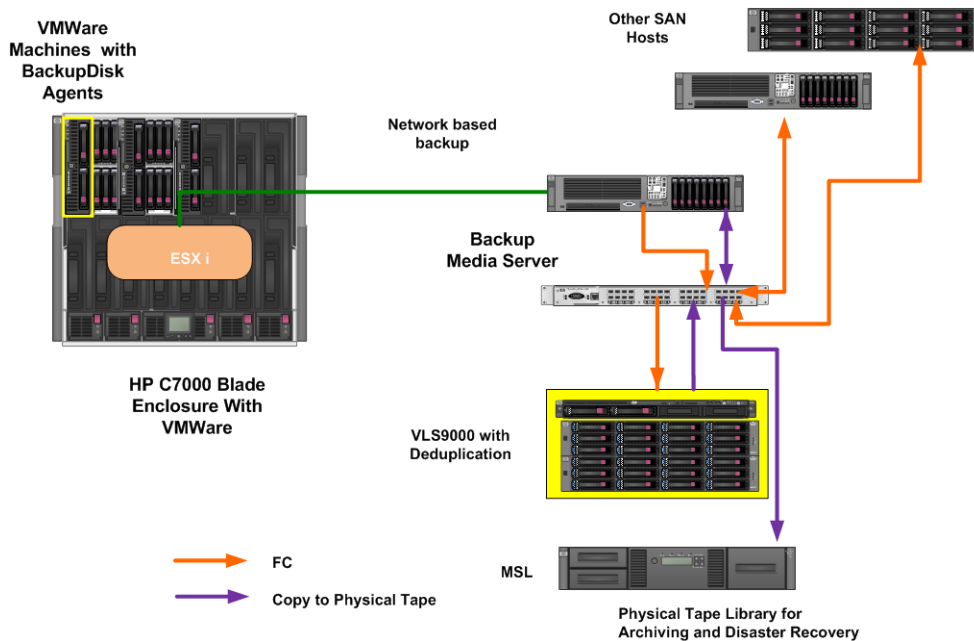
**Solution:** VLS9000 15 TB with deduplication (licensable feature) plus object copy to physical tape for disaster recovery.

**Caveats:** VMWare-based backup with ESXi is over a network which limits the backup performance to an absolute maximum of around 80 MB/sec. Each VMWare machine suffers loading at backup time.

**ISVs:** HP Data Protector, Netbackup and TSM only at this time (VLS supported deduplication ISVs).

**Device configuration highlights:** Configure VLS with dedicated library for VMWare machines, do not multiplex backups, use separate virtual drives for each VMWare machine and backup VMWare machines in parallel.

**Backup application configuration highlights:** None – works as normal network backup, disk agents and online backup agents on VMWare machines, media agents on separate backup server.



**Figure 55 Using VLS with VMWare using Backup Agents**

**Recovery options:** Using the separate server the restore can take place directly from the VLS device over the network to the VMWare machines. If the data no longer exists on the VLS but has been archived to tape then it can be restored directly from tape to VMWare machines.

**Solution advantages:**

- Uses traditional backup methods – well understood
- Significantly improved data retention capability with deduplication
- Online backup of File System and Application Data inside the VMs. Enable single object and application recovery (including point in time)
- Single appliance solution for VMWare and non-VMWare hosts

**Solution trade-offs:** Network Backups can only reach around 80 MB/sec best case. For high volumes of data this may not be fast enough to meet a fixed backup window, but VMWare backups tend to be a lot of small machines with not high volumes of application data.

**Other information:** HP strongly recommends using the backup application to perform the copy from VLS to physical tape. Single management console for all activities and better utilization of physical media through object copy technology.

Further enhancements possible by providing multi-site disaster recovery by using deduplication-enabled replication on VLS devices.

## VMWare Backup Environment with ESX4 and VCB-based Backups

**Company requirements:** Mid-sized VMWare environment, wants VMWare-based snapshot backup with individual windows file-based recovery using VMWare consolidated backup (VCB) in ESX4. Company also wants single device for VMWare backups and other SAN host backups in their infrastructure.

**Solution:** VLS9000 15 TB with deduplication (licensable feature) plus object copy to physical tape for disaster recovery. VCB snapshots created on VCB proxy server then backed up.

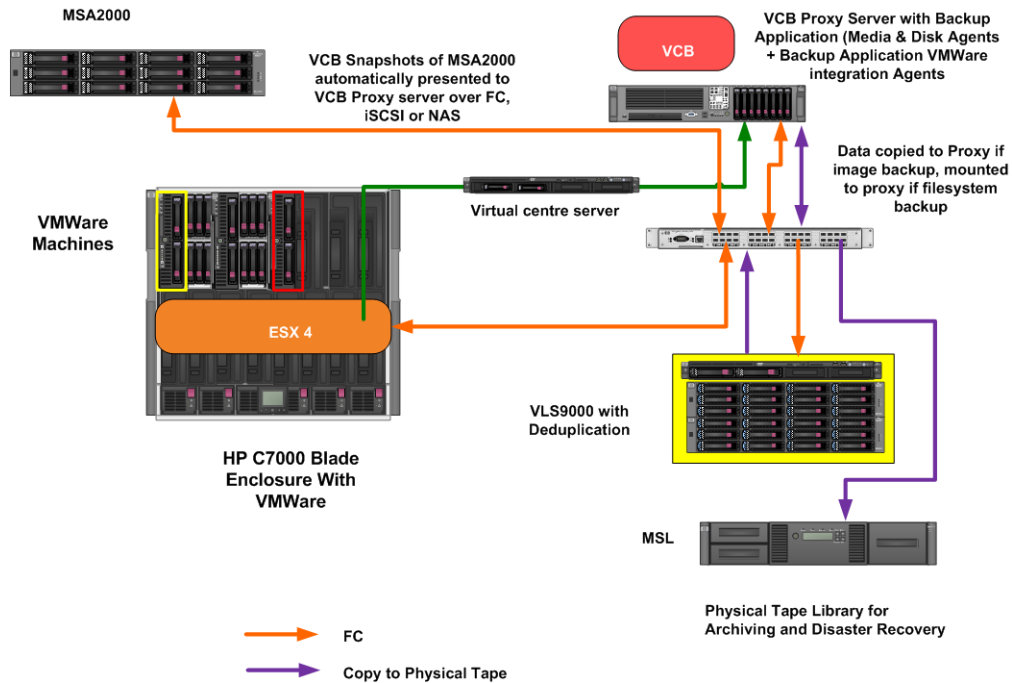
**Caveats:** This is mainly a snapshot-based backup technology where complete images have to be restored. File level recovery from a file system backup is only possible with Windows. The proxy

server has to have sufficient storage to store the snapshots if image-based backups are being used because the images are copied from the primary storage to the proxy server for these types of backup.

**ISVs:** HP Data Protector and Netbackup only at this time. (Tested by HP enterprise backup solutions lab in depth.)

**Device configuration highlights:** Separate virtual library created for each VMWare machine. Backup source is VCB proxy server with snapshot copied to it from the MSA2000 in this case.

**Backup application configuration highlights:** Backup is a combination of VCB doing the snapshots and backup applications writing the snapshots to backup device. VCB is hardware independent hence can accommodate snapshots with all array types.



**Figure 56 Using VLS with VMWare using VCB**

**Recovery options:** Recovery is a two stage recovery:

1. The image is restored to the VCB proxy or the ESX server.
2. The image is restored as a VM using VMware Converter if restored to the VCB proxy, or vcbRestore if restored to the ESX server.

Restart the VM and you are back in business as of the last snap. This is the fastest way to do a complete disaster recovery. For example, HP Data Protector on Proxy can restore directly from VLS or tape onto VM machines.

For Windows OS, you can also mount the file system backup to VCB proxy and restore single files or restore the complete system.

**Solution advantages:**

- VMWare approved method of backing up, takes the load off the ESX console for backup and restore and passes it to the proxy server
- Allows all HP supported Fibre Channel backup devices to be connected to Proxy server
- VCB is disk array independent snapshot technology
- Automation of backup and restore of VM snapshots

- Disaster recovery of ESX environments
- Easy transfer to physical tape if required, if backup application used
- With HP Data Protector and most other backup software the following approaches are possible:
  - Online and offline snapshots of VMs. File level backup and recovery of data in the VMs.
  - VCB Image Backup: Online and offline snapshots - full backup only.
  - VCB File Backup: Full backup. Incremental and differential backups based on per-file modification time. Incremental contains changes since previous backup; differential contains changes since latest full backup.

**Solution trade-offs:**

- Full snapshot-based recovery – single file recovery only on Windows file system backups.
- Distinctly separate proxy server required at additional cost – requires additional storage for copy of snapshots.
- Backup application integrations such as Oracle, SQL, Exchange, etc., cannot be used with VCB.
- VLS cannot deduplicate NetBackup VCB with FlashBackup data type yet which is used for single file restore from a VM image backup.

**Other information:** After snapshot has been copied to proxy and backed up , it is deleted from the primary storage array where it was first taken.

Further enhancements possible by providing multi site disaster recovery by using low bandwidth replication on HP VLS devices.

## Larger VMWare Environment with ESX4 and HP Data Protector ZDB IR Integration with VMWare Consolidated Backup

**Company requirements:** Large VMWare environment, high throughput required to meet backup window (Fibre Channel). High SLAs require "Instant Recovery." Requirement for online application backups using techniques such as HP Zero Downtime backup. Common backup solution required for VMWare and other SAN host in the company's infrastructure. Room for growth, and more backup data online (for SLA) dictates need for deduplication.

**Solution:** HP VLS9000 with deduplication plus HP Data Protector ZDB implementation for SQL, Exchange and Oracle.

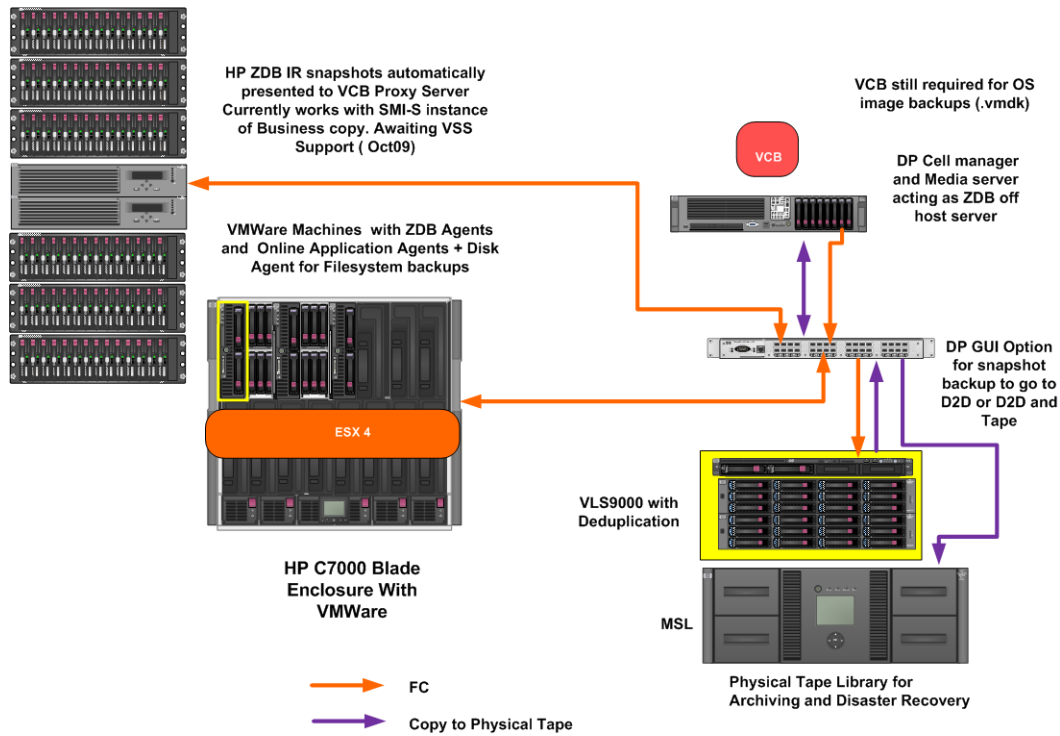
**Caveats:** Limited to arrays that support DP ZDB IR – XP, EVA and LHN (future).

**ISVs:** HP Data Protector and Symantec Netbackup only at this time fully tested by HP.

**Device configuration highlights:** Configure different virtual libraries for SQL, Exchange, Oracle snapshot backups.

**Backup application configuration highlights:** Configuring Zero Downtime Backup within Data Protector and Symantec Netbackup requires some additional scripting work. See the *HP EBS SAN Design Guide* at <http://www.hp.com/go/ebs>.





**Figure 57 Using VLS in Larger VMWare Environments with HP Data Protector ZDB IR Backup**

**Recovery options:**

- Two step restore process that provides full recovery of VM and data
  - Recover the VM image from VCB backup
  - Required only in case the whole VM is lost or restore to a new ESX server
- Recover Application Data
  - For replication-based disc backups (IR), perform Instant Recovery
  - For disc/tape backups, restore directly to the VM
  - Recovery performed automatically by Data Protector as specified by user

**Solution advantages:**

- VCB snapshot-based backup is not really “application aware” whereas HP Data Protector ZDB quiesces the application prior to the hardware-based snapshot. It is much more suitable for live databases and email servers than VCB-based backups.
- Remove the backup load from VM to a backup server with Data Protector Zero Downtime Backup
- Reduce application recovery times from hours to minutes with Data Protector Instant Recovery
- Easy copy to physical tape if required

**Solution trade-offs:**

- Limited to EVA & XP & LHN (Future)
- Snapshots could take up large amount of EVA and XP space depending on frequency of snapshots- at extra cost

**Other information:**

A simple two step backup for complete protection:

1. VM backup

- Data Protector Integration with VMware VCB provides the first layer of protection
- Provides snapshots that can be used for restoring VM
- Needed only when VM configuration changes

## 2. Application backup

- Provides the second layer that enables protection of application data
- Can be performed using ZDB/IR feature of Data Protector by using HP Business Copy for EVA, XP (LHN to follow)

Further enhancements possible by providing multi-site disaster recovery by using low bandwidth replication on HP VLS devices.

## Large Enterprise Cross-site Backup with Deduplication and Inbuilt Disaster Recovery

**Company requirements:** Fully redundant system with automated offsite backup over extended SAN – getting data offsite immediately is the requirement. Manage data growth, keep more data online, regular archiving to tape using object copy.

**Solution:** VLS9000 with deduplication over extended SAN. ESL library for high volume archiving.

**Caveats:** CWDM bandwidth is the critical performance factor in this solution typically using HP Multi-Protocol Routers.

**ISVs:** HP Data Protector, Symantec Netbackup and TSM (deduplication support) only at this time.

**Device configuration highlights:** Normal VLS deduplication requirements – small cartridges, strict backup policy naming.

**Backup application configuration highlights:** RMAN stream naming scripting required for best Oracle deduplication performance.

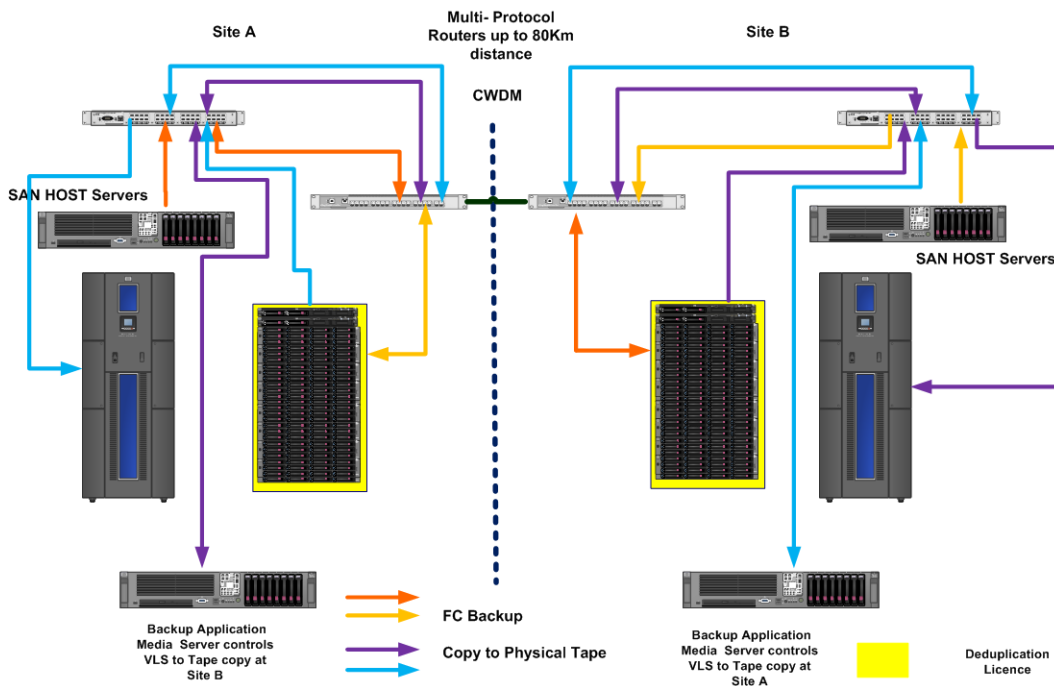


Figure 58 Using VLS in Enterprise Cross Site Backup Scenarios

**Recovery options:** Depending on retention time on the VLS, data can be recovered directly from VLS on the remote site or if that retention has expired then a copy should be archived on physical tape which is again directly available from the remote site over the extended SAN.

**Solution advantages:**

- Data off-sited as part of the initial backup, good resilience solution, no delays
- Part of an Active/Active Data Center disaster recovery solution
- Backup load can be split across sites 50:50 – best use of all available resources
- No data offsite costs for physical tape – it is already offsite with this method

**Solution trade-offs:** Company must be willing to invest in multiprotocol routers and DWDM inter site link technology at some cost.

## Combining D2D and VLS in an End-to-end Solution for ROBO/Regional DCs and Main Data Center using HP Data Protector

**Company requirements:** Protect data anywhere mentality – bring all data home to the main data center. Use the right technology at the right price for the right location.

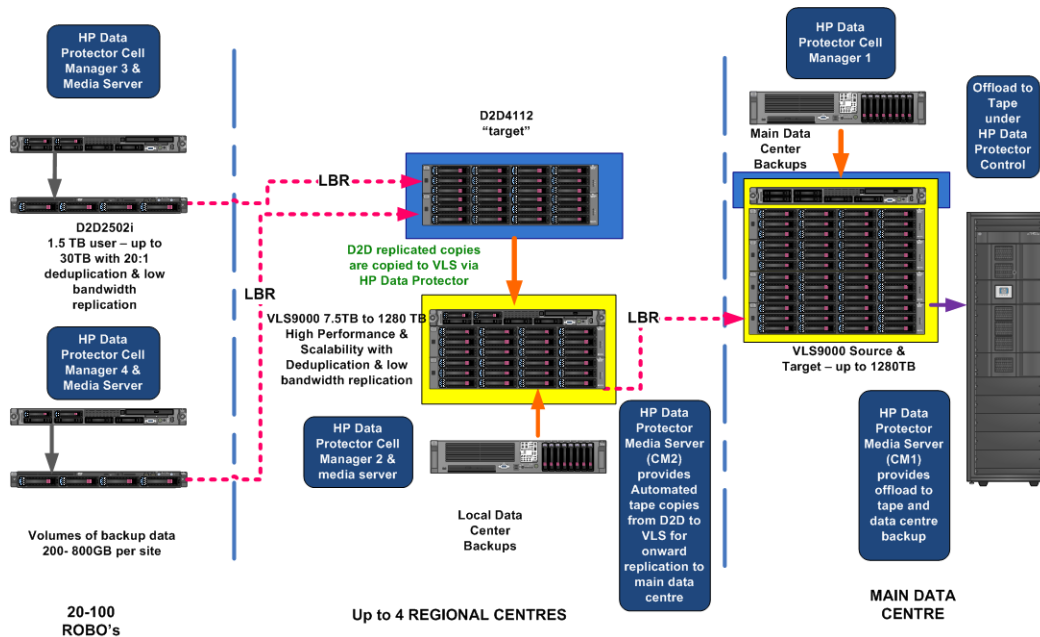
**Solution:** HP D2D in remote offices and regional data centers. VLS in regional data centers and main data center. All with deduplication and replication enabled.

**Caveats:** Sizing replication target size and links is critical; use the HP StorageWorks Backup Sizer. This solution needs a very structured implementation.

**ISVs:** HP Data Protector has the best integration in this end to end scenario with scripted automated importing of the replicated data cartridges at the regional data centers and main data centers.

**Device configuration highlights:** Probably single VTL configured on D2D at remote sites , four slot mappings per D2D virtual library is possible at regional data center. VLS at regional data center should have one library dedicated to accepting data from D2D at regional data center. Regional data center and main data center use of VLS units as local backup resources should be deployed to provide as many parallel backup paths as possible.

**Backup application configuration highlights:** D2D units are polled (scripted) by HP Data protector for recently replicated cartridges, VLS units email Data Protector for recently replicated cartridges. Both of these trigger an automated catalog import at the replication target sites (HP unique feature).



**Figure 59 Using D2D and VLS in ROBO/RDC/MDC Environments**

**Recovery options:** As with all replication solutions using D2D and VLS, under normal situation data is recoverable from local D2D or VLS. If retention period has expired it may be on physical tape at that site. In the event of a major disaster at a site then the replicated copy of data can be accessed in several ways:

- At disaster recovery site using a separate backup application master server data can be recovered to new servers.
- A replacement VLS and servers can be installed at the primary site and data can be reverse replicated from the disaster recovery site over the low bandwidth link, although this is not recommended for high volumes of data (TBs) because of the time to replicate data wholesale over a relatively small link speed.
- Data can be copied to physical tape at the disaster recovery site then transported to the primary site and imported to new servers and arrays using physical tape at the primary site. Tape could also be copied to the VLS and then recovered from there.

**Solution advantages:**

- High Data availability solution, data always held in three locations and on tape
- Total Solution under HP Data Protector control
- Cost effective disaster recovery using low bandwidth links.

**Solution trade-offs:** Some scripting of backup application is required to allow background imports of replicated cartridges to take place

## VLS Automigration

In addition to copying virtual media to physical media via the backup application, another option is to use automigration to perform transparent tape migration. The automigration feature allows the VLS to act as a tape copy engine that transfers data from virtual cartridges on disk to a physical tape library connected to the VLS device. Echo copy is the supported VLS automigration scheme.

## Echo Copy Concepts

Echo copy essentially acts as a transparent disk cache to the physical library, so that echo copy jobs can be performed at times other than the peak backup window. Once automigration is set up, echo copy operations are automatic and seamless to the user. During the normal backup window, the backup application writes to virtual cartridges. Each virtual cartridge that pertains to an echo copy pool has header information that matches a physical cartridge in the tape library. During the user-established automigration window, data migrates from each virtual cartridges to its matching physical cartridge.

Echo copy is managed through the automigration software, not the backup application. For that reason:

- The destination library is not visible to the backup application, so it does not need licensing.
- Copy traffic is no longer routed through the SAN.
- Maximizes performance of high-speed tape drives by freeing bandwidth during the normal backup window.
- All destination tapes created by echo copy are in the native tape format, so they can be restored directly from any tape drive or library that is visible to the backup application.

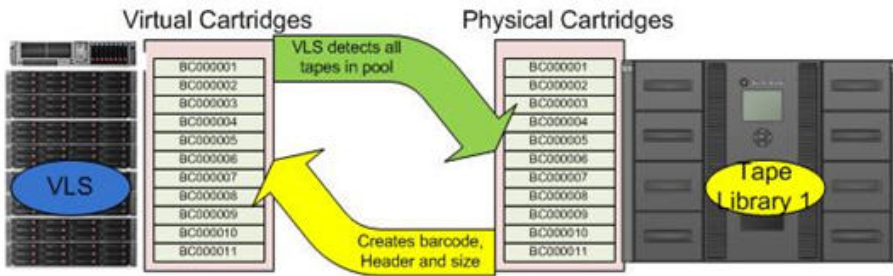
When determining whether to use echo copy, remember:

- The destination library can only be used for copy operations.
- Echo copy is a full tape copy, rather than an incremental change copy, so it can be an inefficient use of media if you are using non-appending media pools in your backup jobs.
- Human error is a factor to keep in mind. With linked media management, any mistakes with the destination library's media will also affect the virtual cartridges. For example, if new tapes are not loaded into the destination library, new, matching virtual tapes will not be created. Subsequent backups will fail because the virtual tapes present are protected.
- The backup application will not be aware of any copy failures.

## Implementation

### Echo Copy Pools

An echo copy pool is used to define which destination library slots are to be echoed by a specified virtual library. Automigration then monitors the echo copy pool slots to detect cartridge loads and ejects in the destination library. Automigration automatically synchronizes the virtual cartridges to the destination cartridges, creating matching virtual cartridges as needed. When new destination tapes are loaded, automigration scans them, and reads the backup application headers. It then creates matching virtual cartridges in the virtual library specified for the echo copy pool the physical tape is in. See [Figure 60](#).



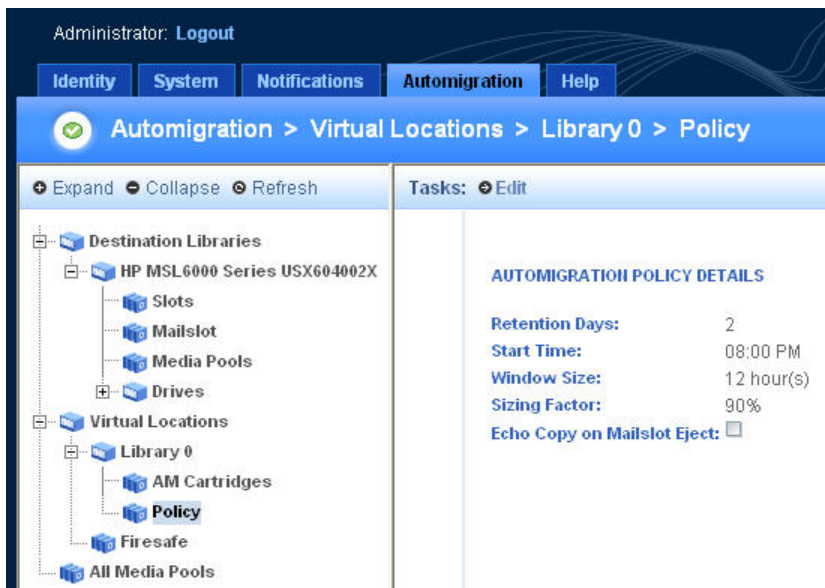
**Figure 60 Linked Media**

When virtual tapes are ejected by a backup application, or when tapes are ejected in the destination library, the matching virtual tapes are moved to the device's firesafe — a temporary holding place. This ensures the device does not fill up its disk space with older cartridges.

When virtual media is ejected, the matching destination tapes are also automatically ejected. If a destination tape is ejected while a copy is pending, notifications in the VLS GUI will inform the user that they have interrupted the copy process. However, if the user selected the Echo Copy on Mailslot Eject option and ejects a tape mid-copy, the copy job increases in priority and completes as soon as possible before the tape is ejected. See [Automigration Policy](#).

## Automigration Policy

Data written to virtual cartridges is automatically copied onto matching destination tapes based on the automigration policy. See [Figure 61](#).



**Figure 61 Automigration Policy Details**

The automigration policy defines the following for each source virtual library:

- Retention days – the number of days information is kept in the firesafe before it is deleted. The default is 2 days.  
While virtual cartridges are in the firesafe, they can be used to access data without having to restore from a physical copy. However, as far as the backup application is concerned, these cartridges

are no longer present. Therefore, to restore data from virtual cartridges in the firesafe, the user must move the virtual cartridges out of the firesafe first. See [Restoring from Automigration Media](#).

- Start time – the time at which the echo copy job begins. HP recommends that copies are scheduled within a different time window from other backup activities.
- Window size – the number of hours available during which the echo copy job can take place. If the migration does not take place during this window, the job is maintained in the queue until the next start time.
- Sizing factor – the percentage of space, compared to the actual size of a disk, allocated on the virtual tape for the copy job. For example, if the sizing factor is set to 90% and the physical tape is 100 GB, the allocated space for the virtual tape will be 90 GB.

The sizing factor is crucial to creating the right size virtual tapes. The sizing factor should be based on the size of the physical tape or the tape type if possible. HP recommends against setting the sizing factor too high (for example 110%). Although the virtual copy can be made, if the drive compression is too low the data might not fit back onto the physical tape. The sizing factor applies to all operations except Load for Restore, in which case the sizing factor is ignored and the virtual copy is exactly the size of the physical cartridge. In general, the virtual cartridge is approximately 90% of the actual cartridge size. Base the size of the copy on the eventual maximum size of the actual cartridge that is to be copied.

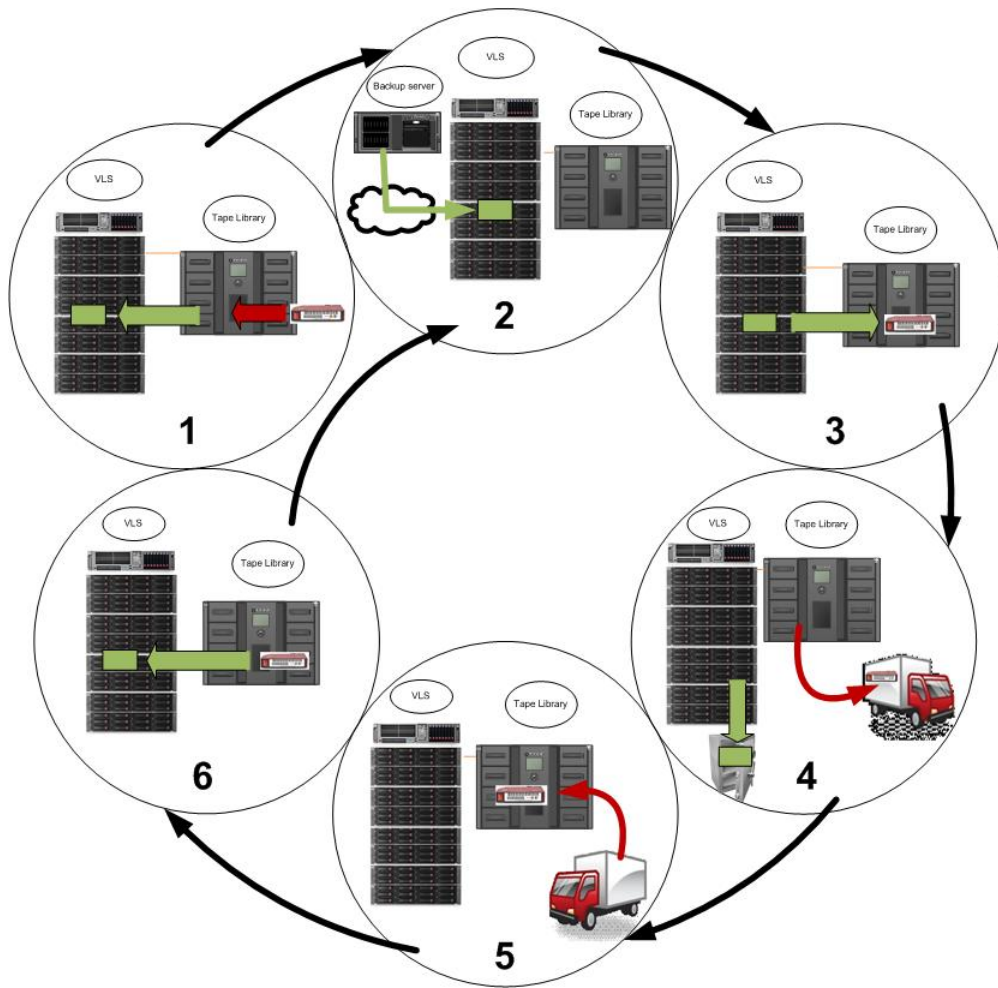
- Echo Copy on Mailslot Eject — when this option is checked, upon mailslot eject, the copy will be started as soon as possible (depending on available drives), regardless of the start time and window size specified in the automigration policy. When the copy is complete, the destination cartridge will be ejected and the virtual copy will move to the firesafe. If this option is not checked, the matching destination cartridge will not be ejected until the copy is complete, which will depend on the start time and window size.

## Automigration Setup

To setup automigration, the first step is to use the GUI to connect a destination library or libraries to the VLS. Once connected, the destination library must be managed, or added as a destination for automigration virtual cartridges. After managing the destination library, the user establishes the automigration policy for each virtual library. See [Automigration Policy](#).

The user then creates echo copy pools. Each echo copy pool defines the single source virtual library, the single destination library, and the slot range to be used on the destination library. Echo copy pools follow the automigration policy of the virtual library they pertain to.

Once echo copy pools have been created, the initial automigration setup is complete.



**Figure 62 Automigration Media Life Cycle**

This section presents an example, in which a user who works for Company X did everything described in the previous section to set up automigration. See [Automigration Setup](#).

The numbered list that follows coincides with [Figure 62](#), which illustrates the stages in the life cycle of automigration virtual and physical media.

1. The user selects **Load Blank Media** in the GUI to load media into the destination library. Automigration automatically creates virtual cartridges with matching header information in the source virtual library.
2. Scheduled backups run to the virtual cartridges in the VLS.
3. During the echo copy window, automigration creates physical copies of the virtual cartridges.
4. The echo copy window ends. One of the following methods is used to eject the physical tapes:
  - The backup application.
  - The automigration tab in the GUI.
  - The library front panel GUI.
  - Manual eject on the physical library.

The ejected tapes from the destination library are transported off-site. Company X has a policy of sending physical tapes off-site for 30 days. When the physical tapes are ejected from the destination library, the matching virtual cartridges are automatically moved to the device's firesafe where they are retained for as long as the user established in the automigration policies. The



user in this example set the retention time for seven days. Therefore, if there is a need to do a restore within the 7-day retention period, the user can simply move a cartridge from the firesafe, and use the backup application to do the restore. When the seven-day retention period passes, the virtual cartridges are deleted, at which point a physical tape is needed for restore operations. See [Restoring from Automigration Media](#).

5. Thirty days have passed. The truck returns with the physical tapes previously transported off-site that are now to be recycled.
6. The media life cycle starts again. The recycled tapes are loaded into the destination library using **Load Media for Overwrite**. After scanning the tape headers, automigration automatically creates virtual cartridges with matching headers. At this point, both the virtual and the physical cartridges are ready to back up data again.


## Design Considerations

### Automigration Use Models

The following sections give examples of backup and restore use cases for automigration.

#### Backup to Non-shared Virtual Libraries

In this use case, the customer has multiple virtual libraries on the VLS device, and each virtual library is mapped to a different backup server. In other words, there is no SAN-sharing of virtual library resources. Instead, there is a dedicated, non-shared virtual library for each backup server. This is a possible scenario for SMBs who use backup applications that cannot be SAN-shared across multiple backup servers (for example, single-server backup applications). This situation could also occur with a large enterprise that has multiple different types of backup applications in the same SAN (for example, due to company acquisitions). In both cases, copying to tape via the backup application is very difficult because without SAN sharing of the virtual libraries, the customer needs a separate tape drive/library for every backup server. With automigration, the customer can use a single tape library and effectively share it across the multiple virtual libraries. See [Figure 63](#).




**Figure 63 Multiple Virtual Libraries**

In this use case, the physical tape library slots are split into separate slot ranges, with each slot range being assigned to a different virtual library. This allows separate backup servers to backup to their own dedicated virtual libraries. During the echo copy window, all the virtual cartridges from the virtual libraries are automatically and transparently copied to the single physical tape library. This scheme provides daily off-site protection and archiving of the physical tapes.

The user can configure additional virtual slots in the virtual libraries. That is, there can be more virtual slots than the number of physical slots mapped to the virtual library, to provide spare slots for restore from disk. These additional slots in the virtual library can also be used to contain purely virtual media (for example, media that has no mapping to any physical tape). The user can then create a separate media pool for the purely virtual cartridges, so that not all backups go to physical tape. For example, the user can direct full backups to the media pool containing the automigration virtual cartridges, and incremental backups to the purely virtual media. Thus, only full backups are going to physical tape. An alternative to this would be to create a separate virtual library for incremental backups which would contain purely virtual cartridges.

### Backup to a Shared Virtual Library

In this use case, the customer has multiple backup hosts on the SAN fabric that are sharing a single virtual library on the VLS device. This is the normal configuration for the majority of enterprise customers. Even though this shared library configuration allows tape copying via the backup application, some customers still wish to use automigration to reduce the licensing costs of the backup application or remove copy data from their SAN. See [Figure 64](#).



**Figure 64 Shared Virtual Library**

In this scenario, there is one or more physical tape libraries that need to be mapped to a single, shared virtual library. The physical library or libraries can be bigger or smaller than the available disk space on the virtual library. If the physical library has more capacity than the available disk of the virtual library, only a subset of its slot range is assigned to the virtual library. This allows multiple backup servers to backup to the shared virtual library, and then during the copy window, all the virtual cartridges from the virtual library are automatically and transparently copied to the one or more physical tape libraries mapped to that virtual library. This scheme therefore provides daily off-site protection of the physical tapes.

The ability to map multiple physical tape libraries to a single virtual library means that customers can start with smaller, single physical libraries. As they upgrade the capacity of their virtual libraries, customers can add additional physical libraries without needing to modify their backup jobs.

The user can configure additional virtual slots in the virtual libraries. That is, there can be more virtual slots than the number of physical slots mapped to the virtual library, to provide spare slots for restore from disk. These additional slots in the virtual library can also be used to contain purely virtual media (for example, media that has no mapping to any physical tape). The user can then create a separate media pool for the purely virtual cartridges, so that not all backups go to physical tape. For example, the user can direct full backups to the media pool containing the automigration virtual cartridges, and incremental backups to the purely virtual media. Thus, only full backups are going to physical tape. An alternative to this would be to create a separate virtual library for incremental backups which would contain purely virtual cartridges.

## Sizing the Tape Library

When sizing the physical tape library (or library partition) for use with automigration, the number of tape drives should be selected depending on the required copy performance. A single automigration copy stream runs at 100-120 MB/sec, so six copy streams per node will saturate the node at 600 MB/second aggregate copy performance. Maximum copy performance would therefore be achieved using six LTO3 or LTO4 tape drives per node (or ten LTO2 tape drives) multiplied by the number of

nodes in the VLS. If the copy window does not need maximum copy performance, you can use even fewer tape drives per node.

For example, a 4-node VLS9000 device that needs to automigrate 64 TB every night within an eight hour copy window (that runs after the backups are complete) requires an aggregate copy performance of 2200 MB/sec. With a 100-120 MB/sec per copy stream, this implies a tape library would need 18-22 LTO3/LTO4 tape drives to achieve this copy performance. Remember these tape drives should be evenly spread across both back-end Fibre Channel ports of every VLS node. In other words, half the drives should be connected to the first back-end Fibre Channel port on every node, and the other half should be presented to the second back-end Fibre Channel port on every node.

## Restoring from Automigration Media

When using automigration to create physical backups, if the destination tape is still loaded in the destination library, its matching virtual cartridge will still be present in the virtual library. In this case, the user can simply restore from the virtual cartridge using the backup application.

If the destination tape has been ejected from the destination library, the user has the following options:

1. If the automigration virtual cartridge still exists in the firesafe (for example, if the policy-defined retention period has not expired), use Command View VLS to manually move the cartridge back into the virtual library. You can then restore from the virtual cartridge using the backup application, and then move the cartridge back to the firesafe.
2. If the automigration virtual cartridge has been deleted from the firesafe, retrieve the physical cartridge by one of the following methods:
  - Load the physical cartridge into any compatible physical drive or library that is visible to the backup application and restore from that tape using the backup application.
  - Load the physical cartridge into the destination library and use Command View VLS to perform a **Load Media for Restore**, which copies the entire destination tape back into a virtual cartridge. You can then restore from the virtual cartridge using the backup application.

---

### NOTE:

Customers may choose to configure their devices such that there are additional, unmapped slots in the virtual library. These additional virtual slots can be used for backups so that even if the physical slots are all full, there are still spare slots available to hold the cartridges needed for restore. Another option is to have a separate virtual library for restore operations.

---

## VLS Accelerated Deduplication

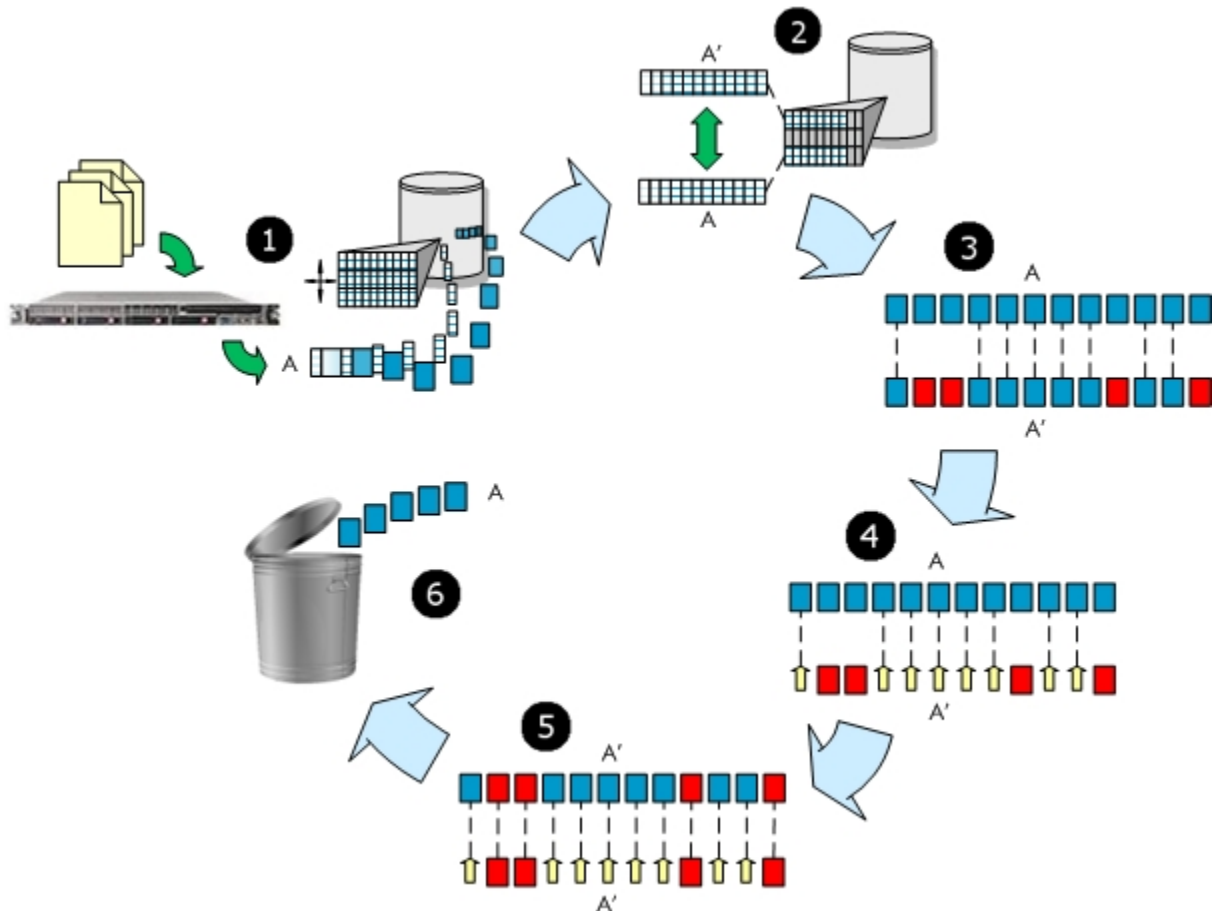
HP Accelerated deduplication technology is designed for optimal performance and scalability to meet the needs of enterprise data centers. It offers the following features and benefits:

- Leverages object-level differencing code, which targets matching backup “objects.” Rather than comparing the current backup with every byte of data stored on the device, Accelerated deduplication can compare backup objects from the current backup to matching objects from the previous backup, where there is likely to be duplicate data. This allows the deduplication to be performed faster, consuming fewer resources.
- Is performed as a post-process. This delivers the fastest possible backup performance, because there is no loss of bandwidth to the normal backup window.
- Allows data to be retained longer because redundant information is removed from previous backups, freeing disk space.

- Delivers fastest restore because it maintains a complete copy of the most recent backup data.
- Older data can be restored faster because it is still on the VLS – there is no need to locate physical media, load it into a drive, and perform the restore in the traditional manner.
- Provides scalable deduplication performance which can be increased simply by adding nodes. However, the multi-node VLS still maintains the manageability of a single device.
- Provides maximum hardware efficiency because the virtual cartridge file system was designed to be deduplication aware from the beginning, and data is not copied to a separate deduplication store.


## How it Works

Figure 65 illustrates the steps in the process that backup data undergoes with Accelerated deduplication. The explanations that follow correspond with the numbers in the figure.



**Figure 65 Steps of Accelerated deduplication**


1. **On-the-fly backup analysis** Figure 66 illustrates the steps in the first phase of Accelerated deduplication. Although Accelerated deduplication uses post-processing technology, backups are analyzed initially while the backup is running. This process has minimal performance impact. The deduplication software uses the metadata attached by the backup application to decode the format of the tape. The resulting “metadata” of the tape’s contents is stored in a database for later use. This database stores a “map” of the logical backup objects on the tape, with tape format headers removed to reveal the actual data.



Item	Description
1	The VLS analyzes the backup as it goes through memory.
2	The backup application metadata is stripped and used to understand the format of the tape.
3	The metadata database creates a map of the locations of the logical backup data.

**Figure 66 Backup Analysis**

In [Figure 67](#), two instances of File A are shown, one from Session 1, and one from Session 2.



At the backup object level, each version of File A is different because the metadata attached by the backup application contains different information. However, at the logical level, these files are the same. Accelerated deduplication strips away the backup application metadata to reveal duplicated data that may be present in the logical data.

**Figure 67 Physical versus Logical Data**

2. **Data grooming** This phase is triggered when a tape is unloaded either inside a drive, or when it is ejected to a library slot. This phase only occurs if backups are complete. The on-the-fly backup analysis has already been performed, and the metadata database is queried to find out:
  - a. what new backups are on the tape


- b. the names of the backup jobs
- c. the server being backed up
- d. the type of backup – full or incremental
- e. the type of data in the backup – files, database, etc.

The deduplication software then queries the metadata database to find an equivalent older version of the same backup job to compare it against the new backup. If the current backup is full, it will be compared against the last equivalent full backup version. If the current backup is differential or incremental, it will be compared against the next equivalent full backup version (TSM "incremental forever" backups are all treated as full backups by deduplication).

3. **Data comparison** This phase is also called delta differencing. After the data grooming phase is complete, and if there is an older version of the backup, either full or incremental, the new backup is compared, or delta-differenced against it. The deduplication software looks for differences at the byte-level between backup objects.

4. **Data reassembly** In this phase, duplicate data is replaced with pointers, and pointers are readjusted so that they point to the most recent instance of data. If the new backup is a full backup, this phase does not touch the new data, so that the most recent backup is complete and available for restores. Instead, a new copy of the *older* version is created. See [Figure 68](#).

If the new backup is incremental, it is reassembled so that duplicate data is replaced with pointers to the last full backup version.



Item	Description
1	Backup session 1 consists of files A and B.
2	Duplicate data from A is replaced with a pointer to A'. The unique data from file A, as well as file B in session 1 one is retained.
3	The pointer from A to A' is readjusted to point to A'' in session 3, the most recent instance of that data.
4	Pointers are placed from files A' and C in session 2 to A'' and C' in session 3.

**Figure 68 Backup Sessions with Duplicate and Unique Data**

5. **Optional integrity check** If enabled, the optional integrity check runs a byte-for-byte read comparison between the original tape and the new reassembled version of that tape. This option is disabled by default.
6. **Space reclamation** After the data reassembly creates a new version of the older backup tape, the original tape is deleted, freeing up disk blocks. The new version of that tape is renamed to the same barcode as the original tape.

 **NOTE:**

A tape must be full before reclamation of its data will occur.

## Accelerated Deduplication Implementation

The following must be performed to implement Accelerated deduplication:

- Upgrade from 4 GB to 16 GB RAM per VLS node
- Upgrade to firmware version 3.1.2 or higher



- Install deduplication licenses
- Configure the available deduplication options using Command View VLS

A complete backup is necessary to initialize Accelerated deduplication, but data in subsequent sessions can then be deduplicated. See the *HP StorageWorks Virtual Library System users guide* for the complete setup procedures.

## Supported Backup Applications and Data Types

Accelerated deduplication can only be used with supported backup applications and operating systems. (In firmware version 3.3.0 and higher, all data types within the supported backup application and operating system matrix will deduplicate.) See <http://www.hp.com/go/ebs> for more information on backup applications and operating system types supported with HP Accelerated deduplication.

## Deduplication Preparation

Before upgrading the device, review your current free disk capacity to ensure that there is enough to hold the deduplication database (1 TB for VLS6000 models, and 2 TB for VLS9000 and VLS12000 models) and deduplication working space (requires a minimum of 5% of your disk capacity). For example, if you have a VLS6600 with 53 TB capacity you need at least 3.65 TB of free capacity. If you have a VLS9000 or VLS12000 with 120 TB capacity, you need at least 8 TB of free capacity. If you do not have enough free disk space (i.e., all the device disk capacity is allocated to virtual cartridges), you must delete some virtual cartridges (those holding the lowest priority backup data) to free up enough disk space for the deduplication database and working space.

You must then upgrade the VLS to the minimum hardware/firmware level that supports Accelerated deduplication:

- For VLS nodes that current have 4 GB of RAM, you must upgrade them with an additional 12 GB RAM so that every VLS node in the device has 16 GB RAM. Nodes with the "B" part number (e.g., VLS9000 node AG310B) will already have 16 GB of RAM pre-installed.

To determine how much RAM you have on each node, log into Command View VLS, select the **System** tab, expand **Chassis** in the navigation pane, and select a node. The Total Memory is listed in the General section of the node details. The GUI reports usable memory and will display memory of 14465MB when 16 GB RAM is installed in the node.

- Multiple storage pools are not currently supported for deduplication. Before upgrading to firmware version 3.x, if your VLS9000 or VLS12000 has multiple storage pools you must destroy the existing pools and change them to a single storage pool. If your VLS has only one storage pool configured, the upgrade will not destroy the existing backup data on the virtual cartridges.

On the VLS9000 you can use the **Edit Policy** task and select the "All" arrays option; this destroys all existing virtual cartridges. On the VLS12000 you can retain Storage Pool1 but delete all other existing storage pools, then add their EVA LUNs to Storage Pool1 (selecting the "Format" option); this retains the virtual cartridges in Pool1 but destroy the virtual cartridges in the other pools.

- Upgrade to firmware version 3.3.0 or higher. The *HP StorageWorks Virtual Library System Firmware Version 3.3.0 release notes* covers the detailed installation process which uses the version 3.3.0 Quick Restore DVDs to reinstall the deduplication-capable firmware on all nodes in the device. The VLS warm failover feature maintains the original device configuration so that the resulting upgraded device will present its original virtual libraries with the same WWPNs and the original virtual cartridges with all their backup data intact.
- Install deduplication licenses and wait for the "Completed deduplication installation" and "Initialization for node0 complete" notifications.
- If deduplication failed to initialize due to lack of free space for the deduplication metadata database (1 TB for VLS6000-series, 2 TB for VLS9000-series and VLS12000), you must free up that disk

space by deleting cartridges. In the 3.2 firmware this is disabled in the GUI, so use the VLS CLI to perform this.

---

 **NOTE:**

Use only the CLI cartridge delete option to delete non-deduplicated cartridges.

---

After deduplication is successfully enabled, ensure virtual cartridges are created correctly for optimum deduplication performance and for correct backup capacity:

- If you are creating new cartridges for deduplication, then sizing should be 50-100 GB up to 300 GB maximum (see [Media Management](#)). If you are sizing cartridges for a VLS that has existing backup data created when the VLS was running 2.x firmware, the Storage Pool created by the 2.x firmware has lower cartridge limits. You should size the virtual cartridges to at least 100 GB or larger.
- If you already have existing non-deduplicated cartridges that exceed these cartridge sizing requirements, you must resize these existing cartridges to match the deduplication requirements. This will not affect any existing backup data until the next backup is performed to the resized cartridge.
- Because deduplication automatically enables cartridge oversubscription, you can create more virtual cartridge capacity than the available physical disk capacity (the deduplication will reduce the amount of physical disk used on the virtual cartridge backup data by eliminating duplicate data). Only create the amount of virtual cartridge capacity that you need for your backup retention policy (see [Tape Oversubscription](#)). This means you need to ensure you create enough virtual library slots to contain the oversubscribed set of virtual cartridges. For existing virtual libraries that do not have enough slots, you must modify the existing virtual library to increase the number of slots (either delete/recreate the virtual library with more slots or contact your HP support representative for assistance in expanding the existing virtual library).
- Remember to update the backup application library configuration to detect any new library slots and new virtual cartridges.
- Use shared media pools (i.e., share the media across backup policies/names and across media servers). See [Media Management](#).
- Review [Detailed Backup Application Guidelines for VLS](#) for backup application-specific deduplication guidelines such as setting 256 KB tape record size, etc.

A complete backup is necessary to initialize Accelerated deduplication, but data in subsequent sessions can then be deduplicated. See the *HP StorageWorks Virtual Library System user guide* for complete setup procedures. See [Migrating your Existing Backup Data](#) for details on how to handle existing non-deduplicated backup data. After deduplication has been running for multiple full backup cycles, you can check the deduplication reports to see how effectively it is running; see [Configuration and Reporting](#). Deduplication configuration options are available through Command View VLS to allow optimization of file server deduplication (see [Optimizing File Server Deduplication](#)) or handling backup application-specific data type configurations (see ???).

## Licensing

Accelerated deduplication is a licensed feature. Licenses are based on the physical disk capacity of the VLS device. You must install enough deduplication licenses for all disk LUNs.

---

 **NOTE:**

You cannot license a partial configuration. The entire capacity of the device must be licensed.

---

There are three licenses available, one for each VLS platform – VLS6000–series, VLS9000–series, and VLS12000 Gateway. See [Table 24](#).

**Table 24 Required Deduplication Licenses by Platform**

Platform	Deduplication licenses required
VLS6000–series	One license per MSA20 for 250 or 500 GB drives Two licenses per MSA20 for 750 GB drives Two licenses per MSA20 for 1 TB drives
VLS9000–series 7.5 and 10 TB systems	One license per base unit One license per expansion kit <sup>1</sup>
VLS9000–series	Three licenses for 30 TB capacity Four licenses for 40 TB capacity
VLS12000–series Gateway	One license per 2 TB EVA LUN

<sup>1</sup>The VLS9000 7.5 TB system does not require a license for a third expansion kit because the first two expansion kit licenses cover the storage provided by the third kit.

## Migrating your Existing Backup Data

If you are enabling deduplication on an already active VLS, the following process can be used to retain your existing backup data (stored before deduplication was enabled) and gradually overwrite it with new backups that will then be deduplicated. This is because deduplication only works on backups that are stored after deduplication was enabled.

After you enable deduplication on your VLS and perform all the preparation steps:


1. Perform your normal backups.
2. Wait one data retention cycle for all the existing virtual cartridges to be expired by the backup application and then rewritten with new backup data. If there are cartridges that would not normally expire, you may need to manually expire them to force them to be reused. At the beginning of the next retention cycle, all of your old backup data should have been overwritten with new backup data that has been deduplicated (freeing up disk space).
3. Once all the existing virtual cartridges have been rewritten and deduplicated, you can extend your retention policy for your backup jobs to take advantage of the deduplication.


## Configuration and Reporting


Accelerated deduplication is aware of the backup contents of tapes. Therefore, you have reporting and policy control options available over the deduplication process. You can:


- Enable and disable deduplication by backup job type or individual backup job.
- Switch from backup to File-level differencing for file server backup jobs.
- View deduplication status and statistics per backup job, by backup job type, or by cartridge.
- View a detailed breakdown of the capacity used by reference data, deduplicated data, data awaiting deduplication, and metadata.


The deduplication reports show the following states:


-  **Waiting for Next Backup** — the backup has nothing to difference against and is therefore waiting for a new backup. Incr/diff backups difference against the new full backup, not the old full backup.

 Delta-diff in Process — the backup has identified another version of itself to difference against and is now running differencing to identify the duplicate data between the two versions. With multi-stream backups, this process may take multiple tries (going back to "Waiting for Next Backup" state each time) until the differencing locates the correct stream.

 Pending Reclamation — the differencing is complete (duplicate data has been identified) and space reclamation is now waiting for full cartridges before it can begin.

 Partially Reclaimed — in a backup job that spanned multiple cartridges, the full cartridges have been space reclaimed but the last tape in the spanned set (which is partially full) is waiting until that tape is full before it can reclaim space.

 Dedupe Complete — deduplication, including all space reclamation, is complete for this backup job.

 No Match Found — indicates backups that cannot be deduplicated because they were unable to match with their previous version due to large-scale modification.

In the detailed Backup and Cartridge reports, the "Dedup Ratio" represents an estimate of the overall compression ratio (including both deduplication and compression) of that backup/cartridge based on the amount of logical backup data stored versus the amount of physical disk space used. If the backup job status is "Waiting for next backup," "Delta-diff in Progress," or "Pending Reclamation," the ratio will only include the compression because deduplication will not have actually space reclaimed that tape and therefore will not have saved any disk space. If the status is "Partially Reclaimed" or "Dedupe Complete," the backup disk space has been partially or fully reduced by deduplication and the ratio represents compression plus deduplication. The "Physical Size" value in GB follows the same logic.

Use the System Capacity report to see how much disk space is saved by deduplication. The physical used space will vary as deduplication creates temporary working data and then eliminates the duplicate data.

## Handling the Device Out of Capacity Condition

In VLS firmware version 3.3.0, there is a graceful shutdown of backup services in the VLS when the device reaches 98% full disk capacity. This ensures that there is no loss of data access when the storage capacity is full. The VLS will reject all attempts to write data from the backup application at 98% full capacity, so you must reformat/erase unnecessary cartridges via the backup application to free storage. All read operations will remain unaffected so you can still restore existing backups on the device or copy backups to tape prior to reformatting them to free up disk space. Follow the recommended sequence to free up disk space when the device is full:

1. Use the VLS GUI to generate a Cartridge Utilization report (via **Notifications Report Setup**). This report lists the barcode of each cartridge, disk space usage in bytes, deduplication status, and dependent cartridges (if any) which can be used to identify the cartridges that you can reformat/erase to create free storage. For example, first reformat/erase cartridges that are consuming a considerable amount of space and are expired, and then reformat/erase non-deduplicated cartridges which have no or few dependent cartridges. When reformatted/erasing non-deduplicated cartridges, the disk space is only freed up if dependent cartridges are also reformatted/erased.

2. Use the backup application to reformat/erase the cartridges identified by the Cartridge Utilization Report as good candidates to free up disk space. For example, in Netbackup you can expire and relabel the cartridges, in Data Protector you can recycle and format the cartridges, and in TSM and NetWorker you can relabel the cartridges. In the backup application, you may need to re-enable tape drives that were offlined or verify media that was marked as Poor because backups were disabled in the VLS.
3. In the VLS GUI, use the **Reclaim Space** task in the Chassis pane to initiate an immediate space reclamation (to free up the disk space on the reformatted/erased cartridges). A "Space Reclamation is scheduled" notification is displayed on the notification screen.
4. If the VLS disk capacity drops below 90% utilization (due to reformatting/erasing cartridges and reclaiming that disk space), backup writes are automatically re-enabled in the VLS. However, if you need to re-enable writes with a disk utilization above 90% (and below 98%), use the VLS GUI option **Enable Writes** in the Chassis pane to do this.

## Design Considerations

Because every enterprise backup environment is different, there is no single answer for optimizing deduplication. This section discusses deduplication design considerations to help you maximize performance in your environment.

For guidelines based on your specific backup application, see ???.

## Performance

Because the various phases of deduplication may be distributed among nodes, you must determine how many nodes are necessary to optimize the performance in your environment. The following concepts can help you determine this.

1. The ingest rate of backups, which is the amount of data that is backed up within the backup window.
2. The number of nodes required for the maximum ingest rate.  
For example, if your nightly backup consists of 40 TB to complete in eight hours at 5 TB per hour (1388 MB per second) ingest rate, each node can ingest 500 MB per second. Therefore, three nodes are required for this backup job.
3. Maximum daily deduplication limit per node.  
For example, if 15 TB can be deduplicated in 24 hours by one node, deduplication of 40 TB per day requires three nodes.

---

 **NOTE:**

If the number of nodes required for backing up data differs from the number of nodes required to perform deduplication, select the larger number of nodes.

---

---

 **NOTE:**

On the VLS6200, do not enable software compression and deduplication together. Enabling software compression will drastically reduce the performance if deduplication is enabled.

---

## Capacity Sizing


Several factors affect how much disk space is required for the desired retention scheme:

- Data compression rate.
- Daily change rate.
- Retention policy, or the number of full and incremental backups retained.
- “Scratch” space required for post-processing.

Download the HP StorageWorks Sizer tool to help determine your capacity needs. The tool is available on the HP website: <http://www.hp.com/go/storageworks/sizer>.

## Node Oversubscription

With Accelerated deduplication on the VLS9000 and VLS12000, you can add more nodes to enhance the deduplication performance of the systems. For every standard node (with virtual library devices configured on its front-end Fibre Channel ports), you can add one extra node dedicated to deduplication processing (with no virtual devices configured on its front-end Fibre Channel ports). You can add up to the maximum of eight total nodes (currently the maximum number of nodes for a deduplication-enabled device is four). For example, if you have a 2-node VLS9000 system licensed for deduplication, you can add up to two additional nodes to increase deduplication performance. See [Figure 69](#). This configuration must have a minimum of two VLS9000 arrays configured to meet the performance sizing rules for the two standard nodes; you do not need to configure additional arrays when adding the extra two nodes for deduplication. See [Table 25](#).



**Figure 69 VLS9000 and VLS12000 node oversubscription**

**Table 25 VLS9000 and VLS12000 Node Oversubscription Options**

Number of arrays	Number of nodes			
	1	2	3	4
1	X	Y		

2	X	X	Y	Y
3		X	X	Y
4		X	X	X
5			X	X
6			X	X
7				X
8				X

X = node for backup and restore, Y = node for deduplication processing only.  
Note: the maximum deduplication configuration currently supports four nodes.

## Optimum Record Sizing

You must configure your backup application to use a record size (i.e., tape block size) of 256 KB when using deduplication.

## Optimizing File Server Deduplication

There are many reasons why you may get a low deduplication compression ratio.

For example, deduplication does not occur on:

- Files less than 32 KB in size for File-level deduplication.
- Files less than 8 KB in size for Backup-level deduplication.
- File names that keep changing (e.g., some database dumps).
- Backup sets less than 1 MB in size.

If your file server contains the above items, it will not deduplicate efficiently.

In these cases, switch from File-level to Backup-level deduplication for more optimal results.

---

### NOTE:

Deduplication statistics and reports only include the data that is eligible for deduplication. For example, deduplication does not occur on unsupported formats. Therefore, that data is not included in the reports.

---



---

### NOTE:

For database flat file backups that have static file names (such as Oracle file backups), with VLS firmware 3.1.2 or higher you do not need to manually change from File-level differencing to Backup-level differencing because optimum performance and deduplication ratios for database file backups are achieved with the default File-level differencing setting. This does not apply to standard file backups or database file backups with varying names (such as database dump files created with a Back Up Database Task within an SQL Maintenance Plan) where you still need to switch to Backup-level differencing in the cases listed above.

---

## Media Management

To make the most of storage savings with Accelerated deduplication, consider the following:

- Reclamation does not occur until a tape is full (and everything on the tape is either delta-differenced, disabled, or unsupported).
- The deduplication software cannot deduplicate versions of a backup that are on the same cartridge; the versions are not deduplicated until a new version is written to a different virtual cartridge.

Use the following guidelines:

- **Appending media pool**

HP recommends using standard, appending shared media pools across multiple tapes. By doing so, tapes fill with backup jobs from different hosts, which allows deduplication to occur and storage to be reclaimed.

You should not use non-appending media pools.

- **Virtual cartridge sizing**

Generally it is always better to create smaller cartridges than larger ones because it frees up the cartridge quicker for use by other processes such as restore, tape copy or deduplication processing. Cartridge sizes of 100-200 GB are common and this size can be used regardless of the tape emulation type because the cartridge is virtual (i.e. you can create a 100 GB cartridge even if you selected LTO4 drive emulation).

However, with Accelerated deduplication there are other sizing considerations that could recommend even smaller cartridge sizes. If the daily amount of backup data written to your cartridges is significantly less than the cartridge size, the deduplication process will only be able to eliminate the duplicate data (i.e., free up disk space) once the tape has been filled up by multiple days of backups.

In other words, if your average backup job size is significantly smaller than the virtual cartridge size:

- Divide the number of jobs by the number of virtual tape drives. This determines approximately how many jobs will write to each cartridge per night and thus how much logical data is written on average to each virtual cartridge.
- Divide this logical data per cartridge by the average data compression ratio to get the native capacity required.
- Compare the native capacity required to the cartridge size. If the native capacity written per night is less than the cartridge size, reduce the cartridge size to roughly match.

For example, if you have 50 backup jobs per night with an average size of 25 GB each, writing to 10 virtual tape drives, this writes 125 GB logical data per drive ( $50/10 * 25$ ). With 2:1 compression you get 62.5 GB native disk capacity per drive per night. In this case, use a cartridge size of 50 GB.

The recommended virtual cartridge sizing for most environments is 50-100 GB. You may need to configure a subset of your cartridges at a larger size if some of your backups are very large and you want to reduce the number of tapes the backup spans. (Spanning too many tapes reduces deduplication performance.) The maximum recommended cartridge size is 300 GB.

## Client and Backup Job Naming

Deduplication relies on the names of the backup job and the client to correlate them.

- If you rename the backup job or client after you have deduplication running on your system, the deduplication for that backup data is reset (as though you just activated it for the first time).



- Do not divide full backups and incremental backups across different job names. They must operate under a common backup job name.
- If a client is removed from the backup application configuration (e.g., the client is decommissioned or renamed), you must set that client's backup policy/name to Deduplication Disabled in the VLS GUI. Otherwise, the last backup of that client will remain in the "Waiting for next backup" state.

## Disabling Deduplication on Specific Backup Policies

To improve deduplication performance, you can disable deduplication in Command View VLS for specific backup policies where the backup data has a very high change rate (and thus does not deduplicate effectively). Examples of such backup policies:

- Pre-compressed or pre-encrypted backups
- Database redo/archive log file backups
- Database block-level incremental/differential backups (because these are already reduced by the database agent to only the blocks that changed and are therefore always unique new data)

---

### NOTE:

Never use an external hardware compression or encryption solution in front of the VLS because this will disable all deduplication and compression inside the VLS. (If all backup data is pre-compressed or pre-encrypted, it will always be 100% unique data.)

---


## VLS Replication

The VLS deduplication-enabled replication technology (which leverages Accelerated deduplication) is designed to maintain performance and scalability to meet the needs of enterprise data centers:

- On a multi-node VLS device the replication is also multi-node capable and thus can scale as the device scales.
- Fast restore performance is preserved on the replicated cartridges on the target device by maintaining the most recent full backup as a complete copy and then deduplicating the older backup cartridges (just like the source device).
- Optimized post-process architecture means that the deduplication analysis of new backup data is only run once (on the source) and the results from this are used for the space reclamation on the source device, the replication delta transfer, and the space reclamation on the target device.

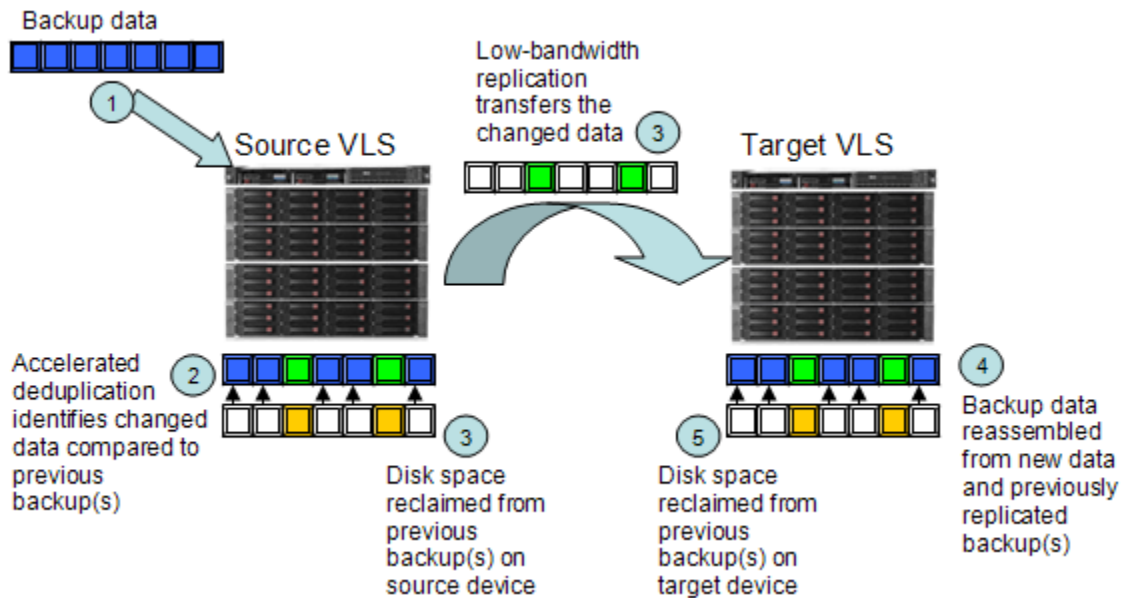
## How it Works

The HP VLS Accelerated deduplication is the key technology enabler for replication. Understand where replication fits into the post-processing deduplication architecture:



**Figure 70 How Replication Fits into the Deduplication Architecture**

1. As the backup takes place, a metadata database is created.
2. Using the metadata compiled in step 1, the deduplication software running on the VLS nodes identifies which backups to compare against which similar backups in order to remove duplicate data.
3. Once the backup job is finished and the tape (which does not need to be full) is unloaded, the system compares the current backup with previous backups identified in step 2. Identifying the differences typically takes about the same length of time as the backup.
4. Once the differences are identified, the backup on the source begins to replicate. The deltas (changes) are sent across the TCP/IP replication link. The space reclamation of the previous backup tape on the source device (not the tape currently replicating) runs in parallel with the replication.
5. Once the backup has successfully replicated, the backup is reassembled using the transferred deltas and the older existing duplicate components on the target device.
6. Space reclamation of the previous backups (not the tape just replicated) on the target device begins. The target runs its own space reclamation to eliminate the duplicate data on the older backups on the target.



**Figure 71 The Replication Sequence**

1. A backup is performed (and its metadata is stored in the database).
2. The new backup is compared to the previous versions and the unique data in the new backup is identified.
3. The source device performs space reclamation on the previous backups to eliminate the duplicate data, and concurrently the replication data (containing just the unique data and metadata) is transferred to the target device.
4. The replication data is reassembled into an intact backup combining the new unique data with the duplicate data already replicated from previous backups.
5. The target device performs space reclamation on the previous backups to eliminate the duplicate data.

HP VLS backup systems with HP Accelerated deduplication allow only unique pieces of data to replicate across the WAN link. This efficiency in understanding precisely which data needs to replicate may result in bandwidth savings in excess of 90% compared to transmitting the full contents of a cartridge at the source site. In addition to the unique data transmitted, there is some overhead of data instructions that also needs to pass across the replication link. This is known as control data, and is estimated to be about 1% of the full backup size. The control data is optionally compressed before it is transmitted. There are also some variances in the amount of data replicated from database backups and file system backups because of the different deduplication algorithms used and differences with incremental backups.

## Replication of Incremental Backups

Most database incremental backups consist of block level changes in the database, while most file level incremental backups consist of changes at a file level. Depending on your recovery point objective (RPO) requirements at the disaster recovery site, you may choose whether or not to replicate incrementals. With the current capacity-optimized comparison algorithms in the HP VLS, the amount of data replicated for file system incrementals is almost the entire size of the incremental itself. (In firmware version 3.3.0 or higher, differential backups will deduplicate against each other so the amount of data replicated is reduced to the size of incremental backups.) Database block-level incrementals will also always be replicated as the entire size of the incremental backup because the data is always unique.

## Multi-node Replication Scaling

The HP VLS9000 and VLS12000 architectures allow formidable scaling in terms of capacity and performance. Supporting multiple nodes, the backup throughput, deduplication comparison, and the replication process can be load balanced across all the available nodes to ensure no bottlenecks to performance. Because the VLS hardware should never be the bottleneck for replication, the LAN/WAN link will be the bottleneck. For example, a 4-node configuration backing up 60 TB per day could require a 1000 Mbit link for the daily replication traffic; this bandwidth is easily achieved by balancing the replication workload across the four nodes.

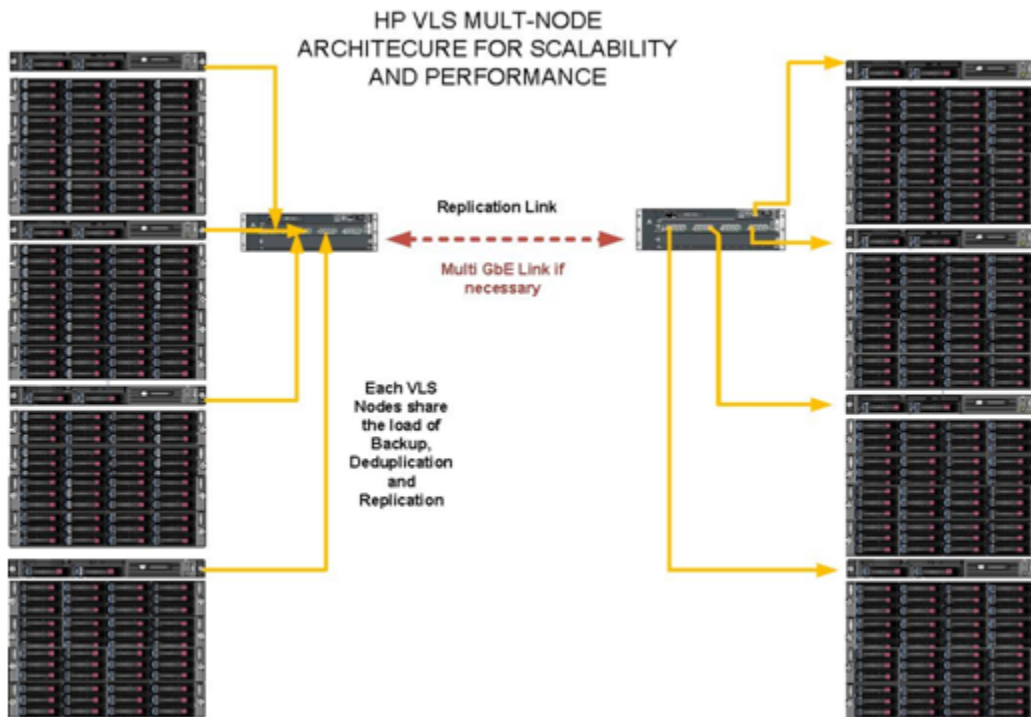


Figure 72 Replication Scaling

The main reason to have replication LAN links on each node is to distribute the replication processing evenly across the nodes (not to increase the bandwidth). Always use all of the nodes for replication or you risk overloading the master node.

## VLS Replication Implementation

You must perform the following to implement replication on VLS systems.

### Licensing

On HP VLS devices, the replication is licensed as follows:

- Requires Accelerated deduplication to be fully licensed and enabled on both the source and target devices.
- Requires VLS replication licenses on the replication target device, with one license for each node configured in the target device. In active-passive deployment, you would install one license per node in the passive target device. In many-to-one deployment, you would install one license per node on the one target device (this then supports up to ten source devices, and additional licenses

are required for more than 10 sources). In the active-active deployment, you would install one license per node in both the source and target devices (for all nodes in those devices).

## Replication Preparation

Prior to configuring replication, ensure that deduplication is licensed and enabled on both the source and target devices and enter the replication licenses on the target device. Create the required virtual libraries on the source and target devices.

You also need to create the required virtual libraries and virtual tape drives on source and target VLS devices. With VLS, the virtual libraries/drives on both source and target device are created as normal (the replication configuration is overlaid onto these standard virtual libraries by creating replication targets, LAN/WAN destination libraries, echo copy pools, etc.). If your source device's backup application has media agents in the target site then when mapping the replication library to the front-end Fibre Channel ports in the VLS, remember to hide the target virtual library from the source backup application (e.g., assign it to dedicated front-end Fibre Channel ports that you can zone away from the source backup application agents). Best practices include creating more library slots than initially needed in the source and target virtual libraries to provide room for future growth.

---

### NOTE:

On a VLS6200 that is acting as a target device, you must enable compression when creating the virtual tape drives in the target virtual library (because all the replicated cartridges on the target device will be software compressed and you will not be able to read these tapes via the target virtual library unless it has compressed tape drives). If the source device is also a VLS6200, this should have software compression disabled on the source virtual library tape drives because deduplication performance will be significantly degraded if software compressed drives are used for backups.

---

After creating the virtual libraries, create the virtual cartridges. If starting from a blank state on both source and target devices, you only need to create virtual cartridges in the replication target library. The matching cartridges are automatically created in the source virtual library when the echo copy pool is created.

However, if you are starting with an existing source device that already has a virtual library containing virtual cartridges with backup data, in order to retain the existing backup data you must use the following process to prepare for replication:

- Create a virtual library on the target device with enough slots to hold the number of current source device cartridges that you want to replicate (plus additional slots for future growth).
- Manually create matching virtual cartridges (with the same barcodes, cartridge sizes, and cartridge type) in the target virtual library that match the source device cartridges that you want to replicate.
- When you create the echo copy pool on the source device, this automatically links the matching cartridges on the source and target.

---

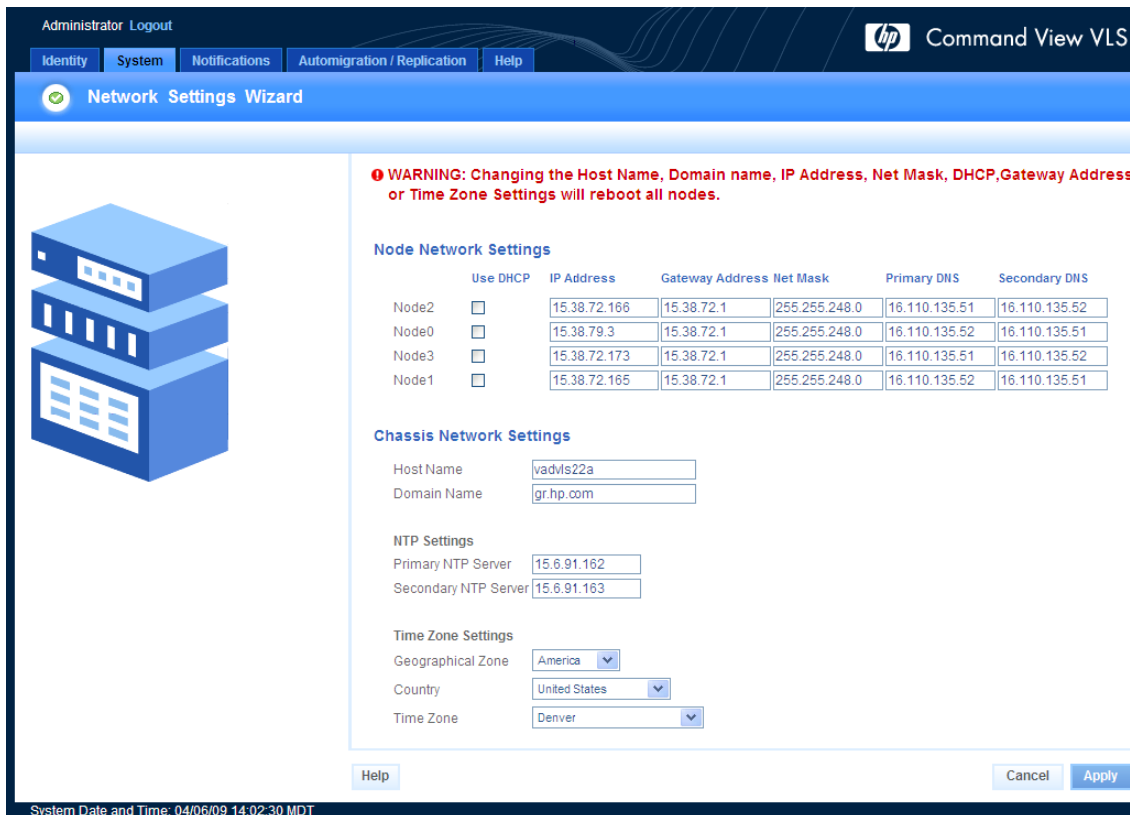
### NOTE:

Only backups performed *after* replication is configured will be replicated. Backups that were performed before replication configuration will not be replicated until the tapes are overwritten by the backup application (for example, after the retention time for the older backups the virtual cartridges will be recycled and reused for new backups which will then replicate).

---

You must also prepare the TCP/IP network connection between the source and target devices. In VLS devices, the node0 will already have a TCP/IP address assigned to its external Gbit LAN connection to allow device management (GUI, CLI, etc.) and this existing connection is also used for node0

replication traffic. However, in multi-node device configurations you must also assign TCP/IP addresses to every node (matching the static or DHCP mode on node0) in both the source and target VLS devices in order to balance the replication load across the entire device. (See [Multi-node Replication Scaling](#) for details.) The VLS will automatically balance the active replication jobs across all the nodes that have an external TCP/IP address configured. The VLS GUI provides the means to configure the secondary node TCP/IP settings in the “Network Settings” task under Maintenance Tasks on the Chassis screen (note that the order of the nodes on this screen is not sequential):



**Figure 73** Preparing the Network Connection for Replication in a VLS

**NOTE:**

If you have any network firewalls between your source and target devices then you must ensure the firewall is configured to allow traffic on the required ports needed for replication to work:

- The source VLS must be able to contact the target VLS on HTTP port (80), SSH port (22), and ICMP.
- The replication data transfers will use the ‘Un-Encrypted Data Transfer Port’ setting with a default value of 5570 (which you can modify on the target).

## Replication Setup

After preparing your source and destination VLS devices, the following steps configure replication. See your VLS user guide for details. [Figure 74](#) shows an overview of the steps outlined below.

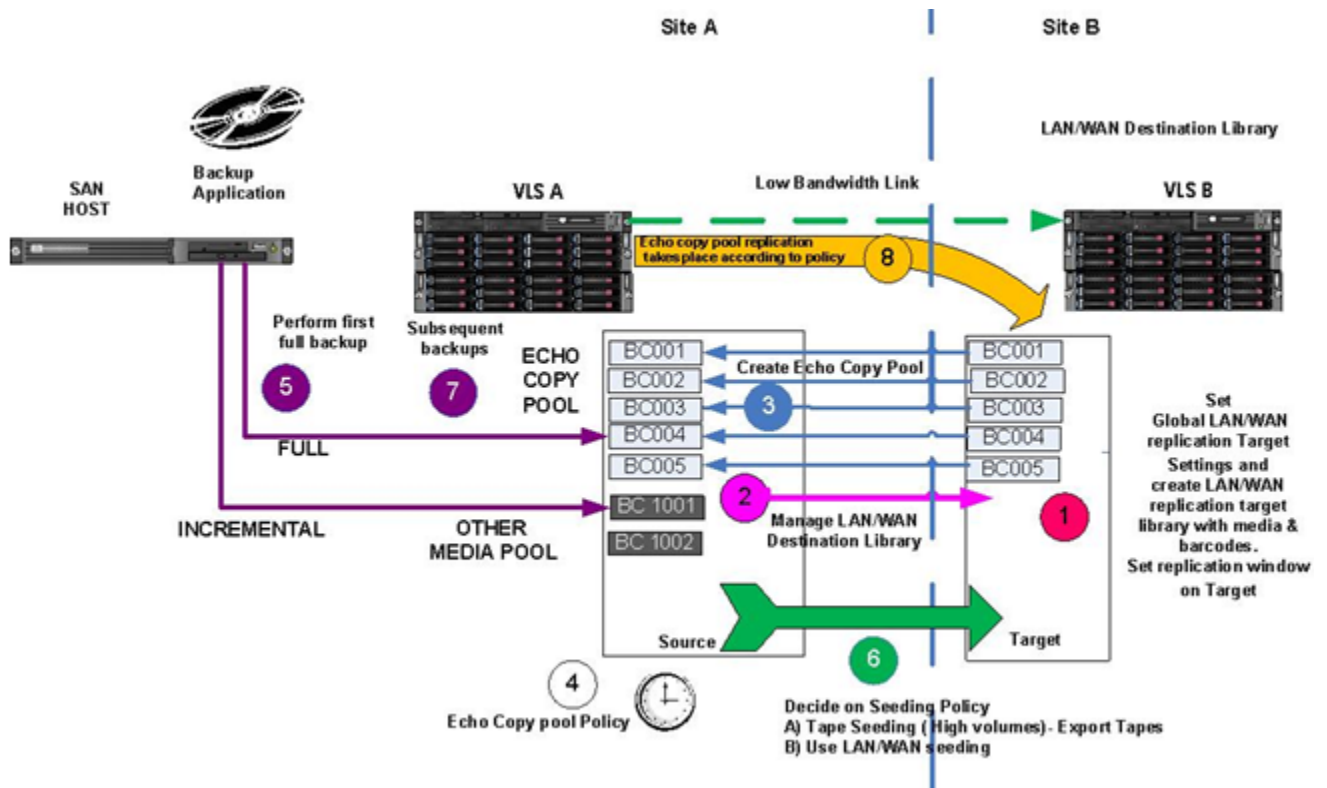


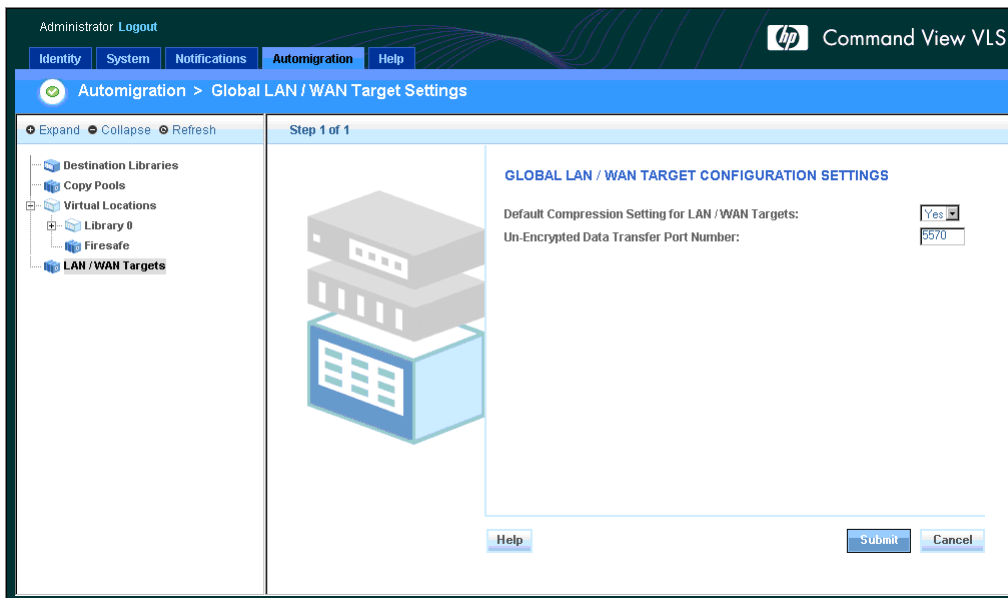
Figure 74 VLS Replication by Numbers

1. Create a replication target on the target device.

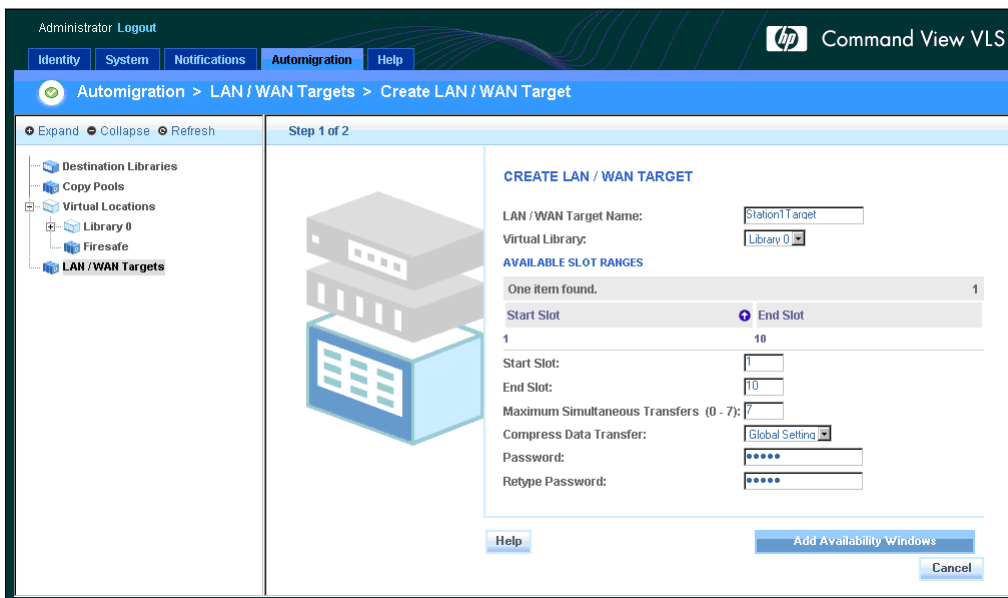
On the target device, set up “Global LAN/WAN Target Settings” and then “Create a “LAN/WAN Target” with:

- Name
- Start and end slots
- Maximum simultaneous transfers
- Password
- Target replication start window
- Target replication duration window (typically 24 hours)

Set up the global LAN/WAN settings by going to the Automigration/Replication tab and setting global settings for compression over the replication link and the TCP/IP port number used for replication data transfers.

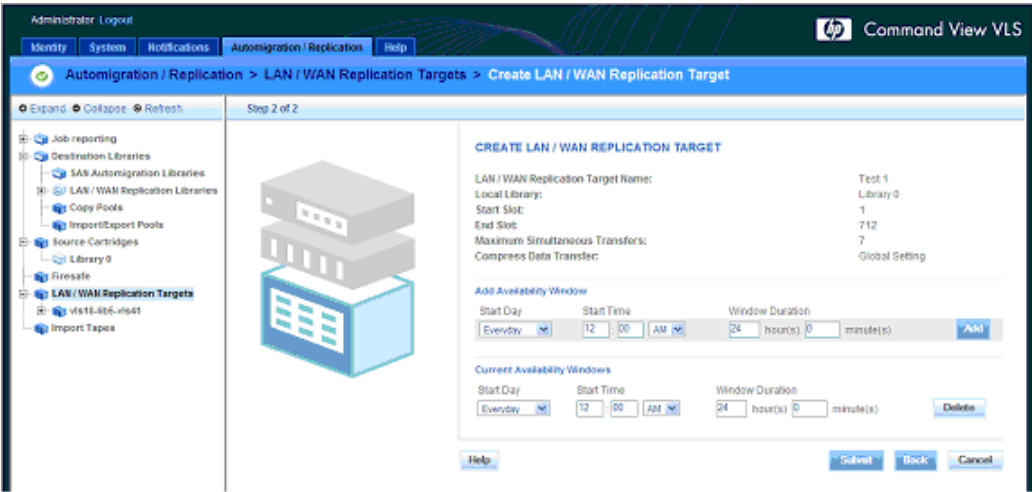


Configure a LAN/WAN target, selecting names, slot mappings, and password for access.

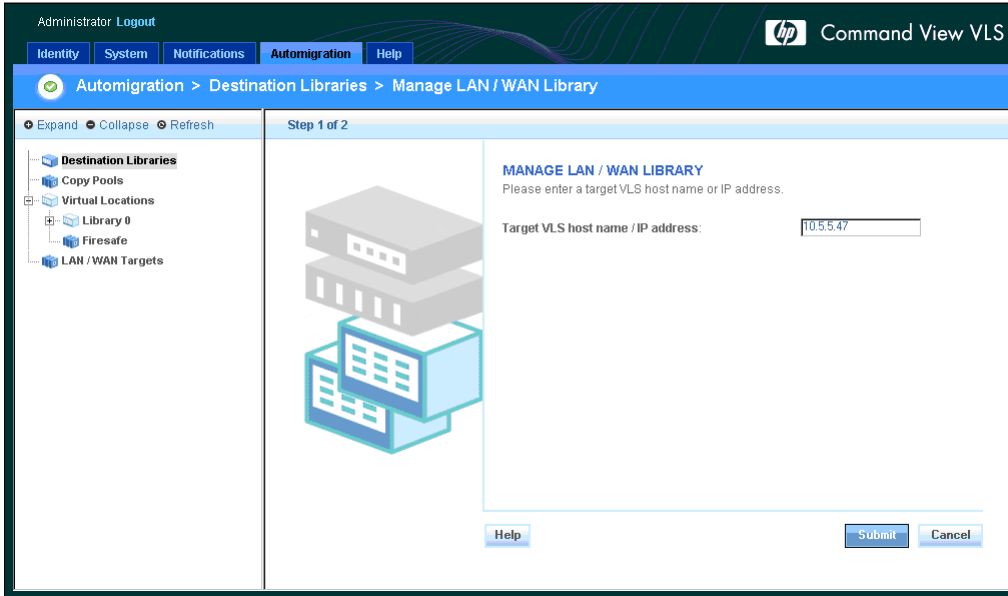




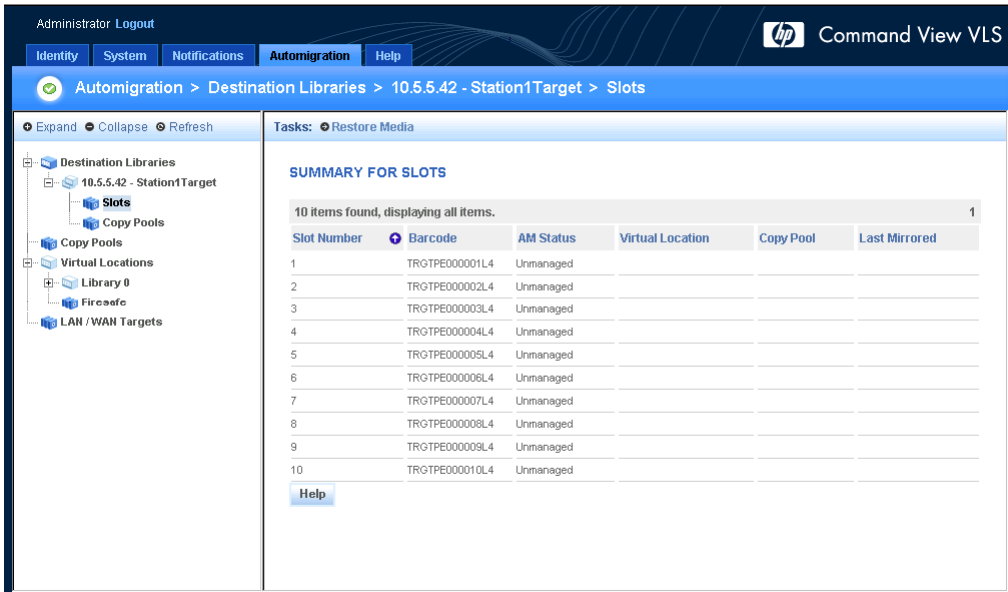
Configure the replication target's window limits — start time and window duration.



2. Add “LAN/WAN Destination” to the source device (to link the source device to target device).  
On the source device, use “Manage LAN/WAN Library” and the relevant password to link the source and target devices together in a replication pair. You can see the target VLS, slots, and copy pools on the target from the source.



Note that the slots in the target LAN/WAN library at this stage are shown with an automigration status of “Unmanaged” because they are not yet associated with a copy pool.



3. Create an echo copy pool (including selecting whether to use tape initialization).

On the source device, select the destination library and “Create echo copy pools.” In this context, the echo copy pool acts as an asynchronous mirror between the selected source virtual library and the target library. The barcodes created initially on the target device are now also “echoed” in the source virtual library.

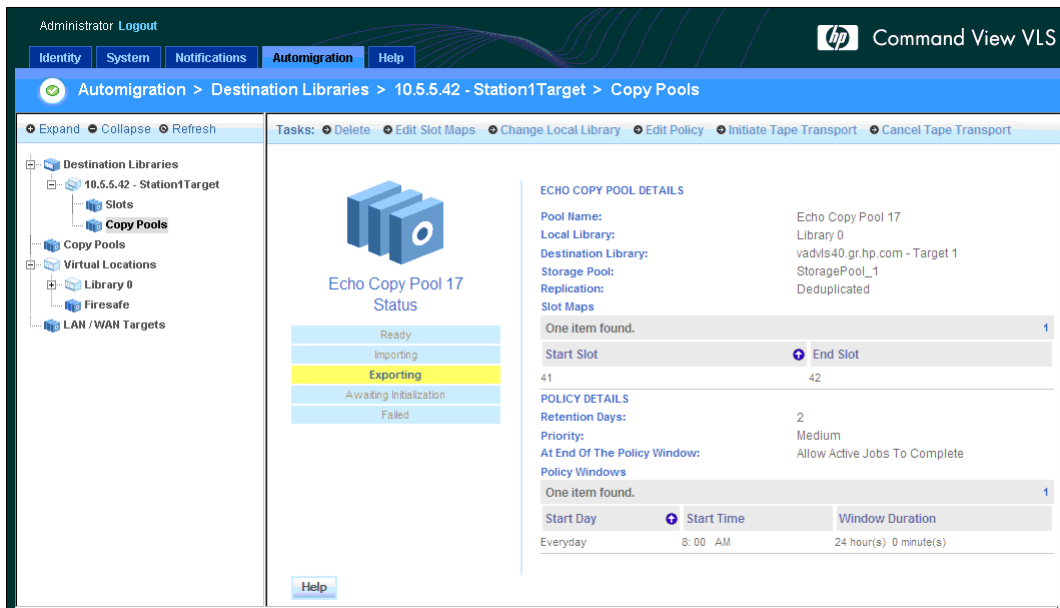
The slot range represents the slots on the replication target (not the source virtual library). Generally you would create an echo copy pool across all the slots in that replication target (because there is only ever one source device linked to a replication target). If you want to create multiple sets of replicating cartridges, for example, with different replication priorities (see [Advanced Replication Job Control](#)), you can assign part of the slot range to one echo copy pool and another part to another echo copy pool, etc.

The option to “Initialize via Tape Transport” is selected when creating the echo copy pool and determines whether tape initialization will be used or not. If not selected, initialization will be performed over the LAN/WAN.

The Deduplicated vs. Non-Deduplicated option is used when you may have a VLS shared between two different backup applications of which one of them is not currently supported by Accelerated deduplication. In this case, you must create a separate non-deduplicated echo copy pool for the virtual cartridges that cannot be deduplicated and thus force non-deduplicated replication for these cartridges (otherwise these cartridges will not be replicated). This will improve performance because the replication system will not have to attempt to perform deduplication-enabled replication on non-deduplicated cartridges.

4. Edit the replication policy (in echo copy pool).

By default the LAN/WAN replication policy (from the source perspective) is set to replicate 24 hours per day every day of the week, and set to continue running jobs at the end of each window. If you want to restrict the replication jobs to only run during specific hours of the day on specific days, you can edit the default replication policy by changing the Start Day, Start Time, Duration; the “At End of the Policy Window” can be changed to cancel replication jobs if they have not completed within the specified replication window (they would be restarted on the next window).



5. Run the first full backup (to replicating media in the source echo copy pool).  
Only backups performed after the echo copy pool is created will be replicated. Any previous existing backups will not be replicated until their cartridges expire in the backup application and are overwritten by new backups. For any backups that you want to replicate, you must ensure your backup application is configured to use the virtual cartridge media in the source device's echo copy pool. Any backups that go to virtual cartridges on the source device that are not in an echo copy pool will not be replicated.
6. Trigger the tape initialization (if tape initialization selected, then trigger after first full backup).  
If you selected "Initialize via Tape Transport" when creating the echo copy pool, after the first backup to the source device has successfully completed you must trigger the tape initialization process. (See [Implementing the Initialization Process](#)). After tape initialization is complete, subsequent replication jobs will transfer over the LAN/WAN.
7. Subsequent backups occur.
8. Run new backups and replicate (over the LAN/WAN to the target device).  
Replication takes place when the source start and duration and target start and duration windows are aligned and there is a subsequent backup to the source cartridges in the echo copy pool. See [Advanced Replication Job Control](#) for design consideration around having different priority jobs on the source device and having job throttling on the target device.

## Implementing the Initialization Process

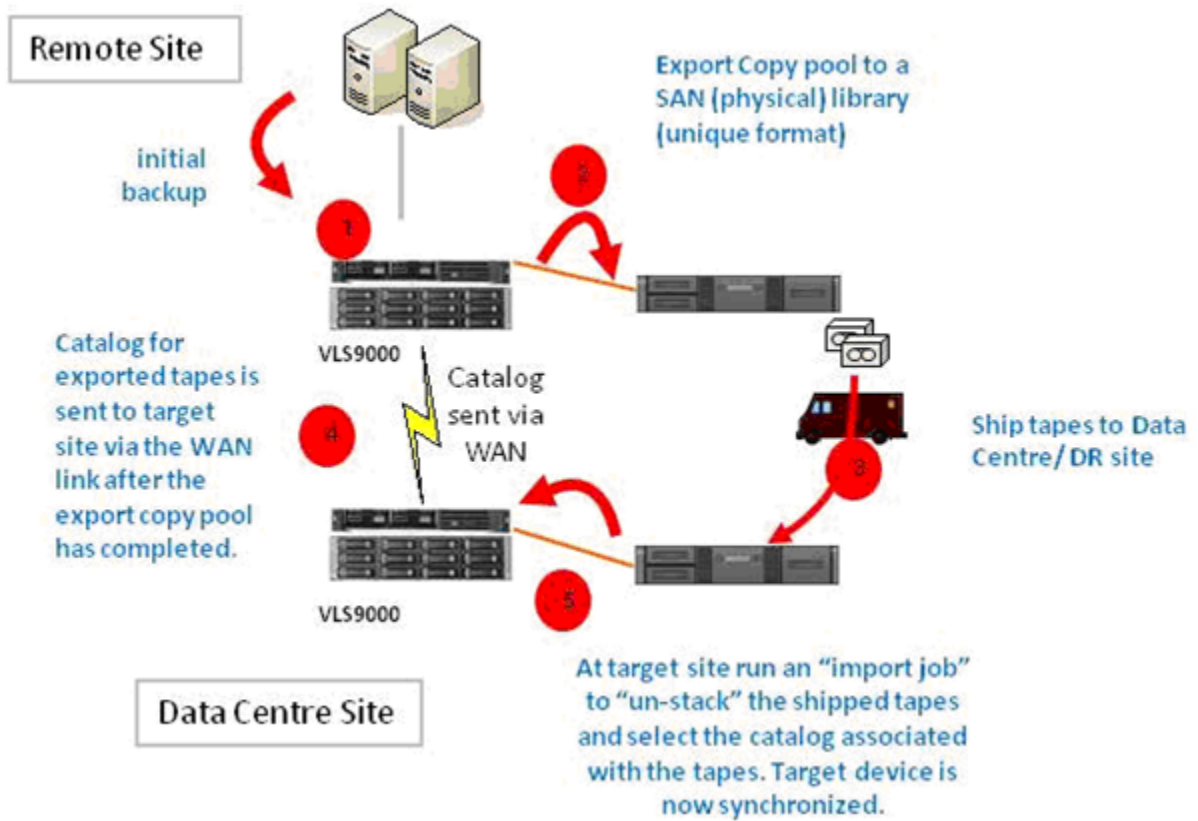
Before replication can take place, the reference data (the first full backup cycle performed with deduplication and replication enabled) needs to be sent to the target. This initial reference copy of the data is large; if the WAN/LAN replication link was sized around the daily deduplicated data to be transferred it could take a significant amount of time to transfer the reference copy (weeks or months). Once the source tape references are transferred to the target (initialized), the daily deduplication/replication can commence over the WAN/LAN link. There are several options for initialization of the reference data:

- Physical tape initialization
- WAN initialization
- Co-location initialization

You must decide whether to use the tape initialization process when initially configuring the VLS replication (when creating the echo copy pool). If tape initialization is not selected, WAN or co-location initialization will be performed over LAN/WAN.

### Physical Tape Initialization

Because in enterprise backup environments the initial backups can be quite large (many TB), in most cases it will be impractical to perform tape transfers using the low bandwidth link (because it would take weeks to replicate the first backup). Tape transfer via co-location at a site might also be impractical because of the size and space requirements of the VLS units and inter-site shipment logistics, or because both devices are always busy (e.g., using active-active or many-to-one deployments). The most popular method for initialization on VLS systems is via physical tape transfer as shown in the following diagram:

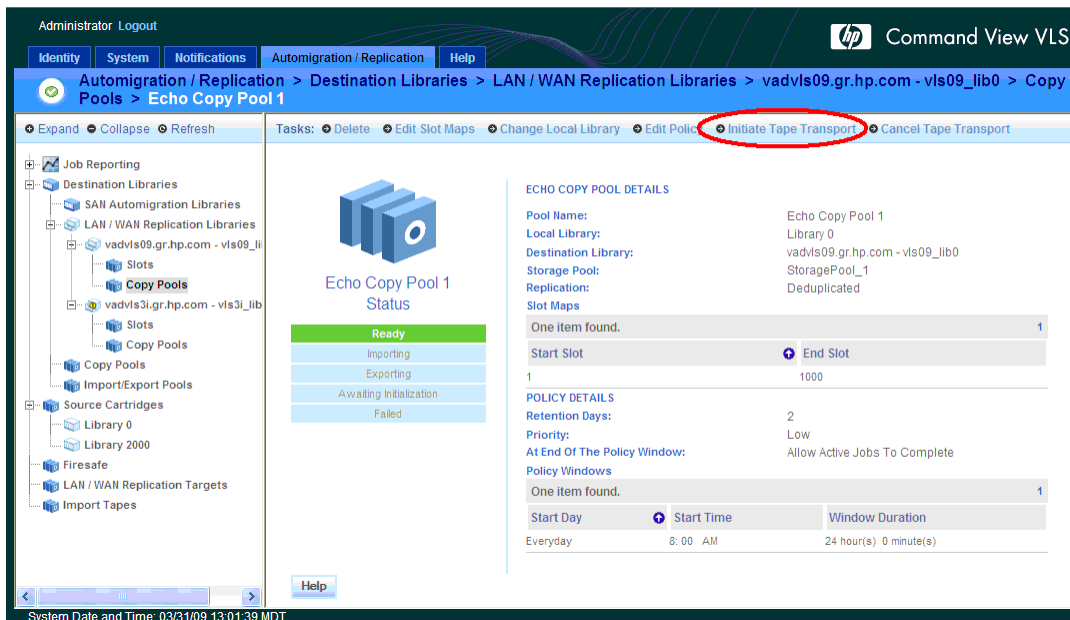


**Figure 75 VLS Initialization using Physical Tape**

This shows how using the export/import pools function within the VLS replication GUI creates a stack of tapes that can be used to populate the deduplication store at the target (destination library). To implement this way of tape transfer, you must select "Initialize via Tape Transport" when creating the echo copy pools.

You must prepare for tape initialization by plugging an HP Fibre Channel-attached tape library (MSL, EML, ESL-E) into the storage ports of the VLS device. (For a VLS9000 this would plug into the internal Fibre Channel switches inside the device, for a VLS12000 this would plug into the same Fibre Channel SAN zones containing the EVA arrays, and for the VLS6200/6600 this would require switching 1-2 of the front-end Fibre Channel ports to be back-end storage ports.) Then add this tape library as a "SAN Destination Library" in the Automigration tab.

The tape initialization process can be started when the first full backup performed after replication configuration has successfully run to the source device (backups that existed before the replication configuration are not replicated). You can then trigger the tape initialization using the "Initiate Tape Transport" option as shown in the following Command View VLS management GUI screen below. This option prompts you to select from the available "SAN Destination Library" devices to use for the tape export:



The tape export will stack the virtual cartridges to replicate (i.e., the cartridges that contain the full backup) onto the available physical tapes in the selected SAN destination library. The tape export GUI will show the export status to the tape handler to determine which physical tapes are ready for removal, whether new tapes need to be loaded, etc.

The format of these export/import tapes is unique via this process and associated with the physical tapes is a catalog which is sent over the WAN to the remote site as soon as the tapes are produced. On the subsequent import at the target site, the tape stack is read into the VLS and referenced using the catalog that was sent via the WAN.


You must plug an HP Fibre Channel-attached tape library into the storage ports of the target device and add it as a "SAN Destination Library." Then you can run the "Import Tapes" task in the GUI to perform the tape import on the target. This will prompt to select the SAN destination library to use for the import and then instruct which tapes to load into this library, show the import status to the tape handler to determine which physical tapes are ready for removal, whether new tapes need to be loaded, etc.

Once tape initialization is complete (all tapes successfully imported into the target device), subsequent replication jobs will be transferred over the LAN/WAN to the target device. In the case where you may need to override tape initialization (for example, a tape could not be imported because it was lost) you can force the switch to LAN/WAN replication using the "Cancel Tape Transport" task.

## WAN Initialization

In certain circumstances where the volume of data required to be transferred for the initialization process is small and the link speed allows it to happen in a reasonable time, the WAN link itself can be used to initialize the target device. See the following diagram. To implement this way of tape transfer you should leave "Initialize via Tape Transport" unselected when creating the echo copy pools.

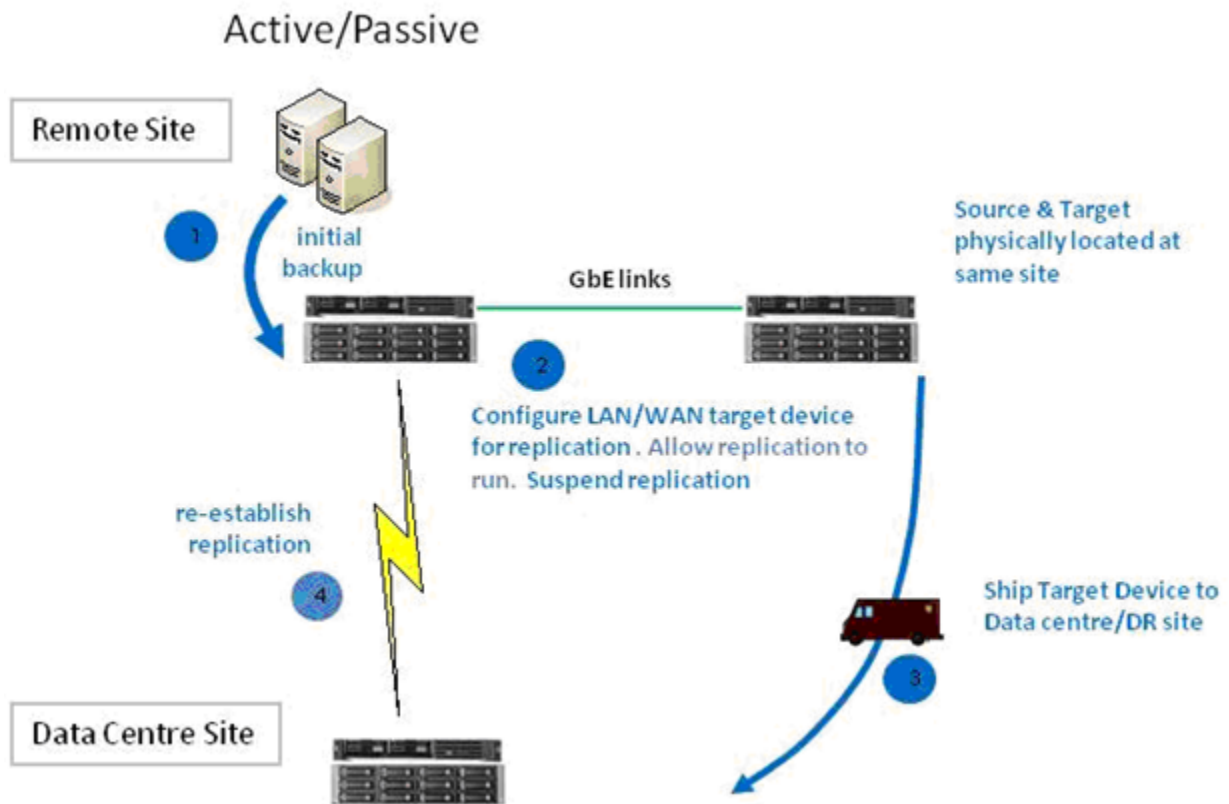
## Active/Passive



**Figure 76 VLS Initialization using the WAN**

### Co-location Initialization

In circumstances where data centers are within a small distance of each other it may be practical to co-locate the two VLS units at the same location, directly connect the GbE links from the VLS nodes together by plugging both source and target device external LAN ports into the same external LAN switch. Then backup to the source VLS and replicate to the target VLS using the high speed GbE links at the co-location. This allows high volumes of data to be initialized in a faster time. To implement this, do *not* select "Initialize via Tape Transport" when creating the echo copy pools and configuring replication.



**Figure 77 VLS Initialization by Co-location**

Co-location is practical if replication is being installed along with VLS technology from day 1. If replication is being added to VLS already installed at different sites, WAN transfer is best for small quantities of data and tape transfer for larger quantities of data. Again, for many-to-one or active-active implementations, co-location is impractical and WAN transfer or tape transfer is best.

**NOTE:**

With co-location initialization, you ship the initialized target device to the other site and will have to change the target device TCP/IP address. In this case, you need to use the "Re-Manage LAN/WAN Replication Target" task in the source device GUI to re-link the source to the target's new address.

### Continuing Role for Physical Tape

HP has insisted all along that virtual tape libraries and physical tape are not mutually exclusive but complementary technologies. Replication makes disaster recovery easier and more automated than ever before, but disk is not infallible and some astute users will still archive data to tape at the disaster recovery site allowing three copies of the backup data. This is prevalent in industries that have compliance or severe audit requirements. Physical tape also has a role to play in replication setup (tape initialization/transfer) and in the recovery process (see [Restore Directly from the VLS Target Device](#)), so do not disregard those tape libraries; they still have a valuable role to play.

### Reporting

You can view the replication status and history from both the source and target VLS devices.



## Status of LAN/WAN Replication Connection

You can check that the replication LAN/WAN connection between the source and target devices is online by viewing “LAN/WAN Replication Libraries” in the source device GUI. If the replication connection is offline this can be because:

- The target device is “Unreachable” (network problem, target device failure, you manually set the replication target to be Offline, etc.).
- “Configuration out of Sync” (policy time window settings on the target have been modified to now conflict with the source replication policy time window).

Also, if tape initialization was selected when configuring the echo copy pool, the replication LAN/WAN connection is not enabled until tape initialization is complete (or cancelled). If you view the echo copy pool in the source device GUI the pool state shows the progress of the tape initialization:

- “Awaiting Initialization” (waiting for first full backup to be completed before the “Initiate Tape Transport” can be triggered).
- “Exporting” (copying the first full backup to tape at the source device).
- “Importing” (copying the exported data from tape into the target device).

## Source Device Replication Status

On the source device GUI you can check the replication status of each source virtual cartridge by viewing “Slots” in a LAN/WAN replication library or by viewing a library in “Source Cartridges”. This lists all the local virtual cartridges being replicated to that target device, with “Last Mirrored” showing the data of the last successful replication and the “Cartridge Status” showing any active replication states for each cartridge as shown in the screen shot below:

The screenshot shows the HP Command View VLS interface. The breadcrumb path is: Automigration / Replication > Destination Libraries > LAN / WAN Replication Libraries > vadvis09.gr.hp.com - vls09\_lib0 > Slots. The main content area is titled 'SUMMARY FOR SLOTS' and shows '1,000 items found, displaying 1 to 25 of items found.' Below this is a table with the following data:

Slot Number	Barcode	Cartridge Status	Local Library	Copy Pool	Last Mirrored
37	T09036	Mirror Broken	Library 0	Echo Copy Pool 1	No Previous Mirroring
325	T09324	Mirror Postprocessing	Library 0	Echo Copy Pool 1	03/28/2009 22:33:50
319	T09318	Mirror Scheduled	Library 0	Echo Copy Pool 1	03/27/2009 16:58:39
327	T09326	Mirror Scheduled	Library 0	Echo Copy Pool 1	03/27/2009 16:58:43
330	T09329	Mirror Scheduled	Library 0	Echo Copy Pool 1	03/27/2009 16:58:45
332	T09331	Mirror Scheduled	Library 0	Echo Copy Pool 1	03/27/2009 16:58:49
336	T09335	Mirror Scheduled	Library 0	Echo Copy Pool 1	03/27/2009 16:58:49
340	T09339	Mirror Scheduled	Library 0	Echo Copy Pool 1	03/27/2009 16:58:50
342	T09341	Mirror Scheduled	Library 0	Echo Copy Pool 1	03/27/2009 16:58:51
343	T09342	Mirror Scheduled	Library 0	Echo Copy Pool 1	03/27/2009 16:58:51
344	T09343	Mirror Scheduled	Library 0	Echo Copy Pool 1	03/27/2009 16:58:53
350	T09346	Mirror Scheduled	Library 0	Echo Copy Pool 1	03/27/2009 16:58:55

The “Cartridge Status” allows you to see which cartridges do not have any replicating pending (i.e., do not need replicating because they have already replicated or have not been modified because replication was enabled), which cartridges are on the queue waiting to replicate within the specified window, and which cartridges are actively replicating over the LAN/WAN link. The “Advanced Search” task can be used to search for a specific set of cartridges based on barcode string or cartridge state

The states for cartridges that do not have pending replication are:

- Adding Cartridge. This is when the source virtual cartridge is being initially synchronized to the target cartridges. This can take a significant time when dealing with a large number of virtual

cartridges. For example, if you create 1000 cartridges in the replication target it can take 40 minutes create the matching 1000 cartridges on the source device and complete the “Adding Cartridge” operations for all 1000 cartridges.

- Up to Date. The source virtual cartridge is fully synchronized with the matching target cartridge (i.e., the mirror is in sync).
- Waiting for Backup Data. Waiting for the backup application to update the source virtual cartridge (append or overwrite with new backup data). Note that replication is only performed on backups that occur after the replication configuration is performed, so even if there are existing backups on the source virtual cartridges when replication is configured they won’t replicate until they are updated.
- In Use. The source virtual cartridge has been loaded into a virtual tape drive (this always overrides any replication job which will be restarted once the cartridge is unloaded from the drive).

Once a source virtual cartridge has been modified with new backup data, then the cartridge is queued for replicating with the following states:

- Backup Postprocessing. The source virtual cartridge has been modified but replication is waiting for the deduplication to finish analyzing it to determine what data is unique in the new backup. (All data is unique if this is a new backup, if deduplication was disabled in the GUI for this backup, or if the backup data type is not supported by deduplication.) Virtual cartridges that have been pre-formatted by the backup application (e.g., Data Protector) but have not been used yet for backups will also appear in this state.
- Out of Sync. The source virtual cartridge has been modified (last mirror date on the source virtual cartridge is newer than its last write date) and is ready for replication (deduplication has finished analyzing it) but is currently outside of the replication window policy defined in the echo copy pool. The “Copy Now” task can be used to manually override the replication window.
- Mirror Scheduled. When the replication window starts, the modified source virtual cartridge is queued up for processing. If a source virtual cartridge is modified within a replication window then it will be added to the queue as soon as the backup postprocessing is complete (it does not have to wait until the next window starts).

The queued cartridges are then processed, running as many concurrent replication jobs as allowed based on the “max jobs” setting in the replication target on the target device (the default is to run four replication jobs concurrently). You can view the active jobs in the replication job reporting screens but the following source cartridge states also show the progress of an active job:

- Mirror Preprocessing. The active replication job is gathering together all the data needed for the replication transfer including the unique data (the deltas) on the cartridge and replication metadata.
- Mirroring. The active replication job is transferring the data (deltas plus metadata) over the LAN/WAN link to the target device.
- Mirror Postprocessing. The active replication job has successfully transferred the delta data over to the target device and is now converting it to a fully intact cartridge on the target device (to provide maximum restore performance).
- Mirror Broken. This state shows an unexpected error occurred while running the replication job. You can use “Restart Broken Mirror” to attempt to recover from the error.
- Restoring. This state is when a Restore over LAN/WAN is being performed.



#### NOTE:

If an active replication job was interrupted (e.g., the cartridge was loaded into a virtual tape drive or the link failed), this will cancel/reschedule the replication job and return it to “Mirror Scheduled” status (if within a replication window) or “Out of Sync” (if outside of a replication window) so that the replication job is run again once the interruption is resolved.

---

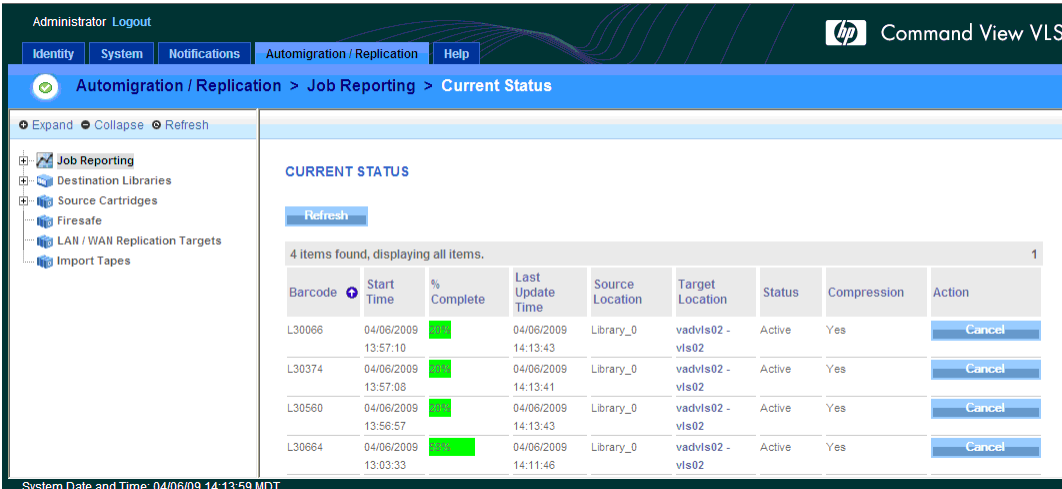
## Target Device Replication Status

On the target device GUI you can check the replication status of each target virtual cartridge by viewing “Slots” in a LAN/WAN replication targets. The “Last Successful Echo Copy” field shows the date/time of the last successful replication operation. If the cartridge has never been replicated or is in the middle of replicating the date/time is shown as “Unknown.”

This target status information is very important in the event of a disaster because the source site (including the source device) is destroyed and therefore the target device is the only available source of information on what cartridges were replicated successfully. This information is needed because when you restore from the target device you do not want to restore from cartridge in an “Unknown” replication state because it could be half way through replicating or it could be obsolete data. You can therefore use the “Last Successful Echo Copy” date/time to determine which target virtual cartridges are safe to restore from.

## Replication Job Status

On both the source and target device GUI, in the Automigration/Replication tab you can click on the “Job Reporting” branch to view the active replication jobs. This allows you to monitor the progress of the active jobs (with percentage complete indicators to monitor the actual replication transfer) and you can also manually cancel any active jobs using the “Cancel” buttons.



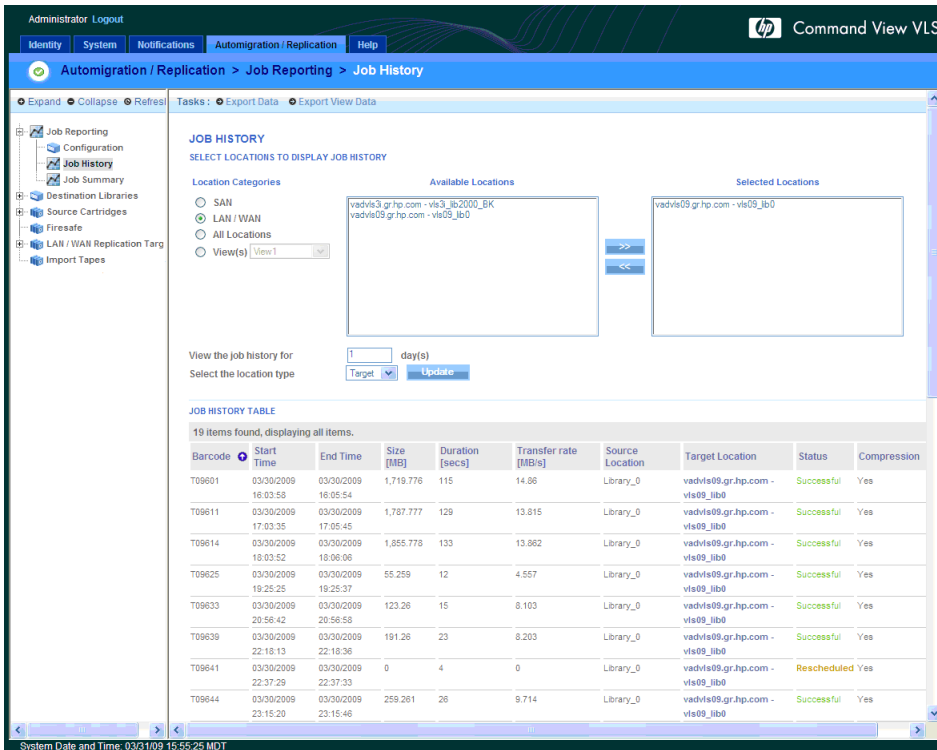
Barcode	Start Time	% Complete	Last Update Time	Source Location	Target Location	Status	Compression	Action
L30066	04/06/2009 13:57:10	100%	04/06/2009 14:13:43	Library_0	vadvis02 - vis02	Active	Yes	Cancel
L30374	04/06/2009 13:57:08	100%	04/06/2009 14:13:41	Library_0	vadvis02 - vis02	Active	Yes	Cancel
L30560	04/06/2009 13:56:57	100%	04/06/2009 14:13:43	Library_0	vadvis02 - vis02	Active	Yes	Cancel
L30664	04/06/2009 13:03:33	100%	04/06/2009 14:11:46	Library_0	vadvis02 - vis02	Active	Yes	Cancel

## Replication Job History

On both the source and target device GUI, you can use the replication “Job History” to view the active and historical replication jobs. This can be used to monitor job successes and failures and also the performance of the replication jobs. The “Start Time” and “End Time” represent the entire replication job including preprocessing and postprocessing (so you can determine whether the entire replication process is finishing on time), but the “Duration” and “Transfer Rate” represents just the delta data (plus metadata) that is actually transferred from source to target device.

Given this provides a historical view, there will generally be multiple replication jobs for each virtual cartridge. For example, each time a source virtual cartridge is appended/overwritten the cartridge will be re-replicated each time, or if a replication was interrupted (e.g., by loading the cartridge into a drive) then each time the job is restarted it will be a new job in the history. Interrupted jobs will have a status of “Rescheduled” to indicate this.

The following screen shows an example of the job history on a target device but you can also view the same information on the source device, and in the case of an active-active deployment you can view both the outgoing and incoming jobs on a device that is both a source and a target:



On both the source and target devices there is also an option to view the “Job Summary” which provides a summary of either all outgoing jobs (source) or all incoming (target) jobs. The summary includes how many replication jobs were successful, how many jobs were rescheduled (e.g., due to cartridges being appended while they were actively replicating), and how many jobs failed. For the incoming target jobs, the summary is shown for each source device. The job summary provides a quick and easy overview of the replication activity on a VLS.

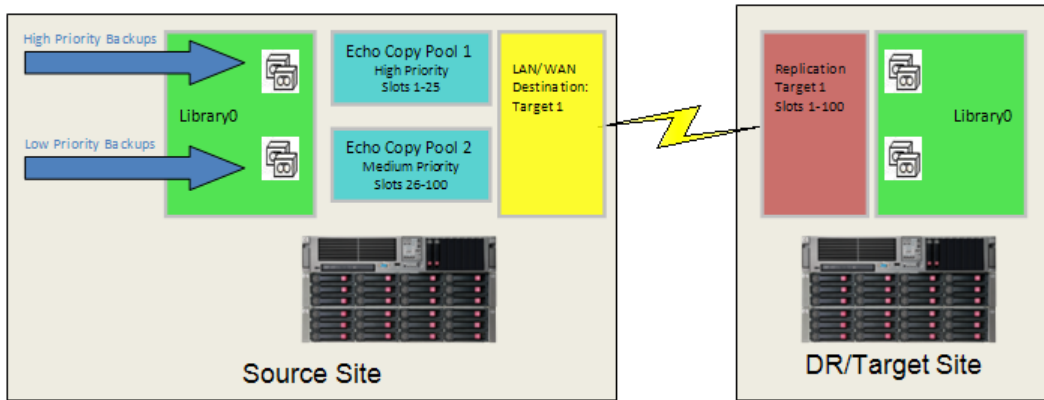
## Design Considerations

Because every enterprise backup environment is different, there is no single answer for optimizing replication. This section discusses replication design considerations to help you optimize replication in your environment.

## Advanced Replication Job Control

The replication windows in the echo copy pool policies and replication targets allow you to restrict replication traffic to specific hours of the days or days or the week, but for the more complex configurations there are additional ways to control the replication process at both the source and destination devices.

If you have different classifications of backups at the source device, you may want to create multiple echo copy pools with different policy priority settings so that high priority backups are replicated before lower priority backups. You can still keep one replication target and configure the multiple echo copy pools to have different slots ranges (and thus different sets of cartridges) within that replication target. On the source device you could either have one source virtual library mapped to the multiple echo copy pools or you could have a separate virtual library for each echo copy pool. In either case you would divide your backup jobs so that the high priority jobs went to the cartridge set or virtual library mapped to the echo copy pool with high priority set in the policies, and the lower priority jobs would go to the other cartridges or other virtual library, etc.



**Figure 78 Dividing the Backup Jobs by Priority Level**

On the target device you also have the ability to limit replication traffic beyond limiting the replication time windows. This can be done by controlling how many concurrent replication jobs are allowed per replication target, and optionally controlling the maximum Mbytes/sec throughput of a single replication job. (This is set as a global configuration value that applies across all replication targets on that device.) In this way, the amount of replication traffic can be controlled to each replication target on a target device.

Setting the maximum number of concurrent jobs per replication target needs to take into account various factors such as VLS device hardware limits and the network latency of the replication link:

- In order to get maximum replication throughput, you must scale the number of concurrent replication jobs per target depending on the number of nodes used for replication in the target device. For example, if you have a 4-node VLS configured as a single replication target, you should set the maximum number of jobs for the replication target to be 28 (based on four nodes each with a maximum of 7 concurrent replication jobs per node).
- There is a hardware limit of seven concurrent replication jobs per VLS node, so if you have multiple replication targets on a VLS (e.g., many-to-one deployment) you need to ensure that you do not exceed the hardware limit when you configure the “Maximum Simultaneous Transfers” on each replication target.
  - For example, if you have four replication targets configured on a 2-node VLS, you do not want to set each Target to allow seven concurrent jobs because this would add up to 28 concurrent jobs across all four targets but you only have 14 concurrent jobs capability in the 2-node configuration. So if all the replication windows were the same across the four source devices, you would potentially starve some of the replication targets. (If you had non-overlapping replication windows then there would not be any contention issues.)
  - Another example would be an active-active configuration where you had a 1-node VLS device, because then you have to share the hardware limit of seven replication jobs across both the outgoing and incoming replication jobs. In this case you could leave the default setting of four concurrent jobs in the target thus leaving at least three jobs for the outgoing replication traffic.
- The network latency can affect how many concurrent replication jobs are needed in each replication target. The VLS replication incorporates latency-acceleration technology, but with longer network latencies this still requires some level of concurrent replication jobs to saturate the replication link as shown in [Table 26](#).

**Table 26 Network Latency**

Job Concurrency Required to Saturate Link	Network Latency	Aggregate Mbytes/sec
2	0ms	80 MB/sec

Job Concurrency Required to Saturate Link	Network Latency	Aggregate Mbytes/sec
4	50ms	80 MB/sec
4	100ms	80 MB/sec
7	200ms	80 MB/sec
7	500ms	52 MB/sec

When you set the global per job Mbytes/sec limit on the target device, you can calculate the Mbytes/sec limit on each replication target by multiplying the per job limit by the maximum number of concurrent jobs set for that replication target. This is only the Mbytes/sec limit but it will only apply if the replication link bandwidth is higher than this limit.

## Link Sizing

This section focuses on understanding how to size the replication link bandwidth required.

The size of the replication link will depend on your usage model for replication.



### NOTE:

For the volumes of data to replicate with HP VLS systems, HP recommends a dedicated replication link wherever possible.

- Backup and replication can take place in a fixed window (up to 12 hours) preferably when the WAN link is not heavily used by applications (that is outside business hours). The replication process has almost the full bandwidth of the WAN link available to it. Within the HP VLS System configuration for replication you can set a “replication start time and window” during which time replication is allowed to operate.
- The business runs 24x7 and there are no fixed windows in which to do replication. In this example the replication can be “limited” and the HP VLS devices can be configured to only use a percentage of the available bandwidth so as not to adversely affect application performance if applications are sharing the same WAN link.
- The most likely scenario for enterprise data centers using HP VLS replication is that it will be enabled to run 24 hours a day on dedicated links used only for replication. HP VLS has the ability to enable windows for replication to further customize the implementation to best suit your needs.
- With HP Accelerated deduplication there is a time delay after the backup completes for the deduplication to run and the replication instructions to be generated before new data can replicate. This means there is probably a stronger need for a dedicated link and the full utilization of that link in order to get the backup deduplication and replication completed within an acceptable service level agreement.

HP also provides a Backup Sizing Tool to help you size any backup solution including deduplication and replication. See <http://www.hp.com/go/storageworks/sizer>. Wherever possible, do a proof of concept before implementing into production.

## Basic Calculations

The time required for replication is derived from the equation:

$$\frac{(\text{New replication data} + (\text{Control data} - \text{approx } 1\% \text{ full backup size})) \times \text{comp ratio}}{\text{Rated WAN link speed} \times \text{Efficiency} \times \text{Throttling Factor}}$$

Rated WAN link speed x Efficiency x Throttling Factor

 **NOTE:**


You should always factor in additional bandwidth to the calculated link size to take possible link downtime or replication initialization delays into account. If the replication falls behind for any reason (e.g., the link is down for 24 hours), you need more than the normal daily bandwidth to allow replication to “catch up.” Also, when new server clients are added to the backup system, their first backup goes over the link wholesale thus causing a temporary increase in daily bandwidth. HP recommends sizing the link with an additional 25% extra bandwidth to handle these cases.

The calculations below are based on several assumptions and represent worst case scenarios.

- A pessimistic compression ratio of 1.5:1.
- Links are only 66% efficient (e.g., 1 Gb/sec link can typically deliver 80 MB/sec).
- Because the VLS is a multi-node architecture with significant performance for backup, deduplication, and replication, the VLS will never be the bottleneck with respect to getting data on the wire for replication.
- Target usage model is backup and replication complete within 24 hours.
- Allow a four hour delay after backup completes for the first set of replication instructions to be produced.
- There is a “typical” split of 75% database, 25% files system data. File system incrementals are used which will result in larger amounts of data going across the replication links. Database fulls are used.

Each example illustrates a unique point:

1. Sizing Scenario: 20 TB of database and file-system replication.



**Figure 79 VLS Sizing Example 1**

- Database Full Backup size 15TB – 1% change rate
- File System Full backup Size 5TB – Typical incremental size of 10%
- Replication Window 24 Hours

- Required Link speed 40% of existing 200Mbit/sec link

This example shows a 15 TB full database backup and 10% incremental backup of file system data to replicate daily between sites. There is already a 200Mbit/sec link between sites. The actual amount of data to replicate is 164 GB (representing 15 TB full database backup and 10% incremental of 5 TB File system backup). Using bandwidth throttling at a rate of 40% of the 200Mbit/sec link, you can easily accomplish the backup and replication within 24 hours without impacting application performance using the same WAN link.

2. Sizing Scenario: 66 TB of database and file-system replication.

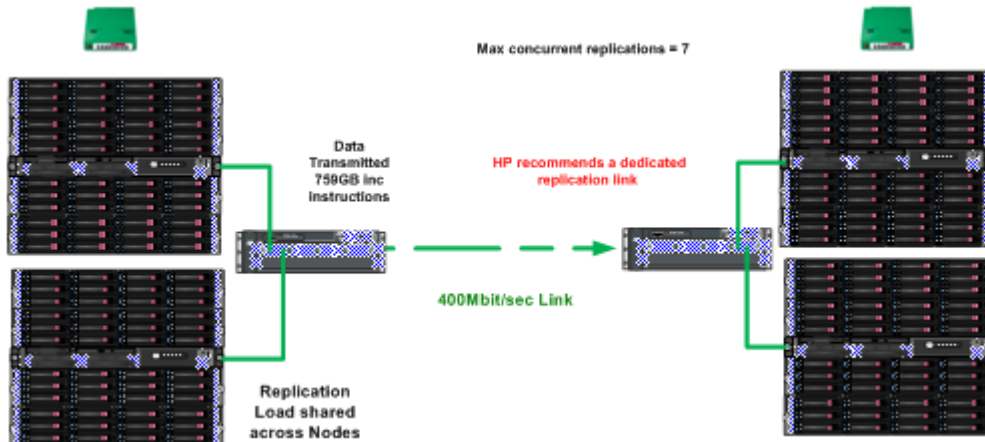


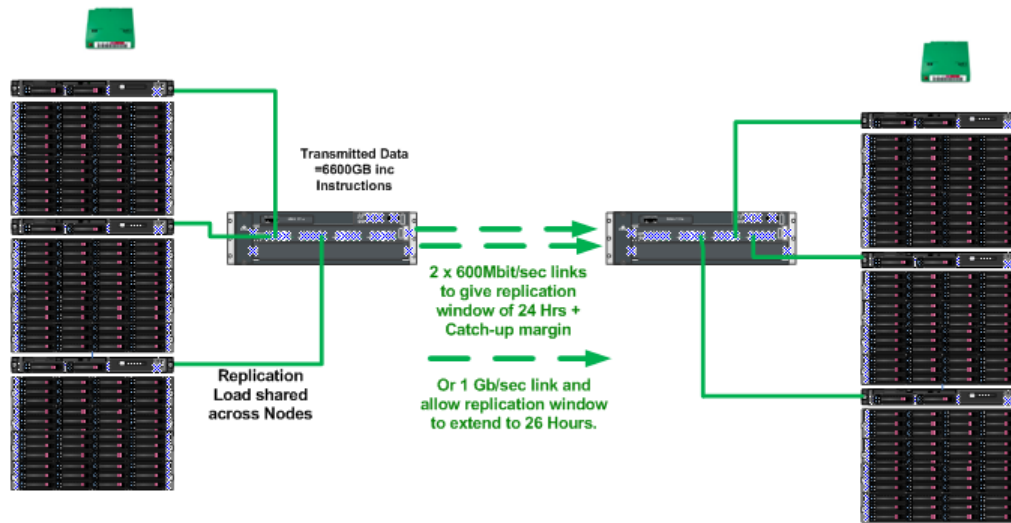
Figure 80 VLS Sizing Example 2

- Database Full Backup size 50TB – 1% change rate
- File System Full backup Size 16.6 TB – This example replicates one of the incremental backups, size of 10%
- Replication Window 24 hours
- Required Link - 400Mbit/sec link made available
- Typical replication time 12.8 hours

Now consider a larger backup requirement with a full database backup of 50 TB and an incremental of 10% of 16.6 TB of file system data. Notice how as the amount of data backed up scales the number of nodes scales, and as the number of nodes scales more processing power is made available to both backup, deduplication, and replication. Here in this 2-node system both nodes work to ensure the replication link is saturated at all times. In this example the Sizer calculations show around 659 GB of data needs to replicate. There is a dedicated 400 Mbit/sec link between sites. The calculations show that this replication can complete within 12.8 hours (even allowing four hours for the first set of deduplication replication instructions to be generated, so it is easy to achieve the requirement of backup and replication within 24 hours).

3. Sizing Scenario: 266 TB of database and file-system replication.





**Figure 81 VLS Sizing Example 3**

- Database Full Backup size 200TB – 1% change rate
- File System Full backup Size 66.6 TB – This example replicates one of the incremental backups, size of 10%
- Initial Replication Window 24 hours
- Required Link 1100bit/sec + some margin for catch-up
- Consider 1 Gb/sec link and relax replication window to 26 hours

The final example shows the trade-off and the sheer capabilities of this technology. This simulates a 266 TB environment being replicated (200 TB database, 66.6 TB file system) with an initially stated 24 hour replication window (allowing four hours until the first replication starts). When you run the Sizer calculations in order to replicate this volume of data within 24 hours, it really needs a 1200 Mbits/sec link (or 2 x 600 Mbits because few telcos deliver greater than a 1Gb/sec link). Again note how, because of the volume of data, the number of nodes has increased and the replication ability has also scaled appropriately. You now have a choice to either invest in the 2 x 600 Mbit/sec links or compromise on the time to data off-site. If you are willing to accept data replicated within 26.2 hours, you can save considerable cost in requiring only a 1 Gb/sec link.

This example replicated the equivalent of 266 TB in just over 26 hours, a task that without deduplication technology would have taken 205 hours on DWDM 4Gb Fibre Channel at significantly higher cost. While some may say 1 GB/sec is not low bandwidth, everything is relative to the task at hand.

**Table 27 Summary of the Calculations Involved in the Sizing Examples**

Data types (3 DB:1 FS) & backup type	Full size in GB	Replicated data (compressed) in GB	Time available to replication/time to replicate	Link speed available or required in Mbit/sec	Combined link speed
Database (Full) 1% change	15000	198	24 hours	83 at 40% utilization	20 Mbits/sec at 40% utilization
File data (Full)	5000	297		125 at 40% utilization	

Data types (3 DB:1 FS) & backup type	Full size in GB	Replicated data (compressed) in GB	Time available to replication/time to replicate	Link speed available or required in Mbit/sec	Combined link speed
Database (Full) 1% change	50000	660	9.5 hours	Assume fixed dedicated link speed of 400 Mbit/sec	Assume fixed dedicated link speed of 400 Mbit/sec
File Data (incremental 10%)*	16666	99	3.3 hours		
Database (Full) 1% change	200000	2640	24 hours	444 at 100%	At greater than 1 Gb/sec+ at high volumes add extra link bandwidth for catch-up
File Data (incremental 10%)*	66666	3960		666 at 100%	

\* Fulls on the weekend.

4. Sizing Scenario: Many-to-one (4:1) replication.

4 to 1 VLS Replication Bandwidth Sizing Example

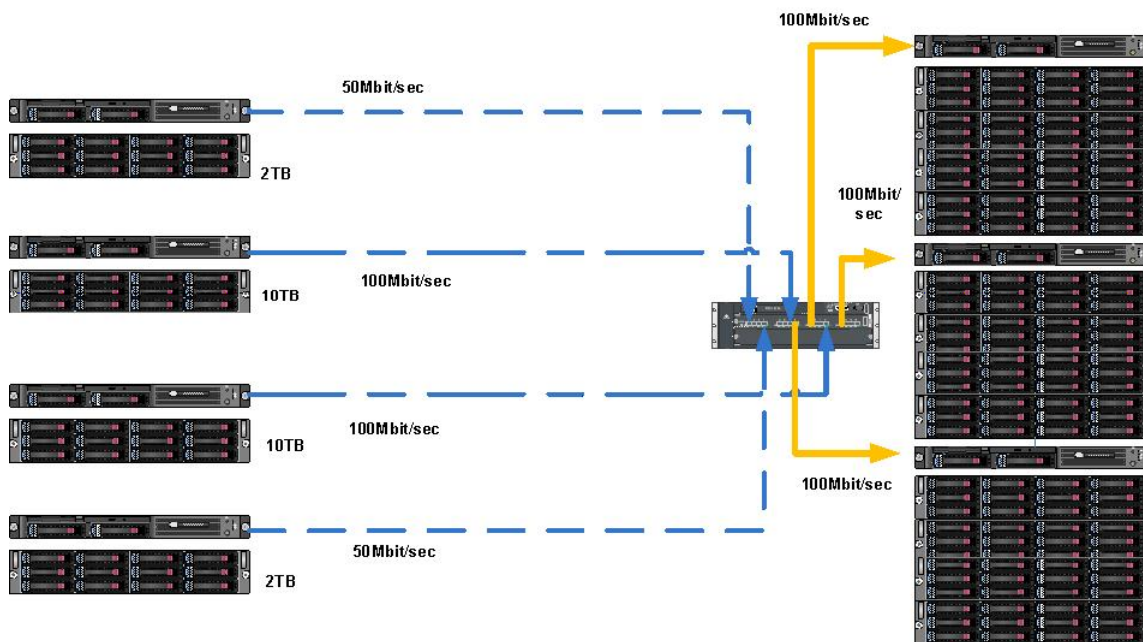


Figure 82 VLS Sizing Example 4

Here are four remote sites replicating 2 TB and 10 TB backups daily into a central VLS9000. The best approach in a many-to-one scenario is:

- Size each of the individual links separately.
- Ensure the links at the target at least match the total remote size replication capacity (in this example 300 Mbits/sec).
- If possible, allow 25% extra capacity for “catch up” in the event of a backlog or paused replications should a link go down.

- Each remote site could have different start time and replication window to ensure target device is not overloaded and best use is made of inter-site link speeds.

## Device Sizing

When sizing a replication target device, the replicated virtual cartridges on the target device take up the same amount of disk space as the original cartridges on the source device, and they also require the same amount of working disk space for deduplication. In other words, if you replicate all of the cartridges in the source device then you need the target device to have the same capacity as the source device. In a many-to-one deployment, add up all the source disk capacity and the total capacity is what you need on the target device.

However, the target device is not just adding capacity; it also need to add sufficient extra nodes to handle the replication target disk I/O. With a 24-hour replication window, you need as much node performance on the target device as on the source devices. For example:

- With an active-passive deployment, if you have a 4-node/4-array VLS9000 as the source device and you are replicating all of its cartridges, you need a 4-node/4-array target device.
- With a many-to-one deployment, if you have four VLS6600/24-disk source devices (which are effectively approximately half a node of performance and half an array of capacity each), you need a 2-node/2-array VLS9000 as the target device.
- If you have an active-active deployment, you must combine the required source performance/capacity and the replication target performance/capacity to give you the overall device. If in each site you have a 2-node/2-array requirement for the source piece (i.e., backups), you need 2-node/2-array for the target piece. Therefore, you would have a 4-node/4-array device in each site.

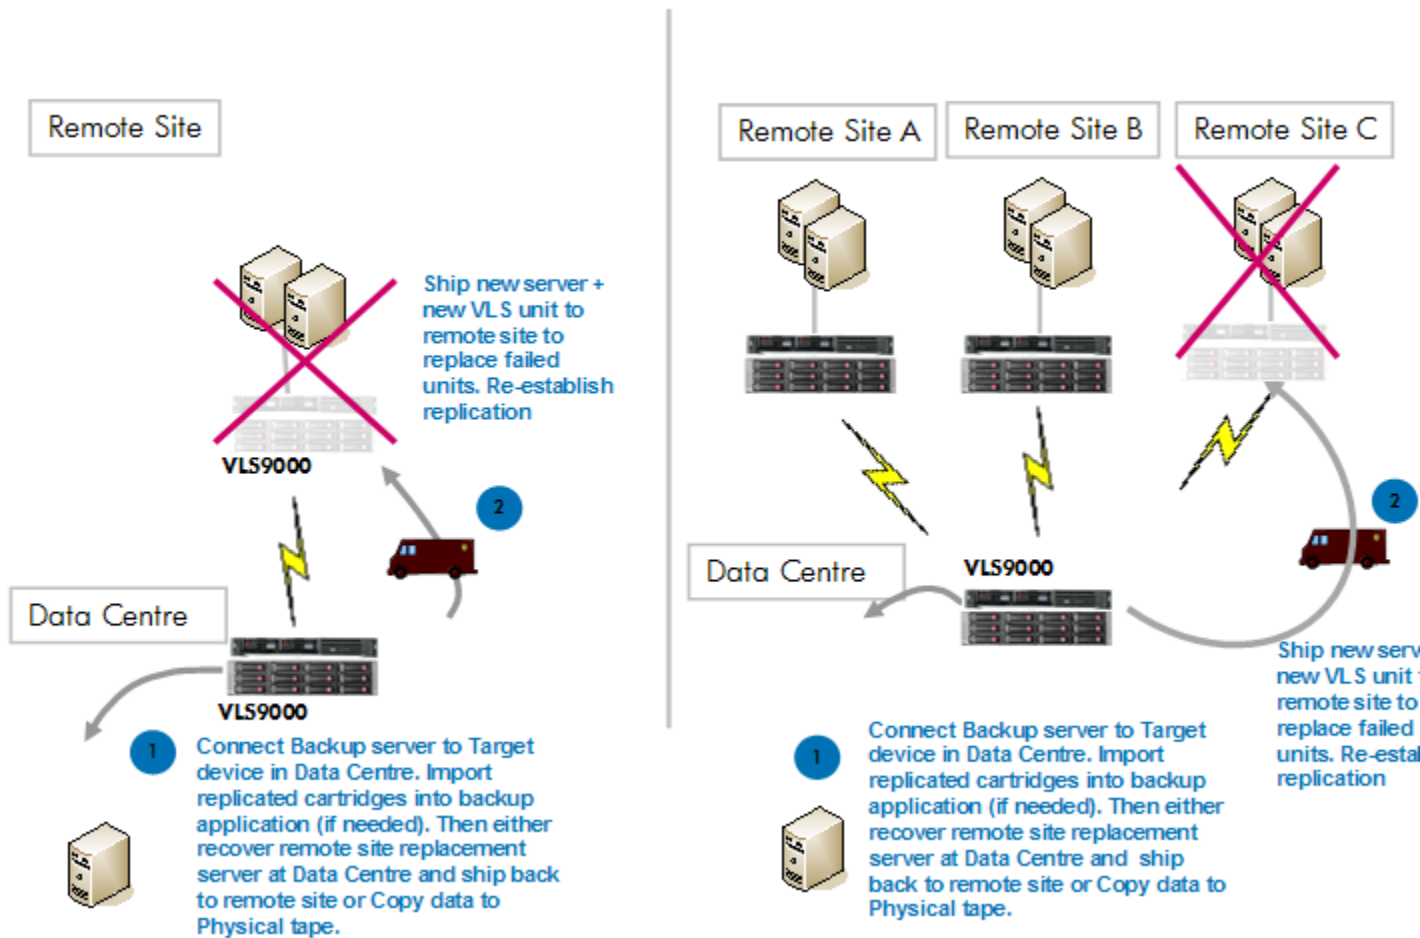
## VLS Replication Data Recovery Options

Replication on VLS enables easier recovery of data in the event of a major site disaster. The data is not instantly available but has to be recovered through a standard restore process using a backup application. These scenarios examine what happens with a total site disaster including losing the server and the VLS appliance that was protecting it.

In the event of a total disaster at a site where the VLS and the servers it is protecting are destroyed or damaged, the replicated data at the disaster recovery (target) site can be accessed by means of a backup server on the remote site and transferred either to an application or onto physical tape for distribution. Alternatively, when the damaged site is repaired, the most recent backups on the virtual tape library at the disaster recovery site can be restored over the LAN/WAN to the virtual tape library at the previously damaged site.

### Restore Directly from the VLS Target Device

The first method of recovery uses the fact that there are probably more trained IT personnel at the data center and that spare servers, etc., are stored there. With VLS the virtual library containing the replication target (on the target device in the disaster recovery site) is always available through the front-end Fibre Channel ports and thus can be presented as a library to a backup application in the disaster recovery site (note the backup application can only be presented to the target device, it cannot also be presented to the source device as well).

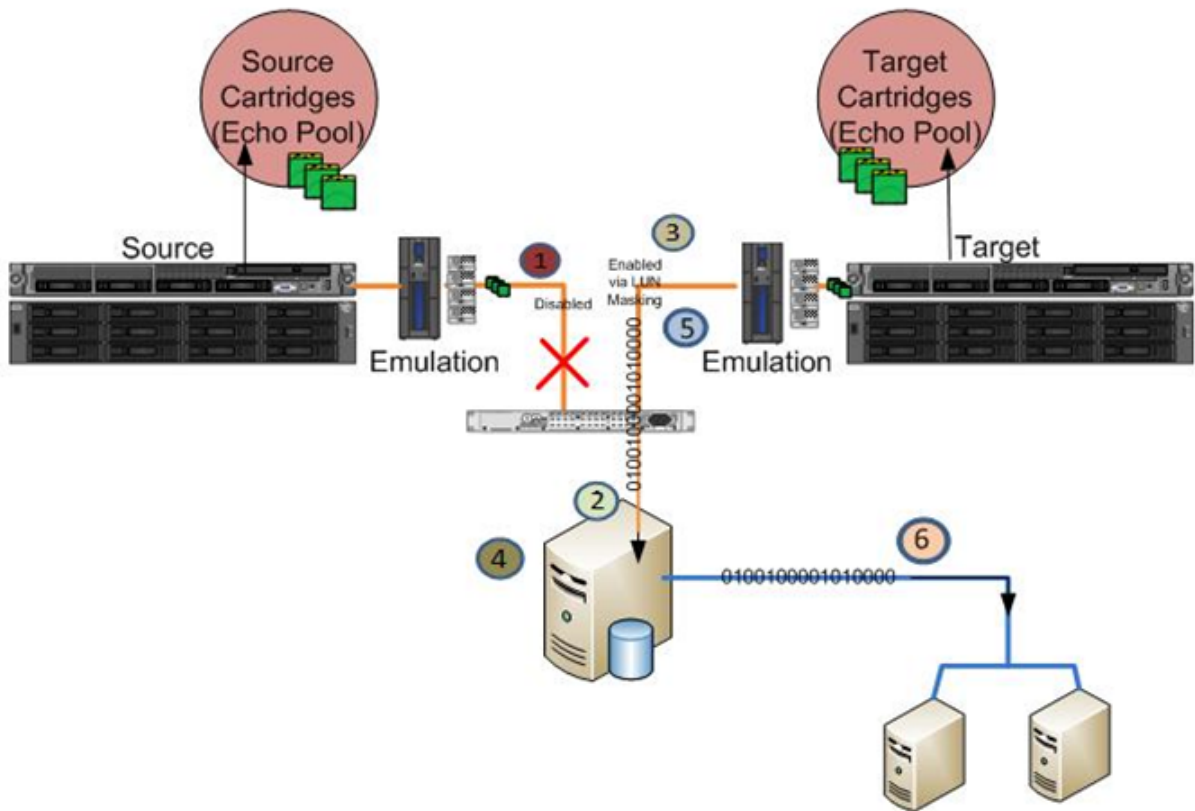


**Figure 83 VLS Recovery Options for a Data Center Rebuild**

This target virtual library can be used to restore directly onto new servers that can then be shipped to the replacement site, or to copy the data to physical tape which can be shipped to the replacement site and used to rebuild the new servers there. To prepare for the restore process there are different scenarios depending on which backup application is used for the restore:

- If you use a different backup application (not the instance that did the original backups on the source site), you must recreate the media database information from the replicated cartridges.
  - If this is a blank backup application just used for recovery, you can restore the replicated backup of the media database and then restore from the replicated cartridges that are fully synchronized. (The most recent backups in the media database may not have been fully replicated and thus may not be usable for restore.)
  - If this is a live backup application (for example, it is performing its own local backups in the target site), you can import the replicated cartridges into the existing backup application database. You can do this manually at the time of the disaster recovery, or you can use an automated import script (see [Creating Archive Tapes from the Target](#)) to keep the media database roughly in sync with the replication cartridges (less than 24 hours out of date in the event of a disaster recovery).
- If you use the same backup application (the instance that did the original backups on the source site), you need to switch the backup application to “see” the target device instead of the source device. When you initially configuring this type of solution, you would have needed to ensure that the target device was *not* seen by the backup application. (You cannot have the backup application see two copies of the same cartridge at the same time; see [Backup Application Interaction with](#)

Replication). One option would be to use the LUN Mapping feature in the VLS to hide the target virtual library from the backup application. But in the event of a disaster recovery, you can perform the sequence shown in Figure 84. (The most recent backups in the media database may not have been fully replicated and thus may not be usable for restore.)



**Figure 84 Data Recovery from the Target VLS**

1. The source VLS becomes inaccessible to the remote VLS and to the backup host.
2. The source tape library and drive are deleted from the backup host (cartridge media pools remain).
3. The target VLS is enabled using the VLS LUN mapping feature in the GUI.
4. The backup application server configures the target library and tape drives and performs a barcode scan.
5. The backup application recognizes data cartridges within the new library.
6. Data is restored using target cartridges (unknown to the backup application) where the most recent backup is forward referenced.

## Restore the VLS over the LAN/WAN

Another method of recovery is to restore the backups over the WAN/LAN from the target device back to a rebuilt source device. The restore over WAN/LAN can be initiated from the source device in the "Restore Media" task in LAN/WAN destination library slots, where you can then select which target cartridges you want to restore back to their matching cartridges on the source device.

The LAN/WAN restore is a whole cartridge non-deduplicated copy so the entire contents of the cartridge are read back over the LAN/WAN, and the resulting restored cartridges on the source device will never deduplicate against any other cartridges on the device (until they are recycled and then overwritten with new backups). Given this, you would be limited by the bandwidth of the replication link in terms of how much data can be restored, and so generally you would restore the

cartridges containing the last full backup and any incremental backups after this point. These can then be used by the backup application to import/restore from the source device's virtual library to replacement servers.

## Rebuilding Source Device and Re-establishing Replication

If the entire source VLS device was destroyed (i.e., you had a site disaster) and was replaced with a brand new blank device, then you must recreate the original virtual library configuration and replication configuration. If you have a previously saved configuration file (created using Save Configuration in the VLS GUI) you can perform a "Restore Configuration" in the VLS GUI (loading in a previous saved configuration). Alternatively you can manually recreate the original virtual library configuration and replication configuration.

There is currently no method of restoring the entire deduplicated set of cartridges from the target VLS to the source VLS. Once your source VLS is rebuilt and reconfigured and you have restored your servers, you must start backups again to the source device and then repeat the initialization process after the first full backup is finished.

---

### NOTE:

If you need to re-initialize the replication system using tape initialization, then you must create your echo copy pools with the "Initialize via Tape Transport" option enabled and then repeat the tape initialization process. This means if you used "Restore Configuration" to recreate your replication configuration you must delete your Echo Copy pools and recreate them.

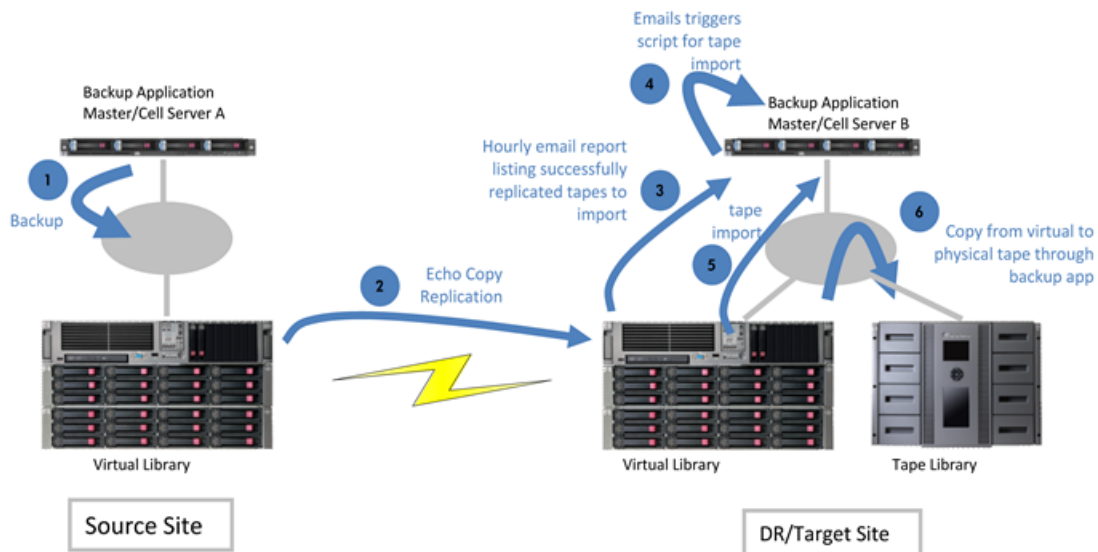
---

## Creating Archive Tapes from the Target

Even with replication removing the need for physical tape off-site, there are still many users who wish to use physical tape for archive or test recoveries, etc. An archive tape will have the following differences from the original backup:

- A different retention time from the original backup (for example, you may keep backups and off-site replicas for three months but keep tape archive for several years).
- Different cartridge contents because it will be different size (optimum virtual cartridge size for deduplication is 50-100 GB but physical tapes such as LTO-4 are 800 GB).
- A different barcode and be tracked separately in the backup application.

The above characteristics mean that the backup application must control the creation of the archive tape and must use an object-copy based system (cannot use a whole-tape copy system). This is easy if the physical library is on the source site because you can use the normal backup application object-copy mechanisms to simply copy from their source virtual library to their physical library (via the backup application media servers). However, if the physical library is at the target site and you want to perform the tape archive from the replicated virtual cartridges in the target device, this requires additional steps. This is because the backup application at the target device that sees the replicated virtual cartridges cannot be the same backup application that performed the backups at the source site (see [Backup Application Interaction with Replication](#)). The target site backup application does not have any information about the replicated cartridges in its media database because it did not backup to those tapes.



**Figure 85 Creating Archive Tapes from the VLS**

In order to have the target backup application copy from the replicated cartridges to a physical tape library it needs to be “taught” what is on the replicated cartridges. This can be done automatically as shown in the diagram above, by using a new “ISV Import” email report on the VLS that sends an hourly email report listing which virtual cartridges have been successfully replicated in the last hour, and this cartridge list can then be fed into a script that automatically triggers tape import jobs in the backup application (which read the new cartridge data and import this content into the media database). The target backup application can then restore from the imported virtual cartridges or copy them to physical tape, etc.

### ISV Import Email Format

The following shows the format of the ISV Import email from the VLS, which lists the virtual cartridges that were successfully replicated into replication targets since the last report. The attachment format is:

ISV~<virtual cartridge barcode>~<virtual library serial number>~<virtual library slot number>

Email body example:

**From:** administrator@devicename.test.com  
**Subject:** Replication Data for ISV Import

**Status** : Successful  
**Report Name** : Replication Data for ISV Import  
**Report Generation Time** : Wed Mar 18 16:00:00 MDT 2009  
**Report Frequency** : Hourly

Email attachment example:

ISV~G231006~01YFOfAM00~44  
 ISV~G231011~01YFOfAM00~55  
 ISV~G231022~01YFOfAM07~112  
 ISV~G231025~01YFOfAM04~12  
 ISV~G231030~01YFOfAM00~45  
 ISV~G231036~01YFOfAM04~19

## Email Processing Example Script

The following section details an example script to process the “ISV Import” email report from the VLS and convert it into a list of cartridges to be imported. The following steps show how this might be done in the target site (this example uses a Linux server client):

- You will need an SMTP gateway server in the network (needed for the VLS SMTP emails).
- Allocate a client server to use for the import script (could be a separate server or the actual master/cell server itself).
  - Ensure the selected client is running an email messaging system (e.g., sendmail).
  - Create a user account on the client.
  - Ensure the selected client has CLI access to the backup application.
- In the user account, create a filter that sequentially processes the incoming emails, and for the ISV import emails it should send the message body (containing the list of replicated cartridges) to the script that actually performs the tape import. One way of doing this is to create a `.forward` that sends your incoming emails to procmail. For example: `| /usr/bin/procmail` and then add the following lines to the `.procmailrc` that copy an ISV import email (containing the list of tapes to import) to another script that runs the ISV-specific import commands:  

```
:0 bc: # send VLS DP import files for processing* ^Subject:. *ISV| /home/importuser1/isv_import/dpimport.sh
```

  
See [Data Protector Import Example Script](#) and [NetBackup Import Example Script](#) for sample import scripts.
- On the target VLS device GUI, setup the SMTP Gateway address and then configure the “ISV Import” email report using the selected email account.

## VLS Non-deduplicated Replication

If you have *not* enabled Accelerated deduplication on your VLS device, non-deduplicated replication can still provide a more reliable method of sending mission-critical data off-site compared to off-site using physical tape (because there is no need to load/unload physical media and ship it to the off-site location).


VLS-to-VLS non-deduplicated replication options:

- Using SAN-based Automigration  
Automigration also supports using a VLS as a destination library, as a device-to-device replication scheme over an extended SAN. To use a VLS as a destination library, designate one VLS as the source and a second VLS as the destination. The destination VLS is presented over an extended SAN so that it is visible to the source VLS. With Command View VLS, establish an automigration policy and echo copy pools in the same way you set up automigration with a physical destination library. See [Echo Copy Pools](#) and [Automigration Policy](#). The source VLS performs normal backups during your regular backup window. During the automigration window you defined, source virtual cartridges automatically migrate to matching virtual cartridges on the destination VLS.  
The VLS with the source library must support automigration; however, it is not necessary for the destination VLS to do so. If you have an earlier generation VLS system and you purchase a VLS6200, VLS6600, or VLS12000 Gateway, you can use the newer models to handle automigration as the source device. Earlier generation VLS devices can be used as destinations.



 **NOTE:**

You cannot use device-to-device replication with the VLS9000 because its back-end storage ports are wired into its private SAN and cannot be connected to a public SAN.



**Figure 86 Site-to-site Replication**

- Using TCP/IP-based non-deduplicated replication  
You can configure an echo copy pool to a LAN/WAN destination library (which connects to a replication target on the target device). This provides an echo copy replication link over TCP/IP between the source and destination devices.
  - You must be running VLS firmware version 2.3 on both the source and target VLS devices.
  - This option only supports one replication link (node0 external LAN connection).
  - The process to set up the non-deduplicated replication is very similar to the deduplication-enabled replication (see [Replication Setup](#)) except that there is no need for initialization step because the data is transferred in its entirety every time.

## Replicating a Subset of Virtual Cartridges - a Use Case

The following is just one possible use case for non-deduplicated replication. In a typical backup environment, companies rarely have network bandwidth surpassing what meets their needs. For this reason, the ability to replicate just a subset of business critical data, thus conserving bandwidth, is advantageous. It is fully supported to have a mixture of standard and automigration tapes in the same virtual library at the same time.

Replicating a subset of backup data to a remote VLS also provides customers with the flexibility and simplicity of not having to create a second library target in the backup application just for replication. Customers can configure one library on the source VLS that performs local backups, and that also replicates a subset of data they designate.

To replicate a subset of backup jobs:

1. On VLS1 (local VLS) configure one library with more slots than what you need for regular backups (to be used for replicating virtual cartridges).
2. On VLS2 (remote VLS) configure a small replication library.
3. Configure the SAN zoning so that the VLS2 replication library is presented over the extended SAN into the VLS1 back-end array zones. That is, the back-end fibre channel ports of VLS1 can see the replication library on VLS2.
4. On VLS1, “manage” or add the VLS2 replication library as the destination library. See [Automigration Setup](#).

5. Create an echo copy pool for the VLS2 replication library, with the library on VLS1 as the source.

Automigration will automatically create the replicating virtual cartridges on VLS1 in the spare slots configured in step 1. During the automigration window you created, backups already copied to replicating virtual cartridges on VLS1 will automatically be copied over to VLS2. If you back up to a non-automigration tape in VLS1, then nothing else happens. To the backup application, there is just one target library, though it contains both replicating and non-replicating tapes. It is easy to put the replicating tapes into a different media pool in the backup application, and simply configure the backup jobs to use either the replication media pool, or the non-replication media pool. In either case, all the backup jobs go to the one, shared destination library.

---

 **NOTE:**

For cross-replication you can create the reverse configuration. For example, create another large library on VLS2 for its backups, and create a small replication library on VLS1 that is visible to the back-end ports of VLS2.

---

## Detailed Backup Application Guidelines for VLS

To achieve full performance benefits from the VLS, you may need to modify your enterprise backup application configuration (for example, to take advantage of the fact that you can create many more virtual tape drives than a physical tape library holds). If you have enabled deduplication, you may need additional changes in your enterprise backup application to take full advantage of this capacity optimization technology. Follow these guidelines *before* you begin deduplicating your data or some data may not deduplicate efficiently.

### HP Data Protector

This section includes both general and deduplication guidelines as well as other useful information for HP Data Protector.

#### Data Protector General Guidelines

The following Data Protector guidelines apply to a VLS regardless of whether deduplication is enabled or disabled:

- Tape block size: set the tape block size to 256 KB in the Drive Properties of the virtual tape drives configured in Data Protector.
- Multiple paths to tape devices: in an environment with shared tape devices, there may be “MultiPath” devices within Data Protector (where there are both direct and indirect paths to the same shared tape drive). To ensure optimal performance, verify that the direct path is used for backups.
- Data Protector media pool policy: the recommended media pool policy is “Loose” and “Appendable.” This prevents dedicating specific cartridges to specific backup jobs which results in inefficient media usage and can also cause significant impact to VLS deduplication because it does not consistently fill tapes.
- Disable multiplexing: set the concurrency to 1 (default is set in the Drive Properties, and can also be set in the backup specifications), because multiplexing is not recommended for virtual libraries due to the performance impact on restores and on deduplication. Because you can create additional drives easily in the virtual library and allocate them to Media Agents as necessary, sufficient backup streams can be concurrently written without requiring multiplexing.
- Disable client backup compression.

## Data Protector Deduplication Guidelines

In addition to the Data Protector general guidelines, the following guidelines apply to a VLS with deduplication enabled:

- Patch all Windows clients: HP Data Protector on Windows clients requires a software fix to eliminate 56 KB file fragments that diminish deduplication performance on VLS and D2D products. The fix depends on your version of Data Protector:
  - DP 6.0: install the patch `DPWIN_00426` (or later patches). Download the patch from the HP website on the Data Protector version 6.0 “Download drivers and software” page.
  - DP 6.1: included in the 6.11 service pack. If this is not yet available, request `SSPNT610_012` (which you can automatically deploy to all Windows Disk Agent clients) through the Data Protector support channel.
- Disable multiplexing in Oracle RMAN backups: In the VLS firmware version 3.3.0 (or higher) the scripting requirements for Oracle data files have been relaxed, so it is no longer mandatory to assign prefix naming conventions to Oracle data files to make sure the data is properly deduplicated. The minimum requirement for RMAN to deduplicate efficiently is to disable multiplexing (where you have multiple RMAN streams concurrently writing to the same drive). You can disable multiplexing by setting concurrency to 1 in Data Protector and setting `FILESPERSET=1` for each database backup in the RMAN scripts (not needed for archive logs or control files). You can define additional channels to increase the number of parallel backup processes running with `FILESPERSET=1` to generate the same RMAN performance as before, so that each stream goes to a different virtual tape drive.

The following RMAN settings will not deduplicate because they always generate totally unique data every backup:

- Incremental backups
- RMAN compressed backups
- RMAN encryption

---

### NOTE:

If you had previously configured the prefix naming conventions in your Oracle RMAN scripts, this will continue to deduplicate efficiently with the firmware version 3.3.0 (or higher), but you should ensure that multiplexing is disabled as well. If you wish to remove the prefix naming conventions, you will need to rename the backup job name in Data Protector. (New backup jobs will not deduplicate against the previous jobs with the different job name.)

---

- Allocate Unformatted Media First: If you have deduplication enabled, you should disable the “Allocate unformatted media first” option in the media pool. You must also set the global variable `InitOnLoosePolicy=1` defined in the global file and the Data Protector services must be restarted for the global variable to take effect. In HPUX the file is located in `/etc/opt/omni/server/options/global`, and in Windows the file is located in `C:\Program Files\OmniBack\Config\Server\Options\global`. These settings prevents Data Protector from using all available blank cartridges for backups before it starts overwriting expired/scratch cartridges; creating too many virtual cartridges results in the device disk capacity filling up.
- Interactive backups will not deduplicate: If you edit a Data Protector backup specification and then manually run it from the GUI without saving it, the backup session will run with a backup specification name of 'Interactive' and therefore will not deduplicate. This is because deduplication correlates backup versions using the backup specification name.
- Full and incremental backups with different names: Some Data Protector configurations have different backup job names for full backups and for incremental backups (for example, to have dif-

ferent pools with different retention times for fulls vs. incrementals), but this means that the incremental backups will not be fully deduplicated because their backup job names will not match the full backup job names. To address this, the VLS firmware version 3.3.0 includes a new GUI option called “Edit Data Protector Configuration” (in the System Maintenance screen). This option allows you to enter suffixes that will be stripped off the end of the Data Protector backup job names when they are deduplicated. For example, if you have a full backup called `BACKUP1_F` and an incremental backup of the same data called `BACKUP1_I`, you can enter `_F` and `_I` in the Edit Data Protector Configuration and the incremental backups will deduplicate against the full backups.

## Data Protector Import Example Script

The following section details an example script to perform the tape import commands for HP Data Protector. The input to the script is the ISV Import email report (containing a list of cartridges by `/barcode/library/slot` that were successfully replicated and thus are ready for import).

This example uses a Linux server client which contains the email processing script (above) and has CLI access to the Data Protector CLI. The virtual library on the VLS target device has been presented to Data Protector over Fibre Channel and configured in Data Protector as a tape library with tape drives (each tape drive will have its own Data Protector drive name). A recommended configuration would be to reserve a subset of the drives in the library just for use by tape import (i.e., do not use these drives for any other jobs) and the remaining drives can be used for restore jobs or copy to physical tape etc. This is because the import script needs to define the Data Protector drive names that it can use for import jobs and it would make the script more complex to handle drives that might be in use for other jobs when you want to run an import. One way of defining the Data Protector drive names would be to create a lookup file that defines the library serial number and the Data Protector drive name (or names) that can be used for import. For example, a “libmap\_file” containing:

```
DEC06150U7~DP_cell12_lib03_drv1~DP_cell12_lib03_drv2
```

The actual script would process the cartridge list from stdin (identified by the “ISV” tag at the beginning of the line) and use the Data Protector CLI to trigger a tape import on the specified drive(s). For example:

```
#!/bin/bash
#read from stdin
CARTLIST=`grep ISV~ /dev/stdin`
DRVBSY=( 0 0 0 0 )

for CART in $CARTLIST
do
    BARCODE=`echo $CART | awk -F"~" '{print $2}`
    LIBSERIAL=`echo $CART | awk -F"~" '{print $3}`
    BCSLOT=`echo $CART | awk -F"~" '{print $4}`

    DRVNAME[0]=`grep $LIBSERIAL $LIBMAP | awk -F"~" '{print $2}`
    DRVNAME[1]=`grep $LIBSERIAL $LIBMAP | awk -F"~" '{print $3}`
    DRVNAME[2]=`grep $LIBSERIAL $LIBMAP | awk -F"~" '{print $4}`
    DRVNAME[3]=`grep $LIBSERIAL $LIBMAP | awk -F"~" '{print $5}`

    if [ -z ${DRVNAME[0]} ]
    then
        DRVBSY[0]=1
    fi
    if [ -z ${DRVNAME[1]} ]
    then
        DRVBSY[1]=1
    fi
    if [ -z ${DRVNAME[2]} ]
    then
        DRVBSY[2]=1
    fi
    if [ -z ${DRVNAME[3]} ]
    then
        DRVBSY[3]=1
    fi

    #find a drive to use
```

```

for i in `seq 0 3`;
do
  if [ ${DRVBSY[$i]} -eq 0 ]
  then
    /usr/omni/bin/omnimm -recycle $BARCODE
    /usr/omni/bin/omnimm -export $BARCODE
    /usr/omni/bin/omnimm -import ${DRVNAME[$i]} -slot $BCSLOT &

    DRVBSY[$i]=1
    break
  fi
done

if [ ${DRVBSY[0]} -eq 1 ] && [ ${DRVBSY[1]} -eq 1 ] && [ ${DRVBSY[2]} -eq 1 ] && [ ${DRVBSY[3]} -eq 1 ]
then
  wait
  DRVBSY=( 0 0 0 0 )
fi

done
wait

```

The above basic example script template can be extended to add error handling or support for multiple libraries, etc.

## Symantec NetBackup

This section includes both general and deduplication guidelines as well as other useful information for Symantec NetBackup.

### NetBackup General Guidelines

The following NetBackup guidelines apply to a VLS regardless of whether deduplication is enabled or disabled:

- **HP VLS Library Emulation:** When configuring the VLS for use with NetBackup, select “HP VLS” as the library emulation. NetBackup requires the use of this emulation in all HP StorageWorks VLS products for device support purposes.
- **Barcode formats:** When configuring the barcode formats for the virtual cartridges, do not exceed eight characters.
- **Disable multiplexing:** Even though multiplexing is supported, HP does not recommend it for virtual libraries because of the performance impact on restores and deduplication. Because you can create additional drives easily in the virtual library and allocate them to Media Servers as necessary, sufficient backup streams can be concurrently written without requiring multiplexing. Be aware of the Netbackup Device Manager limitation of 128 tape drives per media server.
- **Tape block size:** Set the tape block size to 256 KB (maximum supported size) using `<install path>/netbackup/db/config/SIZE_DATA_BUFFERS`.
- **Storage unit fragment size:** Set the storage unit fragment size to default, 1TB, with the box unchecked.
- **Enable multi-streaming:** Multi-streaming (a backup job concurrently writing to multiple virtual tape drives) provides efficient use of bandwidth when data objects reside on different physical disks/arrays.
- **Avoid multiple retention periods per cartridge:** Avoid mixing different retention periods on the same cartridges; this is off by default.
- **Number buffers:** Set the number of data buffers to 32 or 64 using `<install path>/netbackup/db/config/NUMBER_DATA_BUFFERS`.

## NetBackup Deduplication Guidelines

In addition to the NetBackup general guidelines, the following guidelines apply to a VLS with deduplication enabled:

- NetBackup raw/image file backups: NetBackup raw file backups will only deduplicate if the differencing algorithm for the corresponding policy is set to Backup-level. If the policy for these files is set to File-level, these files will not be deduplicated.
- NetBackup catalog backups: NetBackup catalog backups, which consist of the image record files containing the information about the backups, will not deduplicate.
- Volume pools: Use as few as possible to quickly fill up tapes.
- Turn Off NetBackup Checkpoint: The NetBackup Checkpoint option creates a duplicate backup set with a duplicate name. This prevents the VLS from deduplicating data with that policy because it cannot process duplicate backup sets. You must turn off the NetBackup Checkpoint option for deduplication to work.
- Disable multiplexing in Oracle RMAN backups: In the VLS firmware version 3.3.0 (or higher) the scripting requirements for Oracle data files have been relaxed, so it is no longer mandatory to assign prefix naming conventions to Oracle data files to make sure the data is properly deduplicated. The minimum requirement for RMAN to deduplicate efficiently is to disable multiplexing (where you have multiple RMAN streams concurrently writing to the same drive). You can disable multiplexing by setting `FILESERSET=1` for each database backup in the RMAN scripts (not needed for archive logs or control files). You can define additional channels to increase the number of parallel backup processes running with `FILESERSET=1` to generate the same RMAN performance as before, so that each stream goes to a different virtual tape drive.

The following RMAN settings will not deduplicate because they always generate totally unique data every backup:

- Incremental backups
- RMAN compressed backups
- RMAN encryption

---

### NOTE:

If you had previously configured the prefix naming conventions in your Oracle RMAN scripts, this will continue to deduplicate efficiently with the firmware version 3.3.0 (or higher), but you should ensure that multiplexing is disabled as well. If you wish to remove the prefix naming conventions, you will need to rename the backup job name in NetBackup. (New backup jobs will not deduplicate against the previous jobs with the different job name.)

---

## NetBackup Import Example Script

The following section details a script to perform the tape import commands for Symantec NetBackup. The input to the script is the ISV Import email report (containing a list of cartridges by /barcode/library/slot that were successfully replicated and thus are ready for import). This example uses a Linux server client which contains the [Email Processing Example Script](#) and has CLI access to the NetBackup CLI. The virtual library on the VLS target device has been presented to NetBackup over Fibre Channel and configured in NetBackup as a tape library with tape drives. To perform the import you must define the target virtual library in NetBackup terms with a robot type and robot number. For example, create a `libmap_file` containing entries with the library serial number, library type, and library number:

```
DEC06150U7~TLD~2
```

The actual script would process the cartridge list from `stdin` (identified by the "ISV" tag at the beginning of the line) and use the NetBackup CLI to trigger a tape import on the specified library. For example:

```
#!/bin/bash
#set -x

#read from stdin
CARTLIST=`grep ISV~ /dev/stdin`

for CART in $CARTLIST
do
  BARCODE=`echo $CART | awk -F"~" '{print $2}'`
  LIBSERIAL=`echo $CART | awk -F"~" '{print $3}'`
  BCSLOT=`echo $CART | awk -F"~" '{print $4}'`

  ROB_TYPE=`grep $LIBSERIAL libmap_file | awk -F"~" '{print $2}'`
  ROB_NUM=`grep $LIBSERIAL libmap_file | awk -F"~" '{print $3}'`

  vmupdate -rt $ROB_TYPE -rn $ROB_NUM
  bpmedia -unfreeze -m $BARCODE
  bpmedia -unsuspend -m $BARCODE
  bpexpdate -d 0 -force -m $BARCODE
  bpimport -create_db_info -id $BARCODE &
done
wait
```

## IBM TSM

This section includes both general and deduplication guidelines as well as other useful information for IBM TSM.

### TSM General Guidelines

The following TSM guidelines apply to a VLS regardless of whether deduplication is enabled or disabled:

- Configuration of libraries: See the *HP StorageWorks Enterprise Backup Solution with Tivoli Storage Manager implementation guide* for details on how to configure the device drivers for HP libraries and drives, including how to set the library element numbers, install the appropriate "tmscsi" drivers, etc.  
The "HP VLS" library emulation is only supported in TSM version 5.5.2.1 or higher.
- Disable client compression: When configuring clients on the TSM server, HP recommends not enabling client compression. The built-in compression in the virtual library device gives better performance. In addition, the client compression prevents deduplication from operating. Therefore, when registering a new client node on the server, set the client compression setting to NO rather than CLIENT. The default value for the `COMPRESSION` parameter for the `REGISTER Node` command is CLIENT.
- TSM server/client performance: There are a multitude of settings available in the TSM server options file that make TSM very configurable in a wide variety of environments. However, there are a few key recommended settings shown in [Table 28](#) and [Table 29](#) to improve performance with backup and restore to tape libraries. See the *IBM Tivoli Storage Manager Performance Tuning Guide* for more information on server and client performance.

**Table 28 Recommended Settings for each TSM Server's `dsmserv.opt` File or Storage Agent's `dsmsta.opt`**

Command <sup>1</sup>	Value	Notes
TCPNODELAY	YES	Default

Command <sup>1</sup>	Value	Notes
TCPWINDOWSIZE	63	Default
BUFPOOLSIZE	32768	Default
TXNGROUPMAX	256	Default
EXPINTERVAL	0	Default = 24. When setting this to 0, use an administrative schedule to execute expiration at an appropriate time each day. A value of 0 means the expiration must be started with the EXPIre Inventory command.
SELFTUNEBUFPOOLSIZE	NO	Default
MOVEBATCHSIZE	1000	Default
MOVESIZETHRESH	2048	Default
LOGPOOLSIZE	2048 to 8192	Default = 512

<sup>1</sup>These settings have been tested and verified on all EBS-supported platforms. However, there may be additional factors that might require further adjustments of these variables. See the *IBM Tivoli Storage Manager Performance Tuning Guide*, available on the IBM support web site, for more information.

**Table 29 Recommended Settings for each Backup-archive Client's dsm.sys File (dsm.opt File for Windows)**

Command <sup>1</sup>	Value	Notes
TCPBUFFSIZE	32	Default
TCPNODELAY	YES	Default
TCPWINDOWSIZE	63	Default
LARGECOMMBUFFERS	YES	Default = NO
TXNBYTELIMIT	2097152	For backup to tape/virtual tape only.

## TSM Deduplication Guidelines

In addition to the TSM general guidelines, the following guidelines apply to a VLS with deduplication enabled:

- Incremental forever: The TSM (Tivoli Storage Manager) "incremental forever" scheme of backups are all effectively treated as full backups by the deduplication, so every new backup is compared against the previous backup. With file server backups that contain just the modified files, this may result in lower deduplication ratios because most or all of the new backup data is unique, but database backups (which are normally the entire database every time) will give normal deduplication ratios.
- Mount retention: You must set the mount retention to 0 (zero) for your server configuration as follows:  
`dsmadm -id=admin -pass=admin update devclass <vtlclass> mountret=0`  
Substitute the password for "admin" above with your TSM server admin password.



- Keep mount point: The Keep Mount Point setting on each node must be set to Yes using the command `keepmp=yes`. This command only works if there are no client sessions running when enabled.

Update node \* keepmp=yes

Replace the asterisk (\*) with a node name to update an individual node.

- Resource utilization: Use the “resourceutilization” option in your `dsm.sys` file to regulate the level of resources the TSM server and client can use during processing. To limit the number of backup sessions per resource (i.e., per tape drive) to one, use the resourceutilization setting of 2; this is equivalent to limiting the VLS to one stream. If you do not limit the number of resources to one, multiple backup sets will appear in the user interface for a single backup and the deduplication ratios will be significantly reduced.
- Management classes: To assign a management class other than the default to all files, use an include statement in your client systems options file, or by using an include-exclude file specified by the “includ” option. To specify your own default management class for files that are not explicitly included, specify the following as the first include or exclude option defined:

For non-Windows hosts (e.g., HP-UX, Linux, AIX), add the following to the `dsm.sys` file for the appropriate directories being backed up:

- This will backup all directories and files under `/backup/`:

```
include /backup/.../* mgmt_class_name
```

- This will backup all directories and files under `/`:

```
include /.../* mgmt_class_name
```

For Windows hosts, to specify a management class named `mgmt_class_name` for all files to which you do not explicitly assign a management class, add the following to the `dsm.opt` file for Windows systems for the appropriate volumes/directories being backed up:

- This will backup all directories and files under `e:\backup`:

```
include e:\backup\...\\* mgmt_class_name
```

- This will backup all directories and files under `e:\`:

```
include e:\...\\* mgmt_class_name
```

- This will backup all directories and files under `d:\`:

```
include d:\...\\* mgmt_class_name
```

The management class names are not case-sensitive.

- Disk pools: A disk pool is a primary storage pool that stores data on disks. When using TSM, the primary disk pool is typically the initial destination for data being backed up. This storage pool is sized to accommodate the projected amount of daily changed data in your environment, as with an incremental backup. The disk pool is typically configured to migrate to the next storage pool in the hierarchy, the tape storage pool. When TSM migrates data from a disk pool, it determines what data to migrate based on the node that has the most amount of data stored in the disk pool. TSM moves the data for each node in the order in which it was written to the storage pool. The cache feature on the primary disk pool, which is disabled by default, allows a rapid restore of recently backed up data as necessary.

The recommended option is to eliminate the disk pool and back up directly to the VLS. However, if you continue using a disk pool for some of your backups (for example, to provide a very high availability primary pool), the disk pool must be sized to contain the entire night’s backup to ensure good deduplication. At a minimum, the disk pool must be equal to you largest nightly backup.

- Copy storage pools: The copy storage pool provides a way to recover from disasters or media failures. TSM relies on the process of backing up its primary storage pools to an off-site copy pool to recover the TSM environment if a disaster (e.g., destruction of the data center) occurs. With a backup of the TSM database in a remote location, the off-site copy storage pool is a duplicate of

all objects in the primary storage pools. Data is stored in a copy storage pool when you back up active and inactive data in a primary storage pool. A copy storage pool uses only sequential-access storage (e.g., a tape or FILE device class).

The following TSM copy methods will deduplicate successfully in the VLS:

- Using “simultaneous writes” to a primary storage pool and one or more copy storage pools. For example, you could backup to a disk pool and simultaneously copy to one or more VLS devices, or you could backup to a VLS and simultaneously copy to another VLS.
- Copying data from a disk storage pool to a VLS. This requires the disk pool size to be at least equal to your largest nightly backup; see “Disk pools” above.

---

 **NOTE:**

Deduplication does not currently support backing up to one VLS and then copying from that VLS to another VLS. Deduplication is performed on the first VLS but not the second.

---

- Virtual volumes: TSM allows one TSM server to store its data on another TSM server. The TSM server that sends the data to be stored on the other server is the source server, and the server that stores the received data is the target server. With virtual volumes, the source server treats the target server like a tape library, and the target server treats the source server as a special type of client node.

Deduplication does *not* support TSM virtual volumes.

- Reclamation: **Limit the number of concurrent reclamation jobs per pool to one.** TSM reclamation is the process of reclaiming space on sequential access storage pool volumes. Space becomes reclaimable when files expire or are deleted (as identified by expiration processing). Reclamation processing consolidates the remaining active data from multiple sequential access volumes onto fewer volumes.

Reclaimed cartridges, those which were the source of reclaimed data, are essentially untouched by reclamation but are made available (scratched) for later reuse. The TSM reclaimed data is copied to a new cartridge, but none of this data will be deduplicated. The TSM reclaimed data is shown in the user interface as “Disabled, Unsupported and Overhead Data” on the **System Capacity** screen (**System > Chassis > Deduplication > System Capacity**).

The TSM default reclamation threshold should be set to at least 90% or higher. If you run reclamation, reclaimed backups are expanded and subsequent backups are not deduplicated against them. For example, with the default reclamation threshold at 90%, 10% of the data will not be deduplicated.

You should run reclamation outside of the backup window because it reduces processing bandwidth. The reclamation process is very device-intensive and requires at least two available drives in the library (to read and to write). You must increase the amount of available capacity for the retention period of the reclaimed backup data.

- Co-location: Co-location minimizes the number of tape volumes that need to be written to and is activated on the storage pool level. If the storage pool that you send backups to is co-located, then backups are co-located as specified by that storage pool. The greatest degree of co-location is filespace co-location, followed by client co-location, then co-location groups. The more data you collocate, the fewer tape mounts are required during restores. The downside is that you will significantly increase the amount of tapes you use.  
Deduplication *does not* support co-location because this feature is unnecessary when using a disk-based backup solution.
- NDMP backup: Deduplication *does not* currently support LAN-free NDMP with TSM. This means you can only backup an NDMP filer over the LAN to the VLS.

- Disable multiplexing in Oracle RMAN backups: In the VLS firmware version 3.3.0 (or higher) the scripting requirements for Oracle data files have been relaxed, so it is no longer mandatory to assign prefix naming conventions to Oracle data files to make sure the data is properly deduplicated. The minimum requirement for RMAN to deduplicate efficiently is to disable multiplexing (where you have multiple RMAN streams concurrently writing to the same drive). You can disable multiplexing by setting `FILESERSET=1` for each database backup in the RMAN scripts (not needed for archive logs or control files). You can define additional channels to increase the number of parallel backup processes running with `FILESERSET=1` to generate the same RMAN performance as before, so that each stream goes to a different virtual tape drive. The RMAN script defines the number of channels configured (drive mounts). The number of channels must be equal to the client node's maximum number of mount points.

The following RMAN settings will not deduplicate because they always generate totally unique data every backup:

- Incremental backups
- RMAN compressed backups
- RMAN encryption

---

 **NOTE:**

If you had previously configured the prefix naming conventions in your Oracle RMAN scripts, this will continue to deduplicate efficiently with the firmware version 3.3.0 (or higher), but you should ensure that multiplexing is disabled as well. If you wish to remove the prefix naming conventions, you will need to rename the backup job name in NetBackup. (New backup jobs will not deduplicate against the previous jobs with the different job name.)

---

- TSM image files: TSM image file backups will only deduplicate if the differencing algorithm for the corresponding policy is set to Backup-level. If the policy for these files is set to File-level, these files will not be deduplicated. To simplify configuration, define a specific management class for TSM image backups and configure that class in the VLS GUI for a backup-level deduplication policy.

## TSM Useful Queries

The following calls can help you better understand the current TSM environment and help with sizing the virtual library:

- This select statement will list all primary storage pools that are not DISK and show the amount of backup data in those pools:
 

```
select node_name, stgpool_name, physical_mb from occupancy where stgpool_name in (select stgpool_name from stgpools where pool-type='PRIMARY'and devclass<>'DISK') and type='Bkup'
```
- This select statement will show the totaled sizes listed in the output from the previous select statement:
 

```
select sum(physical_mb) as "Primary Pool Storage (MB)" from occupancy where stgpool_name in (select stgpool_name from stgpools where pool-type='PRIMARY' and devclass<>'DISK') and type='Bkup'
```
- Display GB's data backed up broken down by "nodename" for a 7 day period:
 

```
select entity, cast(float(sum(bytes))/1024/1024/1024 as dec(8,4)) as "Backed up data in GB" from summary where activity='BACKUP' and start_time>timestamp(current_date)-(7)days group by entity
```
- Display backup totals in TB's for a 7 day period:

```
select cast(float(sum(bytes))/1024/1024/1024/1024 as dec(8,4)) as
"Backed up data in TB" from summary where activity='BACKUP' and
start_time>timestamp(current_date)-(7)days
```

- Display backup volume in Bytes by node for the past 24 hours:  

```
select start_time, end_time, entity, bytes from summary where activ-
ity='BACKUP' and start_time>timestamp(current_date)-(1)days order by
end_time
```
- Display value of mount retention for the device class:  

```
Query device class q devc f=d
```
- Display details on the TSM options:  

```
Query TSM Options q opt
```
- Display detailed node information:  

```
Query q n f=d
```
- Display if client side compression is being used (should be disabled if you want to run deduplication or even compression in the virtual library):  

```
q ac begint=-24 search="compressed by"
```
- Display if collocation is turned on – it should be turned off for a virtual library:  

```
Query storagepool f=d
```

## EMC NetWorker

This section includes both general and deduplication guidelines as well as other useful information for EMC NetWorker.

### NetWorker General Guidelines

The following NetWorker guidelines apply to a VLS regardless of whether deduplication is enabled or disabled:

- Tape block size: Set the tape block size to 256 KB (maximum supported size, 512KB is not supported). Otherwise, NetWorker uses a block size based on the device type used. For example, LTO-2 is set by EMC to 64 KB, assuming that the operating system is configured for variable block size mode. You can experience metadata errors when setting a fixed block size on your initiator if the block size in the `stinit.def` file does not match the NetWorker block size, so set a fixed block size on your initiator in the `stinit.def` file that matches the NetWorker block size. Use the full number in the `stinit.def` file with a lowercase “k.”
- Disable multiplexing: Even though multiplexing is supported, HP does not recommend it for virtual libraries because the performance impact on restores and deduplication. Because you can create additional drives easily in the virtual library and allocate them to Storage Nodes as necessary, sufficient backup streams can be concurrently written without requiring multiplexing. To disable multiplexing, do one of the following:
  - Default pool: Set the NetWorker device “target sessions” for any virtual tape library tape drives that will be used by the Default pool to 1.
  - All other pools: Set the pool “parallelism” parameter to 1, which enforces one stream per virtual tape library tape drive when writing to a pool configured with this parameter. HP recommends using a non-default pool with the virtual tape library with the “parallelism” parameter, because NetWorker autochanger reconfigurations may delete and recreate devices, and re-setting the “target sessions” parameter to the default of 4.

- Windows device configuration: The following guidelines apply to NetWorker on Windows:
  - Disable Plug and Play Test Unit Ready (TUR) for any tape drivers in use. A SCSI TUR command may rewind a mounted NetWorker tape unexpectedly, causing data loss.
  - Enable “verify header on eject” in the NetWorker autochanger. This marks any tapes corrupted by a SCSI TUR reset/rewind as suspect.
  - For versions of NetWorker newer than 7.4.1, use Windows Tape Drive Persistence to ensure that tape device names do not change when Windows is rebooted. This will enable device paths other than the standard `\\.\Tape0` to be used, such as `\\.\Tape2147483646` or `\\.\tape4801101`. For Windows Tape Drive Persistence, see <http://support.microsoft.com/kb/873337>.
  - Confirm that the Fibre Channel device supports a maximum block size larger than that being used in NetWorker. This can most easily be done in Windows via the NetWorker `mt` command. If the maximum supported block size found via the `mt` command is too small (typically 64 or 96 KB), contact the Fibre Channel HBA Vendor for assistance in modifying the SCSI Scatter/Gather registry setting in the Windows Registry for the Fibre Channel HBA. The following is an example `mt` command:

```
C:\>mt -f \\.\Tape4801110 status
\\.\Tape4801110:
Media Capacity = 292.97GByte
Media Remaining = 289.95GByte
Media Blocksize = 0
Media Partition Count = 1
Media is not write protected
default blocksize = 65536
maximum blocksize = 524288
```

- Linux device configuration: The following guidelines apply to NetWorker on Linux:
  - Configure `stinit.def` to enable variable block size. This requires the `st-mt` package to be installed and `/etc/stinit.def` to be configured. See the Linux man page for configuration details and examples. HP recommend setting `stinit` to load via an init file (for example, `/etc/rc.local` on Redhat) to ensure it is loaded properly.
  - NetWorker supports Linux `udev` device naming. See the EMC whitepaper *Persistent Binding and udev Changes for EMC NetWorker* for details (document h5795, available at <https://powerlink.emc.com>).
- UNIX device configuration: The following guidelines apply to NetWorker on UNIX:
  - AIX – Ensure that variable block mode is enabled. AIX does not have a persistent naming option, although the ODM database generally tracks device links properly across reboots.
  - Solaris – Ensure that the tape drive is supported in the kernel. This is most easily done by running `strings /kernel/drv/sparcv9/st |grep -i LTO` (for LTO) and comparing this to the tape drive string found in `/var/adm/messages`. If tape drive is unsupported by the kernel, configure `st.conf` to allow variable block mode. For Solaris Persistent Naming, see Section 8, *Persistent Binding for Tape Devices* available at <http://docs.sun.com/source/819-0139/>.
  - HP-UX – An issue exists when EMC NetWorker uses tape library SCSI commands via activation of a Common Device Interface (CDI) with the HP-UX 11.31. Operating CDI is activated by default in EMC NetWorker. When CDI is activated, HP-UX `estape` driver has a problem distinguishing between EOT and BOT, thus causing job failures. The EMC messages log may contain the following as symptom:

```
-----
Oct 22 13:30:29 nairobi: NetWorker media: (warning)
rd=suzanne.ebs.net:/dev/rmt/c8t2d0BESTnb writing: continuation, No space left on
device, at file 2 record 0Oct 22 13:30:29 nairobi: NetWorker media: (emergency)
```

Media is marked as full and the application selects another tape and progresses through the same until all tapes are listed as "full." The recommended action is to obtain and install HPUX 11.31 patch PHKL\_36312.

## NetWorker Deduplication Guidelines

In addition to the NetWorker general guidelines, the following guidelines apply to a VLS with deduplication enabled:

- Pools: Use as few pools as possible to quickly fill up tapes. You should create a separate pool for each projected retention period. For example:
  - Daily (except for the last day of the month)
  - One Month
  - One Year
- Enable auto-ejection: Enable auto-ejection by setting the `idle device timeout` parameter in the NetWorker autochanger (tapes will not be deduplicated until they are ejected from the virtual tape drive).
- Disable multiplexing in Oracle RMAN backups: In the VLS firmware version 3.3.0 (or higher) the scripting requirements for Oracle data files have been relaxed, so it is no longer mandatory to assign prefix naming conventions to Oracle data files to make sure the data is properly deduplicated. The minimum requirement for RMAN to deduplicate efficiently is to disable multiplexing (where you have multiple RMAN streams concurrently writing to the same drive). You can disable multiplexing by setting `FILESERSET=1` for each database backup in the RMAN scripts (not needed for archive logs or control files). You can define additional channels to increase the number of parallel backup processes running with `FILESERSET=1` to generate the same RMAN performance as before, so that each stream goes to a different virtual tape drive.

The following RMAN settings will not deduplicate because they always generate totally unique data every backup:

- Incremental backups
- RMAN compressed backups
- RMAN encryption

---

### NOTE:

If you had previously configured the prefix naming conventions in your Oracle RMAN scripts, this will continue to deduplicate efficiently with the firmware version 3.3.0 (or higher), but you should ensure that multiplexing is disabled as well. If you wish to remove the prefix naming conventions, you will need to rename the backup job name in NetBackup. (New backup jobs will not deduplicate against the previous jobs with the different job name.)

---

- Disable multiple streams per data file for SQL database backups: Multiple streams from a given data file will reduce deduplication effectiveness. For some database backups, NetWorker can create multiple savestreams per database file. Disabling this feature is critical when using deduplication. Multiple savesets from a database is not an issue, only multiple streams from a data file. Disable the "Striped Backup" feature in the NetWorker Module for Microsoft SQL Server because this creates multiple streams from each data file. To disable this for a scheduled backup, edit the NetWorker client object for the SQL backup and remove the `-sn` option from the backup command (`n` is the number of parallel streams per data file; removing this argument from the backup com-

mands sets it to the default of 1 stream per data file). The default setting for this is disabled, and most sites do not use Striped Backups with SQL Server.





---

# 6 Support and Other Resources

## Related Information

### Documents

- *HP StorageWorks 9000-series virtual library system user guide*
- *HP StorageWorks 300 and 12000 Gateway virtual library system user guide*
- *HP StorageWorks 6000-series virtual library system user guide*
- *HP StorageWorks Enterprise Backup Solution design guide*
- *HP StorageWorks D2D Backup System HP D2D4004 and HP D2D4009 User Guide*
- *HP StorageWorks D2D Backup System HP D2D2502, HP D2D2503 and HP D2D2504 User Guide*
- *HP StorageWorks D2D Backup System HP D2D4112 User Guide*

### Websites

- HP website:  
<http://www.hp.com>
- HP Enterprise Backup Solutions Overview and Features:  
<http://www.hp.com/go/ebs>
- HP StorageWorks Sizer:  
<http://www.hp.com/go/storageworks/sizer>
- HP Data Storage:  
<http://www.hp.com/go/storage>
- HP Library and Tape Tools:  
<http://www.hp.com/support/tapetools>
- HP Technical Documentation:  
<http://www.docs.hp.com>
- HP Partner Locator:  
[http://www.hp.com/service\\_locator](http://www.hp.com/service_locator)

## Contacting HP

### Before you contact HP

Be sure to have the following information available before calling HP:

- Technical support registration number (if applicable)

- Product serial numbers
- Product model names and numbers
- Applicable error messages
- Operating system type and revision level
- Detailed, specific questions

## HP contact information

For the name of the nearest HP authorized reseller:

- See the Contact HP worldwide (in English) webpage ([http://welcome.hp.com/country/us/en/wwcontact\\_us.html](http://welcome.hp.com/country/us/en/wwcontact_us.html)).
- In the United States, call 1-800-345-1518.
- In Canada, call 1-800-263-5868.

For HP technical support:

- In the United States, for contact options see the Contact HP United States webpage ([http://welcome.hp.com/country/us/en/contact\\_us.html](http://welcome.hp.com/country/us/en/contact_us.html)). To contact HP by phone:
  - Call 1-800-HP-INVENT (1-800-474-6836). This service is available 24 hours a day, 7 days a week. For continuous quality improvement, calls may be recorded or monitored.
  - If you have purchased a Care Pack (service upgrade), call 1-800-633-3600. For more information about Care Packs, refer to the HP website (<http://www.hp.com/hps>).
- In other locations, see the Contact HP worldwide (in English) webpage (<http://welcome.hp.com/country/us/en/wwcontact.html>).

## Subscription Service

HP strongly recommends that customers register online using the Subscriber's choice web site: <http://www.hp.com/go-e-updates>.

Subscribing to this service provides you with e-mail updates on the latest product enhancements, newest driver versions, and firmware documentation updates as well as instant access to numerous other product resources.

After subscribing, locate your products by selecting **Business support** and then **Storage** under Product Category.

## Document Conventions and Symbols

**Table 30 Document conventions**

Convention	Element
Blue text: <a href="#">Table 30</a>	Cross-reference links and e-mail addresses
Blue, underlined text: <a href="http://www.hp.com">http://www.hp.com</a>	Web site addresses
<b>Bold text</b>	<ul style="list-style-type: none"> <li>• Keys that are pressed</li> <li>• Text typed into a GUI element, such as a box</li> <li>• GUI elements that are clicked or selected, such as menu and list items, buttons, tabs, and check boxes</li> </ul>

Convention	Element
<i>Italic text</i>	Text emphasis
Monospace text	<ul style="list-style-type: none"> <li>• File and directory names</li> <li>• System output</li> <li>• Code</li> <li>• Commands, their arguments, and argument values</li> </ul>
<i>Monospace, italic text</i>	<ul style="list-style-type: none"> <li>• Code variables</li> <li>• Command variables</li> </ul>
Monospace, bold text	Emphasized monospace text

---

**⚠ WARNING!**

Indicates that failure to follow directions could result in bodily harm or death.

---



---

**⚠ CAUTION:**

Indicates that failure to follow directions could result in damage to equipment or data.

---



---

**📝 NOTE:**

Provides additional information.

---

The following equipment symbols may be found on hardware to which this guide pertains. They have the following meanings:

---

**⚠ WARNING!**



These symbols, which mark an enclosed surface or area of the equipment, indicate the presence of electrical shock hazards. The enclosed area contains no operator serviceable parts.

**WARNING:** To reduce the risk of injury from electrical shock hazards, do not open this enclosure.

---



---

**⚠ WARNING!**



These symbols, which mark an RJ-45 receptacle, indicate a network interface connection.

**WARNING:** To reduce the risk of electrical shock, fire, or damage to the equipment, do not plug telephone or telecommunications connectors into this receptacle.

---

---

⚠ **WARNING!**



These symbols, which mark a surface or area of the equipment, indicate the presence of a hot surface or hot component. Contact with this surface could result in injury.

**WARNING:** To reduce the risk of injury from a hot component, allow the surface to cool before touching.

---

---

⚠ **WARNING!**



Power supplies or systems marked with these symbols indicate the presence of multiple sources of power.

**WARNING:** To reduce the risk of personal injury from electrical shock, remove all power cords to completely disconnect power from the power supplies and systems.

---

---

⚠ **WARNING!**



Any product or assembly marked with these symbols indicates that the component exceeds the recommended weight for one individual to handle safely.

**WARNING:** To reduce the risk of personal injury or damage to the equipment, observe local occupational health and safety requirements and guidelines for manually handling material.

---

---

# Glossary

<b>automigration</b>	The feature in which the virtual tape library acts as a tape copy engine that transfers data from virtual cartridges on disk to a physical tape library connected to the virtual tape device.
<b>D2D</b>	The HP StorageWorks D2D Backup Systems product line.
<b>deduplication</b>	The feature in which only a single copy of a data block is stored on a device. Duplicate information is removed, thereby reducing the amount of storage used by a given data block.
<b>NAS</b>	Network attached storage.
<b>replication</b>	Automigration between two VLS devices via LAN/WAN. <i>See also</i> automigration
<b>SAN</b>	Storage area network.
<b>VLS</b>	The HP StorageWorks Virtual Library System product line.
<b>VTL</b>	Virtual tape library.
<b>WAN</b>	Wide area network.
<b>WORM</b>	Write once, read many. A characteristic of data storage media.



---

# Index

## A

- Accelerated deduplication
  - backup job naming, 160
  - capacity sizing, 158
  - client naming, 160
  - defined, 148
  - design considerations, 157
  - disabling on specific backup policies, 161
  - how it works, 149
  - implementing, 152
  - licensing, 154
  - media management, 160
  - migrating existing data, 155
  - node oversubscription, 158
  - optimizing file server deduplication, 159
  - performance, 157
  - preparing for, 153
  - record sizing, 159
  - reporting, 155
  - supported backup applications and data types, 153
- active-active replication, 27
- active-passive replication, 27
- application-based disk backup
  - alternative to virtual tape, 18
- archive tapes, creating, 44
- authorized reseller, 210
- automigration
  - copying to tape, 41
  - defined, 140
  - design considerations, 145
  - implementing, 141
  - restoring from, 148
  - setting up, 143
  - use models, 145

## B

- backup solutions, 31
  - analyzing environment, 31
  - application guidelines, 50
  - backup considerations, 31
  - blocksize, 37
  - copy considerations, 38
  - future growth, 38
  - identifying performance bottleneck, 45
  - LAN-free, 37
  - multiple library
    - See multiple library backup
  - restore considerations, 33
  - retention planning, 37
  - single library
    - See single library backup
  - single vs. multiple library, 33
  - speed considerations, 33
  - transfer size, 37
- backup technologies, 16
- business copy
  - alternative to virtual tape, 19

## C

- capacity
  - See storage capacity
- cartridge sizing
  - D2D, 59
- conventions
  - text symbols, 211
- conventions, document, 210
- copying data
  - considerations, 38
  - to physical tape using backup application, 39
  - using D2D tape offload, 43
  - using replication, 43
  - using VLS automigration, 41

## D

### D2D

- cartridge sizing, 59
  - defined, 55
  - design considerations, 56
  - Dynamic deduplication, 69
  - housekeeping, 72
  - implementing replication, 73
  - multiple backup streams, 57
  - offload performance, 58
  - port optimization, 56
  - portfolio, 16
  - reducing housekeeping, 59
  - replication, 73
  - replication data recovery, 95
  - replication design, 84
  - replication reporting, 82
  - restoring data, 71
  - setting up replication, 76
  - tape library emulation, 59
  - tape offload, 98
  - technical specifications, 56
  - typical environment, 17
- D2D2500
- technical specifications, 56
- D2D4000
- technical specifications, 56
- D2D4112
- technical specifications, 56
- data
- backing up to VLS or D2D, 31
  - copying considerations, 38
  - future growth, 38
  - restoration, 44
    - from disk backup, 45
    - from replication target, 45
    - from tape copy, 45
  - retention planning, 37
- data protection
- before replication, 25
  - with replication, 25
- Data Protector
- See HP Data Protector
- data recovery
- D2D replication, 95

### deduplication

- See Accelerated deduplication (VLS) or Dynamic deduplication (D2D)
  - Accelerated, 148
  - backup application guidelines, 194
  - defined, 20
  - Dynamic, 69
  - HP solutions, 20
  - ratios, 21
  - tape oversubscription, 23
  - target-based, 22
- design considerations
- D2D, 56
  - D2D replication, 84
  - for backup system, 31
  - of automigration, 145
  - of the VLS, 123
- disk-based backup, 15
- document conventions, 210
- Dynamic deduplication
- defined, 69
  - how it works, 69
  - implementing, 70
  - licensing, 74

## E

- echo copy
- benefits, 42
  - considerations, 42
- Ethernet ports on the D2D, 56

## H

- hash-based chunking, 69
- help, obtaining, 210
- housekeeping
- See D2D housekeeping
- HP
- storage web site, 209
  - StorageWorks, deduplication, 20
  - StorageWorks, replication, 24
  - Subscriber's choice web site, 210
- HP Data Protector
- deduplication guidelines, 195
  - general guidelines, 194
  - import example script, 196
  - with HP Accelerated deduplication, 194
  - with HP virtual tape libraries, 51

## I

- IBM TSM
- with HP Accelerated deduplication, 199
  - with HP virtual libraries, 52



## L

- LAN-free backups, 37
- licensing
  - Accelerated deduplication, 154
  - capacity on the VLS, 117
  - D2D replication, 74
  - Dynamic deduplication, 74
  - VLS replication, 164
- link sizing
  - D2D replication, 84
- LUN mapping, 49
- LUN masking, 49

## M

- many-to-one replication, 27
- mapping
  - See LUN mapping
- masking
  - See LUN masking
- media server considerations, 40
- multipathing, 36
- multiple library backup
  - benefits, 35
  - use model, 35
- multiplexing, 36
- multistreaming, 36

## N

- n-way replication, 27
- NAS
  - alternative to virtual tape, 18
- NetBackup
  - See Symantec NetBackup
  - deduplication guidelines, 198
  - general guidelines, 197
  - import example script, 198
- NetWorker
  - deduplication guidelines, 206
  - general guidelines, 204
  - with HP Accelerated deduplication, 204
- node oversubscription, 158

## O

- object-level differencing, 148
- operating system
  - tape configuration, 47
- OS
  - See operating system
- oversubscription, node, 158
- oversubscription, tape, 23

## P

- performance bottlenecks, 45
- physical tape
  - alternative to virtual tape, 18
- port optimization, D2D, 56

## R

- replication
  - active-active, 27
  - active-passive, 27
  - benefits, 44
  - considerations, 44
  - creating archive tapes, 44
  - D2D, 73
    - configuring, 78
    - data recovery options, 95
    - design considerations, 84
    - how it works, 73
    - implementing, 73
    - initialization, 74
    - licensing, 74
    - link sizing, 84
    - reporting, 82
    - setup, 76
  - defined, 23
  - deployment options, 26
  - HP solutions, 24
  - licensing on the VLS, 164
  - limitations, 29
  - many-to-one, 27
  - n-way, 27
  - restoring data from target, 45
  - time estimates, 26
  - VLS, 161
    - data recovery, 187
    - design considerations, 180
    - device sizing, 187
    - how it works, 161
    - implementing, 164
    - incremental backups, 163
    - initialization, 172
    - job control, 180
    - link sizing, 182
    - multi-node scaling, 164
    - preparation, 165
    - reporting, 176
    - setup, 166
    - with backup applications, 28
  - reporting
    - Accelerated deduplication, 155
    - replication on the VLS, 176

- restoring data
  - from disk backup, 45
  - from replication target, 45
  - from tape copy, 45
- retention planning, 37

## S

- SAN
  - backup guidelines, 46
  - design considerations, 46
  - zoning, 46
- script examples
  - HP Data Protector, 196
  - NetBackup, 198
- single library backup
  - benefits, 34
  - considerations, 34
  - use model, 33
- snap and clone
  - See business copy
- specifications
  - See technical specifications
- storage capacity
  - saved by deduplication, 21
- Subscriber's choice, HP, 210
- Symantec NetBackup
  - with HP Accelerated deduplication, 197
  - with HP virtual libraries, 51
- symbols in text, 211
- symbols on equipment, 211

## T

- tape
  - See physical tape
- tape offload
  - See D2D tape offload
- tape oversubscription, 23
- technical specifications
  - D2D2500, 56
  - D2D4000, 56
  - D2D4112, 56
  - VLS12000, 108, 109, 111
  - VLS6000-series, 107
  - VLS9000-series, 105, 106
- technical support, 210
- Telco
  - cost examples, 92
  - terminology, 93
- text symbols, 211

- TSM
  - See IBM TSM
  - deduplication guidelines, 200
  - general guidelines, 199
  - useful queries, 203

## V

- virtual tape libraries, 15
  - alternatives to, 18
  - integration, 15
  - problems addressed, 15
- VLS
  - Accelerated deduplication, 148
  - automigration, 140
  - benefits of, 102
  - capacity licensing, 117
  - defined, 101
  - design considerations, 123
  - how it works, 111
  - implementing, 115
  - portfolio, 16
  - scalability, 111
  - sizing, 123
  - technical specifications, 104
  - typical environment, 17
  - warm failover, 114
- VLS12000
  - technical specifications, 108, 109, 111
- VLS6000-series
  - technical specifications, 107
- VLS9000-series
  - technical specifications, 105, 106

## W

- warm failover, 114
- web sites
  - HP documentation, 209
  - HP storage, 209
  - HP Subscriber's choice, 210
- write-to-disk, compared to virtual tape, 19