

# HP Data Protector 6.20

## StoreOnce Software Deduplication

Introducing the StoreOnce library

White Paper

First published December, 2011. Latest version, V2, December 2011

### Table of contents

Summary .....	3
Overview.....	3
When to use a B2D device .....	4
About StoreOnce software deduplication .....	4
Advantages of B2D devices and deduplication .....	4
Deduplication performance .....	5
How target-side deduplication integrates with Data Protector.....	5
Backup to Disk device concepts .....	6
B2D device operation .....	8
Device locking .....	9
Object consolidation .....	9
Backing up data using B2D devices .....	9
Gateways .....	11
StoreOnce library (deduplication store) .....	11
Deleting expired backup data from the deduplication store .....	11
Clearing redundant data from the deduplication store .....	12
Store robustness .....	12
Deduplication statistics.....	12
Deduplication ratio.....	13
Limitations.....	14
Installation .....	14
Prerequisites .....	14
Installation procedure .....	14
Installation package contents .....	17
Licensing.....	18
Configuration .....	18
Example configuration using a B2D device .....	18
Adding a B2D device .....	21
Creating a store.....	23
Configuring the root directory of the deduplication stores.....	23
Backup.....	25
Restore .....	25
Troubleshooting .....	25
Before you begin .....	25
Low disk space warning.....	25



Backup of the system.db file .....	26
Problems .....	26
Tutorial: Adding a Backup to Disk device and creating a backup specification .....	28
Appendix A: StoreOnceSoftware utility .....	33
Appendix B: Command line interface changes to support B2D devices .....	38
Addition to omnimm .....	38
Addition to omnidownload .....	38
omniupload .....	39
Appendix C: Omnirc variables related to B2D devices .....	40
Appendix D: System requirements and performance .....	41
Supported platforms .....	41
Deduplication server requirements .....	41
Glossary .....	43
Index .....	45
For more information .....	46

## Summary

This document describes how HP Data Protector integrates with the HP StoreOnce software deduplication engine. By supporting software deduplication, several new concepts are introduced to Data Protector including the StoreOnce library and a new device type, the Backup to Disk device. Both are discussed in detail in this document.

**HP Data Protector software deduplication** provides the ability to deploy target-side deduplication on virtually any industry-standard hardware, offers greater flexibility than existing solutions as it can be deployed in a wider range of hardware set-ups, and provides enterprise-class scalability.

Because of the way Data Protector makes use of the extremely efficient HP StoreOnce engine, Data Protector software deduplication uses memory very efficiently. As a result, you can deploy deduplication on application or backup servers without lowering application performance. Data Protector software deduplication can even be deployed on a virtual machine. In addition, Data Protector software deduplication delivers very high throughput.

**Backup to Disk devices** use the StoreOnce software deduplication interface to back up data to the StoreOnce library (the deduplication store). This document also describes how the Backup to Disk device integrates with Data Protector and the basic principles behind deduplication technology. A short tutorial is included to guide you through the process of adding a Backup to Disk device, creating a deduplication store, and running a backup session.

Backup to Disk devices are supported by Data Protector 6.20 with the installed patch bundle set 6.21 or superseding updates. For full details of supported systems, see the latest HP Data Protector support matrices at <http://www.hp.com/support/manuals>. For general Data Protector procedures, see the online Help.

## Overview

The Backup to Disk device (abbreviated to B2D throughout this document), together with the StoreOnce software deduplication interface, uses deduplication technology to backup data to disk. Data deduplication is a data compression technology which reduces the size of the backed up data by not backing up duplicate data.

The deduplication process splits the data stream into manageable chunks (or blocks) of data. The contents of these data chunks are then compared to each other. If identical chunks are found, they are replaced by a pointer to a unique chunk. In other words, if 20 identical chunks are found, only one unique chunk is retained (and backed up) and the other 19 are replaced by pointers. The backed up data is written to a disk-based destination device called a *deduplication store* (with StoreOnce software deduplication, this is the StoreOnce library). When a restore operation is done, the unique chunk is duplicated and inserted in the correct position as identified by the pointer. With deduplication-type backup and restore operations, the restore process is sometimes referred to as *rehydration* of the backed up data.

## When to use a B2D device

Typically, you would use a B2D device (and by default, data deduplication) when backing up an e-mail filesystem which may contain 100 instances of the same 1 MB graphic file attachment. If the system is backed up using a conventional backup technique, all 100 instances of the attachment are backed up. This requires approximately 100 MB of storage space. However, if the backup is done through a B2D device, only one instance of the attachment is actually stored. All other instances are referenced to the unique stored copy. In this example, the deduplication ratio is approximately 100 to 1. Although this example is referred to as *file-level deduplication*, it serves to demonstrate the benefits of B2D devices and StoreOnce software deduplication.

Other points to consider when deciding to use deduplication technology:

- Some data is not a good deduplication candidate! Data that is automatically created by a computer does not deduplicate well, for example, database files. Photos, video, audio, imaging, seismic data are all examples of data that do not deduplicate very well.
- Do not compress data before sending it to the deduplication server. It will impact on the deduplication ratio (see also [Deduplication ratio](#)) and is unnecessary as compression is done following deduplication.
- Do not encrypt data before sending it to the deduplication server. This produces a deduplication ratio 1:1, basically, no deduplication.

## About StoreOnce software deduplication

There are several deduplication technologies available in the marketplace. They are generally grouped into hardware-based and software-based solutions. These solutions can be further sub-grouped, for example, into file-level (single-instancing) or block-level deduplication. Data Protector's StoreOnce software duplication offers a software-based, block-level deduplication solution.

When using StoreOnce software deduplication, note the following:

- Deduplication backs up to disk-based devices only. It cannot be used with removable media such as tape drives or libraries.
- Because Data Protector uses a software-only approach to deduplication (that is, when using StoreOnce software deduplication), no specific hardware is required other than standard hard disks to store the backed up data.
- In the deduplication process, duplicate data is deleted, leaving only one copy of the data to be stored, along with reference links to the unique copy. Deduplication is able to reduce the required storage capacity since only the unique data is stored.
- StoreOnce software deduplication uses hash-based chunking technology to split the data stream into sizeable chunks of data.
- Specifying a *Backup To Disk* target in the backup specification tells Data Protector to do a deduplication-type backup.

## Advantages of B2D devices and deduplication

Generally, data deduplication increases the speed of the backup service as a whole and reduces overall storage costs. Data deduplication significantly reduces the amount of required disk storage space. Because data deduplication is a disk-based system, restore service levels are significantly higher and tape (or other media) handling errors are reduced. Additional benefits of deduplication include:

- Data deduplication is more appropriate with large volumes of data.
- Data Protector uses well-proven deduplication algorithms to guarantee data integrity (StoreOnce software deduplication uses deduplication technology developed by HP Labs for HP StoreOnce Backup Systems. These systems use hardware-based deduplication. For more information about HP StoreOnce products, see [For more information](#)).
- Disk-to-disk (D2D) storage with deduplication is rapidly becoming the preferred method for backup and recovery in both local and remote applications.
- The total cost of recovery for duplication-enabled D2D systems is significantly lower than with tape-based systems. Data deduplication backups can provide considerable capacity and cost savings compared to conventional disk backup technologies.

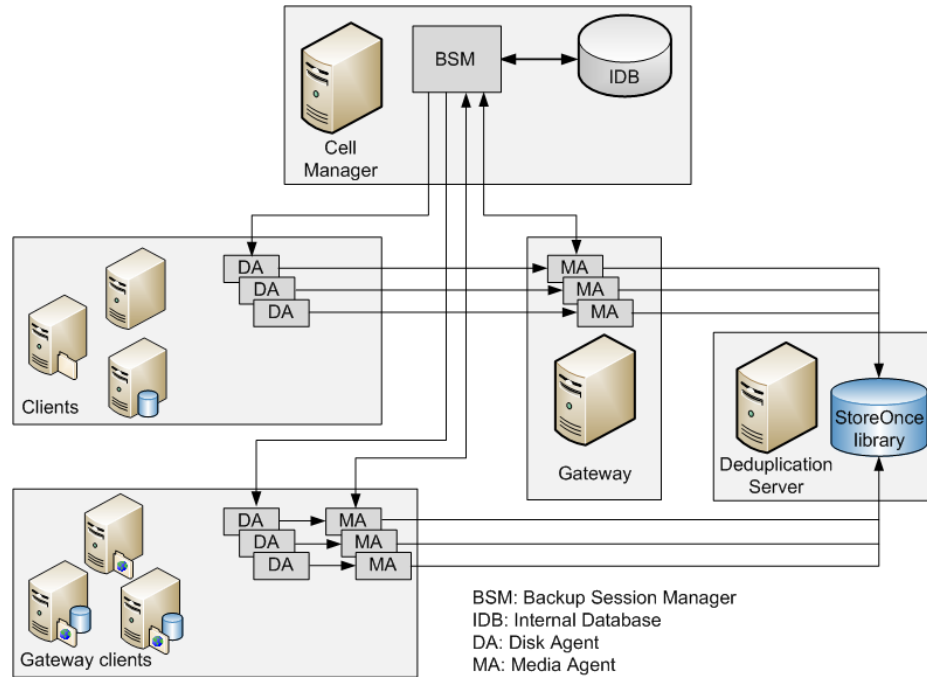
## Deduplication performance

There are many factors that can affect deduplication performance. These include hardware and network speed, how the storage disk is setup, the size of the store, the deduplication ratio of the data, and how many concurrent backups are running. Using multiple streams can significantly improve backup performance. Data Protector supports up to 30 parallel streams of writing and reading data to each store. For performance-related details, see [Appendix D: System requirements and performance](#).

## How target-side deduplication integrates with Data Protector

Data deduplication integrates with Data Protector as shown in Figure 1. The deduplication process takes place on the deduplication server. It receives data to be backed from Media Agents on the gateway (serving the clients group) and from Media Agents installed on clients in the *gateway clients* group. The deduplication server then writes the deduplicated data to the StoreOnce library (this is the physical store and is sometimes referred to as the deduplication store). This type of deduplication is referred to as *target-side deduplication* because the deduplication process takes place near to where the data is stored. (The role of *gateway clients* is discussed in [Gateways](#).)

Figure 1: How target-side deduplication integrates with Data Protector

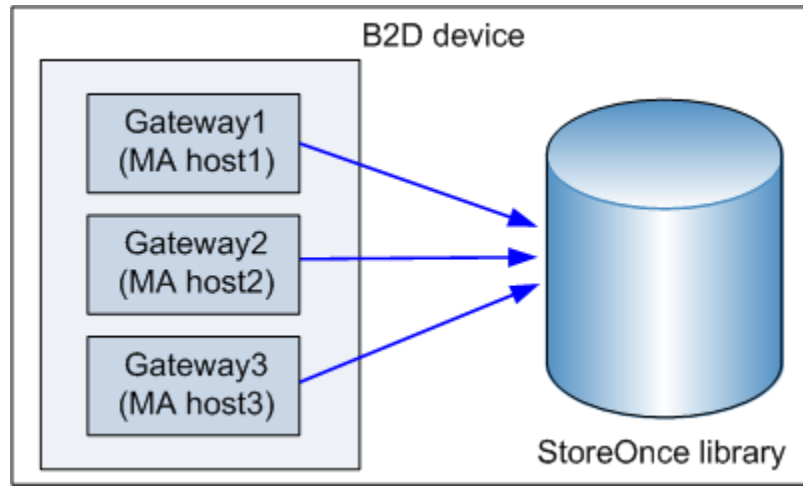


The deduplication server allows connections from several Media Agents, locally or remotely. It also provides synchronization mechanisms to enable multiple Media Agents to work with the StoreOnce library at the same time. The Media Agent reads or writes data in terms of object versions to or from the StoreOnce library. Each object version is represented as an item in the StoreOnce library. To optimize deduplication performance, Disk Agent concurrency is not supported (this means, one Disk Agent *talks* to one Media Agent – there is no multiplexing of streams). B2D device configuration is stored in the IDB. An example configuration showing a basic local and remote office deployment is given in [Example configuration using a B2D device](#).

## Backup to Disk device concepts

A Backup to Disk (B2D) device backs up data to the StoreOnce library (a physical storage disk). The B2D device supports multi-host configurations. This means that a single physical storage disk can be accessed through multiple hosts called *gateways*. Each gateway represents a Data Protector client with the Media Agent component installed. A B2D device is a logical device and consists of gateways and a store. Figure 2 shows the relationship between a generic B2D device with multiple gateways and the StoreOnce library.


Figure 2: B2D device (logical view)



The StoreOnce library can also be divided into individual stores representing specific storage sections, similar to partitioning a hard disk. A store is represented by a network path and is used by the backup application. These parameters, along with any other device-configuration information, are stored in the device configuration in the IDB.

Each individual store on the physical storage disk can be accessed by one B2D device only. However, several B2D devices can access different stores on the same physical storage. Figure 3 shows three individual B2D devices accessing three individual stores on the same physical disk.

Figure 3: Multiple B2D devices accessing multiple stores on the same physical storage (physical view)



When configuring a B2D device, note the following:

- It is possible to configure multiple stores on a single deduplication-server node. These stores share resources, such as, CPU, memory, disk I/O, and the number of connections per


deduplication system (up to a maximum of 30). However, each store represents its own deduplication domain. Deduplication does not happen across different stores.

- Each store must have its own dedicated B2D device configured. It is not possible to have two stores configured to the same B2D device.
- Each B2D device must use a store exclusively. Accessing the same store by more than one B2D device is not supported.

## B2D device operation

The B2D device type is similar in structure to other library-based devices, that is, the device is comparable to a library and the gateway is comparable to a drive in the library. However, the operation is significantly different. Whereas library drives could be considered as *unique* because they could only be accessed by one Media Agent at a certain point in time, gateways behave differently and allow more flexibility. Each gateway represents a host on which multiple Media Agents can be started simultaneously, either in single or multiple sessions.

Figure 4: Gateways and Media Agents (physical view)



The number of Media Agents that can be started on a specific gateway is defined by:

- Gateway limits. Each B2D gateway is limited to a maximum number of parallel streams. This limit is specified in the GUI.
- Connection limits to the store. Each B2D device is limited to a maximum number of connections per store. This limit is specified in the GUI. If the value is left unchecked, Data Protector uses the maximum available. A StoreOnce library supports a maximum of 30 connections per store.



- The physical connection limits of the physical storage disk. This value is retrieved from the physical store (see below).
- Depending on the current operation, each Session Manager attempts to balance the number of Media Agents on a gateway with regards to the following input parameters:
  - The number of objects being backed up
  - Object location
  - Physical connection limits.

The physical connection limit (the maximum number physically possible) is verified during the session. The value entered in the GUI is checked against the number of available connections. If the value entered exceeds the physical limit, the physical limit is used. The physical connection limit cannot be configured in the GUI. (Note: To use the maximum number, uncheck the option.) If a large physical store has been partitioned into smaller stores (Store1, Store2, Store3 as previous), each of these stores has a limit to the number of connections.

## Device locking

The purpose of locking is to ensure that only one system at a time communicates with a device that is shared between several systems. With B2D devices, certain connection limits must be obeyed. These connection limits are the maximum number of parallel streams per gateway and the maximum number of connections per store. Data Protector keeps a lock count for both these resources. When the limit is reached, the lock is denied. If the locking request is successful, the lock counts for both the gateway and the store are increased. When the gateway is unlocked, the lock counts are decreased. This ensures that B2D connection limits are considered *cell-server-wide* and not just during a specific session.

## Object consolidation

To accommodate gateway and gateway/store/device connection limits, object copy and consolidation functionality makes sure that:

- When B2D devices are used as sources, at least one connection is available for object copy and at least  $n$  connections for object consolidation – where  $n$  is the number of source media used for consolidation (see the next paragraph for details).
- When B2D devices are used as targets, at least  $m$  connections must be available – where  $m$  is the minimum device setting in the copy/consolidation specification. If other types of devices are used in parallel, the CSM (Copy and Consolidation Session Manager) tries to balance them such that the minimum setting is reached, otherwise it terminates the session.

When consolidating backed up data (full and incremental backups), make sure there are enough available connections to the store (a maximum of 30 connections per store is allowed). This is easier to explain by considering an example consolidation session of six incrementals. In this case, the number of connections = 1 (Full) + 6 (Incr) + 1 (target) = 8 connections. It is recommended to run a weekly consolidation session for between 6 to 10 incrementals.


## Backing up data using B2D devices

Making a backup to a B2D device is similar to backing up to tape-based devices. The notable differences are that the Session Manager dynamically spawns Media Agents on defined gateways and the Media Agent communicates with the devices through a device-specific API.

The following is an example of a backup session utilizing one B2D device with two gateways (Gateway1 and Gateway2). Five objects are being backed up (DA 1 ... DA 5), two objects are local to first gateway and three are remote to both gateways. The number of connections to the physical storage is six. The backup specification is configured with:

- Load balancing (Max) parameter set to 5 (meaning that up to five Media Agents can be used during this session)
- The connection limit for the B2D device is set to 10.
- The connection limit for both gateways is 5.

Figure 5: Sample backup configuration using two gateways (local and remote objects)



The Session Manager dynamically starts five Media Agents based on the above configuration. Since there are two gateways involved, the five Media Agents are distributed between the two gateways. The load balancing algorithm distributes the Disk Agents between the Media Agents such that Disk Agents local to Gateway1 are assigned to Media Agents on that gateway. Other Disk Agents are load-balanced between the two gateways since they are remote to all Media Agents.

When creating a backup specification, a B2D device can be selected as a target. It is also possible to select specific gateways. If a B2D device is selected as a target device, then all gateways will be chosen during the backup process using up-to-date device configuration information. However, this only works for load-balanced backups. When configuring a static backup (non-load balanced), you can only assign each object to a gateway and not to a B2D device.

B2D devices use a special data format for fast read/write access and to improve the deduplication ratio. The format splits the metadata from the actual data to be backed up. The data format is automatically set when you select a B2D device and is only used for B2D devices.

## Gateways

Backup to Disk (B2D) devices must be configured to access pre-defined gateways. A gateway, or rather a *gateway client*, is a client with the Media Agent component installed (the client must be a 64-bit system, see below). It can be backed up as with any other client in the cell. Gateways are identified by a unique name. A gateway name includes a network pathname which points the device to the physical store. The default naming convention is similar to file libraries: *device\_name\_gwnumber*. Gateways are selected from a drop-down list in the GUI. Clients which cannot act as gateways are not listed.

Gateway connectivity can be verified (checked) to make sure the device is able to communicate with the gateway. If the gateway is unavailable for any reason, an error state is displayed.

Additionally, gateway properties and features include:

- Gateway properties such as the gateway name and advanced options can be modified in the Gateway Properties dialog. The properties of multiple gateways can be modified simultaneously (see [Viewing gateway properties](#)).
- Gateways can be enabled or disabled. To enable or disable a gateway, right-click the gateway you want to enable or disable in the Device list and select Enable Gateways or Disable Gateways as appropriate.
- The Media Agent component can only be installed on 64-bit client systems. This means that if a client is to be designated as a gateway, it must be a 64-bit system (see [Supported platforms](#)).
- Each gateway represents a host on which multiple Media Agents can be started at the same time, either in a single session or in multiple sessions. For this reason, a gateway is sometimes referred to as a Media Agent host.
- If you are familiar with library-based devices, the B2D device is comparable to a library and the gateway is comparable to a drive in the library.

## StoreOnce library (deduplication store)

The StoreOnce library (or deduplication store), is the physical storage disk used by the StoreOnce software deduplication interface. (StoreOnce software deduplication uses HP StoreOnce Backup System technology developed by HP Labs.) The supported capacity of the physical disk is 10 TB (of deduplicated data). Typically, with a deduplication ratio of say 20:1, this is equivalent to backing up 200 TB of data. If multiple stores are used, the total supported capacity is still 10 TB. To view the properties of a store, use the CLI commands (see [Appendix A: StoreOnceSoftware utility](#)).

One deduplication server can host multiple deduplication stores providing the stores share the same root directory. Although Data Protector supports up to 32 stores per volume, best performance (with respect to the deduplication ratio) is achieved with one store only. Configuring the deduplication store is completed in a single stage (see [Configuring the root directory of the deduplication stores](#)).

## Deleting expired backup data from the deduplication store

Periodically, Data Protector automatically triggers a cleanup session to delete backed up data residing in the physical store. Several methods are employed for removing unprotected data (for details, see [Addition to omnimm](#)):

- Manual removal of unprotected B2D backup objects  
Data Protector builds a list of unprotected backed-up objects residing in the store. Data Protector first removes them from the store and then deletes the media (object's) information from the Data Protector database. Note that removing media from the store does not free-up disk space, it merely instructs the store to obsolete the data.
- Automatic removal of unprotected B2D backup objects  
This method does the same as the above but is performed automatically by Data Protector at regular intervals. The interval can be configured in the global options file.
- Immediate removal when a slot is deleted  
Deleting a slot deletes the slot from the IDB, deletes the object in the slot, and the slot itself is removed from the store. This is the same action as recycling and deleting.

Removing unprotected B2D backup objects deletes the associated slots immediately. Deleting items does not free-up disk space immediately. During the next housekeeping job, expired files and unreferenced chunks are deleted, and possibly freeing up some disk space.

---

#### Note

Redundant data is data that is no longer referenced in the store. With expired data, the protection date has expired.

---

## Clearing redundant data from the deduplication store

Data Protector provides a space-management (housekeeping) utility to optimize storage space. The housekeeping utility is started by default and runs in the background.

A data chunk becomes redundant when it is no longer referenced by the indexing table. Data is not deleted automatically from the store. This only occurs when the housekeeping utility runs and frees-up disk space.

### Store robustness

StoreOnce software deduplication has a built-in mechanism to verify the integrity of the store. To minimize or prevent data loss, take note of the following:

- Use an uninterruptible power supply (UPS). It enhances the fault tolerance of the deduplication server. A UPS allows your computer to keep running for a short time when the primary power source is lost. It also provides protection from power surges.
- The store must be configured as a RAID array. Due to the directory structure of deduplication stores, if one disk corrupts, the whole store becomes unusable. Hardware RAID is preferable.
- For critical data, it is recommended to do object copy operations from the deduplication store to a tape. Do not write to the store while backing it up.

### Deduplication statistics

For backup sessions using StoreOnce software deduplication, Data Protector displays the backup statistics after each object version is complete, for example:

```
Deduplication Statistics for Name:
```

```
Mbytes Total:          52 MB
Mbytes Written to Disk:  3 MB
Deduplication Ratio:   17.3 : 1
```

The statistics includes:

- `Mbytes Total`: The original size of the object version (the data to be backed up).
- `Mbytes Written to Disk`: The actual size written to the disk after deduplication. (If less than 1 MB, 1 MB is displayed).
- `Deduplication Ratio`: The 'Mbytes Total' divided by the 'Mbytes Written to Disk'. (See notes below.)

When interpreting the deduplication ratio, note the following:

- If the value for 'Mbytes Written to Disk' is less than 1 MB, it is rounded to 1 MB (otherwise, the calculation produces an unrealistic result).
- Typically, you can expect a deduplication ratio of the order 10 - 20 : 1. Ignore erroneous ratios (for example, 4435 : 1). This can occur when the denominator (Mbytes Written to Disk) is extremely small.

The ratio displayed in the backup statistics applies to the current session. The ratio displayed in the CLI, applies to the store as a whole.

## Deduplication ratio

The storage capacity saved by using deduplication is typically expressed as a ratio. The sum of all pre-deduplicated backup data is compared with the actual amount of storage the deduplicated data requires. For example, a ratio of 10:1 means that 10 times more data is stored than would have been if deduplication had not been used.

The most significant factors affecting the deduplication ratio are:

- How long the data is retained.
- How much the data changes between backups.
- Small file sizes may result in a low deduplication ratio.

However, many factors influence how much storage is saved in your specific environment. The ratio is reported in the summary screen (after adding a device), in the Devices context (**Devices > Stores**), and in the backup statistics following a backup operation (for a typical output, see [Deduplication statistics](#)).

## Limitations

- Data deduplication is not suitable for archiving of data.
- The Data Protector Java GUI is not supported.
- It is not supported for more than one B2D device to access the same store. This means that each B2D device must be configured to a dedicated store. Do not configure a second device to use the same store.
- If the number of connections required for consolidating backed up data (full and incremental backups) exceeds the maximum number of connections, the restore chain which could not be consolidated, terminates. A maximum of 30 connections per store is allowed. See also [Object consolidation](#).

## Installation

This section provides an overview of the main patch installation tasks and specific requirements when installing the StoreOnce software deduplication package.

### Prerequisites

Make sure that the HP Data Protector 6.20 Cell Manager, user interface client, and Installation Server are installed on supported systems. Make sure the appropriate patch bundle is installed (for example, Windows, Linux).

For an overview of supported platforms, see [Supported platforms](#). For details, see the latest HP Data Protector support matrices at <http://www.hp.com/support/manuals>. See the *HP Data Protector Installation and Licensing Guide* on how to install Data Protector in various architectures.

### Firewall configuration

Ensure that the following ports are open for incoming connections on the StoreOnce library:

- 5555/tcp – on all Data Protector clients (this is a prerequisite for the general Data Protector installation procedure)
- 9387/tcp – command port (see also [Appendix C: Omnirc variable](#))
- 9388/tcp – data port (see also [Appendix C: Omnirc variable](#)).

Ports 9387 and 9388 must be open in a firewall separating the deduplication server from any gateways. (Windows systems: Ports are opened during the installation process, Linux systems: Ports must be manually opened.) For details on Data Protector ports, see the online Help index: "port range".

## Installation procedure

The installation procedure consists of two main parts:

1. Install Data Protector patch bundle as described in the patch readme file.
2. Install Data Protector components on appropriate clients and servers. For details, see [Installing Data Protector StoreOnce software deduplication components](#).

The StoreOnce software deduplication package must be installed locally and remotely.

## Installing Data Protector StoreOnce software deduplication components

1. Connect to any client with the updated Data Protector User Interface component.
2. Open the Data Protector GUI and, in the Context List, select **Clients**.
3. Add the Data Protector StoreOnce software deduplication component to a backup host:
  - If the backup host is not part of the Data Protector cell, use the Data Protector **Add Clients** functionality.
  - If the backup host is already part of the Data Protector cell, use the Data Protector **Add Components** functionality.

The StoreOnce software deduplication component can be installed from the Installation Server.

Note the following:

- Windows systems: The StoreOnce Software deduplication component can also be selected from the Components list to install on clients for both a remote installation and a local installation.
- Linux systems: For local installations, use `omnisetup.sh -install StoreOnceSoftware` for the StoreOnce Software deduplication component.

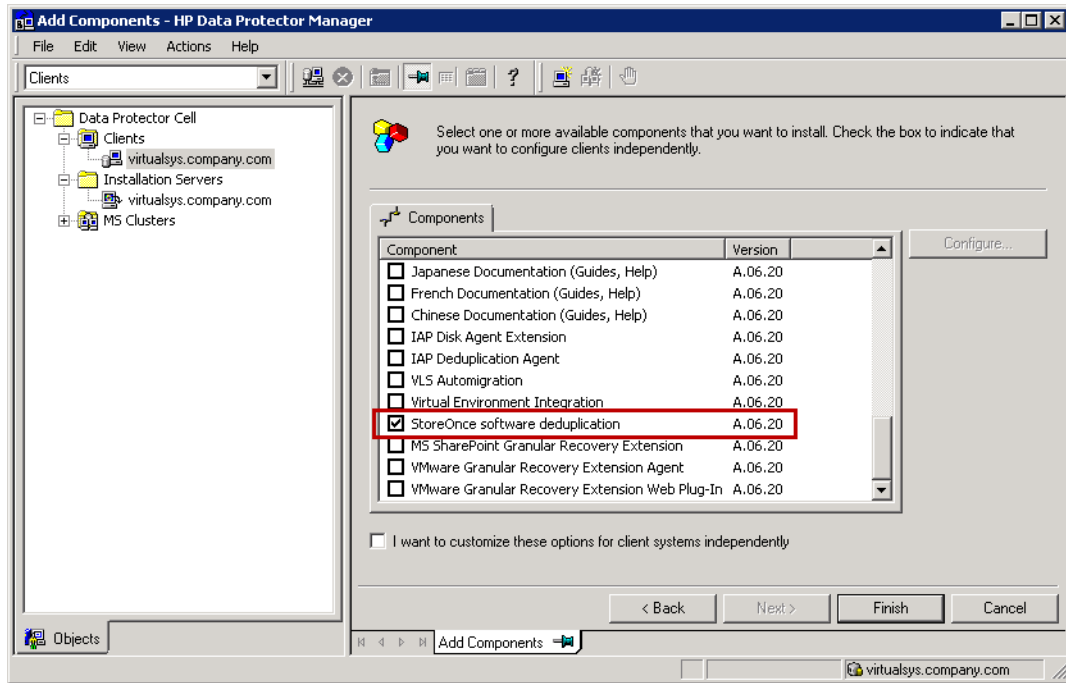
Linux systems only:

To install the StoreOnceSoftware daemon such that it starts automatically after a reboot, copy the file `StoreOnceSoftwared` to the `/etc/init.d` directory and include it in startup scripts. The daemon can also be started or stopped manually using the commands:

```
/opt/omni/sbin/StoreOnceSoftwared start
and
/opt/omni/sbin/StoreOnceSoftwared stop
```

Removing the StoreOnce software deduplication component from the system automatically stops the process and removes the file `StoreOnceSoftwared` from the `/etc/init.d/` directory.

Figure 6: Selecting the StoreOnce software deduplication component

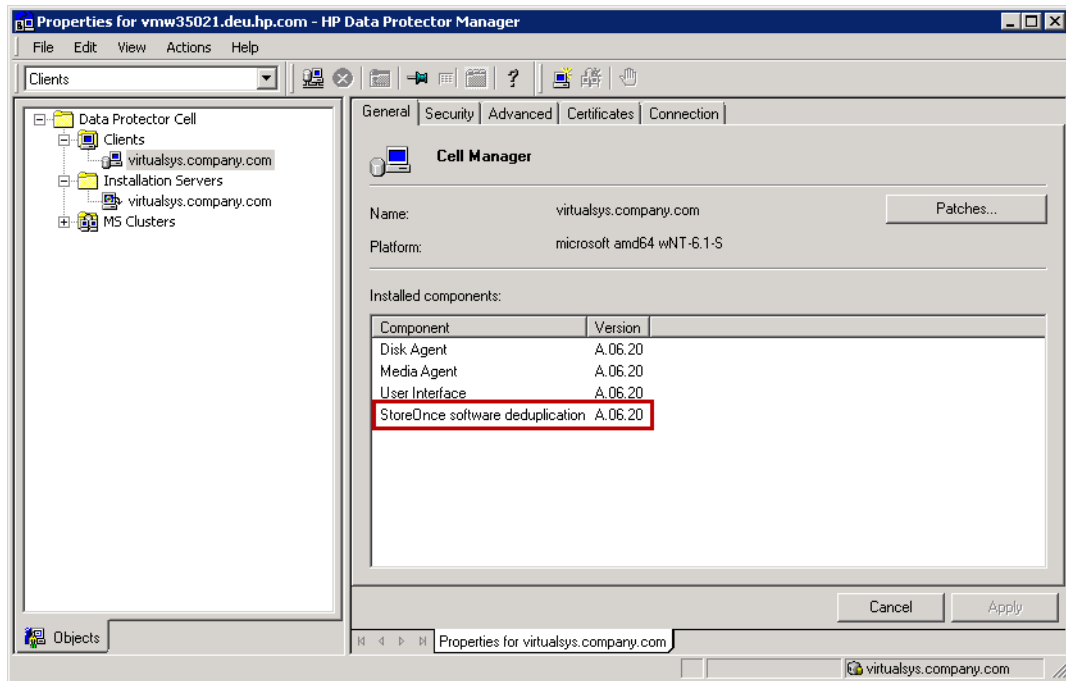


Following successful installation, the StoreOnce software deduplication component is listed in the Installed components list.

Before you can use StoreOnce software deduplication, the root directory of the stores must be configured. To find out how to do this, see [Configuring the root directory of the deduplication stores](#).



Figure 7: Viewing the installed components



## Installation package contents

The package names are:

- Windows systems: StoreOnceSoftware.msi
- Linux systems: StoreOnceSoftware

### Windows systems

Following successful installation, the StoreOnceSoftware executable is started as a service (see Services tab in the Task Manager). The service name is Data Protector StoreOnceSoftware, the description is StoreOnce Software Deduplication, and the startup type is automatic.

The installation package includes the following files:

Filename	File location
StoreOnceSoftware.exe	<Data_Protector_home>\bin
system.db	<Data_Protector_program_data>\Config\client\StoreOnceSoftware

## Linux systems

Following successful installation, StoreOnceSoftware is started as a background process (daemon). Following a reboot, it can be started automatically.

The installation package includes the following files:

Filename	File location
StoreOnceSoftware	/opt/omni/lbin
StoreOnceSoftwared	/etc/init.d/ /opt/omni/lbin
system.db	/etc/opt/omni/client/StoreOnceSoftware

## Licensing

B2D device licensing uses the capacity-based licensing (similar to File Library licensing). During the session, the Media Agent gathers statistical information from the physical store (that is, the deduplicated data in the deduplication store) and sends it to the Session Manager, which in turn updates the licensing database. When a B2D device is deleted or modified in the IDB, the licensing database is updated. The B2D device requires a capacity-based *Advanced Backup to Disk* LTU which is available as follows (the used capacity is based on the usage of deduplicated data on disk):

- B7038AA – 1 TB
- B7038BA – 10 TB
- B7038CA – 100 TB

For more information about capacity-based licensing, see the *HP Data Protector Installation and Licensing Guide*.

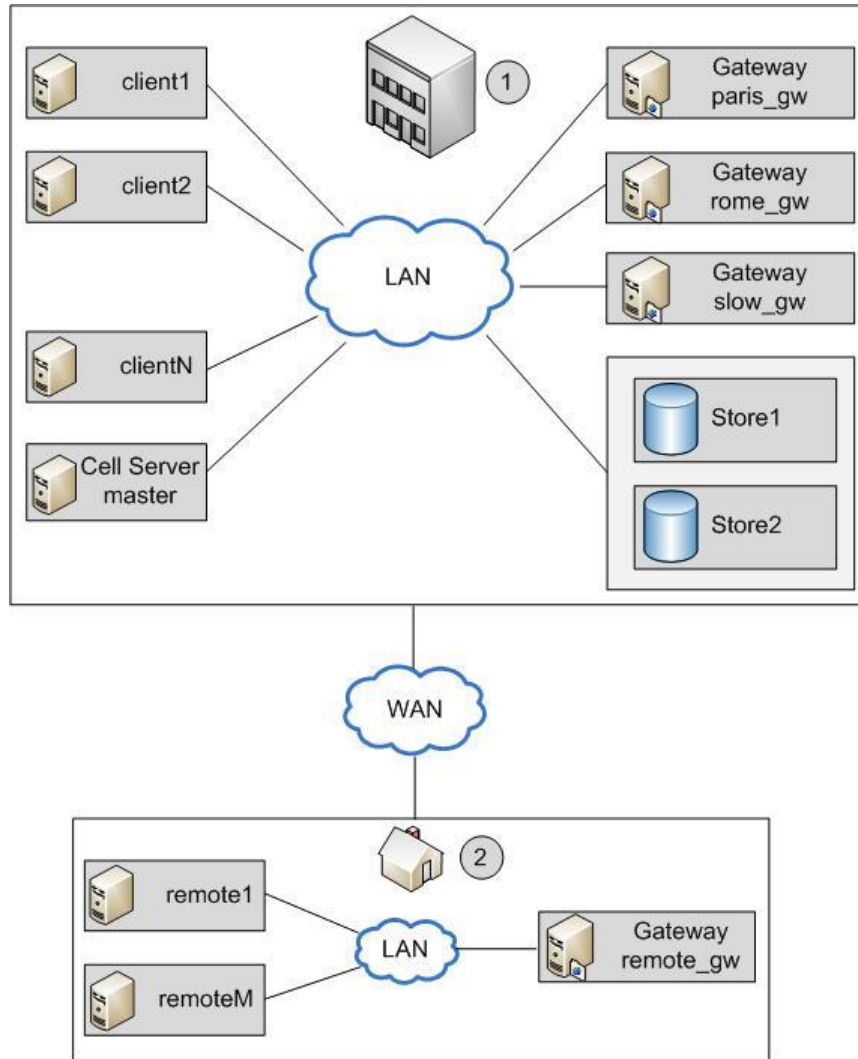
## Configuration

This section includes an example environment and configuration procedures.

### Example configuration using a B2D device

A typical use-scenario with a B2D device is to have the Media Agent running on the client, that is, the gateway for a client is the client itself. In effect, this means that almost all clients become gateways. The following provides a typical usage model of a central office/remote office configuration.

Figure 8: Deployment example within a central office/ remote office environment



Item	Description
1	Central office. This LAN is located in the central office. It is connected to the LAN in the remote office over a WAN.
2	Remote office. This LAN is located in the remote office.

The Data Protector Cell Server is installed in the central office on the host *master*. There are several clients in the central office: *client1* to *clientN* (non-gateway clients), *paris\_gw*, *rome\_gw*, and

*slow\_gw* (gateway clients). Additionally, two object stores (*Store1* and *Store2*) are configured in the central office.

The remote office includes clients *remote1* to *remoteM* and *remote\_gw*. All clients in the remote office are part of the same Data Protector cell as the clients in the central office. The remote office is connected to central office over a slow WAN network.

---

**Note**

Gateways are simply clients with the Media Agent component installed. Think of them as *gateway clients*. For a client to become a gateway, it must be a 64-bit system (see [Supported platformsPrerequisites](#)).

---

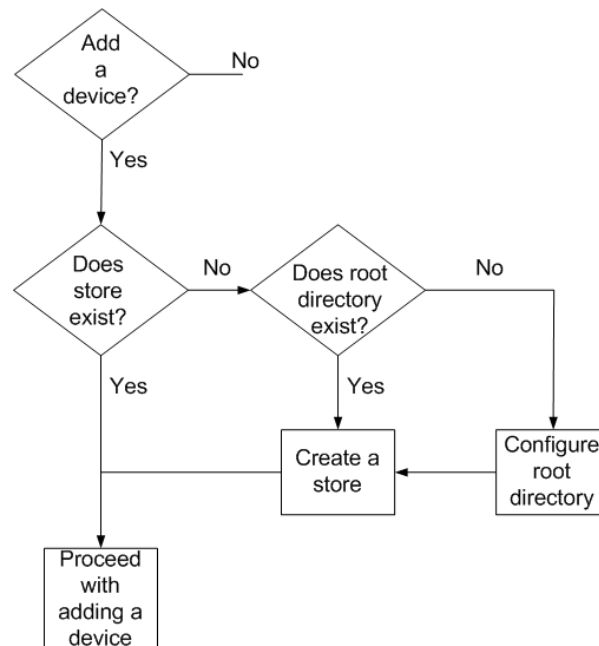
When you configure a B2D device, you must specify certain parameters such as the name and location of the store, gateways, and network paths. In the above example, you want to use the store *Store1* (which is accessed by StoreOnce software deduplication) for backup of clients in your environment. To do this, you configure the B2D device to use *Store1* as the repository. You also decide that clients *paris\_gw*, *rome\_gw*, and *slow\_gw* are to be used as gateways for other DP clients in the central office. Additionally, note the following:

- Concurrency specifies a number of Disk Agents writing to the device in parallel. Multiple Disk Agents read data in parallel (from disks) to provide a constant data stream to the Media Agent. With StoreOnce software deduplication, Disk Agent concurrency for each Media Agent is set to 1 (this improves the deduplication ratio).
- Only one store can be configured per device.
- Stores are represented by network paths (UNC) that contain information about the deduplication system and the store name. (Note: In the context of B2D devices, the deduplication system refers to name of the *hosting* machine where the deduplication store is located.)

## Adding a B2D device

The procedure for adding a B2D device is the same as for other device types. However, you must first configure a root directory and then create a store (see [Configuring the root directory of the deduplication stores](#) and [Creating a store](#)). Figure 9: Overview to adding a deviceFigure 9 provides an overview to adding a device.

Figure 9: Overview to adding a device



To add a B2D device (which targets an existing store), proceed as follows:

1. In the **Devices & Media** context, right-click **Devices > Add Device**.
2. Specify a device name, add a description, select the device type Backup to Disk, and select the interface StoreOnce software deduplication.
3. Click **Next**.
4. Select the Deduplication System (the hosting machine where the deduplication store is located) and specify the actual store name (this assumes that the store already exists).
5. (Optional) Check Max. Number of Connections per Store and specify a value (1 to 30).
6. Select a gateway and click **Add** to display the properties dialog.
7. If necessary, change any gateway properties and then click **OK** to add the gateway.
8. Click **Check** to verify the connection. (Clicking **Next** also verifies the connection.)
9. Click **Next** to display the Summary screen.
10. Click **Finish**. The B2D device is shown in the Scoping Pane.

### Setting additional options

The following options and settings are available when adding a B2D device. In the Specify the store and list of gateways dialog:

- **Max. Number of Connections per Store:** Specifies the maximum number of connections that can be opened to the deduplication store. The value specified is checked against the number of available connections (that is, the value set in the configuration file). If the value specified exceeds the physical limit, the physical limit is used (you can specify up to a maximum of 30 connections).
- **Check:** Verifies the connectivity to the specified gateways system. First highlight the gateway to be checked, and then click **Check**. If Data Protector is able to connect (communicate) with the store through the gateway, the status OK is displayed in the Status field. If the gateway-store connection is unavailable for any reason, or the store does not exist, an error state is displayed. A dash (-) indicates that the gateway-store connection has not yet been checked. You must run the check operation on all added gateways. Clicking **Next** has the same action as clicking **Check**.
- **Properties:** Displays various gateway properties, see [Viewing gateway properties](#).

### Viewing gateway properties

To view gateway properties, highlight a gateway and click **Properties** in the Specify a list of gateways dialog or select the gateway in the Scoping Pane (**Devices & Media > Devices > device\_name > Gateways**). Gateway properties are grouped under the following tabs:

- **General:** Displays information about the gateway including name and description.
- **Settings:** In addition to media type and media pool, a checkbox is provided to disable the gateway (Disable Gateway). The Advanced options dialog includes the option Max. Number of Parallel Streams per Gateway. This value specifies the maximum number of Media Agents that can be started per gateway. Each gateway is limited to a maximum number of parallel streams (you can specify up to a maximum of 30 parallel streams). This means that multiple objects can be backed up in parallel. The default is the maximum available, that is, up to 30 (in this case, a dash (-) appears in the gateway list). Check the checkbox and specify a value (1 to 30).
- **Policies:** Specify gateway policies for the gateway as follows:
  - **Gateway may be used for restore:** A gateway with this option selected may replace any gateway with the same gateway tag. If the original gateway is not available for a restore session, Data Protector automatically selects an alternative gateway of the same type. Default: OFF
  - **Gateway may be used as source gateway for object copy:** A gateway with this option selected may replace any gateway with the same gateway tag. If the original gateway is not available for an object copy session, Data Protector automatically selects an alternative gateway of the same type. Default: OFF
  - **Gateway Tag:** Specify a name for the gateway tag. Gateways with the same gateway tag name can replace each other if needed. The name can be a maximum of 80 characters including spaces.

---

## Note

To modify the properties of multiple gateways simultaneously, first select the gateways, click **Properties**, make any changes, and click **OK**. When multiple gateways are selected, the individual gateway name is deselected.

---

For a step-by-step procedure to add a B2D device, see [Tutorial: Adding a Backup to Disk device and creating a backup specification](#).

## Creating a store

Before creating a store, make sure the root directory of the stores has been configured and the physical storage disks (LUN devices) are formatted and mounted on the deduplication server. The LUN devices may be on local disks or a disk array (SCSI or Fiber Channel interface) or on a NAS device in the same LAN (iSCSI interface). When using iSCSI interface, the reliable network connection must provide a latency of at most 2 ms and a throughput of at least 1 Gbit/s.

A store can be created from:

- The GUI: Follow the procedure for adding a device and when prompted, specify the name of the store (see below for details).
- The CLI: Use the command `StoreOnceSoftware --create_store` (see [Appendix A: StoreOnceSoftware utility](#) for details).

The procedure for creating a store is similar to adding a device but includes a few additional steps. To create a store, proceed as follows:

1. Follow the procedure for adding a device:
  - a. In the **Devices & Media** context, right-click **Devices > Add Device**.
  - b. Specify a device name, add a description, select the device type Backup To Disk, and select the interface StoreOnce software deduplication.
  - c. Click **Next** to display the screen where you specify a store and a list of gateways.
2. Select the Deduplication System and specify a name for the store. The maximum length of the store name is 80 characters (alphanumeric characters and underscore (\_) only).
3. Select a gateway, click **Add** to display the properties dialog, then click **OK** to add the gateway.
4. Click **Check** to verify the connection. If the store does not exist, it is created. (Note: Clicking **Next** also verifies the connection.)
5. Continue with the procedure for adding a device.

If you specify the store name incorrectly, you cannot change it through the GUI. Run through the procedure again and create the store with the correct name. Use the CLI to delete the incorrectly-named store (assuming data has not been written to it).

## Configuring the root directory of the deduplication stores

This section describes how to configure the root directory of the stores. This must be done after installing the software and before creating the first deduplication store.

One deduplication server can host multiple deduplication stores providing the stores share the same root directory. Each store operates independently of the other, that is, deduplication only occurs

within one store and each store has its own index table. Although all stores run under the same process, they can be started/stopped individually (this does not mean to physically start/stop a store, see [Appendix A: StoreOnceSoftware utility](#) for details). Operations cannot be done on a store if it is stopped (offline).

Stores sharing the same root directory cannot be separated physically. This design guarantees uniform loading on all disks and provides better performance.

Following successful installation, the StoreOnceSoftware utility starts in a mode where it is running but waiting for the root directory of the stores to be configured. A B2D device cannot be added and a store cannot be created until the root directory is configured.

The root directory of the stores can be configured from:

- The GUI: Follow the procedure for adding a device and when prompted, specify the root directory (see below for details).
- The CLI: Use the command `StoreOnceSoftware --configure_store_root` (see [Appendix A: StoreOnceSoftware utility](#) for details).

---

### Note

The root directory must already exist (on the server) and you must have write permissions before it can be configured. This is because the (GUI) configuration process asks you to specify its location.

---

The procedure for configuring the root directory using the GUI is similar to creating a store but includes a few additional steps. Once the root directory has been configured, these additional steps are no longer necessary. To configure the root directory (and create a store at the same time), proceed as follows:

1. Follow the procedure for adding a device:
  - a. In the **Devices & Media** context, right-click **Devices > Add Device**.
  - b. Specify a device name, add a description, select the device type Backup To Disk, and select the interface StoreOnce software deduplication.
  - c. Click **Next** to display the screen where you specify a store and a list of gateways.
2. Select the Deduplication System and specify a name for the store.
3. Select a gateway, click **Add** to display the properties dialog, then click **OK** to add the gateway.
4. Click **Check**. The message `Root directory not configured` is displayed.
5. In the dialog, specify the root directory path (for example, `C:\Volumes\StoreOnceRoot`) where all the stores are to reside and click **OK**. (Note: Browsing to the valid root directory is not possible).
6. If the root directory exists, the dialog closes and device configuration continues. The StoreOnceSoftware utility creates a subdirectory (the store) in the specified root directory. If the root directory does not exist, an error message is displayed.
7. Continue with the procedure for adding a device.

Note the following points when configuring the root directory and creating stores:

- Do not use the same disk where the operating system (OS) is installed.
- Use dedicated (exclusive) storage disks.
- Data Protector supports a maximum of 32 stores per volume.



## Backup

Specifying a B2D device in the backup specification tells Data Protector to do a deduplication-type backup. The deduplication process runs in the background and the deduplicated data is written to the StoreOnce library (the deduplication store).

You make a data deduplication-type backup in the same way as with a conventional backup. First, add a new B2D device (in this case, by specifying StoreOnce software deduplication) and then create a backup specification which targets this device. See the online Help index: "creating, backup specification".

## Restore

You restore backed up data in the same way as with a conventional restore operation. Although, the background process of retrieving data from the deduplication store is significantly different when compared to conventional restore processes, there are no special tasks to be performed. The main operations in the retrieving process includes loading data to be restored into memory, reading reference information from index tables, and using this information to *rehydrate* the backed up data. See the online Help index: "restore".

## Troubleshooting

This section provides log and event reporting, warnings, diagnostics, and problem-solving information when using Data Protector B2D device integration. For general Data Protector troubleshooting information, see the *HP Data Protector Troubleshooting Guide*.

### Before you begin

- Ensure that the latest official Data Protector patch bundles are installed. On how to verify this, see the online Help index: "patches".
- For general Data Protector limitations, as well as recognized issues and workarounds, see the *HP Data Protector Product Announcements, Software Notes, and References*.
- For an up-to-date list of supported versions, platforms, and other information, see <http://www.hp.com/support/manuals>.

### Low disk space warning

To avoid running out of disk space where the stores reside, a warning message is written (Event Log on Windows systems or Syslog on Linux systems) when a pre-defined threshold is reached. The default value for the threshold is 10% of the store's capacity. The default value can be modified using the `omnirc` variable (see [Appendix C: Omnirc variables related to B2D devices](#)). The warning message is issued before any further read/write operations are done to the store, once per day, or if the StoreOnceSoftware utility is restarted. A warning is also shown in the backup session message at the beginning and at the end of the session. The low disk-space warning message is:

```
You are running out of disk space on Deduplication Store root directory:  
[path]. The threshold x% is reached. Please free space or add more disks.  
[warning].
```

## Backup of the system.db file

The `system.db` database file contains root directory information and information about the stores. It is located under `..\ProgramData\OmniBack\Config\client\StoreOnceSoftware`. If this file is deleted or lost, the store and the backed up data cannot be accessed. To avoid this situation, every time a change is made to the database, a backup copy of the `system.db` file is made to `..\Store_Root\StoreOnceLibrary\system.db.bak`. The `system.db` file can be restored by copying the backup file to the original location, renaming it, and restarting the StoreOnceSoftware utility.

Make sure that files under the root directory are protected (RAID or backup).

## Problems

The following lists common problems and errors reported by the StoreOnceSoftware utility. Errors generally relate to the operating environment and the directory structure of the deduplication store.

The following problems may occur during startup.

### Problem

**Accessing the system.db file: The system.db file is inaccessible (for example, permission denied, or disk full).**

The StoreOnceSoftware utility fails to find the root directory of the store.

### Action

Change permissions, free disk space, or make the database accessible. The database file (`system.db`) contains an empty or no value for the root directory of the deduplication store. To recover the file, see [Backup of the system.db file](#). To reconfigure the root directory to use another location, see `configure_store_root` in [Appendix A: StoreOnceSoftware utility](#).

### Problem

**Accessing the system.db file: The system.db file in the root directory of the store is missing.**

The StoreOnceSoftware utility fails to start.

### Action

Restore or recreate the `system.db` file. See previous problem.

### Problem

**Starting stores: Store directory is inaccessible.**

During startup of the store, an error is logged. The store cannot be accessed.

### Action

Make the store's directory accessible, check permissions, and verify the root directory exists.

### Problem

#### **Starting stores: Store directory is missing.**

Starting the store is successful, but no items are found.

### Action

Restore the root directory and the stores below the root directory.

### Problem

#### **Starting stores: Store is dirty and cannot be recovered.**

An error is logged. The store cannot be accessed.

### Action

Restore the root directory and the stores below the root directory.

The following problems may occur during shutdown.

### Problem

#### **Stopping stores: Items are open (for example, backup or restore sessions are running).**

Stopping the store reports an error. Store is closed in a *dirty* state (items, such as backup or restore, are open or running). Recovery may happen on next startup.

### Action

Check that all operations are finished before stopping the StoreOnceSoftware utility.

### Problem

#### **Stopping stores: Housekeeping utility cannot be stopped.**

An error is logged during shutdown. Recovery may happen on next startup.

### Action

Check that all operations are finished and then stop the StoreOnceSoftware utility. Recovery may happen on next startup.

## Tutorial: Adding a Backup to Disk device and creating a backup specification

This procedure demonstrates the process of adding a Backup to Disk device-type and then creating a backup specification which targets this device. It assumes that the hardware has been correctly installed and configured.

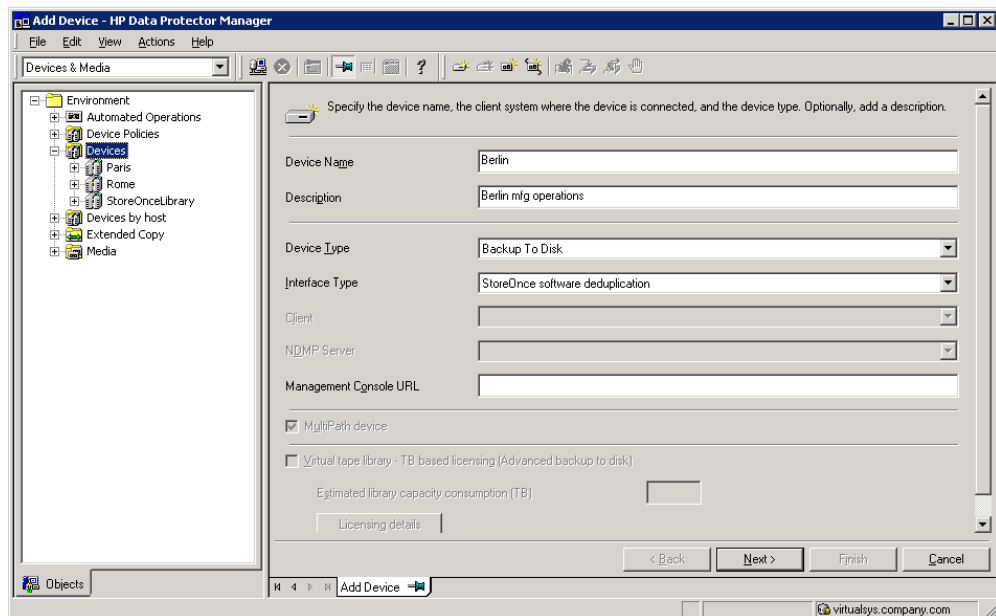
### Prerequisites:

- Data Protector Manager running on a Windows system.
- At least one gateway and a deduplication server (where the store is located) must have been configured at time of installation (including the root directory of the store).

### Add a Backup to Disk device

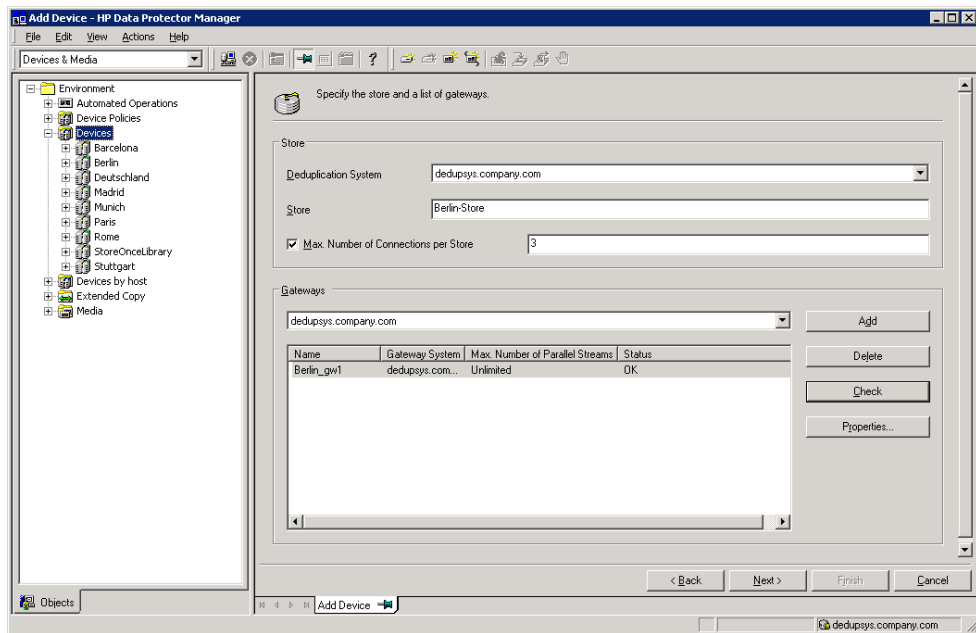
If necessary, start Data Protector Manager. In the **Devices & Media** context:

1. Right-click **Devices** and select **Add Device**.
2. In the Device Name field, specify a name (for example, *Berlin*).
3. In the Description field, specify a description (for example, *Berlin mfg operations*).
4. Select **Backup To Disk** in the Device Type drop-down list.
5. Make sure **StoreOnce software deduplication** is selected in the Interface Type drop-down list.



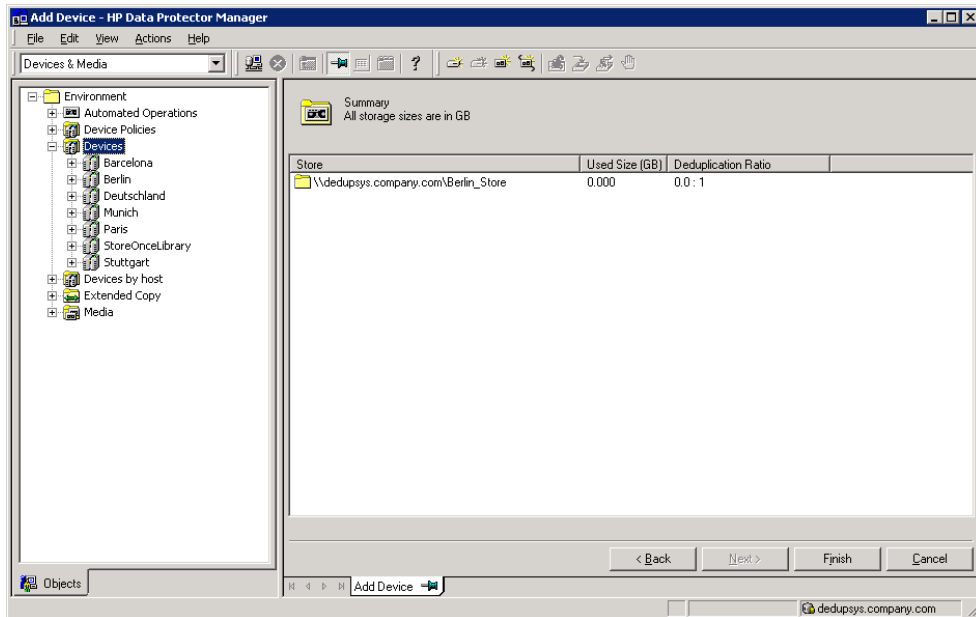
6. Click **Next**.
7. In the Deduplication System field, select the machine where the deduplication store is located. (This is the network path and address of the *hosting* machine for the deduplication store, for example, *virtualsys.company.com*).
8. In the Store field, specify a name (for example, *Berlin\_Store*). The maximum length of the store name is 80 characters (alphanumeric characters and underscore (\_) only).

9. (Optional) Check Max. Number of Connections per Store and specify a value between 1 to 30 (for example, 3).
10. Select a gateway from the drop-down list and click **Add** to add the selected gateway to the list (since you are not going to change any gateway settings, click **OK** in the Add Gateway dialog).
11. Highlight the gateway entry and click **Check** to verify connectivity.
12. Click **Yes** to confirm that you want to create the store. Following a successful connection, the Status field changes to OK.

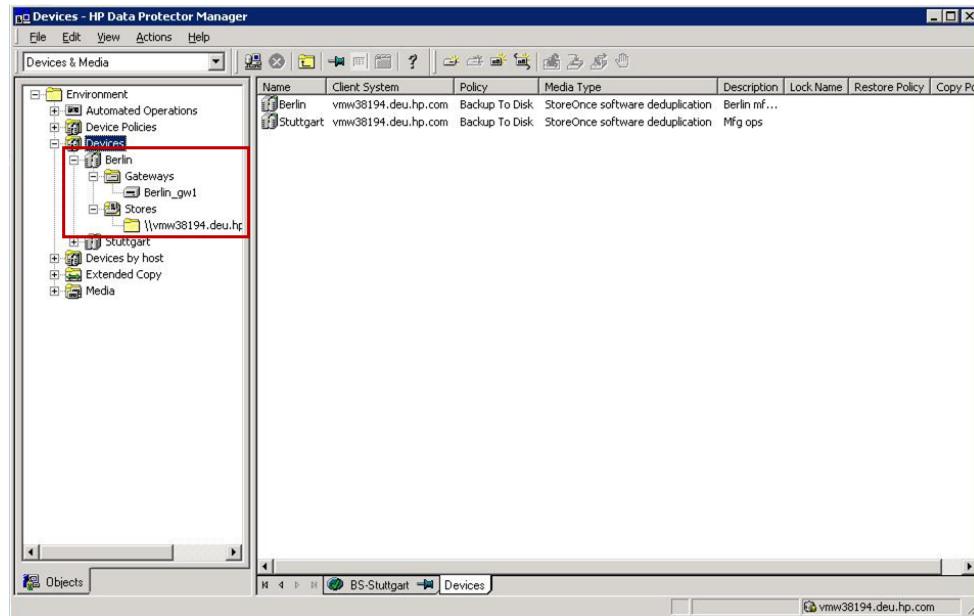


13. To view or modify gateway properties, click **Properties** (for more information on properties and advanced options, see [Setting additional options](#)).

14. Click **Next** to display the Summary screen.



15. Click **Finish**. The B2D device you just added is shown in the Scoping Pane. Note the suffix `_gw1` is appended to the gateway name.




This completes the procedure of adding a Backup to Disk device-type. The following procedure describes how to create and run a backup specification which targets this device. Note: Creating a backup specification for a B2D device is the same as with any other backup specification.

### Create and run a backup specification

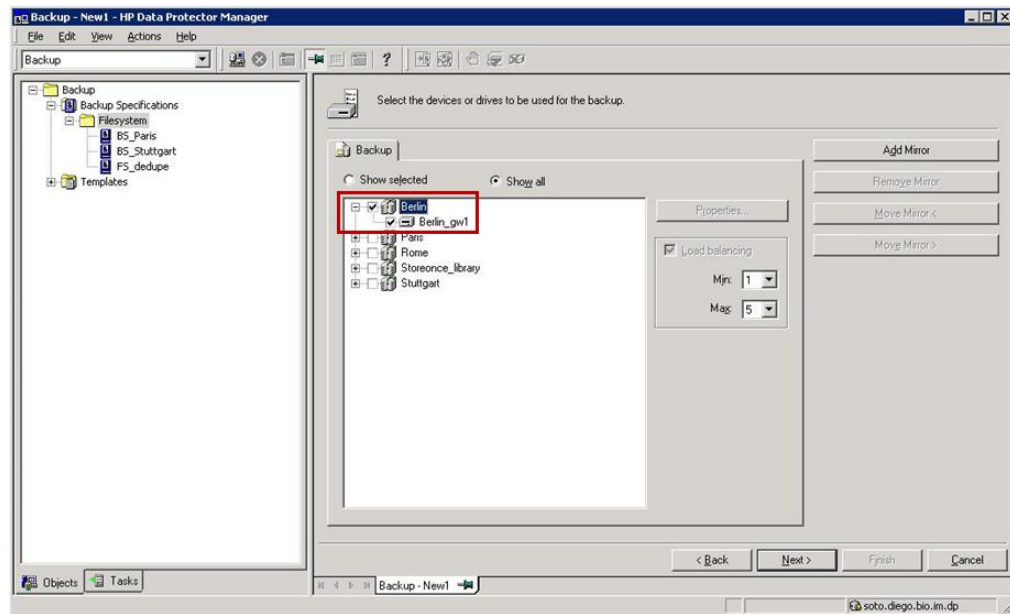
In the **Backup** context:

1. Expand **Backup Specifications**.

2. Right-click **Filesystem** and select **Add Backup**.
3. Click **OK** to accept the default Blank Filesystem Backup template.
4. Expand the filesystem in the Results Area and select a folder to backup, for example, MyTestfiles > Testfiles01.

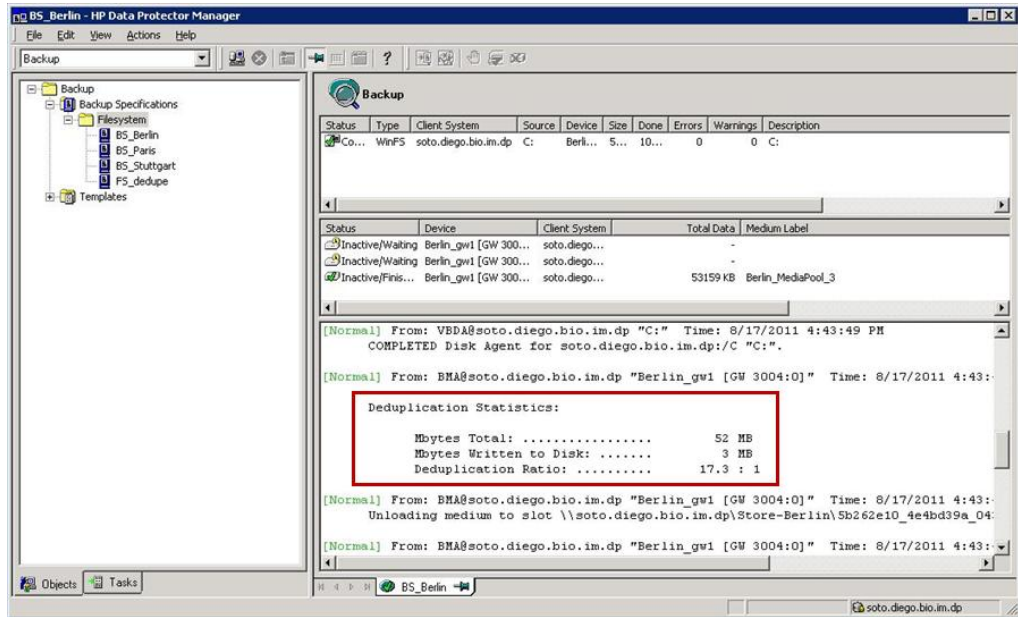


5. Click **Next**.
6. Select the device Berlin (the device you added earlier).



7. Click **Next**.
8. Add a description, for example, Backup Spec for folder Testfiles01.
9. Select **None** in Filesystem Options.
10. Click **Next**.
11. Select a date and click **Add** to specify a time.
12. Click **OK**.
13. Click **Next** to review the summary, and click **Next** again.
14. Click **Save As**, specify a name (for example, BS\_Berlin) and click **OK**.

15. To run the backup immediately, click **Actions > Start Backup > OK**.



The results of the backup operation (total backed up data, data written to disk, and deduplication ratio) are shown in the backup statistics.



## Appendix A: StoreOnceSoftware utility

The StoreOnceSoftware utility is a service/daemon and maintenance tool for performing general administration tasks on the StoreOnce library (the deduplication store). Note: On Windows systems, the term *utility* refers to the `StoreOnceSoftware.exe` service. On Linux systems, `StoreOnceSoftwared` is a script which utilizes the `StoreOnceSoftware` CLI (the options are: {start | stop | status}).

### Synopsis

```
StoreOnceSoftware --help | -h
StoreOnceSoftware [--log_path=LogPath] [--log_level=no_log|fatal|critical|
                  error|warning|notice|information|debug|tracing]
StoreOnceSoftware --configure_store_root --path=RootDirectory [--force]
StoreOnceSoftware --create_store --name=StoreName
                  [--description=StoreDescription]
StoreOnceSoftware --delete_store --name=StoreName [--force]
StoreOnceSoftware --start_store --name=StoreName
StoreOnceSoftware --stop_store --name=StoreName [--force]
StoreOnceSoftware --set_autostart=ON|OFF --name=StoreName
StoreOnceSoftware --list_stores [--name=StoreName]
StoreOnceSoftware --get_server_properties
StoreOnceSoftware --daemon
```

---

### Note

A standard format is used for both Windows and Linux platforms:

- `--option` or `--option=Value`
  - The short notation only works for the help option: `--help` or `-h`
- 

Note the following situations:

- Before any administrative store commands can be used, the `StoreOnceSoftware` utility must be running, otherwise, an error message is displayed.
- If the `StoreOnceSoftware` utility is *not* running and you do not specify a parameter, the following message is displayed: `The daemon is stopped. Restart the utility with the command: net start StoreOnceSoftware.`
- When starting the utility, you can define a path for logs and a logging-level with the options `--log_path` and `--log_level`.
- If the command-line option is not recognized (for example, the option is specified without the leading characters `--`), the following message is displayed:  
`Unknown option specified: unknown_option.`

### Description

```
--help | -h
```

Displays a list of CLI options with descriptions.

```
--configure_store_root --path=RootDirectory [--force]
```

Configures the root directory of the store(s). The root directory must be configured before you can create a store. When StoreOnce software deduplication is first installed, it is

running in a *non-configured* mode and cannot be used until the root directory of the store has been set. The option `--path` specifies the path to the root directory. The path must not be empty, must be a valid directory which already exists on the system, and not previously used as a store-root path. If successful, Data Protector displays the path in the CLI. If the root directory cannot be configured (there may be several reasons), you are prompted to stop/start the utility (daemon). Use the stop/start commands given below.

When the root directory is configured, Data Protector automatically creates the subdirectory `StoreOnceLibrary` below the root directory.

Use the `--force` option to reconfigure the location of the root directory in cases where the actual data has been moved to another location. Make sure the data is already at the specified location before using the `--force` option (see below for details).

Once the root directory is configured, it cannot be reconfigured through the GUI. To move the store to a new, never previously used, location (in case of a disaster recovery or being unable to use the existing mount point), proceed as follows:

1. Stop the StoreOnceSoftware utility.  
(On Windows systems: Use `net stop StoreOnceSoftware` or use the Service Manager, on Linux systems: Use `/opt/omni/sbin/StoreOnceSoftwared stop`)
2. Manually move the data from the `Old_Path` to the `New_Path`. This means the subdirectory `StoreOnceLibrary` and all its contents.
3. Run the command:  
`StoreOnceSoftware --configure_store_root --path=New_Path --force`  
(make sure `New_Path` includes the full path, for example, `--path=C:\Volumes\NewRoot`)
4. Start the StoreOnceSoftware utility.  
(On Windows systems: Use `net start StoreOnceSoftware` or use the Service Manager, on Linux systems: Use `/opt/omni/sbin/StoreOnceSoftwared start`)

See also [Backup of the system.db file](#) for information about the `system.db` file and the root directory, and [Configuring the root directory of the deduplication stores](#) for the GUI-based procedure to configuring the root directory.

`--log_path=LogPath`

Defines a path where logs are to be stored.

`--log_level=no_log | fatal | critical | error | warning | notice | information | debug | tracing`

Defines the logging detail as defined by:

`no_log` = Logging is disabled (default value).

`fatal` = A fatal error. The application will most likely terminate. This is the highest severity.

`critical` = A critical error. The application might not be able to continue running successfully.

`error` = A non-critical error. An operation did not complete successfully, but the application continues.

`warning` = A warning. An operation completed with an unexpected result.

`notice` = A notice, which is information with a higher severity.

`information` = An informational message, usually denoting the successful completion of

an operation.

debug = A debugging message.

tracing = A tracing message. This is the lowest severity.

`--create_store --name=StoreName [--description=StoreDescription]`

Creates a deduplication store with the specified name. Depending on the success of the operation, one of the following messages is displayed:

- If successful: The store: *StoreName* has been created successfully.
- If the store exists: The store *StoreName* has already been created.
- If an error occurs: Failed to create the store: *StoreName*.
- If a name is not specified:  
The options `--start_store`, `--stop_store`, `--create_store`, `--delete_store`, and `--set_autostart`, require the name option.

`--delete_store --name=StoreName [--force]`

Deletes the store with the given store name. The store is stopped before it is deleted.

If the `--force` option is used, the operation will try to close all the current activities, stop the store and then delete the store. One of the following messages is displayed:

If successful: The store: *StoreName* has been deleted successfully.

If an error occurs: Failed to delete the store: *StoreName*.

If name of the store is not specified:

The options `--start_store`, `--stop_store`, `--create_store`, `--delete_store` and `--set_autostart` require the `--name` option.

`--start_store --name=StoreName`

Starts the store specified by the store name. Depending on the success of the operation, one of the following messages is displayed:

- If successful: The Store *StoreName* has been successfully started.
- If already started: The store: *StoreName* has already been started.
- If an error occurs: Failed to start the store: *StoreName*
- If a name is not specified:  
The options `--start_store`, `--stop_store`, `--create_store`, `--delete_store`, and `--set_autostart` require the name option.

`--stop_store --name=StoreName [--force]`

Stops the store specified by the store name. If the `--force` option is used, the store attempts to close all current connections and then stops the store. Depending on the success of the operation, one of the following messages is displayed:

- If successful: The store: *StoreName* has been stopped successfully.
- If already stopped: The store: *StoreName* has already been stopped.
- If an error occurs: Failed to stop the store: *StoreName*
- If a name is not specified:  
The options `--start_store`, `--stop_store`, `--create_store`, `--delete_store`, and `--set_autostart` require the `--name` option.

`--set_autostart=ON|OFF --name=StoreName`

Sets a store to be started automatically (ON) or not automatically started (OFF). Depending on the success of the operation, one of the following messages is displayed:

- **If successful:** Autostart option for the store *StoreName* successfully set to *ON/OFF*.
- **If an error occurs:** Failed to set Autostart option for store: *StoreName*
- **If a name is not specified:**  
The options `--start_store`, `--stop_store`, `--create_store`, and `--set_autostart` require the name option.

#### `--list_stores`

Displays a list of stores that are configured on the deduplication server. It includes the following information:

- Store Name
- Store ID
- Store Description
- Store Status
- Store Autostart status
- User Data Stored (original, not deduplicated data)
- Store Size on Disk (resulting, deduplicated data)
- Deduplication Ratio

Depending on the success of the operation, one of the following messages is displayed:

- **If successful:** Listing of the stores succeeded.
- **If an error occurs:** Failed to list the stores.
- **If no stores can be identified:** There are no stores in database.

`--list_stores` supports the use of the `--name` option, which lists only the information about the store specified with the `--name` option. The following is a typical output showing three stores:

```
C:\Users\Administrator>StoreOnceSoftware --list_stores

Store Name:          StoreOnceLibrary
Store Id:            1
Store Description:   Data protector store
Store Status:        started
Store Autostartable: ON
User Data Stored:    1305 MB
Store Size on Disk:  770 MB
Deduplication Ratio: 1.7 : 1

Store Name:          Berlin_Store
Store Id:            2
Store Description:   Data protector store
Store Status:        started
Store Autostartable: ON
User Data Stored:    20 MB
Store Size on Disk:  1530 KB
Deduplication Ratio: 13.7 : 1

Store Name:          Lisbon_Store
Store Id:            5
Store Description:   Data protector store
Store Status:        started
Store Autostartable: ON
User Data Stored:    4549 MB
Store Size on Disk:  1174 MB
Deduplication Ratio: 3.9 : 1

Listing of the stores succeeded.
```

#### `--get_server_properties`

Displays a list of server/store related properties:

- Path and name of the root directory (`Store Root:`)

- Existing stores
- Available stores
- Space capacity
- Disk space free

```
C:\Users\Administrator>StoreOnceSoftware --get_server_properties
```

```
Store Root:          c:\StoreOnceRoot
Existing Stores:     3
Available Stores:   29
Disk Capacity:      49 GB
Disk Space Free:    29 GB
```

```
Listing of the daemon properties succeeded.
```

--daemon

Applies to Linux systems only. Run StoreOnceSoftware as a daemon for debugging purposes.

If the option --name is used in combination with non-valid options, the following message is displayed: The option --name can only be specified with the --start\_store, --stop\_store, --create\_store, --delete\_store, --set\_autostart or the --list\_stores command.

The option --description can only be used when creating a store (--create\_store).

## Appendix B: Command line interface changes to support B2D devices

The following CLI command options are added or updated.

### Addition to omnimm

#### Synopsis

```
omnimm -delete_unprotected_media [Library | -all]
```

**Description** Manual removal of unprotected B2D backup objects

When invoked, Data Protector builds a list of unprotected backed-up objects residing in the B2D device. Data Protector first removes them from the store and then deletes the media information from the Data Protector database. Note that removing media from the store does not free-up disk space, it merely instructs the store to obsolete the data. You can optionally specify the name of the library. In this case only unprotected media from the specified library are deleted.

**Description** Automatic removal of unprotected B2D backup objects

This method does the same as the above but is performed automatically by Data Protector at regular intervals. The interval can be configured in the global options file.

```
# DeleteUnprotectedMediaFreq=TimesPerDay
# default: 1
# limit: 1 <= DeleteUnprotectedMediaFreq <= 24
# If set to 1, operation is performed once per day (00:00),
# set to 2 two times per day (00:00,12:00), set to 3 three times
# per day (00:00, 08:00, 16:00), set to 4 four times per day
# (00:00, 06:00, 12:00, 18:00). If maximum (24) is specified,
# operation will be started on every hour.
```

This functionality is implemented by the `omnitrig` command which calls `omnimm` at specified intervals.

**Description** Immediate removal of objects when a slot is deleted

Deleting a slot deletes the slot from the IDB, deletes the media in the slot, and the slot itself is removed from the store. This is the same action as recycling and deleting. The commands are:

```
omnimm -recycle MediumID
omnimm -remove_slots Library Slot
```

### Addition to omnidownload

Use this command and options to list devices, device information, and libraries.

#### Synopsis

```
omnidownload -list_devices
```

The following example shows a typical device-list output:

```

Device name           Host                Device type          Pool Name
=====
DeDup Device1_gw1    paris_gw.gateway.com BackupToDisk StoreOnce DeDup
Device1_MediaPool
DeDup Device1_gw2    rome_gw.gateway.com  BackupToDisk StoreOnce DeDup
Device1_MediaPool
DeDup Device1_gw3    tercus.gateway.com    BackupToDisk StoreOnce DeDup
Device1_MediaPool

```

### Synopsis

```
omnidownload -dev_info
```

The following example shows a typical device-information output:

```

Device name           Host                Device type          Pool Name
=====
DeDup Device1         device.box.host1    BackupToDisk StoreOnce
DeDup Device1_gw1     paris_gw.domain.com StoreOnce (in library) DeDup Device1_MediaPool
DeDup Device1_gw2     rome_gw.domain.com  StoreOnce (in library) DeDup Device1_MediaPool
DeDup Device1_gw3     tercus.domain.com    StoreOnce (in library) DeDup Device1_MediaPool
=====
Together : 3 configured device[s].

```

### Synopsis

```
omnidownload -list_libraries -detail
```

The following example shows a typical device-library output:

```

NAME "DeDup Device1"
DESCRIPTION ""
POLICY BackupToDisk
TYPE StoreOnceSoftware
VTLCAPACITY 4
REPOSITORY
    "\\device.box.host1\Store1\7fccb8b7_4d63bb99_01b8_0004"
    "\\device.box.host1\Store1\7fccb8b7_4d63bb99_01b8_000f"
    "\\device.box.host1\Store2\7fccb8b7_4d63bb99_01b8_00a1"
    "\\device.box.host1\Store2\7fccb8b7_4d63bb99_01b8_00e0"
DIRECTORY
    "\\device.box.host1\Store1" "<encoded_client_id>"
"<max_store_connections>"
MGMTCONSOLEURL ""

```

## omniupload

Use the `omniupload` command to upload information about a B2D device from an ASCII file to the Data Protector IDB. For command options, see the *HP Data Protector Command Line Interface Reference*.

## Appendix C: Omnirc variables related to B2D devices

The *omnirc* file is enhanced with additional options. Use this file to set such parameters as port number and disk-space threshold warnings. These changes apply to the client only.

### Synopsis

```
OB2_STOREONCESOFTWARE_COMMAND_PORT=port_number
```

### Description

This variable changes the port which is used for command communication between the Media Agent and the StoreOnceSoftware utility.

For example: `OB2_STOREONCESOFTWARE_COMMAND_PORT=12345`

Default: 9387

### Synopsis

```
OB2_STOREONCESOFTWARE_DATA_PORT=port_number
```

### Description

This variable changes the port which is used for data communication between the Media Agent and the StoreOnceSoftware utility.

For example: `OB2_STOREONCESOFTWARE_DATA_PORT=12346`

Default: 9388

### Synopsis

```
OB2_STOREONCESOFTWARE_SESSION_IDLE_TIMEOUT=s
```

### Description

The StoreOnceSoftware daemon periodically checks for idle connections and terminates them. This variable specifies the number of seconds of inactivity after which a connection is considered idle.

Default: 300 (Range: Minimum: 10)

### Synopsis

```
OB2_STOREONCESOFTWARE_DISK_SPACE_THRESHOLD=%
```

### Description

This variable is to set a threshold for the free disk space (see [Low disk space](#) warning).

Default: 10% (Range: 1% - 95%)

### Synopsis

```
OB2_STOREONCESOFTWARE_MINIMUM_DISK_SPACE=n
```

### Description



This variable controls the minimum disk space (in MB) the StoreOnceSoftware needs to reserve. If this minimum is reached, writing data to any stores will fail. Default: 1000 (Minimum: 500)

## Appendix D: System requirements and performance

This section lists the supported platforms and the minimum and recommended requirements for the deduplication server.

### Supported platforms

Supported platforms for StoreOnce software deduplication component are:

- Windows Server 2008 R2 Enterprise (64-bit, x64)
- Windows Server 2008 (64-bit, x64)
- SUSE Linux Enterprise Server 10.x (64-bit, x86-64)
- SUSE Linux Enterprise Server 11.x (64-bit, x86-64)
- Red Hat Enterprise Linux (RHEL) 5.x (64-bit, x86-64)
- Red Hat Enterprise Linux (RHEL) 6.x (64-bit, x86-64)

StoreOnce software deduplication is not supported on Linux x64 gateways which have the NDMP Media Agent installed.

StoreOnce software deduplication component can be installed on physical systems as well as on virtual machines. When choosing a virtual machine as your deduplication server, note that the performance of deduplication stores may be lower than expected.

For details on supported systems, see the latest HP Data Protector support matrices at <http://www.hp.com/support/manuals>. See the *HP Data Protector Installation and Licensing Guide* on how to install Data Protector in various architectures.

### Deduplication server requirements

	<b>CPU speed / number of cores<sup>1</sup></b> (dedicated to StoreOnce software)	<b>Physical memory<sup>2</sup></b> (dedicated to StoreOnce software)	<b>Number of disks<sup>1</sup></b> (dedicated to the store)
<b>Minimum requirements</b> (for 1 TB store)	2.2 GHz / 1 core	300 MB of RAM	1
<b>Recommended requirements</b> (for 10 TB store with 5 parallel connections)	2.2+ GHz / 4 cores	3 GB of RAM	4 or more using RAID5
(1) For maximum performance, consider 0.7 cores, 0.8 disks, and additional 50 MB of RAM per parallel stream, and use a dedicated physical system as deduplication server. (2) 300 MB of RAM are needed for each 1 TB of store capacity.			

For data security and performance reasons, only local host filesystems are supported. This includes disks or volumes connected through SCSI, Fiber Channel, or iSCSI interfaces; NFS and CIFS are

not supported. On Windows systems, only NTFS volumes are supported and on Linux systems, volumes with the 'ext4' or 'xfs' filesystems are supported.

HP recommends RAID5 with a limited stripe size, which achieves good performance in conjunction with the ability to recover from a single disk failure.

## Glossary

This glossary reflects terms relating to this document.

Term	Description
Backup to Disk (B2D) device	Disk-based backup device. The StoreOnce software deduplication interface is accessed through the device type <i>Backup To Disk</i> .
chunking (chunks)	The process of dividing data into blocks (chunks), where each chunk gets a unique content address. Internal storage unit of between 1.6 kB – 10 kB variable size (4 kB on average). Unique data chunks are stored in the deduplication store. A list of (references to) chunks makes up a portion.
D2D	Disk-to-disk storage. Generic name for disk-to-disk backups. For example, HP StoreOnce Backup systems offers D2D storage (and includes built-in deduplication).
deduplication	Process of eliminating duplicate data from an incoming data stream. Only unique data is written to the storage disk. During a restore operation, the original data stream is reconstructed (sometimes referred to as <i>rehydrated</i> ).
deduplication daemon	See <a href="#">StoreOnceSoftware</a> utility.
deduplication ratio (-rate)	Ratio between size of the source data to be backed up and the size of the actual data which is written to the store. For example, if 10 MB of source data results in only 1 MB of backed up data being written to store, the deduplication ration is 10 : 1.
deduplication store	See <a href="#">StoreOnce library</a> .
deduplication system	Enables multiple Media Agents to work with the StoreOnce library simultaneously. Name of the <i>hosting</i> machine where the deduplication store is located (name of the server where the deduplication store is installed).
gateway	Gateways are clients with the Media Agent component installed. Sometimes referred to as a Media Agent host or gateway client. Each gateway represents a host on which multiple Media Agents can be started at the same time. Gateways are selected from a drop-down list of Data Protector clients that have the Media Agent component.
hash	A short fingerprint of a sequence of bytes, typically of a <i>chunk</i> . It uniquely identifies this chunk with a very high probability.
item	Externally addressable unit of storage in the deduplication store. It can represent Tape (for Virtual Tape personality), File (NFS) in the D2D appliance, and consists of <i>portions</i> . An item contains one or more object versions.
LSU	Logical Storage Unit. Analogous to the D2D Object Store but for Data Domain Boost devices.
LUN	Logical unit number. A number which is used to identify a logical unit such as a device addressed by the SCSI protocol.
portion	Description of input data as list of hashes, variable size (10 MB on average), and limited by input data size or number of hashes. A portion is part of an Item. Portions do not overlap. The sequence of all portions of an item makes up the item.
server-side deduplication	Deduplication is done on the machine of the deduplication store only, so all the data needs to be transferred over the network (this is sometimes referred to as target-side deduplication or high-bandwidth transfer).
SHA-1	Cryptographic hash function. Secure hash algorithm.

source-side deduplication	Deduplication is done on the client machine, so only new, unique data needs to be transferred over the network (this is sometimes referred to as client-side deduplication or low-bandwidth transfer). Note: Data Protector does not support source-side deduplication.
StoreOnce	All HP StoreOnce products feature HP StoreOnce deduplication software. Full product name: <i>HP StoreOnce Backup system</i> .
StoreOnce library	Physical store used by StoreOnce software deduplication.
StoreOnceSoftware utility	The StoreOnceSoftware utility runs as a service ( <i>StoreOnceSoftware.exe</i> ) on Windows systems and as a background process (daemon) on Linux systems. A maintenance tool for performing general administration tasks on the deduplication store.
StoreOnce software deduplication	Interface type used by the Backup to Disk device. Provides software-based deduplication.
target-side deduplication	See server-side deduplication.
UPS	Uninterruptible power supply (or source). Battery backup when the mains electricity fails.

## Index

- backing up, 9
- Backup To Disk
  - concepts, 6
  - device type, 3
  - multiple devices, 7
- capacity, maximum storage, 11
- CLI
  - omnidownload, 38
  - omnimm, 38
  - omnirc, 40
  - omnupload, 39
  - StoreOnceSoftware utility, 33
- concurrency, 20, 22
- consolidation, and object copy, 9
- data format, 10
- deduplication
  - block-level, 4
  - installation package, 17
  - target-side, 5
- delete
  - expired backup data, 11
- environment
  - example scenario, 19
- Gateway Tag, 22
- gateways, 6, 20
  - Advanced options, 22
  - identifying, 11
  - limits, 8
    - multiple, 6
    - multiple Media Agents, 8
    - properties, modifying, 22
- global options, 38
- installation
  - prerequisites, 14
  - procedure, 14
- interface type
  - StoreOnce software deduplication, 11
- licensing, 18
- load balanced, 10
- Media Agent
  - component, 6, 11
  - distribute between gateways, 10
- multi-host configurations, 6
- object copy and consolidation, 9
- package names (Windows, linux), 17
- ratio
  - deduplication, 13
- rehydration, 3
- statistics, deduplication, 12
- store name, length, 28
- StoreOnce software deduplication
  - interface type, 11
- StoreOnceSoftware utility, 33
- striping mechanism, 12
- supported platforms, 41
- target-side deduplication, 5

## For more information

Visit the following Data Protector online resources to get more information:

<http://www.hp.com/go/dataprotector>

<http://www.hp.com/go/imhub/dataprotector>

<http://www.hp.com/go/d2d>



© Copyright 2011 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.

Microsoft and Windows are U.S. registered trademarks of Microsoft Corporation. UNIX is a registered trademark of the Open Group. Linux is a U.S. registered trademark of Linus Torvalds.

4AA3-7696ENA, Created December 2011

