HP Data Protector Software Cell Manager Planning and Sizing

Planning and Configuring a HP Data Protector Software Cell Manager and Managing the Internal Database Growth

Table of contents

Executive summary

This white paper provides complementary information on how to plan, size, and maintain a HP Data Protector Cell Manager in a Data Protector cell.

Solution description

HP Data Protector software is a backup and disaster-recovery software that provides reliable data protection and high accessibility for your fast growing business data. Data Protector offers comprehensive backup and restore functionality specifically tailored for enterprise-wide and distributed environments. Data Protector can be used in environments ranging from a single system to thousands of systems on several sites. Due to the network component concept of Data Protector, elements of the backup infrastructure can be placed in the topology according to user requirements. The numerous backup options and alternatives to setting up a backup infrastructure allow the implementation of virtually any configuration you want.

HP Data Protector enables you to perform backup to a large number of backup devices simultaneously. It supports high-end devices in very large libraries. Various backup possibilities, such as local backup, network backup, online backup, disk image backup, synthetic backup, backup with object mirroring, and built-in support for parallel data streams allow you to tune your backups to best fit your requirements.

As Data Protector supports heterogeneous environments, most features are common to the UNIX and Windows platforms.

The Data Protector cell is a network environment that has a Cell Manager, client systems, and devices. The Cell Manager is the central control point where Data Protector software is installed. After installing Data Protector software, you can add systems to be backed up. These systems become Data Protector client systems that are part of the cell. When Data Protector backs up files, it saves them to media in backup devices. The Data Protector internal database (IDB) keeps track of the files you back up so that you can browse and easily recover the entire system or single files. The HP Data Protector Cell Manager also runs session manager processes that start and stop backup and restore sessions and write session information to the IDB

The UNIX and Windows Cell Managers can control all supported client platforms (UNIX, Windows, and Novell NetWare). The Data Protector user interface can access the entire Data Protector functionality on all supported platforms.

It is crucial to the operation of a cell that the HP Data Protector Cell Manger system is properly planned and setup in order to be able to cope with the various performance aspects in such a complex environment.

Cell Manager software topology

The Cell Manager is the key component of a Data Protector cell. It contains the Data Protector internal database (IDB), and is responsible for the start of backup, restore, copy/consolidation, and media management sessions.

Primary Data Protector processes

The Data Protector services (Windows) and daemons (UNIX) run on the Cell Manager.

Table 1: Data Protector Cell Manager services/daemons

Service	Data Protector 6.0	≥Data Protector 6.1
crs Cell Request Server	☑	☑
rds Raima Database Server	☑	☑
mmd Media Management Daemon	☑	☑
Uiproxy Java GUI Server		☑
Kms Encryption Key Management Server		☑
omnitrig Triggers Data Protector scheduled backups. (on UNIX platform only)	☑	☑
omniinet Inter process communication (on Win Platform only)	☑	☑

In a Windows environment, the media management daemon (MMD) runs as an application process, (`mmd.exe`) rather than a service. The MMD is started by the CRS service.

Run the `omnisv –status` command to check whether services/daemons are running.

```
C:\Program Files\OmniBack\bin>omnisv -status
    ProcName  Status  [PID]
===============================
    rds     : Active  [2992]
    crs     : Active  [2296]
    mmd     : Active  [2504]
    kms     : Active  [2728]
    uiproxy : Active  [2900]
    omniinet: Active  [2376]
    Sending of traps disabled.
===============================
Status: All Data Protector relevant processes/services up and running.
```

Session managers

The Cell Manager CRS service listens for session requests and starts the appropriate session managers, which in turn starts the required clients. A dedicated session manager controls the clients for each operation. If a new session is started, an additional session manager session is generated.

bsm	Backup Session Manager
rsm	Restore Session Manager
csm	Copy Session Manager (used for object copy and object consolidation)
dbsm	Database Session Manager
msm	Media Session Manager
asm	Administration Session Manager

Figure 1: Cell manger topology

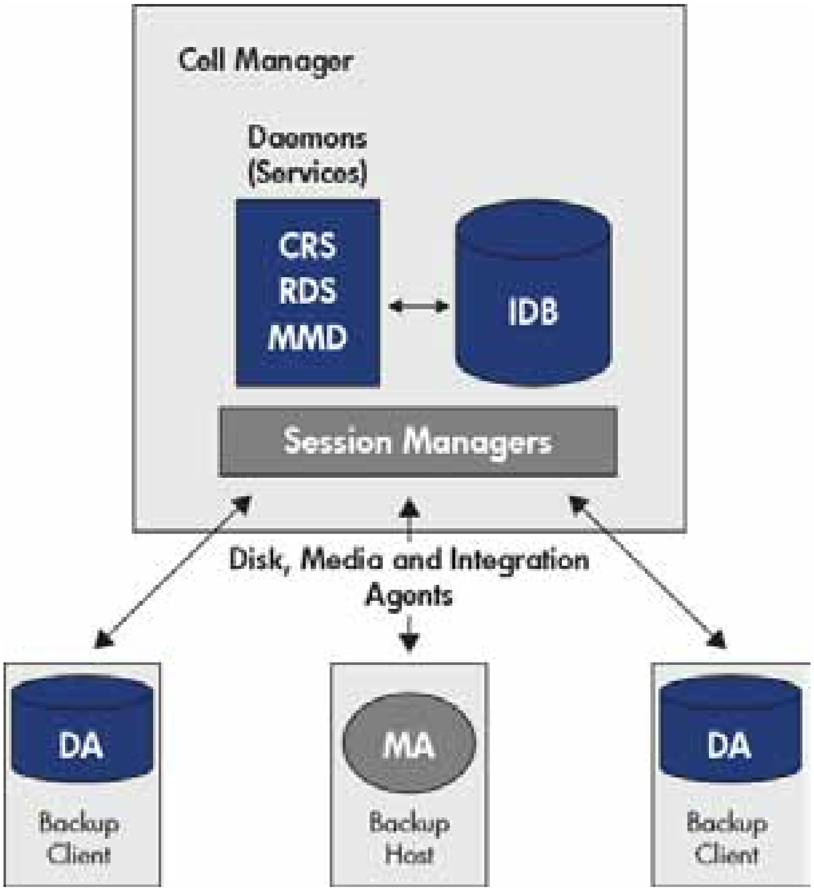

Note:
Disk Agent (DA) and Media Agent (MA) components can be installed on the Data Protector Cell Manager host as well on clients belonging to that Cell Managers cell to fulfill their designated roll.

IDB architecture

The IDB is an internal database located on the Data Protector Cell Manager that keeps information regarding what data is backed up; on which media it resides; the result of backup, restore, copy, object consolidation, and media management sessions; and which devices and libraries are configured.

What is the IDB used for?

The information stored in the IDB enables the following:

- **Fast and convenient restore:** You are able to browse the files and directories to be restored. You can quickly find the media required for a restore and therefore make the restore much faster.
- **Backup management:** You can verify the result of backup sessions.
- **Media management:** You can allocate media during backup, copy, and object consolidation sessions, track media management operations and media attributes, group media in different media pools, and track media location in tape libraries.

Figure 2: IDB parts

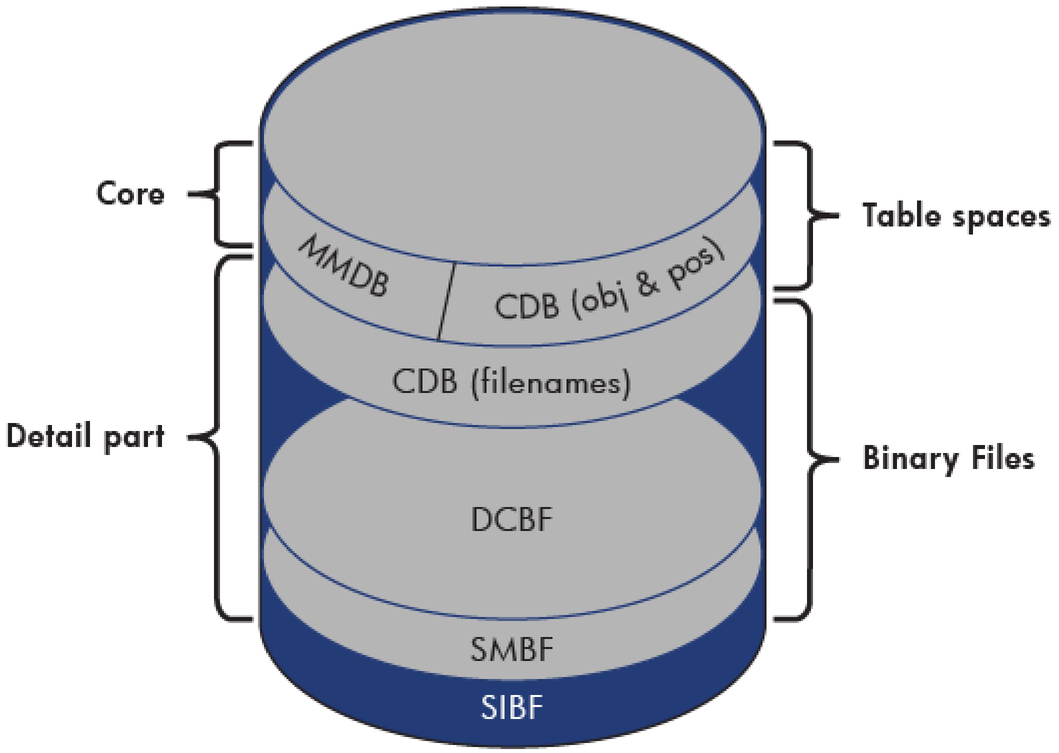

The internal database consists of the following parts:

- MMDB (Media Management Database)
- CDB (Catalog Database)
- DCBF (Detail Catalog Binary Files)
- SMBF (Session Messages Binary Files)
- SIBF (Server-less Integrations Binary Files)

Each of the IDB parts stores certain specific Data Protector information (records), influences the IDB size and growth in different ways, and is located in a separate directory on the Cell Manager.

The MMDB and CDB parts are implemented using an embedded database consisting of tablespaces. This database is controlled by the rds database server process. All changes to the MMDB and CDB are updated using transaction logs. CDB (objects and positions) and MMDB present the core part of IDB.

The DCBF, SMBF, and SIBF parts of the IDB consist of binary files. Updates are direct (no transactions).

In the Manager-of-Managers (MoM) environment, the local MMDB can be merged to a central Cell Manager system to create the Central Media Management Database (CMMDB).

Media Management Database (MMDB)

MMDB records

The Media Management Database stores information about the following:

- Configured devices, libraries, library drives, and slots
- Data Protector media
- Configured media pools and media magazines

MMDB size and growth

The MMDB does not grow very big in size. The largest part of the MMDB is typically occupied by information about the Data Protector media. Space consumption is in range of 30 MB.

Using a large amount of media can result in high memory consumption of IDB and cause a backup to fail. You can calculate the memory usage on HP-UX using the following formula:

```
RDSSize=2048 KB + n*(0.7KB*m+1.5KB*a)
```

Where:

- `n` is the minimum value of the number of parallel strings or the number of RPC threads set in `rdmserver.ini`. Default is 3.
- `m` is the number of media in a selected pool,
- a is the number of media in IDB.

The average sizes are the following:

- 0.7 KB is the size of a medium record in IDB (per pool)
- 1.5 KB is the size of a medium record in a binary file (all media)

MMDB location

Information about configured backup devices is stored in the `devices.dat` file in the

<Data_Protector_program_data>\db40\datafiles\mmdb directory (on Windows Server 2008),
<Data_Protector_home>\db40\datafiles\mmdb directory (on other Windows systems) or
/var/opt/omni/server/db40/datafiles/mmdb (on UNIX systems).

Catalog Database (CDB)

CDB records

The Catalog Database stores information about the following:

- Backup, restore, copy, object consolidation, and media management sessions
  This is the copy of the information sent to the Data Protector Monitor window.
- Backed up objects, their versions, and object copies
- Positions of backed up objects on media
  For each backed up object, Data Protector stores information about the media and data segments used for the backup. The same is done for object copies and object mirrors.
- Pathnames of backed up files (filenames) together with client system names
  Filenames are stored only once per client system. The filenames created between backups are added to the CDB.

Figure 3 depicts the properties of the catalog database of an existing IDB.

Figure 3: IDB catalog database properties

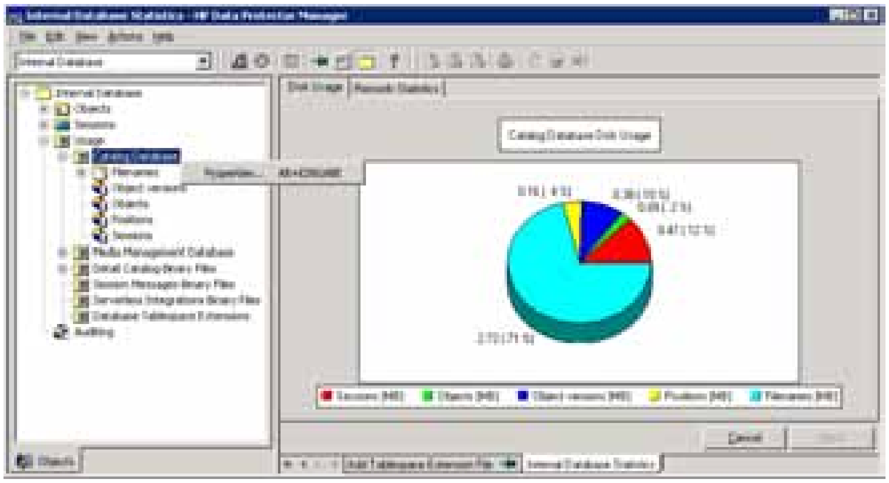
Internal Database Statistics - HP Data Protector Manager
Internal Database
Objects
Sessions
Usage
Catalog Database
Filenames
Properties... Alt+ENTER
Object versions
Objects
Positions
Sessions
Media Management Database
Detail Catalog Binary Files
Session Messages Binary Files
Serverless Integrations Binary Files
Database Tablespace Extensions
Auditing
Disk Usage
Records Statistics
Catalog Database Disk Usage
Sessions [MB]
Objects [MB]
Object versions [MB]
Positions [MB]
Filenames [MB]
Cancel
Objects
Internal Database Statistics

CDB (filenames) size and growth

The biggest and fastest growing part of the CDB is the filenames part. It typically occupies 20% of the entire database.

The growth of the filenames part is proportional to the growth and dynamics of the backup environment and not to the number of backups.

A file or directory in the IDB occupies approximately 50–70 bytes on the UNIX Cell Manager and 70–100 bytes on the Windows Cell Manager.

Filenames are stored in the `fnames.dat` file and in some other files, depending on the filename length. The maximum size of each of these data files is 2 GB. You are notified when one of these files starts running out of space, so that you can add new files to extend the size of the filenames part of the IDB.

Note:
The filenames of object copies are not added to the IDB. Object copying does not produce any impact on the filename part of the CDB if its logging level is the same or less detailed than the logging level of its source object.

Table 2 lists the maximum configurable sizes and limits for the HP Data Protector catalog datafiles.

Table 2: Catalog datafiles and extensions - maximum settings and limits

	Data Protector 6.0	≥ Data Protector 6.1
Size of any IDB datafile (base file or extension file)	2 GB	2 GB
Max size of fnames: (base file and all extensions)	32 GB	48 GB
Max size of fn1…fn4 (base file and all extensions)	8 GB (16 GB with patch)	16 GB
Max size of dirs.dat (base file and all extensions)	8 GB (16 GB with patch)	16 GB

CDB (objects and positions) size and growth

The CDB also records object names and media positions. It occupies minor share of space in the IDB. Space consumption is in range of up to 1 GB for a backup environment.

Prior extending the IDB the backup administrator should check the currently configured tables and their location and size.

Figure 4 depicts a summary of currently configured tables, their size, and location.

Figure 4: Summary of IDB tablespace location and size

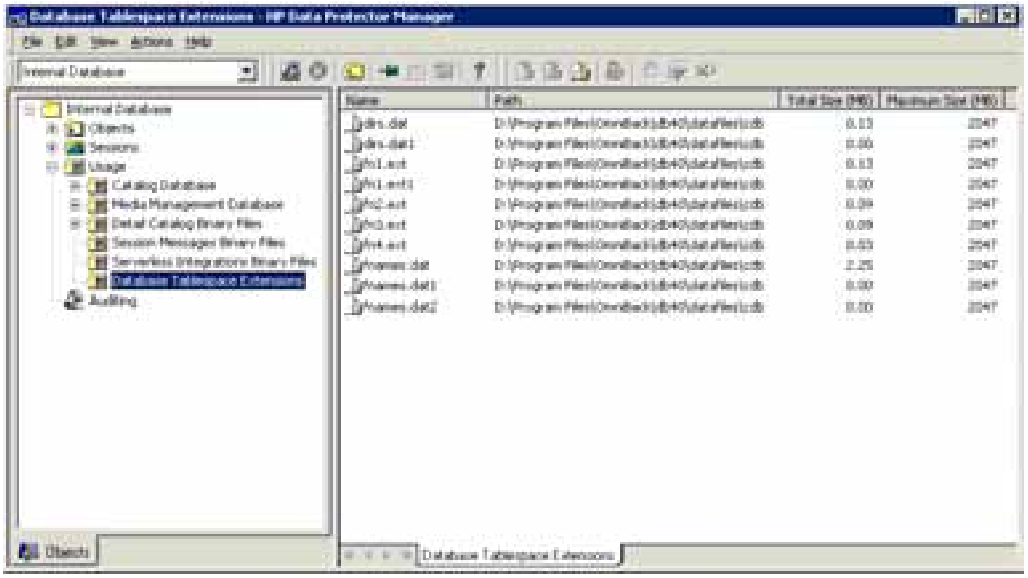
Database Tablespace Extensions - HP Data Protector Manager
File Edit View Actions Help
Internal Database
Internal Database
Objects
Sessions
Usage
Catalog Database
Media Management Database
Detail Catalog Binary Files
Session Messages Binary Files
Serverless Integrations Binary Files
Database Tablespace Extensions
Auditing
Objects
Database Tablespace Extensions

Name	Path	Total Size (MB)	Maximum Size (MB)
dirs.dat	D:\Program Files\OmniBack\db40\datafiles\cdb	0.13	2047
dirs.dat1	D:\Program Files\OmniBack\db40\datafiles\cdb	0.00	2047
fn1.ext	D:\Program Files\OmniBack\db40\datafiles\cdb	0.13	2047
fn1.ext1	D:\Program Files\OmniBack\db40\datafiles\cdb	0.00	2047
fn2.ext	D:\Program Files\OmniBack\db40\datafiles\cdb	0.09	2047
fn3.ext	D:\Program Files\OmniBack\db40\datafiles\cdb	0.09	2047
fn4.ext	D:\Program Files\OmniBack\db40\datafiles\cdb	0.53	2047
fnames.dat	D:\Program Files\OmniBack\db40\datafiles\cdb	2.25	2047
fnames.dat1	D:\Program Files\OmniBack\db40\datafiles\cdb	0.00	2047
fnames.dat2	D:\Program Files\OmniBack\db40\datafiles\cdb	0.00	2047

CDB location

The CDB is located in the following directory:

> <Data_Protector_program_data>\db40\datafiles\cdb (Windows Server 2008),
> <Data_Protector_home>\db40\datafiles\cdb (other Windows systems), or
> /var/opt/omni/server/db40/datafiles/cdb (UNIX systems).

When adding an extension file, you can specify a different path.

Note:
You might consider redirecting the CDB portion of the IDB to a dedicated and fast disk subsystem as this helps to maintain IDB performance in large and/or busy backup environments.

Example: Extending other tablespaces

By default, the IDB Tablespace Space Low notification is triggered when 85% of the space allocated for a specific tablespace is used. Perform the following steps to extend the specified tablespace.

Steps

1. Check if the tablespace is visible in the GUI.
   i. In the Context List, click **Internal Database**.
   ii. In the Scoping Pane, expand **Usage** and click **Database Tablespace Extensions**.

The names of the tablespaces are listed in the Results Area. If the tablespace that has run out of space is listed, perform the remaining steps. Otherwise, contact technical support for further instructions.

2. Extend the tablespace.
   i. In the Scoping Pane, right-click **Database Tablespace Extensions** and click **Add Database Tablespace Extension File**.
   ii. In the **Type** drop-down list, select the tablespace that has run out of space. Specify other options as desired.
   iii. Click **Finish** to exit the wizard.

Figure 5 depicts the database table extension wizard.

Figure 5: Database table space extension wizard

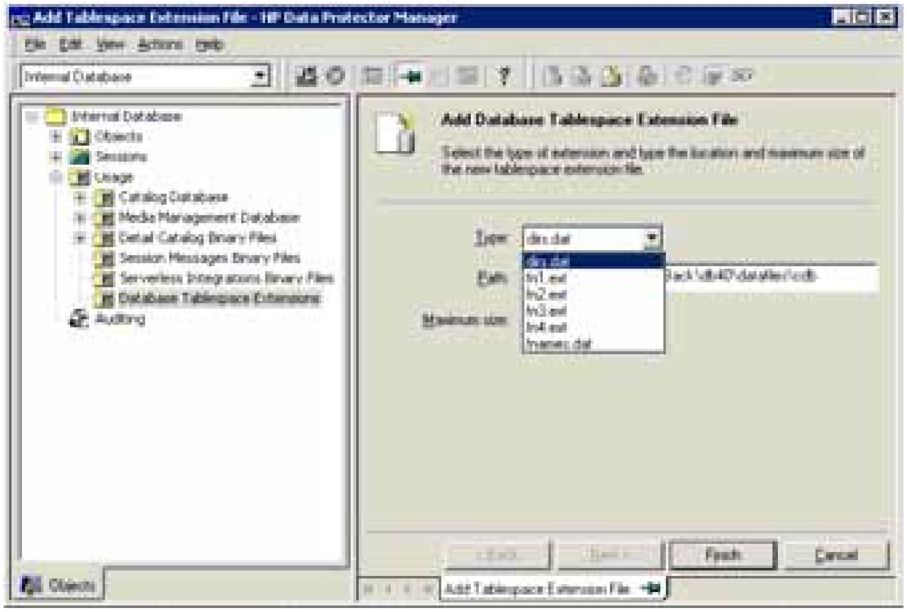
Add Tablespace Extension File - HP Data Protector Manager
File Edit View Actions Help
Internal Database
Internal Database
Objects
Sessions
Usage
Catalog Database
Media Management Database
Detail Catalog Binary Files
Session Messages Binary Files
Serverless Integrations Binary Files
Database Tablespace Extensions
Auditing
< Back
Next >
Finish
Cancel
Objects
Add Tablespace Extension File

Add Database Tablespace Extension File

Select the type of extension and type the location and maximum size of the new tablespace extension file.

Type: dirs.dat

dirs.dat
fn1.ext
fn2.ext
fn3.ext
fn4.ext
fnames.dat

Path: ...ack\db40\datafiles\cdb

Maximum size:

Detail Catalog Binary Files (DCBF)

DCBF information

The Detail Catalog Binary Files part stores file version information. This is information about backed up files, such as file size, modification time, attributes/protection, and so on.

One DC (Detail Catalog) binary file is created for each Data Protector medium used for backup. When the medium is overwritten, the old binary file is removed and a new one is created.

DCBF size and growth

In an environment where filesystem backups using the **Log all** option is typical, the DCBF occupies the largest part (typically 80%) of the IDB. Approximately 30 bytes are used for each version of each backed up file. Logging level and catalog protection can be used to specify what is actually stored in the IDB and for how long.

By default, one DC directory is configured for the DC binary files, the db40\dcbf directory. Its default size is 4 GB (16 GB for Data Protector 6.1), the maximum configurable size is 32 GB. You can create more DC directories and have them on different disks on the Cell Manager, thus extending IDB size. The maximum number of supported directories per cell is 50.

DCBF location

By default, the DCBF is located in the following directory:

<Data_Protector_program_data>\db40\dcbf (Windows Server 2008),
<Data_Protector_home>\db40\dcbf (other Windows systems), or
/var/opt/omni/server/db40/dcbf (UNIX systems).

Consider the disk space on the Cell Manager and relocate the DC directory, if necessary. You can create more DC directories and locate them to different disks. Create several DC directories only if the number of media/DC binary files grows very big (several thousands) or you have space issues.

Example: Creating new DC directories

Create new DC directories to provide additional space for file versions and attributes. If possible, locate them on different disks.

Steps

1. In the Context List, click **Internal Database**.
2. In the Scoping Pane, expand the **Usage** item.
3. Right-click the **Detail Catalog Binary Files** and click **Add Detail Catalog Directory**.
4. In the Allocation sequence text box, specify the order in which Data Protector will choose the DC directory to write information to it.
5. In the Path text box, specify the path for the directory.
6. In the Maximum size text box, specify the maximum size for the directory.
7. In the Maximum files text box, specify the maximum number of DC binary files that can be created in the directory.

One DC binary file is created per one Data Protector medium used for backup. When the medium is rewritten, the old binary file is removed and a new one is created.

8. In the Low space text box, specify the size of the DC directory considered to be close to the maximum size of the directory.

When the difference between used size and maximum size is smaller than the size specified, Data Protector uses the next DC directory specified in the allocation sequence.

9. Click **Finish** to exit the wizard.

Figure 6: DC extension wizard

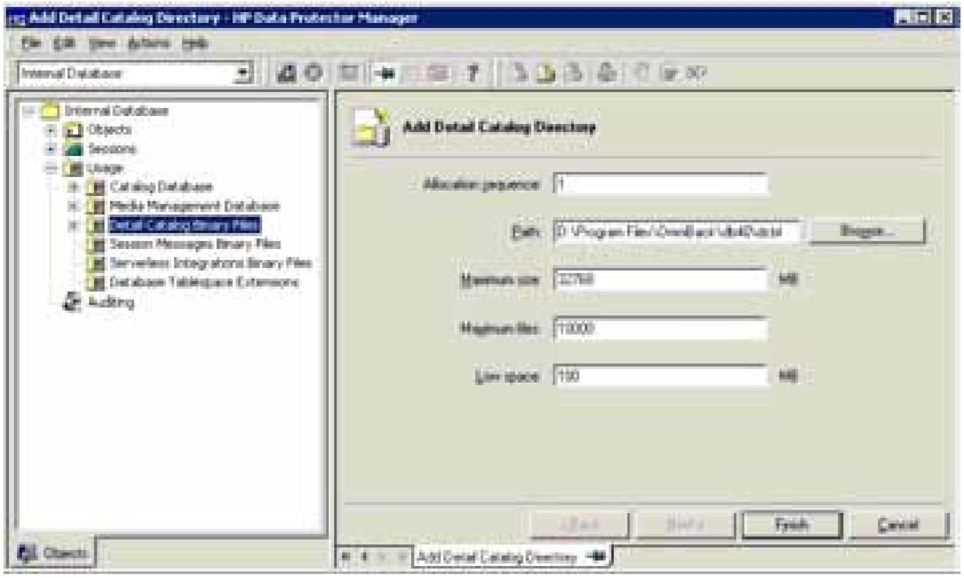

Add Detail Catalog Directory

Allocation sequence: 1

Path: D:\Program Files\OmniBack\db40\dcbf

Maximum size: 32768 MB

Maximum files: 10000

Low space: 150 MB

Add Detail Catalog Directory - HP Data Protector Manager
File
Edit
View
Actions
Help
Internal Database
Internal Database
Objects
Sessions
Usage
Catalog Database
Media Management Database
Detail Catalog Binary Files
Session Messages Binary Files
Serverless Integrations Binary Files
Database Tablespace Extensions
Auditing
Browse...
< Back
Next >
Finish
Cancel
Objects
Add Detail Catalog Directory

Session Messages Binary Files (SMBF)

SMBF records

The Session Messages Binary Files part stores session messages generated during backup, restore, copy, and media management sessions. One binary file is created per session. The files are grouped by year and month.

SMBF size and growth

The SMBF size depends on the following:

- Number of performed sessions
- Number of messages in a session
  One session message occupies approximately 200 bytes on Windows and 130 bytes on UNIX systems. You can change the volume of messages displayed when backup, restore, and media management operations are performed by changing the **Report level** option. This influences the amount of messages stored in the IDB.

SMBF location

The SMBF is located in the following directory:

<Data_Protector_program_data>\db40\msg (Windows Server 2008),

<Data_Protector_home>\db40\msg (other Windows systems), or

/var/opt/omni/server/db40/msg (UNIX systems).

You can relocate the directory by editing the `SessionMessageDir=FullPathToTheMessageDir` variable in the global options file.

Serverless Integrations Binary Files (SIBF)

SIBF records

The Serverless Integrations Binary Files part stores raw NDMP restore data. This data is necessary for restore of NDMP objects.

SIBF size and growth

The SIBF does not grow very big in size. For NDMP backups, the SMBF part grows proportionally to the number of objects backed up. Approximately 3 KB are used for each backed up object.

SIBF location

The SIBF is located in the following directory:

<Data_Protector_program_data>\db40\meta (Windows Server 2008),

<Data_Protector_home>\db40\meta (other Windows systems), or

/var/opt/omni/server/db40/meta (UNIX systems).

Why should you configure the IDB?

Configuring the IDB helps you verify that you have covered all the important aspects of the IDB: locality of the IDB components to meet high availability needs, policies that influence the IDB size and growth, IDB backup configuration, and configuration of all IDB notifications and reports that inform you about the IDB changes and alert conditions. Once you configure the IDB, maintenance is reduced to a minimum, and consists of mainly acting on notifications and reports.

Note:

Additional IDB management details and configuration about the IDB can be found in the HP Data Protector Concepts Guide, which is part of the documentation set of Data Protector software.

Regular IDB backups

An essential part of the Internal Database configuration is configuring the backup of the IDB itself. Once the IDB backup is performed regularly, the most important preparation for recovery in case of a disaster is done. The IDB recovery is essential for restore of other backed up data in case of a Cell Manager crash.

Consider the following when configuring the IDB backup:

- Create a separate backup specification for the IDB backup. This simplifies scheduling and restoring in case of a disk crash. To create an IDB backup specification, select the **IDB** object and follow the standard backup procedure.
- Schedule the IDB backup to be performed once per day. This verifies that you always have an almost up-to-date backup of the IDB. Schedule it to run when there is low activity on the Cell Manager.
- Perform the IDB backup using a separate media pool, on separate media, on a specific device. Use a real tape device or file jukebox—no file library. Media from a file library cannot be imported in case of a Cell Manager crash. Make sure you know which media you use for the IDB backup. You can configure a **Session Media Report** to be informed about the media used for the backup. This greatly simplifies eventual restore. If possible, use a device locally connected to the Cell Manager.
- Set data protection and catalog protection to a few days only. Set these options such that at least the last two IDB backup versions are protected.
- Always have the **Check the Internal Database** option enabled (default).
- Do not overwrite the previous IDB backup with the new one (keeping several copies is suggested).

What Happens During the IDB Backup?

During the IDB backup, Data Protector:

- Checks the consistency of the IDB to prevent backing up and later restoring a corrupted IDB

  For this to happen, you need to have the **Check IDB** option enabled (default). The check operation takes approximately 1.5 hours for a 10-GB database with the `fnames.dat` file size of 1 GB.
- Backs up the IDB online

  Other backup sessions or restore sessions can run during the IDB backup, however, this is not recommended.
- Backs up all the Data Protector configuration data, including the data on backup devices, backup specifications, and schedules

  This simplifies recovery in case of a disaster.
- The IDB backup process creates an `OmniDB` entry in the `media.log` file which allows identifying what tape was used for last IDB backup

  This is extremely useful in case the IDB is lost and needs to be recovered from tape.

Only one IDB backup can run at a time.

Note:
Additional information about IDB backup tasks and IDB recovery methods can be found in the Data Protector software online help. Start the online help and search for "IDB, recovery."

IDB notifications

Data Protector allows you to send notifications from the Cell Manager when specific events occur. For example, when a backup, copy, or consolidation session is completed, you can send an e-mail with the status of the session.

You can set up a notification so that it triggers a report.

You can configure notifications using the Data Protector GUI or any Web browser with Java support.

Input parameters let you customize notifications. Some input parameters allow multiple selections. All other input parameters depend on the type of the notification. Depending on the send method, the recipient can be any of the following:

- A system
- An e-mail address
- An SNMP trap
- A script
- A file
- A configured report group
- The Data Protector Event Log

By default, notifications are configured with default values and are sent to the Data Protector Event Log. To send additional notification using some other sending method and/or other input parameters values, the configuration values must be changed.

To access the Data Protector notification functionality, you either have to be added in the Admin user group or granted the Reporting and notifications user rights.

Example: Configuring an IDB Notification

This is an example of configuring an **IDB Space Low** notification. Once configured, the notification will be triggered when the majority of the allocated space for the IDB is used.

Prerequisite
You have to be either added to the Admin user group, or granted the Reporting, Notifications, and Event Log user right.

Steps:

1. In the Context List, select **Reporting**.
2. Right-click Notifications and then click **Add Notification** to open the wizard.
3. In the Name text box, type a name for the notification, for example, *MyIDBSpaceLow*.
4. In the Event drop-down list, select **IDB Space Low**.
5. In the Send Method drop-down list, select a send method, for example, **Data Protector Event Log**.
6. In the Level drop-down list, select a severity level for the notification, for example **Major**. In the Filename Tablespace Size Limit Threshold [MB] text box, type *200*. In the Disk Free Threshold [MB] text box, type *40*. In the DCBF Size Limit Threshold [MB] text box, type *200*.
7. Click **Finish** to exit the wizard. The notification is listed in the Scoping Pane, under Notifications.

As a result, the notification will be triggered if either the difference between the current and maximum size of all CDB extension files drops below 200 MB, or free space on any of the disks containing the IDB drops below 40 MB, or the difference between the current and maximum size of all DC directories drops below 200 MB.

Existing IDB related notifications

Health Check Failed

Event/notification name	HealthCheckFailed
What triggers the notifications?	A non-zero value returned by the `omnihealthcheck` command. The command returns zero if the following is true: • The Data Protector services (RDS, CRS, MMD, omnitrig, and Inet) are active. • The Data Protector media management database is consistent. • At least one backup of the IDB exists. For more information on this command, refer to the `omnihealthcheck` man page. By default, Data Protector starts the Health Check (which runs the `omnihealthcheck` command) once a day.
Default message level	Critical
Message displayed	Health check message: `<healthcheck_command>` failed.

IDB Corrupted

Event/notification name	IDBCorrupted
What triggers the notification?	Corruption of a part of the IDB
Default message level	Critical
Message displayed	Corruption in the `<IDB_part>` part of the Data Protector Internal Database has been detected (`<error_message>`).

IDB Filename Conversion Needed

Event/notification name	IDBFilenameConversionNeeded
What triggers the notifications?	A need for the IDB filename conversion that may arise during an upgrade from Data Protector A.05.10 or earlier. If the conversion is not performed, you may experience problems when browsing old files for restore. For details, refer to the HP Data Protector Installation and Licensing Guide. By default, Data Protector checks the IDB Filename Conversion Needed condition once a day.
Default message level	Minor
Message displayed	Filenames in the IDB have not been converted since upgrade to A.05.50.

IDB Purge Needed

Event/notification name	IDBPurgeNeeded
What triggers the notifications?	By default, Data Protector checks the IDB Purge Needed condition once a day as a part of checking and maintenance mechanism and triggers the notification if: • The time elapsed since the last IDB filename purge exceeds **Days Last Purge [days]** (default, *180 days*). At the same time, either the number of filename records likely to be purged exceeds **Num. Estimated Filenames [mio]** (default, *6 millions*), or the estimated time needed for the purge exceeds **Estimated Time Purge [min]** (default, *120 minutes*). • The number of filenames in the IDB exceeds **Num. Filenames [mio]** (default, *100 millions*).
Default message level	Warning
Message displayed	Filename purge should be run for Data Protector Internal Database.

IDB Space Low

Event/notification name	IDBSpaceLow
What triggers the notifications?	One of the following events: • The difference between the maximum and current size of all CDB extension files drops below **Filename Tablespace Size Limit Threshold [MB]** (default maximum for `filenames.dat` file, 250 MB) • Free space on any of the disks containing the IDB drops below **Disk Free Threshold [MB]** (default, 50 MB). • The difference between the maximum and current size of all DC directories drops below **DCBF Size Limit Threshold [MB]** (default, 250 MB). By default, Data Protector checks the IDB Space Low condition once a day.
Default message level	Major
Message displayed	Data Protector Internal Database is running out of space.

IDB Tablespace Space Low

Event/notification name	IDBTablespaceSpaceLow
What triggers the notifications?	A lack of free space for the IDB tablespaces. More than **Tablespace Used Threshold [%]** (default, *85%*) of the allocated space is already used. By default, Data Protector checks the IDB Tablespace Space Low condition once a day.
Default message level	Major
Message displayed	Tablespace `<Tablespace_name>` is running out of space.

Note:
It is mandatory to check the Data Protector software event log on a regular basis and check for eventual IDB events. An administrator might consider setting up notifications sent by e-mail allowing prompt action on incoming notifications.

Note:
Additional information about IDB maintenance tasks is available in the white paper "HP Data Protector IDB Purge Best Practices," which can be downloaded from http://www.hp.com/go/dataprotector

Limitations

Internal Database Size

		Data Protector 6.0	**≥ Data Protector 6.1**
CDB	CDB max. size	32 GB	48 GB
	Max # File names Unix (est.)	700 Million	1050 Million
	Max # File names Windows (est.)	450 Million	674 Million
	Text format of all strings including filenames of Unix	ASCII, single- or multi-byte format	ASCII, single- or multi-byte format
	Text format of all strings including filenames of Windows	UNICODE, double-byte format	UNICODE, double-byte format

DCBF **200 GB–1.6TB** **1-50 Stores**	Max number of File Versions	10x # of file names	10x # of file names
	Max number of DCBF directories (containing binary files)	Default: 10 Maximum: 50	Default: 10 Maximum: 50
	Max size per DCBF directory	Default: 4 GB Maximum: 32 GB	Default: 16 GB Limited by file system settings
	Max number of files per DCBF directory	10.000 files	10.000 files
	Max size per DCBF file	2 GB	Limited by file system settings
	Min size per DCBF directory	100 MB	2 GB

MMDB **20-50MB**	Max media per pool	40.000	40.000
	Max media per Cell Manager	500.000	500.000
	Max sessions per cell manger	1.000.000	1.000.000
	Max sessions per day	2.000	9999
	Max parallel backup sessions Unix	100	100
	Max parallel backup sessions Windows	60	60

Number of Backups Scheduled at One Time

The maximum total number of backup sessions running in parallel is 100 on UNIX systems and 60 on Windows systems. The default value is set to five. This can be increased using the `MaxBSessions` global option. When the number of parallel sessions is larger than 50 (recommended maximum) the probability of hitting one of the system limits on the Cell Manager increases significantly (number of file descriptors, TCP/IP limitations, memory limitations).

Concurrent Activities

Each backup session can by default use up to 32 devices at the same time. The upper limit for this parameter is controlled by the `MaxMAperSM` global option (default = 32).

- By default, up to 32 Disk Agents (depending on the concurrency of a device) can write to the same device at the same time. This number can be controlled using the `MaxDAperMA` global option.
- Up to 10 media can be imported in the IDB at the same time.

Note:
It is advisable to not increase the default number of 32 concurrent devices as high memory utilization on the backup host might occur.

Number of Cells in a MoM (Manager of Manager) Environment

There can be up to 50 cells in a MoM environment.

IDB Growth and Performance

For Internal Database configuration and maintenance, you must understand the key factors and parameters that influence IDB growth and performance.

The IDB can grow very big and can have a significant impact on backup performance and the Cell Manager system. The Data Protector administrator has to understand the IDB and decide which information to keep in the IDB and for how long. It is the administrator’s task to balance restores time and functionality with the size and growth of the IDB. Data Protector offers two key parameters, logging level and catalog protection that assist you in balancing your needs.

The data given here is applicable for filesystem backups and illustrates the worst-case scenario (largest or fastest growing IDB). If you perform disk image, online database, or NDMP backup, a smaller amount of data is stored in the IDB.

IDB key growth factors

IDB growth depends on your environment and on Data Protector settings that define how much history and detail you want Data Protector to keep to allow for browsing and search of files.

Key factors	Impact on IDB growth
Details about files and size of the environment	Data Protector can keep track of each file and each version of the file. This means that during each backup, one file version record (up to 30 bytes) will be stored to the DCBF part for each backed up file. Filenames are stored only once (approximately 50–80 bytes per file) and if the filesystem dynamics are low, then the filenames part of the CDB will only grow in proportion to the number of files in your environment. After the first full backup, the size of the filenames part does not grow significantly.
Frequency of (full) backups	The more often you do a backup, the more information is stored in the IDB. If the file system dynamics are low then only the DCBF part will grow.
File system dynamics	The number of files created and removed between backups can have a significant impact on the growth of the filenames part of the IDB. During each backup, all new filenames will be stored in the filenames part. The Report on System Dynamics can give you information about this.
Number of object copies	The more object copies and object mirrors you create, the more information is stored in the IDB. For object copies and object mirrors, the IDB stores the same information as for backed up objects, except for filenames.

IDB key performance factors

Key factors	Impact on IDB load/performance during backup
Number of parallel drives	The number of (tape) drives running in parallel impacts the load on the IDB. If, for example, 10 drives are running in parallel in 10 backup sessions or 10 drives are running in parallel in five sessions, there is almost the same load on the database. Each new drive means another source of file catalogs that must be stored in the database.
Average file size	If small files are backed up, file catalogs are generated faster and load for the IDB is consequently higher.
IDB disk performance	The main Data Protector activity during backup is reading and writing from disk. Therefore, the speed of the disk (subsystem) on the Cell Manager used for the IDB can influence the performance.

IDB key growth and performance parameters

Key parameters	Impact on IDB growth	Impact on IDB performance
Logging level	Defines how much data about files and directories is written to the IDB.	Influences the backup speed and the convenience of browsing data for restore.
Catalog protection	Defines how long information about backed up data (such as filenames and file versions) is kept in the IDB. If the catalog protection expires, data is not removed from the IDB immediately. It is removed on the same day when all the catalog protection for data on the entire media expires.	

The simplified graph in Figure 3 presents the difference in IDB growth when catalog protection is set for a relatively short period of time (one month) versus when the catalog protection is the same as data protection (3 years). Also, the difference in usage of the **Log all** or **Log directories** options is shown. The major growth of the IDB lasts until the catalog protection has been reached. After that, the growth is low and determined by the growth of the backup environment.

Figure 7: IDB growth

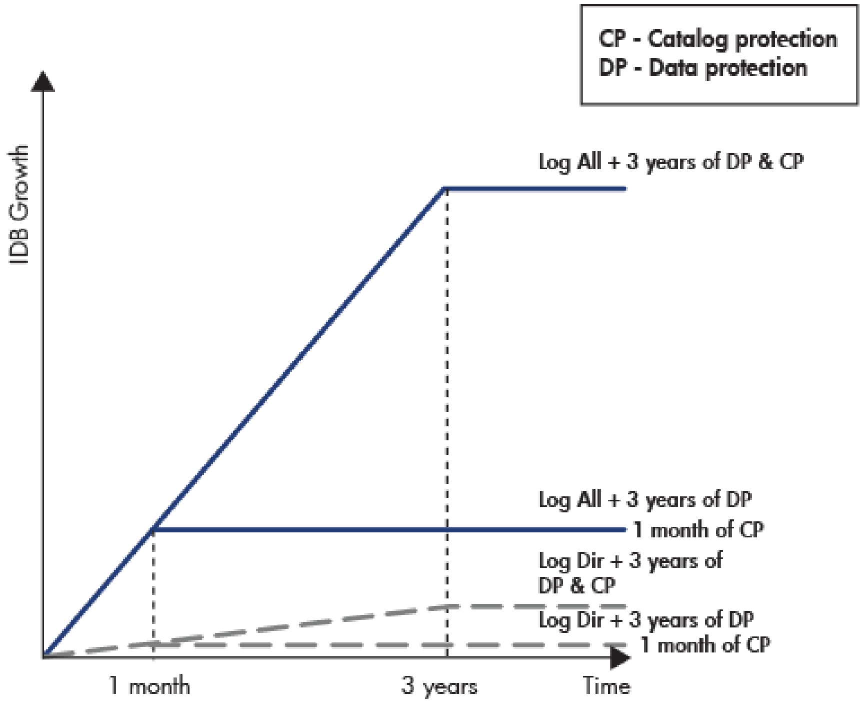

Influence of logging level on IDB

The different logging level settings influence the Internal Database growth, the convenience of browsing filesystems for restore, and, in minor cases, backup performance.
The data provided below applies to filesystems backups. If you perform disk image, online database or NDMP backup, a small amount of data is stored in the IDB.

No log	Only object information is stored, typically 2 KB per filesystem object.
Log directories	Same as **No log**, and in addition, 30 bytes per backed up directory are stored.
Log files	Same as **Log directories**, and in addition, 12 bytes per backed up file are stored.
Log all	Same as **Log files**, and in addition, 18 bytes per backed up file are stored.

Option value: Log all

This is the default logging level. All detailed information about backed up files and directories (names, versions, and attributes) is logged to the IDB.
You can browse directories and files before restoring and in addition look at file attributes. Data Protector can fast position on the tape when restoring a specific file or directory

Option value: Log files

When this logging level is selected, detailed information about backed up files and directories (names and versions) is logged to the IDB.

You can browse directories and files before restoring, and Data Protector can fast position on the tape when restoring a specific file or directory. The information does not occupy much space, since not all file details (file attributes) are logged to the database.

Option value: Log directories

When this logging level is selected, all detailed information about backed up directories (names, versions, and attributes) is logged to the IDB.

You can browse only directories before restoring. However, during the restore Data Protector still performs fast positioning because a file is located on the tape near the directory where it actually resides. This option is suitable for filesystems with many auto-generated files, such as news and log files.

Option value: No log

When this logging level is selected, no information about backed up files and directories is logged to the IDB. You will not be able to search and browse files and directories before restoring

Example: Changing of logging level for filesystem backup

The logging level determines the volume of detail on files and directories written to the IDB during backup, object copy, or object consolidation sessions, and the convenience of browsing data for restore.

Note:
You can restore your data regardless of the logging level used.

The default logging level is **Log All**

Note:
Logging levels are not available for disk image backup.

You can modify an already-configured and saved backup specification.

Steps:

1. In the Data Protector software GUI Context List, click **Backup**.
2. In the Scoping Pane, expand Backup Specifications and then expand the appropriate type of backup specification (for example, *Filesystem*). All saved backup specifications are displayed.
3. Click the backup specification that you want to modify.
4. In the Options property page click on **Advanced** in the Filesystem Options section.
5. In the Filesystem options box, click on **Other**.
6. Select your desired logging level in the **Logging** drop-down menu
7. Click on **Ok** and then on **Apply**.

Figure 8: Filesystem backup logging level

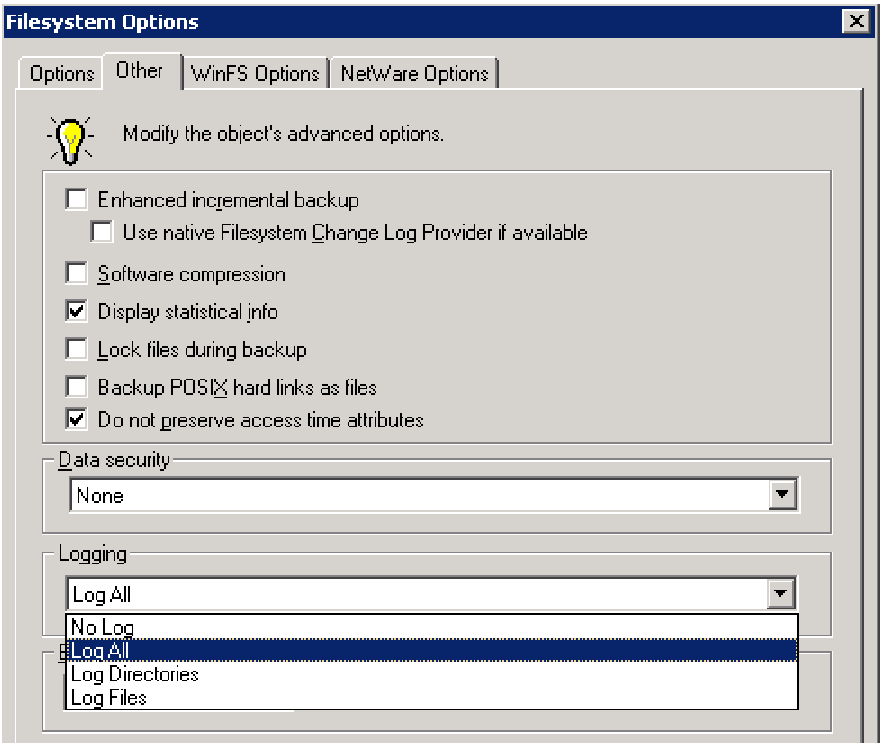

Modify the object's advanced options.

☐ Enhanced incremental backup

☐ Use native Filesystem Change Log Provider if available

☐ Software compression

☑ Display statistical info

☐ Lock files during backup

☐ Backup POSIX hard links as files

☑ Do not preserve access time attributes

Data security: None

Logging: Log All

No Log
Log All
Log Directories
Log Files

Influence of catalog protection on IDB

The largest part of the Internal Database is proportional to the catalog protection period multiplied by the chosen logging level. The more backups are performed within the catalog protection period, the more data accumulates in the IDB. In other words, it multiplies the data needed to store each file version by as many file versions as are backed up during the catalog protection period.

Once the catalog protection expires, the information is not immediately removed from the IDB. Data Protector removes it automatically once per day. Since the information in the IDB is organized on a per-medium basis, it is removed only when the catalog protection expires for all objects on the medium. If so, the entire space occupied by the specific DC binary file becomes free.

You should set the catalog protection such that it includes at least the last full backup. For example, you can set a catalog protection of eight weeks for full backups and one week for incremental backups.

Example: Changing of catalog protection of a backup

Catalog protection determines how long the information about the backed up data is kept in the IDB.

Note:
If there is no catalog protection, you can still restore your data, but you cannot browse for it in the Data Protector GUI.

- **None:** Provides no protection.
- **Until:** Means that the information in the IDB cannot be overwritten until the specified date. Protection for the information stops at noon on the chosen day.
- **Days:** Means that the information in the IDB cannot be overwritten for the specified number of days.
- **Weeks:** Means that the information in the IDB cannot be overwritten for the specified number of weeks.
- **Same as data protection:** Means that the information about backed up data in the IDB is protected as long as the data is protected.

You can modify an already configured and saved backup specification.

Steps:

1. In the Data Protector software GUI Context List, click **Backup**.
2. In the Scoping Pane, expand Backup Specifications and then expand the appropriate type of backup specification (for example, *Filesystem*). All saved backup specifications are displayed.
3. Click the backup specification that you want to modify.
4. In the Options property page click on **Advanced** in the Filesystem Options section.
5. In the Filesystem options box, click on **Options**.
6. Select your desired Catalog Protection duration in the **Catalog Protection** drop-down menu
7. Click on **Ok** and then on **Apply**.

Figure 9: Filesystem backup catalog protection

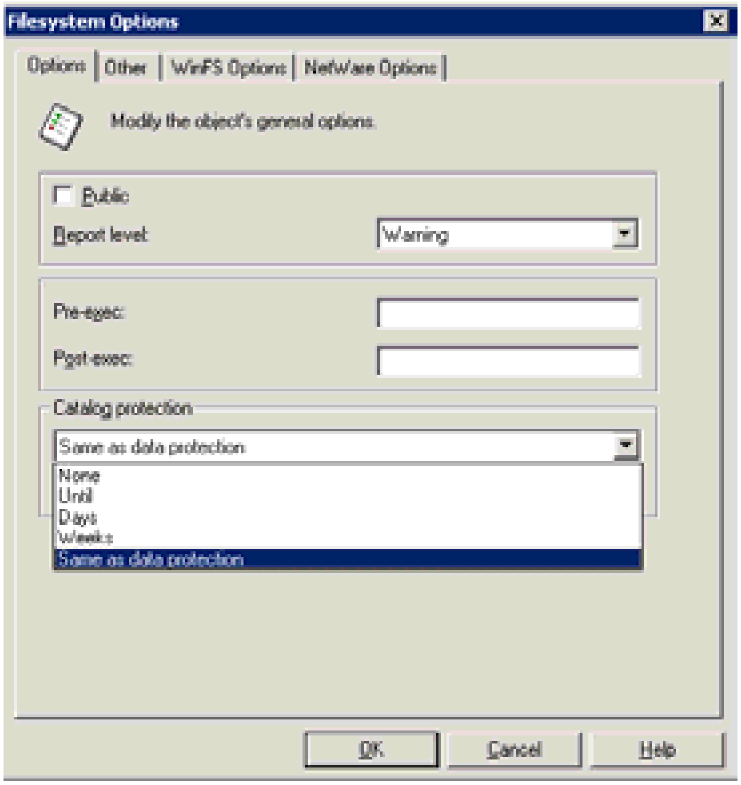

Modify the object's general options.

□ Public

Report level: Warning

Pre-exec:

Post-exec:

Catalog protection: Same as data protection

Recommended usage of logging level and catalog protection

Always set a reasonable level of catalog protection. The only exception is if the Log None option is set (in this case, catalog protection does not apply anyway).

If you set the catalog protection to Permanent, the information in the IDB is removed only when media are exported or sessions are deleted. In this case, the size of the IDB grows linearly until the data protection period is reached, even if the number of files in the cell does not change. For example, if the data protection period is one year and media are recycled, then significant growth of the IDB stops after one year. The addition of new catalogs is approximately equal to the removal of old ones. If catalog protection is set for four weeks, significant growth of the IDB stops after four weeks. Therefore, in this case, the IDB is 13 times larger if the catalog protection is set to one year.

It is recommended that catalog protection includes at least the last full backup. For example, you can set a catalog protection of eight weeks for full backups and one week for incremental backups.

Use different logging levels in the same cell

A cell often consists of mail (or similar) servers that generate a large number of files on a daily basis, database servers that store all information in a handful of files, and some user workstations. Since the dynamics of these systems are rather different, it is very difficult to prescribe one setting that suits them all. Therefore, it is recommended to create several backup specifications with the following logging level settings:

- For database servers, no logging is necessary as they have their own restore policies. Therefore, use the **No Log** option.
- For workstations/file servers, the **Log All** or **Log Files** options allow for searching and restoring different versions of files. For backups with the **Log Directories** or **No Log** options set, you can import catalogs from the media, which, in a reasonably short time, allows the possibility to browse for the selected object. For information on importing catalogs from media, refer to online help, index keyword “importing, catalogs from media.”

Different logging levels for ObjectCopies

Backed up objects and object copies or mirrors of these objects can have the same or different logging levels. Depending on your backup policy, the selected logging level of object copies can be more or less detailed than that of the source objects.

For example, you can specify the **No Log** option for object mirrors if you create these mirrors just to provide a successful completion of a backup session. Or, you can specify the **No Log** option for a backup object to increase the backup performance, and then specify the Log All option for this object in a subsequent object copy session.

Specifics for small cells

If the number of files in a cell is small and will remain small (a million files or less) and the systems in the cell perform usual business activities, you can always use the Log All option, which is the Data Protector default. However, you need to take care of IDB growth and set a reasonable level of catalog protection.

Specifics for large cells

If the number of files grows into the tens of millions, or there are tens of thousands of files generated on a daily basis, and you use the Log All option, then backup speed and IDB growth will become a problem in a relatively short period of time.

In this situation, you have the following options:

- Reduce the logging level to the smallest acceptable level. Setting the Log Files option can reduce the IDB size to a third, and setting the Log Directories option to almost a tenth. This, of course, depends on the nature of the file systems in the cell.
- Reduce the catalog protection to a minimum.
- Split the cell in two. As a final solution, you can always introduce another IDB and redirect half of the systems into it.

You can configure report on system dynamics, which informs you about dynamics of the growth of filenames on a particular client.

Maintenance of DCBF directories

The IDB allows several directories to be registered where DC binary files are stored. The purpose of this is to allow the binary files to be distributed over more disks/filesystems. By default, there is only one directory: `../db40/dcbf`

Each DCBF directory has several configuration attributes. Those are:

MaxSize: specifies what amount of disk space can be used for DCBF in this directory. When this size is reached, new DCBF files will not be created any more.

MaxFiles: specifies how many DCBF files can be stored in the directory.

SpaceLow: Low disk space required before an additional DCBF file can be created.

Whenever there is a need to create a new binary file, the “DCBF allocation procedure” needs to be performed by Data Protector.

First, from the list of all possible DCBF directories, Data Protector will eliminate all that are de-activated or missing. Note that in the case of missing DCBF directory, an `OBDB_Corrupt` event will be generated.

Then, all full DCBF directories will be not considered. DCBF directory is considered full if at least one of the following conditions is true:

```
MaxSize-CurrentSize<SpaceLow
```

```
DiskSpace<SpaceLow
```

```
MaxFiles<=CurrentFiles
```

Then a set of user selectable algorithms will select the actual DCBF directory (global option per cell).

1. Fill in sequence

   HP Data Protector always tries to create the new DCBF file in the first DCBF directory according to the sequence.

2. Balance size

   HP Data Protector will select the DCBF directory that contains (proportionally to size limit) the least DCBF file size. Minimum for the following value is selected:
   `(MaxSize-CurrentSize-SpaceLow)/(MaximumSize-SpaceLow)`

3. Balance number

   Data Protector will select the DCBF directory that contains (proportionally to size limit) the least DCBF file number. Minimum for the following value is selected: `(CurrentFiles/MaximumFiles)`

Variables influencing DCBF behavior are located in the global options file

```
DCDirAllocation=0, 1, 2
```

Default: `0` (1 for Data Protector 6.1)

This global option controls which algorithm will be used to select the directory for the creation of the new DCBF file.

0: Fill in sequence **(default)**

1: Balance size **(recommended)**

2: Balance number

```
MaxDCDirs=NumberOfDirectories
```

Default: `10`

Minimum: 1

Maximum: 50

This option specifies maximum number of configured DCBF directories.

Note:
The default minimum and maximum size for existing DCBF directories that were created with a Data Protector version up to the version 6.0 are not changed automatically to the new supported size with an upgrade to HP Data Protector 6.1.

But it is possible/recommended to change them manually with the `omnidbutil` command:

```
omnidbutil –modify_dcdir <directory> -maxsize 16384
```

`omnidbutil –modify_dcdir <directory> -spacelow 2048`

A manual change of existing dc directories' size is needed in case drives with larger capacities are used (for example, LTO-4) and lots of files are backed up to tape (>10 million)

The `spacelow` value is enforced by HP Data Protector and is checked during every backup session. If available space is less than the sum of all the configured `spacelow` values, the logging is switched to `nolog`.

For example: Configured `spacelow` value = 2048 MB, configured DCBF directories = 10 → required minimum free space = 20 GB

Cell Manager hardware aspects to consider

MS Windows, Red Hat Linux, Suse ES, and Sun Solaris10

Organizations are under growing pressure to support non-stop business operations without increasing the IT budget. Given their position at the center of the IT infrastructure, servers play a critical role in determining both the overall availability as well as the total cost of ownership of that infrastructure. This certainly applies also to an important infrastructure component such as a HP Data Protector Cell Manager server.

Today's business requirements demand IT organizations deploy servers that are designed to be powerful, reliable, and highly available. Leading technologies like redundant Ethernet connections, network adapter teaming, memory management, and Smart Array controllers enable improved flexibility and broad software compatibility. These innovations offer IT professionals the tools to address high-availability platform challenges effectively, both now and in the future.

To enable high availability, a server must proactively respond to failures, both within the server and the network to which it is attached. High availability requires that a system administrator manage a server without having to be co-located with that server, and that the server remains available while a failed component is repaired.

The first choice for a Cell Manager server on MS Windows, Red Hat Linux, Suse ES, and Sun Solaris10 combining all of the above mentioned characteristics is a server from the HP ProLiant server family.

At the time this white paper was compiled a HP ProLiant DL380 G5 (Generation 5) server was providing all of the necessary features providing enough computing power and high availability.

A non-racked tower version (ML370 G5) of this model also exists.

The features in detail are:

- Intel® Xeon™ Processors
- Advanced Memory Protection
- Redundant Cooling and Power
- RAID Technology
- Redundant Ethernet Connections
- Network Adapter Teaming
- ProLiant Essentials
- Intelligent Networking Pack
- Integrated Lights-Out Technology

QuickSpecs for the HP ProLiant DL380 G5 server can be found at:
http://h18004.www1.hp.com/products/quickspecs/12477_na/12477_na.pdf

Performance pack

The HP ProLiant DL380 G5 performance packs are necessary to configure a HP ProLiant DL380 G5 Performance Model. HP ProLiant DL380 G5 Performance model includes the latest performance technologies and enterprise class availability features pre-installed for convenience and value. The performance pack includes a second processor and a redundant power supply.

Processor

It is advisable to equip the Cell Manager server minimum with two Quad-Core Intel® Xeon® Processors E5345 (2.33 GHz, 80 Watts, 1333 FSB) providing required processing power at relatively low power consumption.

Memory

ProLiant servers use a variety of techniques to protect against memory errors. It is advisable to equip the Cell Manager server with Fully-Buffered DIMM technology.

To improve memory protection even further, HP introduced Advanced ECC technology. Advanced ECC technology is capable of correcting a multi-bit error that occurs within one dynamic random access memory (DRAM) chip.

The ProLiant server online spare memory determines if an active DIMM (dual inline memory module) exceeds a predefined error threshold. The error will be corrected and the data from the entire bank that contains the failed DIMM will be copied to online spare memory. The failed bank is deactivated, but the server will remain available until the failed DIMM is replaced during a scheduled shutdown.

Whereas online spare memory mode protects against single-bit errors and entire DRAM failure, mirrored memory mode enables full protection against single-bit and multi-bit errors. In mirrored memory mode, the same data is written to both system memory and mirrored memory banks, but data is read only from the system memory banks. If a DIMM in the system memory banks experiences a multi-bit error or reaches the pre-defined error threshold for single-bit errors, the roles of the system and mirrored memory banks are reversed.

HP is one of the first companies to support hot plug RAID memory, which allows the memory subsystem to operate almost continuously even in the event of a complete memory device failure. In this context, RAID stands for Redundant Array of Industry-standard DIMMs.

Hot plug RAID memory generates parity for an entire cache line of data during write operations and records the parity information on a dedicated parity cartridge. However, hot plug RAID memory does not have the mechanical delays of seek time, rotational latency, and bottlenecks associated with disk drive arrays.

To achieve maximum memory throughput, it is advisable to follow specific memory bank population rules.

Further information can be found in the White Paper “Fully-Buffered DIMM technology in HP ProLiant servers” (http://bizsupport.austin.hp.com/bc/docs/support/SupportManual/c00913926/c00913926.pdf).

Disk technology

It is advisable to use Enterprise class disc technology. Enterprise-class drives are designed for high reliability, high performance, scalability, and error management under heavy 24x365 I/O workloads. They are intended for mission-critical applications. Enterprise-class drives are the only class of drives intended for unlimited I/O loads.

SAS-Serial attached SCSI is usually considered the most cost-effective solution for mission critical, high I/O workload applications, such as business critical databases.

Further information can be found in the White Paper “Disk drive technology overview” at http://bizsupport.austin.hp.com/bc/docs/support/SupportManual/c01071496/c01071496.pdf

HP was instrumental in developing the SAS standard. For a more detailed discussion of SAS, please see the Technical Brief entitled, “Serial Attached SCSI technology” at http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00302340/c00302340.pdf

I/O subsystem

Advanced controllers (such as the HP Smart Array) decouple the logical disks seen by applications from the physical devices used to implement the disk subsystem. A single logical disk (as seen by an application) may be mapped onto an array of multiple physical disks. These controllers include both hardware and software.

This approach provides greatly enhanced flexibility, expandability, maintainability, and performance. Smart Array controllers are available for SAS, SATA, and SCSI interfaces.

For further details, see the technology brief entitled "HP Smart Array Controller Technology," at http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00687518/c00687518.pdf.

HP Smart Array controllers are integrated on ProLiant servers, and support a variety of RAID types including RAID 1+0 and RAID 5. Developed and patented by HP, RAID 6 Advanced Data Guarding (ADG) is further supported on Smart Array controllers. This technology creates two sets of parity striped data across the disks to help enable the system can withstand multiple disk failures without data loss. RAID ADG enables high levels of fault tolerance in a cost-effective manner.

It is advisable to separate operating system from the Cell Manager application and IDB by configuring separate c:\ and d:\ partitions on different drives running in RAID1+0 mode. This configuration provides mirroring and striping, redundancy, and performance improvement.

Network adapter teaming

Network adapter teaming is software-based technology used to increase a server's network availability and performance. Teaming enables the logical grouping of physical adapters in the same server (regardless of whether they are embedded devices or Peripheral Component Interconnect (PCI) adapters) into a virtual adapter. This virtual adapter is seen by the network and server-resident network-aware applications as a single network connection.

Most HP ProLiant Storage Servers are equipped with the HP Network Configuration Utility (NCU). The utility allows administrators to configure and monitor Ethernet network interface controller (NIC) teams in a Windows-based operating system. These teams provide options for increasing fault tolerance and throughput. The NCU is also used for configuring and monitoring individual network adapters.

Fault tolerance provides automatic redundancy. If the primary NIC fails, the secondary NIC takes over. Load Balancing provides the ability to balance transmissions across NICs.

HP-UX, Itanium

The first choice for a Cell Manager server in a Data Protector cell on HP-UX 11i v2/v3 combining the same/similar characteristics than a HP ProLiant server is a server from the HP Integrity server family featuring Intel® Itanium® processor.

At the time this white paper was compiled a HP Integrity rx3600 server was providing all of the necessary features providing enough computing power and high availability.

An optional stand-alone pedestal mount, field installation only, also exists.

The features in detail are:

- Intel® Itanium® Processors
- Advanced Memory Protection
- Redundant Cooling and Power
- RAID Technology
- Redundant Ethernet Connections
- Automatic de-configuration of memory and processors
- Service processor to monitor system status
- Integrated Lights-Out Technology
- HP Integrity Essentials

QuickSpecs for the HP Integrity rx3600 server can be found at:
http://h71028.www7.hp.com/ERC/downloads/4AA0-6529ENW.pdf

Cluster support for Data Protector software Cell Manager

As a part of its high-availability, Data Protector provides integration with a number of cluster technologies. For details on supported operating system versions, level of cluster support and for supported configurations, refer to the HP Data Protector Product Announcements, Software Notes, and References and the HP Data Protector Platform and Integration Support Matrices at: http://www.hp.com/go/dataprotector.

Note:
Complementary information how to configure and maintain HP Data Protector software in a Microsoft cluster environment can be found in the white paper " OpenView Storage Data Protector 5.5 in Microsoft Windows Server 2003 Cluster Server" at: http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00669082/c00669082.pdf

Security

Cell Manager security

The Cell Manager security is important because the Cell Manager has access to all clients and all data in the cell.

Security of the Cell Manager can be enhanced through the Strict IP Checking functionality. However, it is important that the Cell Manager is also secured as a client and that Data Protector users are configured carefully.

While it may not always be necessary to secure each and every client in the cell, it is important that the computers that other clients will trust are secured themselves. These are besides the Cell Manager also the Installation Server and Media Agent clients.

Client security

After you have installed the HP Data Protector clients and imported them into a cell, it is highly recommended to secure them.

HP Data Protector agents installed on the clients in the cell provide numerous powerful capabilities, like access to all the data on the system. It is important that these capabilities are available only to the processes running on cell authorities (Cell Manager and Installation Servers), and that all other requests are rejected.

HP Data Protector allows you to specify from which cell authorities a client will accept requests on the HP Data Protector port (default 5555). For activities such as backing up and restoring, starting pre- and post-exec commands, or importing and exporting clients, the client checks if the computer, which triggers one of these tasks through the HP Data Protector port, is allowed to do so. Other computers are not able to access such a client.

User security

HP Data Protector Users is another security-critical layer of Data Protector. The configuration of users must be carefully planned and tested.

Some user rights are very powerful and therefore represent a security issue. For example, the User configuration and Clients configuration user rights enable a user to change the security settings. The Restore to other clients user right is also very powerful, especially if combined with either the Back up as root or Restore as root user rights.

Even less powerful user rights bear an inherent risk associated with them. Data Protector can be configured to restrict certain user rights to reduce these risks.

Firewall support

You can configure Data Protector in an environment where the HP Data Protector processes communicate across a firewall.

Communication in Data Protector
Data Protector processes communicate using TCP/IP connections. Every Data Protector system accepts connections on port 5555 by default. In addition, some processes dynamically allocate ports on which they accept connections from other Data Protector processes.

To enable Data Protector processes to communicate across a firewall, Data Protector allows you to limit the range of port numbers from which dynamically allocated ports are selected. Port ranges are defined on a per-system base. It is possible to define a port range for all Data Protector processes on a specific system, as well as to define a port range for a specific Data Protector agent only.

Configuration mechanism
You can configure the port allocation behavior through two `omnirc` variables:

- `OB2PORTRANGE`

This option limits the range of port numbers that Data Protector uses when allocating listen ports dynamically. This option is typically set to enable the administration of a cell through a firewall. Note that the firewall needs to be configured separately and that the specified range does not affect the Inet listen port.

- `OB2PORTRANGESPEC`

This option allows you to specify a range of port numbers for every binary. This mechanism gives you more control over the ranges and helps to keep their sizes smaller. Note that the firewall needs to be configured separately and that the specified range does not affect the Inet listen port.

By default, both variables are not set and ports are assigned dynamically by the operating system.

Note:
Details about how to configure the different security aspects are available in the HP Data Protector software online help.
Open the online help menu and enter:
About Security or About Firewall
in the search field to get to the relevant topics.

IDB space consumption example

Due to the complex nature of IT environments in which backup is a requirement it is quite difficult to find a basic formula allowing to precisely estimating the size of the IDB of a Data Protector Cell Manager.

The complexity derives from the following most influencing details:

- Filesystems dynamics
- Large number of object copies
- Filesystem versus online database backup
- Long-term backup data retention policies
- Long-term catalog retention policies
- High detail logging policies
- Number of backups

A basic formula could be:

IDBsize = MMDB+CDB (obj+pos)+CDB (Fnames) +DCBF+SMBF+SIBF

Note:
CDB (obj + pos) is the size of the CDB part without filenames.

To make calculations simpler and more transparent, the backup environment in this example is a bit simplified. The estimation assumes that there is only one backup specification and that this specification is created for backing up filesystems.

Once you estimate the size of the IDB in the simplified environment, you can repeat the estimation for other backup specifications and then sum up your results.

In the following scenario, a backup specification (datalist) with the configuration details listed in table 1 captures data from twenty objects.

Note:
A backup object is a backup unit that contains all items backed up from one disk volume (logical disk or mount point). The backed up items can be any number of files, directories, or the entire disk or mount point. Additionally, a backup object can be a database entity or a disk image (rawdisk).

A backup object is defined by:

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

Table 3: Backup datalist details and backup specifics

Sum of Files to backup:	5,0	Million
Files per directory:	500	
Total Backup Capacity:	10	GB
Objects:	20	
% of incremental Backup Change	5,00%	
Device concurrency:	16	
Log level:	All	
Number of copies	0	
Data protection:	52	Weeks
Catalog protection:	4	Weeks
Full backups/week:	1	
Incr backups/week:	40	

Figure 10 is a graphical representation of the above-described scenario. It shows in the blue box to the left which parameter in the discussed datalist influences the IDB size/growth the most.

To estimate the size of the IDB use the Internal Database Capacity Planning Tool located at:

- On the UNIX Cell Managers:
  /opt/omni/doc/C/IDB_capacity_planning.xls
- On the Windows Cell Manager:
  <Data_Protector_home>\docs\IDB_capacity_planning.xls

You can also use this tool to estimate the size of the IDB in environments with online databases (Oracle, SAP R/3).

Figure 10: Backup scenario

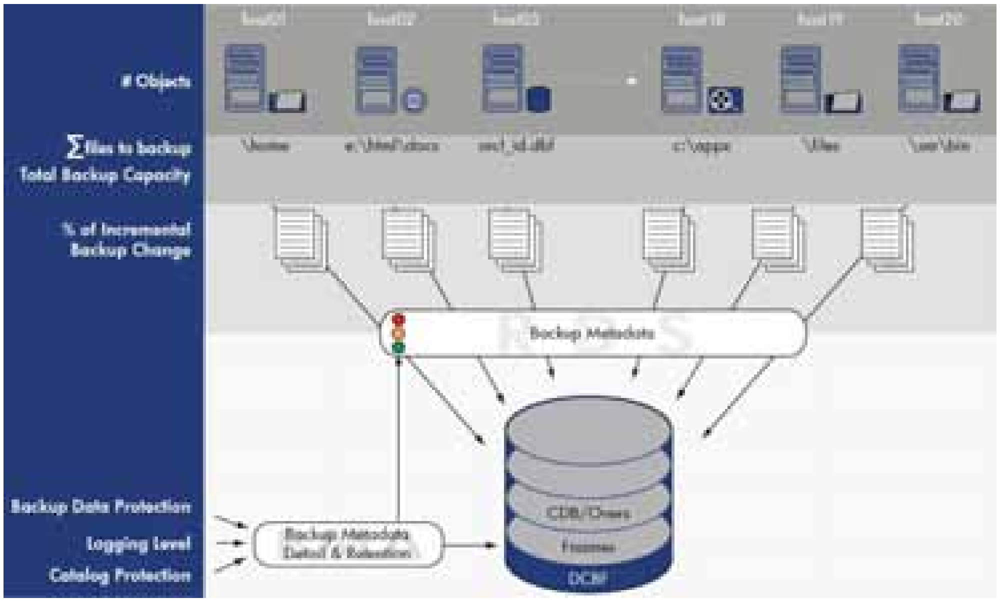

Using the Internal Database Capacity Planning Tool allows to calculate the estimated IDB growth by making changes to the discussed parameters.

Figure 11 shows the IDB growth after changing the Logging level from "None" over "Dir" and "File" and at last "All."

Figure 11: IDB Growth

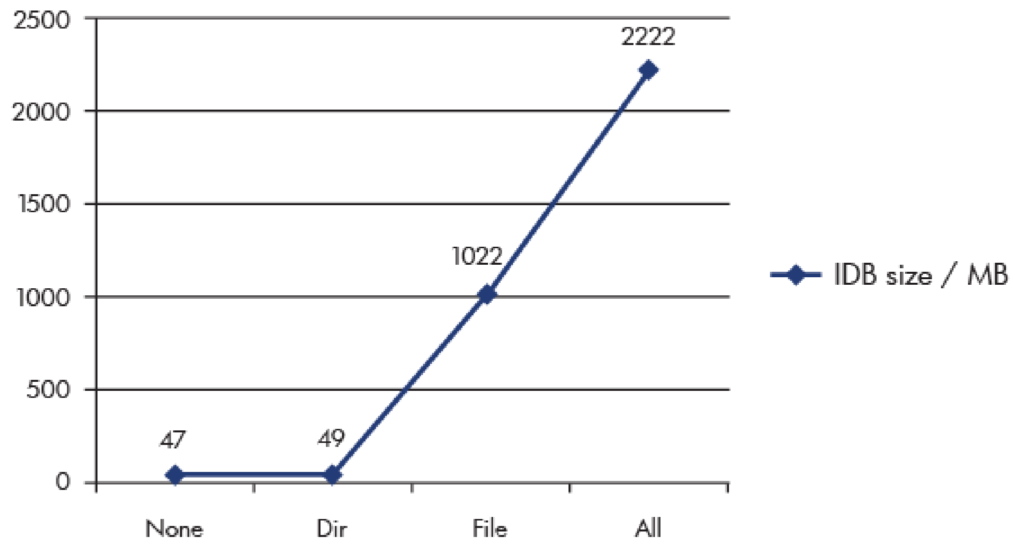

It becomes quite obvious that a careful selection of the datalist configuration parameters/details is recommended to manage the database growth effectively. Please follow the recommendations given in the "IDB key growth and performance parameters" paragraph.

Cell Manager requirements

On systems running HP-UX

The Cell Manager must meet the following minimum requirements:

- The Soft File Limit (maxfiles) per Process on the Cell Manager should be at least 1024.
- 256 MB of RAM (512 MB recommended)

For each parallel backup session, 40 MB of RAM are required and 5–8 MB per data segment size. This means that, for example, if you want to run 60 parallel backup sessions 3 GB of RAM plus 512 MB for data segments are needed.

- 300–425 MB of disk space + approximately 2% of planned data to be backed up (for use by the IDB).
- It is recommended to modify the kernel parameters as follows:
  - set `maxdsiz` (Max Data Segment Size) or `maxdsiz_64` (for 64-bit systems) to at least 134217728 bytes (128 MB).
  - set `semmnu` (Number of Semaphore Undo Structures) to at least 256.

On systems running Solaris

The Cell Manager must meet the following minimum requirements:

- 256 MB of RAM (512 MB recommended)

For each parallel backup session, 40 MB of RAM are required and 5–8 MB per data segment size. This means that, for example, if you want to run 60 parallel backup sessions 3 GB of RAM plus 512 MB of data segments are needed.

- 300–425 MB of disk space + approximately 2% of planned data to be backed up (for use by the IDB)
- The following values of kernel parameters are recommended: SEMMNI (maximum number of semaphore sets in the entire system) = 100 SEMMNS (maximum semaphores on the system) = 256

A system restart is necessary for kernel changes to take effect.

On systems running Windows 2000 or Windows XP

The Cell Manager must meet the following minimum requirements:

- 256 MB of RAM (512 MB recommended).

For each parallel backup session, 40 MB of RAM are required. This means that, for example, if you want to run 60 parallel backup sessions 3 GB of RAM are needed.

- On Windows 2000 systems, Service Pack 3 or later must be installed.
- On Windows XP Professional systems, Service Pack 1 must be installed.
- 190 MB of disk space + approximately 2% of planned data to be backed up (for use by the IDB)
- 2 × size_of_the_biggest_package_to_be_installed + 5MB of disk space needed on system drive

On systems running Windows Server 2003 or Windows Server2008

The Cell Manager must meet the following minimum requirements:

- 256 MB of RAM (512 MB recommended).

For each parallel backup session, 40 MB of RAM are required. This means that, for example, if you want to run 60 parallel backup sessions 3 GB of RAM are needed.

- 190 MB of disk space + approximately 2% of planned data to be backed up (for use by the IDB)
- 2 × size_of_the_biggest_package_to_be_installed + 5 MB of disk space needed on system drive
- On Windows Server 2008 systems, the firewall must be configured to accept “Remote Service Administration” (NP) connections (port 445) additionally.
- On Windows Server 2008 systems, administrative privileges are required to install Data Protector A.06.10.

On systems running Linux

The Cell Manager must meet the following minimum requirements:

- 256 MB of RAM (512 MB recommended)

For each parallel backup session, 40 MB of RAM are required and 5–8 MB per data segment size. This means that, for example, if you want to run 60 parallel backup sessions 3 GB of RAM plus 512 MB for data segments are needed.

- 300–425 MB of disk space + approximately 2% of planned data to be backed up (for use by the IDB).

Further details on Cell manager requirements are documented in the document "Product announcements, software notes, and references," which is part of the HP Data Protector documentation package included in the product DVD.

For more information

www.hp.com/go/dataprotector

Technology for better business outcomes

4AA2-5036ENW, March 2009