Odd notes on DataProtector

Greg Baker (gregb@ifost.org.au)

July 30, 2013

```
Id: dataprotector.tex 581 2013-01-11 04:20:34Z gregb
```

1 About this document

I am a consultant and trainer on HP DataProtector, with experience dating back to 1997. I've been described a few times as one of Australia's most knowledgeable people on DataProtector.

Somewhere around 1999 I started writing up interesting questions I was asked in classes, and jotting down ideas, bugs and information that I found in my work. Nowadays this is the kind of thing you would put in a blog, but back then you couldn't rely on people having access to the web. So instead I published this "odd notes" document, and I've been keeping it updated ever since.

I hope you find it helpful.

I run my own company and I am available for consulting, support, implementations and any other services you might need. Feel free to email (`gregb@ifost.org.au`) or phone (+61 408 245 856) me if you are interested.

Contents

2 Security Issues

To be completed. . .

- Both the MS-Windows and Unix installs of DataProtector by default let anyone walk in with a laptop and be an **admin** in any cell
- Unix installation has lax permissions in `/var/opt/omni/log`
- Don't forget to run Cell Secure – otherwise anyone can perform any restore they like to any computer in the cell (by setting up their own cell manager)

- Remember: the cell server doesn't authenticate, it just trusts the other end of the connection to tell the truth. And it's all sent as plain text.

- The javareporting user doesn't need to have `admin` rights. Create a new group for it, and just give it "user configuration" and "reporting and notifications" – the `classspec` file should say 133120 for its numbering.

- Unless the user `DRM$ADMIN` is allowed to restore as root, most bare-metal disaster recovery won't be able to run.

- Are you concerned about collection of debug logs? Do you need to create a `dls_hosts` file?

3 Media Pools

Here is my methodology for working out what media pools to create.

Geography Make sure each site has separate pools for everything. Otherwise an operator might be told to place a tape currently in Melbourne into a Sydney tape drive.

Virtual tape versus real If you are using a VLS6000 (or similar) to emulate DLT tapes, then you will need different pools for the real tapes and the virtual tapes. Otherwise you might get asked to put a virtual tape into a real tape drive!

Barcoded versus non-barcoded In an ideal world, every tape has a human and machine-readable barcode. In reality, tapes that aren't in a tape library tend just to get handwritten labels. These would have to be in a different media pool, so that DataProtector doesn't start asking for a non-barcoded tape to be put into a tape library device.

Media Generations DDS1 tapes can't go into DDS4 tape drives, so separate the pools of these tapes. This gets tricky for compatible generations of tapes.

WORM vs non-WORM Some tape technologies support special write-once media. Obviously these will need to be treated differently to your standard multi-write tapes.

Block size You can have tapes with different block sizes in one pool, but if you have a backup with one block size on it, no backups with a different block size can be appended to it. DataProtector will load the tape, reject it, and then try the next best tape until it finds a tape that can be used. You could have a lot of tapes each with a tiny backup on it with a 64k block size, giving you a pool with vast amounts of free space which is unuseable for any other backups. If you want to be very sure that you know your tape capacities and usage, you might want to create a pool for each different block size in use. Watch out for MS-SQL backups in 6.2 where you can set the block size to $2^n + 4$kB instead of the 2^nkB block size for file systems and any other kinds of backups.

Archiving reasons If tapes have to be left in a tape library, and the pool has to be appendable because you have multiple backups configured to go into it, but you don't want the next day's backups on the same tapes, then there's little alternative to creating another media pool. Or beg HP to add a 4th usage option: "appendable within 24 hours of the first protected write". I tend to create a media pool called "Tapes for Restore" which I move tapes into when an off-site tape is brought back into the pool. This way it won't get written to accidentally. (The write-protect tab is good for this as well, but watch out that it doesn't get marked as "bad" by a tape drive failing to write to it.)

Financials Maybe it's because of chargeback costing, control, or security. Usually these reasons are quite pathetic, but if it's too hard to battle against the bureaucracy in the name of common sense, well . . . maybe the easiest way to solve the problem is to let each division pay for and manage its own pool of media.

Isolation If you have different tape retention cycles for different data – e.g. backup X must stay on-site for 4 weeks, but backup Y must go-offsite the next day, then you don't want these two backups on the same tape. You can create separate pools to keep these backups isolated. However, usually, the problem is better solved by keeping all originals on-site and creating copies to send off-site where necessary.

4 Using Subversion with DataProtector

One of the most common causes of backups failing is somebody changing something in a backup specification. DataProtector doesn't keep history of configuration by default, so I like to use a "proper" version control system to do this.

The instructions here are for MS-Windows based cell managers.

1. Install TortoiseSVN from `tortoisesvn.tigris.org/`. It will require a reboot.
2. If you don't otherwise have a subversion repository somewhere else in your organisation, and you don't want to use subversion

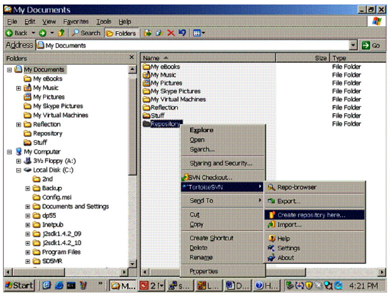

Figure 1: Creating a subversion repository

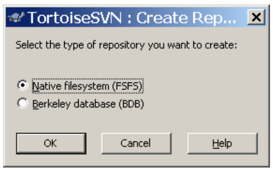

TortoiseSVN : Create Rep...

Select the type of repository you want to create:

● Native filesystem (FSFS)

○ Berkeley database (BDB)

OK Cancel Help

Figure 2: Format options for the subversion repository

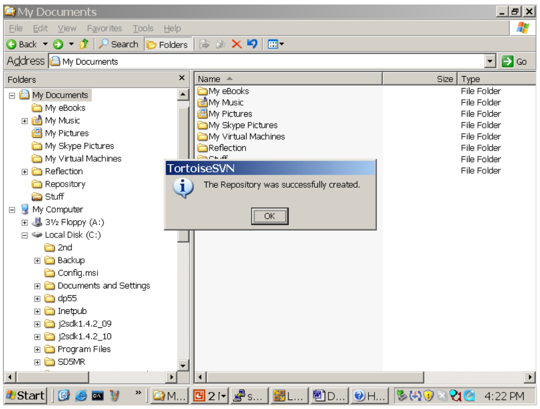

TortoiseSVN

The Repository was successfully created.

OK

Figure 3: Subversion repository created

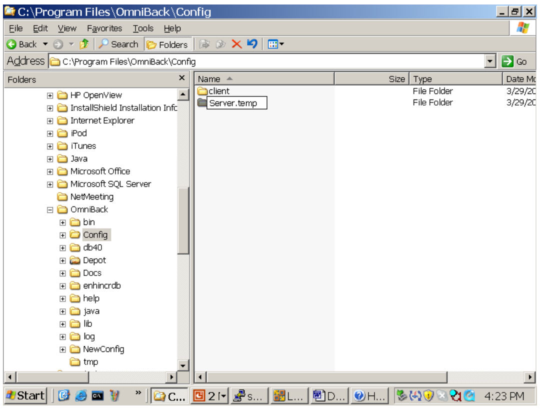

Figure 4: Renaming the DataProtector configuration directory

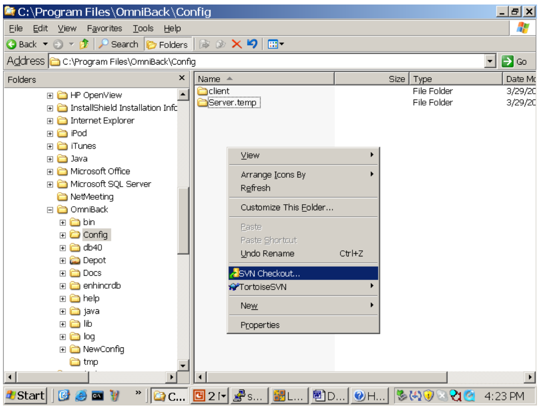

Figure 5: Performing a subversion checkout

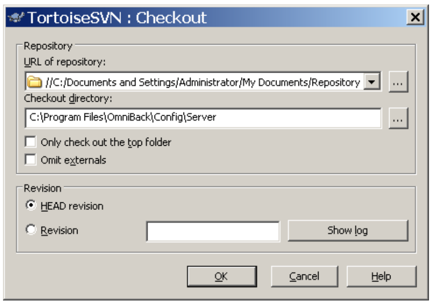

TortoiseSVN : Checkout

Repository

URL of repository: //C:/Documents and Settings/Administrator/My Documents/Repository

Checkout directory: C:\Program Files\OmniBack\Config\Server

☐ Only check out the top folder

☐ Omit externals

Revision

● HEAD revision

○ Revision

Figure 6: Parameters for the checkout

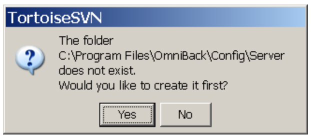

The folder
C:\Program Files\OmniBack\Config\Server
does not exist.
Would you like to create it first?

Figure 7: A checkout will create the directory if needed

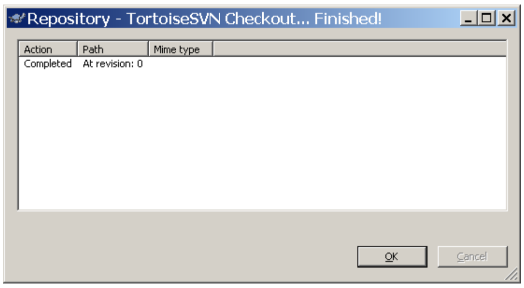

Action	Path	Mime type
Completed	At revision: 0	

Figure 8: The first revision has revision number zero

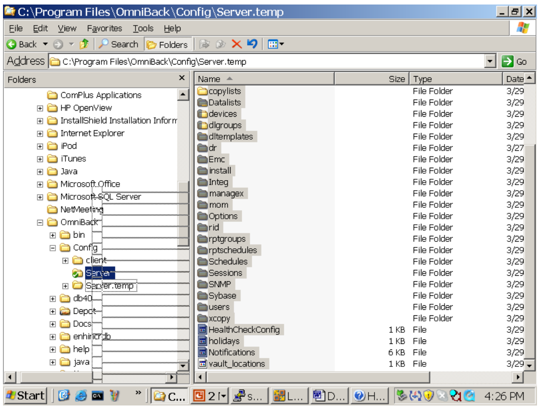

Name	Size	Type	Date
copylists		File Folder	3/29
Datalists		File Folder	3/29
devices		File Folder	3/29
dlgroups		File Folder	3/29
dltemplates		File Folder	3/29
dr		File Folder	3/27
Emc		File Folder	3/29
install		File Folder	3/29
Integ		File Folder	3/29
managex		File Folder	3/29
mom		File Folder	3/29
Options		File Folder	3/29
rid		File Folder	3/29
rptgroups		File Folder	3/29
rptschedules		File Folder	3/29
Schedules		File Folder	3/29
Sessions		File Folder	3/29
SNMP		File Folder	3/29
Sybase		File Folder	3/29
users		File Folder	3/29
xcopy		File Folder	3/29
HealthCheckConfig	1 KB	File	3/29
holidays	1 KB	File	3/29
Notifications	6 KB	File	3/29
vault_locations	1 KB	File	3/29

Figure 9: Moving the configuration copy into the version controlled area.

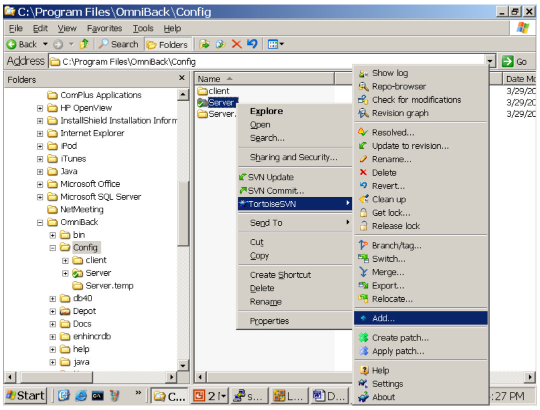

Figure 10: Files and folders need to be explicitly added.

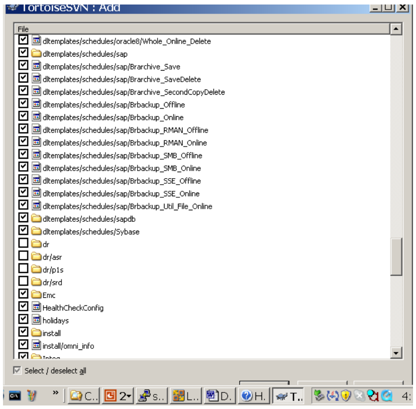

☑ dltemplates/schedules/oracle8/Whole_Online_Delete
☑ dltemplates/schedules/sap
☑ dltemplates/schedules/sap/Brarchive_Save
☑ dltemplates/schedules/sap/Brarchive_SaveDelete
☑ dltemplates/schedules/sap/Brarchive_SecondCopyDelete
☑ dltemplates/schedules/sap/Brbackup_Offline
☑ dltemplates/schedules/sap/Brbackup_Online
☑ dltemplates/schedules/sap/Brbackup_RMAN_Offline
☑ dltemplates/schedules/sap/Brbackup_RMAN_Online
☑ dltemplates/schedules/sap/Brbackup_SMB_Offline
☑ dltemplates/schedules/sap/Brbackup_SMB_Online
☑ dltemplates/schedules/sap/Brbackup_SSE_Offline
☑ dltemplates/schedules/sap/Brbackup_SSE_Online
☑ dltemplates/schedules/sap/Brbackup_Util_File_Online
☑ dltemplates/schedules/sapdb
☑ dltemplates/schedules/Sybase
□ dr
□ dr/asr
□ dr/p1s
□ dr/srd
☑ Emc
☑ HealthCheckConfig
☑ holidays
☑ install
☑ install/omni_info
☑ Integ

☑ Select / deselect all

Figure 11: Checkboxes control what is and isn’t going to be added.

as a mechanism for restoring configuration in a disaster, then you can just create a folder in “My Documents” called “Repository”. (See figure 4 on page 6.) It doesn’t much matter what type (BDB or FSFS). (See figure 4 on page 6.) It should quickly report success. (See figure 4 on page 6.)

3. Locate your DataProtector installation directory (it defaults to `C:\Program Files\Omniback`) then go into `Config`. Rename `Server` to `Server.temp`. (See figure 4 on page 7.)

4. Right click in some empty space, and select SVN Checkout. (See figure 4 on page 7.) If you already have a subversion server in your organisation, you subversion administrator will tell you what to put in the URL of Repository field. Otherwise, click on the triple dot button, and navigate to the “Repository” folder you created. Set the Checkout Directory: to be `...Config\Server`. (See figure 4 on page 7.)

   Answer “yes” when asked if you want to create it. (See figure 4 on page 8.)

   Another dialog box will appear, reporting a checkout of revision zero. (See figure 4 on page 8.)

5. Move everything from `Server.temp` into `Server`. (See figure 4 on page 8.)

6. Right click on `Server`, and select Add... from the TortoiseSVN submenu. (See figure 4 on page 9.) You might want to deselect `Server\dr`, because that directory gets modified on a regular basis by backups. (See figure 4 on page 9.) Press OK. Messages will fly past reporting on the communication with the server. (See figure 4 on page 11.)

7. Right click on `Server` again, and select SVN Commit. (See figure 4 on page 11.)

   Type in a message, such as “Initial import.” (See figure 4 on page 11.)

   Press OK. Committing takes a little longer than addition, because the actual content of the files needs to be sent. (See figure 4 on page 12.)

You now have your DataProtector configuration stored within version control. You might want to explore how the GUI interacts with the configuration – for example, if you add a new user, you will see a large red exclamation mark over the `Users` folder (and if you look in the folder you will see a red mark over the `userlist` file. (See figure 4 on page 12.) After you make a change, right click on the folder (or file) and select SVN Commit – see figure 4 on page 12. In the comment box you can write a message about the change you are making. (See figure 4 on page 13.) The commit should be quite quick. (See figure 4 on page 13.)

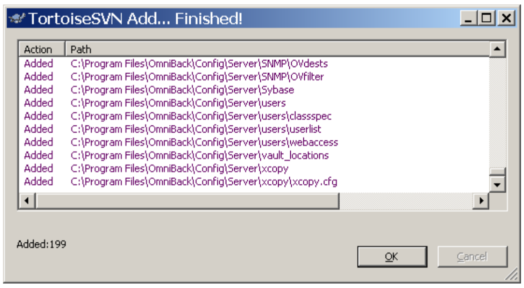

TortoiseSVN Add... Finished!

```
Action   Path
Added    C:\Program Files\OmniBack\Config\Server\SNMP\OVdests
Added    C:\Program Files\OmniBack\Config\Server\SNMP\OVfilter
Added    C:\Program Files\OmniBack\Config\Server\Sybase
Added    C:\Program Files\OmniBack\Config\Server\users
Added    C:\Program Files\OmniBack\Config\Server\users\classspec
Added    C:\Program Files\OmniBack\Config\Server\users\userlist
Added    C:\Program Files\OmniBack\Config\Server\users\webaccess
Added    C:\Program Files\OmniBack\Config\Server\vault_locations
Added    C:\Program Files\OmniBack\Config\Server\xcopy
Added    C:\Program Files\OmniBack\Config\Server\xcopy\xcopy.cfg
```

Added:199

Figure 12: Results from a subversion add operation.

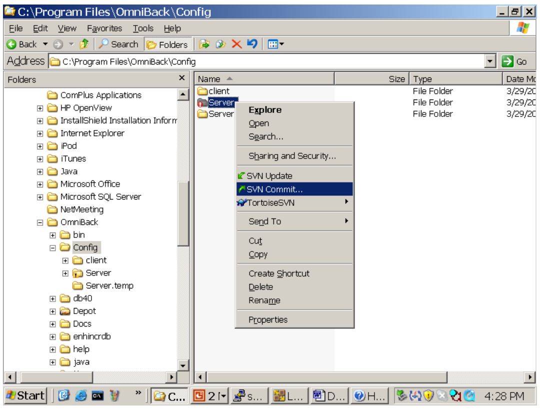

Figure 13: Committing changes.

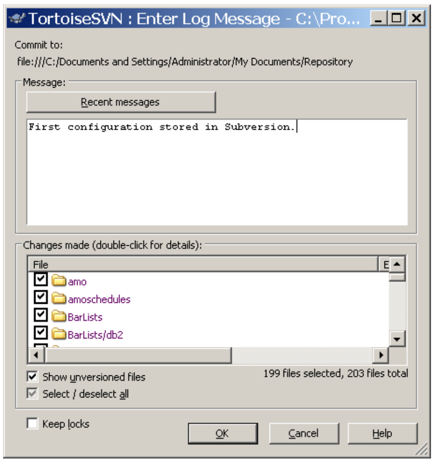

TortoiseSVN : Enter Log Message - C:\Pro...

Commit to:
file:///C:/Documents and Settings/Administrator/My Documents/Repository

Message:

Recent messages

```
First configuration stored in Subversion.
```

Changes made (double-click for details):

File	E
☑ amo	
☑ amoschedules	
☑ BarLists	
☑ BarLists/db2	

☑ Show unversioned files
☑ Select / deselect all

199 files selected, 203 files total

☐ Keep locks

Figure 14: Best practice is to write coherent commit messages.

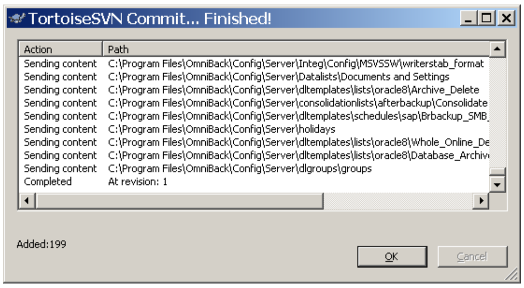

Action	Path
Sending content	C:\Program Files\OmniBack\Config\Server\Integ\Config\MSVSSW\writerstab_format
Sending content	C:\Program Files\OmniBack\Config\Server\Datalists\Documents and Settings
Sending content	C:\Program Files\OmniBack\Config\Server\dltemplates\lists\oracle8\Archive_Delete
Sending content	C:\Program Files\OmniBack\Config\Server\consolidationlists\afterbackup\Consolidate
Sending content	C:\Program Files\OmniBack\Config\Server\dltemplates\schedules\sap\Brbackup_SMB_
Sending content	C:\Program Files\OmniBack\Config\Server\holidays
Sending content	C:\Program Files\OmniBack\Config\Server\dltemplates\lists\oracle8\Whole_Online_De
Sending content	C:\Program Files\OmniBack\Config\Server\dltemplates\lists\oracle8\Database_Archiv
Sending content	C:\Program Files\OmniBack\Config\Server\dlgroups\groups
Completed	At revision: 1

Added:199

Figure 15: The conclusion of a commit.

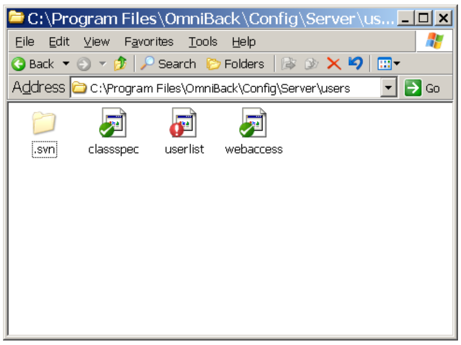
C:\Program Files\OmniBack\Config\Server\us...
File Edit View Favorites Tools Help
Back Search Folders
Address C:\Program Files\OmniBack\Config\Server\users
Go
.svn
classspec
userlist
webaccess

Figure 16: TortoiseSVN provides good visual feedback for what has changed and what is correctly committed to the repository.

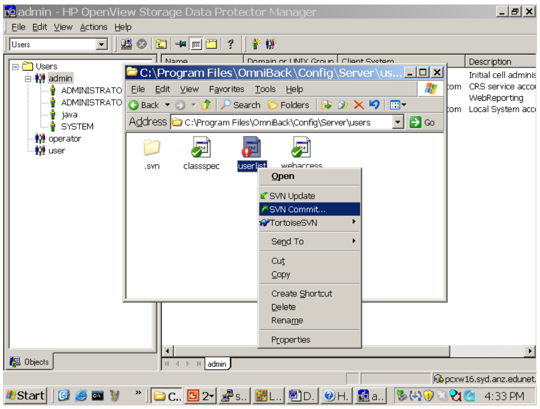

Figure 17: After each change (even from the GUI), commit your changes.

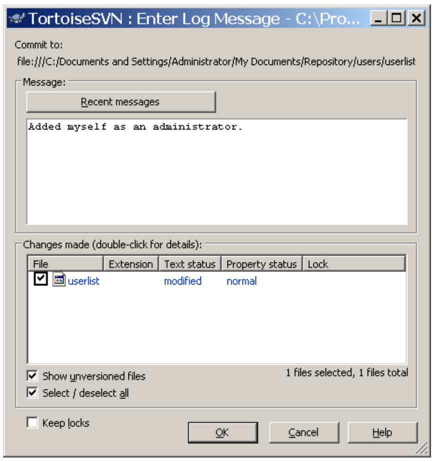

Commit to:
file:///C:/Documents and Settings/Administrator/My Documents/Repository/users/userlist

Message:

Added myself as an administrator.

Changes made (double-click for details):

File	Extension	Text status	Property status	Lock
☑ userlist		modified	normal	

☑ Show unversioned files
☑ Select / deselect all
1 files selected, 1 files total
□ Keep locks

Figure 18: Good log messages help resolve problems later.

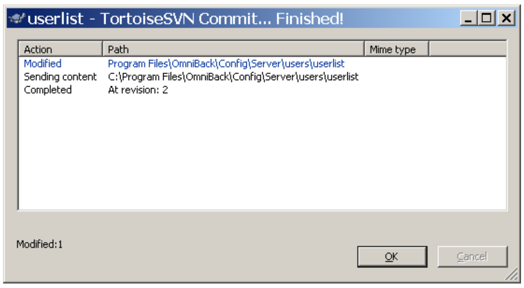

Action	Path	Mime type
Modified	Program Files\OmniBack\Config\Server\users\userlist	
Sending content	C:\Program Files\OmniBack\Config\Server\users\userlist	
Completed	At revision: 2	

Modified:1

Figure 19: Subversion only transfers relevant portions of the changed files.

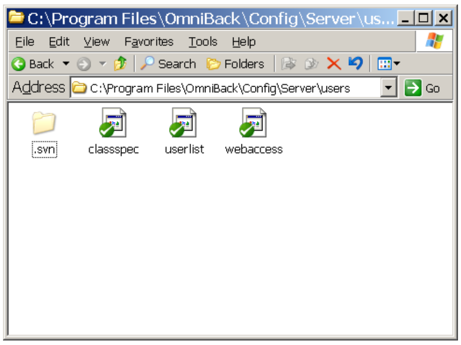

Address C:\Program Files\OmniBack\Config\Server\users

Figure 20: Nice to see everything committed correctly.

The red marks will change to reflect its now-up-to-date status. (See figure 4 on page 13.)

If a backup fails one night, you can see what has changed (either by looking for uncommitted changes, or by looking through the logs of committed changes), and revert back to a working configuration very easily.

5 MacOS X

MacOS X is now finally supported in DataProtector 6.11 (with the latest patches) and DataProtector 6.2.

6 Thoughts on cell managers

To be completed . . .

- Put it in the disaster recovery centre.
- Use a standalone tape drive for backing it up if possible, rather than sitting in a big tape library. Even better, use a different format of tapes.
- Cell manager shouldn't be the robotics controller. But they should be the same operating system if possible, so that in a disaster, the cell manager can recreate an OS disk for the controller, and bootstrap the other machines up from tape.
- Simplicity is key!

Cell manager requirements:

- Enough disk space to store full DR images, plus the database itself
- Sufficient I/O on the database that it's not a bottleneck.
- If you are using the cell manager as a media agent, then roughly dual-channel gigabit ethernet for each tape drive it is writing to. (Otherwise the network will be the bottleneck.)

7 Enabling SSH-based installs from Unix systems

To be completed . . .

For DataProtector 5.5, install PHSS_32831 and its cousins. Set the option OB2_SSH_ENABLED.

DataProtector 6.X, needs no patch, but still needs OB2_SSH_ENABLED to be set.

8 What you want to have around before a disaster

You probably want to email or `rsync` these on a regular basis to somewhere safe.

media.log So that you know what tapes you should look at for the most recent database backup.

mcf files From the tape on which you put the most recent internal database backup. This will speed up the import process.

omnidownload output For every device and library; at the very least you want the ones you will use in recovering the cell manager.

ASR / EADR disks of the cell manager If your cell manager is on MS-Windows 2003 or later.

Ignite-UX of the cell manager If your cell manager is on HP-UX.

mkcdrec of the cell manager If your cell manager is on Linux.

9 Installing on Centos 6

A fresh-out-of-the-box Centos 6 install will fail to install the DataProtector cell manager because there is neither an `inetd` nor `xinetd`. You can fix this with `yum install -y xinetd`.

The install will then proceed, but you will receive the following error:

```
Cannot start "uiproxy" service, system error:
[1053] Unknown error 1053
```

Fix this with `yum install -y compat-libstdc++-33`.

10 Dealing with firewalls

The most common scenario is that you have servers in the DMZ which you want to back up.

The dumbest and simplest solution is to put a tape drive in the DMZ as well, and install a media agent there. Then the only firewall traversal required will be from the cell manager to the DMZ hosts on port 5555.

If you want data to come back in from your DMZ back to a media agent on the inside of your network, you will need to put something like the following on the media agent:

```
OB2PORTRANGESPEC=xMA-NET:5950-5960
```

Put this line into `/opt/omni/.omnirc` if this is a Unix box, or `C:\Program Files\Omniback\omnirc` on MS-Windows (or whatever path you installed DataProtector into).

This will mean that the media agent will still randomly pick a port number, but will pick in the range 5950 to 5960. Obviously, this limits that server to run no more than 11 media agent processes, which would be an issue if you had a very large number of tape drives.

Note also that some firewall admins might misread “5950 to 5960” as meaning that the source will be port 5950, particularly since Cisco IOS does not seem to provide a syntax for specifying a port range. In the above example, a Cisco firewall admin would create 11 lines of configuration: the first would allow port 5950 through, the next 5951, the next 5952, and so on.

There is no way of controlling DataProtector to specify what source port number to use.

11 Performance Tweaks

The worst possible performance you will ever get out of DataProtector is when you are backing up millions of small files. The bottleneck will almost always be the internal database. If there is any way that you can "Log Directories" instead of logging every file you will see improvements. If necessary, run the backup with no logging and then re-scan the media after the backup window is over.

DataProtector has a documented limitation of 10,000 files in each directory but it is not clear if this still applies.

NTFS and UFS and VxFS filesystems all suffer performance problems with large numbers of files in a directory as their lookup of filenames to inodes (or equivalent) is a linear search. This is an operating system limitation rather than a DataProtector limitation, which can be seen by transferring the folder structure over to a Linux machine (e.g. to ReiserFS) where filename to inode lookup is an $O(\log N)$ operation instead of $O(N)$.

Sometimes performance can be improved by creating additional readers in a filesystem. In the GUI you can do this you by right-clicking on a server in the datalist section. Alternatively, just edit the file in the `datalists` folder and copy the relevant stanza – but remember to change the description as DataProtector does not like two objects in the same datalist to have identical names. Whichever method you choose, be careful to set up one writer to include some folders, and the other to include everything else except for those folders.

Other performance tricks:

- On Linux systems, make sure that filesystems are mounted with the `relatime` option (you want this on anyway, even if it weren't for DataProtector); on other Unix systems confirm that you don't need access time auditing, and remount all filesystems with `noatime`. Turn off the filesystem option "Preserve atime". In this way DataProtector doesn't cause any write activity when it is reading data.
- HP-UX has only two filesystems that use hard links – `/` and `/usr`. Linux systems have quite a few hard links, but none of them are essential. So turn on "back up hard links as files" and it should pose few problems.
- Make sure that you are not sending anything over a network accidentally – check that the preferred multi-path host makes sense for each backup.
- Turn on asynchronous backup on MS-Windows boxes. Except

on very slow iSCSI-based SANs, this will almost always improve performance.

Of course, the best possible performance improvement is to use incrementals instead of fulls, and particularly if you can use the Native Change Log provider.

12 Three models of datalists

12.1 Zillions of Tape Drives

- Virtual tape library or file library
- Every object has its own tape drive
- Concurrency is irrelevant

12.2 Network deluge

- Lots of net-attached hosts
- A few fast tape drives
- The bigger the concurrency the better. But check disk agent buffers.

12.3 Free-for-all

- SAN attached servers
- SAN attached tapes
- One big backup job for everything
- Concurrency should be the number of objects per host
- Make sure that each host object in the data list is preceeded by a device stanza with that host as the preferred multi-path host (or else data might be sent across the network).

13 The things you must do when you first install DataProtector

1. Install Subversion, and set it up to use a remote repository. Actually, you don't really have to do that, but it just makes life so much better if you do.

2. Edit the list of users. Make sure there is no user with special rights (such as admin) which can connect from `<Any>` IP address. Since there is no authentication on DataProtector user interface sessions (it just relies on trust) this is a necessary step to maintain security. Note that such entries do exist by default.

3. Set the web password. Otherwise any user can set up a notification which runs a command.

4. In the clients tab, run "Cell Secure" – otherwise anyone who can access port 5555 on any machine in your network can replace any file they want.

5. After you have created a device which will be used for backing up your cell console, run `omnidownload -device ...` (and, if necessary, `omnidownload -library ...`) and copy the resulting output somewhere safe. I'd recommend emailing it, printing it out, and putting it on the USB flash drives of every administrator who might be involved in a disaster recovery.

6. Create a pool called "Tapes used in a current restore". Then, whenever you bring a tape back into the tape library, move it into that pool. In this way you won't accidentally over-write your precious restore data.

7. If your cell manager is a real MS-Windows system (poor you), then schedule Drive Snapshot to take regular backups. If your cell manager is a VMware image (much better), then schedule a regular copy of a vm-snapshot disk image to some safe place. Linux users – regularly schedule `mkcdrec`. HP-UX users – regularly schedule `make_XXXX_recovery` (particularly now that Ignite can create bootable DVD images).

8. Schedule your `media.log` to be copied on a regular basis to somewhere. It's not large, so emailing it to yourself each day is not that silly. Or automatically commit it into subversion each day.

9. Create a new set of notifications, which emails or raises SNMP traps (or whatever is the usual way of alerting an issue in your organisation) for each of the IDB events, the HealthCheckFailed event, the NotEnoughMedia event and UnexpectedEvents. If you are using File Libraries (which you would be doing if you do disk-staged backups) then also alert yourself about FileLibraryDiskUsage. The default is just to write them into the Data Protector event log which nobody reads.

10. Add `omnicheck -dns -full -update` to `$OMNICONFIG/HealthCheckConfig`
11. Seriously think about whether to turn on OB2CRSSTRICTHOSTNAMECHECKING – without it, hostnames are just delivered based on trust.

14 Enhanced Incremental Database on MS-Windows

This is normally found in `C:\Program Files\OmniBack\enhincrdb` unless you have installed DataProtector into a different directory, or used shortcuts to put it somewhere else (as you would with a clustered system).

If you are on DataProtector 6.11 or earlier, try to make sure it is formatted with the smallest block size possible – i.e. a 1k block size. A typical initial `enhincrdb` uses 100-200 bytes per file. On a filesystem with a 4k block size, this means 96% wastage.

In DataProtector 6.2, it's a SQLite database and doesn't waste much space at all.

15 Tape Zap

If you want to quickly and easily let operators update the location of tapes with a barcode, I've written a small python program called `tape-zap.py` which lets you zap barcodes with a barcode reader and update the location field in DataProtector. Collect it from: `http://www.ifost.org.au/Software`

Usage: `python tape-zap.py` *location*

Barcode readers generally act like PS2 or USB keyboards. They enter whatever they zap, and then send a new-line character. `tape-zap.py` runs `omnimm -modify_medium` for each barcode label it reads in from the keyboard.

It needs to run on a machine which has the cell console software on it, and as a user who has at least the rights to run `omnimm`, such as a typical operator account.

16 Options for Cell manager disaster recovery

16.1 Procrastinator's Delight

The only preparation is to remember to copy the `media.log` file off to somewhere else on a regular basis, or have some way of knowing what tape was used for the OMNIDB backup. If you can also replicate the output from an `export_to_mcf` somewhere, that will speed things up.

The disaster recovery procedure is as follows:

- Install an appropriate operating system on to a spare machine. Set its IP address and hostname to be the same as the failed cell manager. (2 hours or less)
- Install DataProtector (15 minutes or less, depending on whether you need to download it from the HP website or not).
- Remember to install any DataProtector patches (15 minutes, depending on whether you needed to download them from the HP website as well).
- Re-read the last OMNIDB backup (2 hours if you have to scan the whole tape, or a few minutes if you have the MCF files).
- Restore the internal database. (Usually well less than 2 hours)

16.2 Cold spare

This is one of the most popular options. Install DataProtector on a spare machine, and then turn off all services. Create a post-exec job on your primary cell manager which runs `omnir` to restore the internal database to the cold spare machine. See section 17 on page 23 for details.

If you have run a cell secure operation (and you should!) you will need to make sure that the cold spare machine is listed in allow_hosts for all systems in the cell.

The disaster recovery procedure is:

- Run `omnidbutil -change_cell_name`
- Run `omnisv start`

- Edit the OMNIDB backup datalist and change it to run the `omnir` on the original cell manager.

Then start using the cold spare as the primary machine. Repeat the procedure in reverse in order to fail back to the original cell manager.

16.3 Hot spare

This only works with DataProtector 6.11 and onwards. It can be used when the bandwidth between the production and DR sites does not support regular scheduled restores as required by the "Cold Spare" scenario.

Install DataProtector on the hot spare, including license keys. At the end of each day (or even several times per day), export_to_mcf every tape that was used that day in the primary production cell. Replicate the MCF files to the hot spare server and import them (this might involve exporting the tapes from the hot spare server first).

There is now no disaster recovery procedure required – the hot spare has full knowledge of all the tapes used in the production system.

16.4 Clunk like it's last century

Buy a one-button-disaster-recovery tape drive. Attach it to the cell manager. Configure OBDR backups which write to that tape drive. In the event of a disaster, insert the latest tape and boot from it.

This is appropriate for very small sites that wouldn't end up buying a tape library anyway. It is also appropriate for very remote and isolated sites where no-one on-site knows much about backup and restore.

16.5 Just make an exception

You could always backup the cell manager with some sort of instant-snapshotting technology. DriveSnapshot only costs EUR95, and can back up to a USB hard disk.

You wouldn't sanely do this for every client in your cell, but if you want to avoid a boot-strapping operation within DataProtector, this approach would make sense for the cell manager.

16.6 It's all virtual anyway

If you can live without "Log File" and "Log All" granularity information in your backups, then you can probably get away with running your cell manager on a virtual machine.

In which case, take a snapshot every day, and the disaster recovery procedure is simply to revert to the previous snapshot.

16.7 Be prepared

Buy a stack of blank CDs, and burn the output from an EADR on a regular basis.

When a disaster strikes, take the most recently burned CD and boot from it. If necessary, restore the internal database as well. Watch as everyone is amazed at your powers of pre-planning.

17 Setting up a Cold Spare Cell manager

17.1 Setting up the backup job on Windows

Create a file `idbcopy.bat` in `%DP_HOME_DIR\bin`. It should contain just one line:

```
omnir -omnidb cellmgr-server:/ "[Database]:  cellmgr-server"
-session $SESSIONID -tree / -into "E:\IDBrestore" -target secondary-server
```

Replace `cellmgr-server` and `secondary-server` with the fully-qualitified domain name of the servers in your domain.

Create a backup job for the IDB and put `idbcopy.bat` into the post-exec definition for the job.

17.2 Setting up the backup job on Unix or Linux

Create a file `idbcopy.sh` in `/opt/omni/bin`

```
#!/bin/sh
CELLMGR=$(hostname)
OTHERHOST=othername.goes.here

/opt/omni/bin/omnir -omnidb \
$CELLMGR:/ "[Database]:  $CELLMGR" \
-session $SESSIONID \
-tree / \
-into "/var/opt/omni" \
-target $OTHERHOST
```

Create a backup job for the IDB and put `idbcopy.sh` into the post-exec definition for the job.

18 Troubleshooting Guide

18.1 Agent failing to start during a backup

1. Have a look at the inet log. Is there a connection being refused for some reason? e.g. does `allow_hosts` refuse this connection?
2. Try running a packet capture on port 5555. This won't be many packets because data gets sent on a different port.
3. If you are seeing no packets at all, then the session manager isn't even starting a connection with the agent. Is there a firewall in the way? Try running `telnet` *agent* `5555` from the cell manager.
4. Try running `omnicheck -dns -full` and see if there could be name resolution problems.

18.2 Push-based installation fails

1. First dumb question: is the username and password correct?
2. What operating system are you trying to push this to?

   HP-UX/Linux/Solaris or other Unix Confirm that the installation server has `OB2SSHENABLED` turned on in `/opt/omni/.omnirc`. No-one leaves `rsh` open any more.

   WinXP Is the Windows box in simple share mode?

   Win2k8 Have you set up `omniinetpasswd` so that the installation server knows what username toinstall with?
3. Is there a firewall blocking something?

18.3 Install completes, but client not imported

1. Is something firewalling port 5555?
2. Is there a naming mismatch? Perhaps DNS is returning a different hostname or domain-name to what the client knows itself as?
3. Is it a Unix system which has neither `inetd` or `xinetd` installed? Then there will be nothing listening on port 5555 to receive the connection.

19 Wish List

- A way of saying "the Media Agent connects to the Disk Agent for this backup specification" which would make backing up machines in a DMZ simpler. (And it wouldn't be much code to add.)
- A new policy for a media pool – "Appendable within 24 hours" – which would use a tape if the oldest protected object on the tape was less than one day old. This would make it much easier to have one pool from which all backups are drawn. Otherwise, how do you get your MS-Exchange data backed up onto the same tapes as the operating system during a weekend backup, but not end up having several days' worth of data on it?
- `/usr/omni/bin/.util` should have support for NSS file systems.
- A speedometer dashboard which shows in real time the bandwidth throughput of each disk agent. It should keep historical statistics and alert when the current speed is several standard deviations away from the mean previous speed.
- Optimisation assistant. As part of the statistical recording, DataProtector could regress the number of files, the average file size and number of folders backed up against the bandwidth throughput. If the bandwidth is decaying close to linearly with the number of files, then DataProtector could suggest turning off file logging.
- Reports on how rapidly files are changing, so that it can suggest RAID5 for data that is hardly ever modified (for example). At the very least, sorted reports of what filesystems have the biggest churn rates could be helpful.
- A report or view to show tapes which are marked FULL, or ones which are unappendable because they have been copied. These are both useful for operators to know what tapes to take out. Also, a report or view of all tapes by location (not just sorting tapes within one media pool by location).

- Autonomous backups – if the media agent and disk agent are on the same machine, perhaps they could be scheduled to run without initiating from the cell manager. (Nearly done!)
- Restartable backup session manager processes. Perhaps when the cell manager starts up it could query every machine in the cell to ask what backups are running and start backup session managers accordingly.
- A script to automatically download any relevant DataProtector patches when they become available.
- MacOS X support (Done!)
- A way of specifying the preferred multi-path host for a device during a restore. Otherwise a restore could data over the network when you might want to constrain it only to run over fibe.
- Ubuntu and Debian support.
- PostgreSQL and MySQL integration agents. They must be close to having as many enterprise users as say, Informix.
- Subversion integration agent.
- Can you do a DR from a file library if the file library's name has a space in it? It seems like you can't.