NETBACKUP 7.1 FEATURE BRIEFING

NetBackup 7.1 Feature Briefing

Auto Image Replication

This document is for internal and partner use only. Do not distribute copies of this document to customers.

If you have any feedback or questions about this document please email them to **IMG-TPM-Requests@symantec.com**

Feature Description

The growth in the use of disk storage (particularly deduplicating disk storage) for backup has highlighted a limitation of disk storage when it comes to site disaster recovery. Tapes can easily be sent to offsite storage and then to a disaster recovery site in the event of a site outage.

Disk doesn't offer this flexibility and, while many disk technologies do have the ability to replicate their contents to a compatible array at a remote location accessing the backups with NetBackup is not always a simple matter. The use of Storage Lifecycle Policies and optimized duplication allows disk based backups to be replicated between devices under NetBackup control but, until now, has been limited to the NetBackup domain in which they originated. Thus to achieve a site disaster recovery capability with optimized duplication requires a NetBackup domain that spans at least two geographically remote sites. (This may range from two rooms in the same building/complex to two data centers in separate cities – obviously the greater the separation between the sites the greater the type and severity of disaster you can protect against.)

Some OpenStorage technologies, including the PureDisk Deduplication Option, support 'out of band' replication where the contents of the storage is replicated between devices without NetBackup's knowledge. The problem with this approach is that the way in which NetBackup accesses the disk based backups means there is no guarantee the replicated data can be accessed without first recreating the NetBackup catalog importing the entire content of the disk storage using the bpimport command.

The nbcatsync utility, introduced in NetBackup 6.5.6 and 7.0.1 goes some way to addressing this problem but relies on being able to restore the catalog from a catalog backup and then 'post-process' it to reconcile the disk device mappings. While faster than importing the contents of the storage this is still a time consuming process.

The Auto Image Replication feature introduced in NetBackup 7.1 seeks to address this issue by allowing Storage Lifecycle Policies (SLPs) to duplicate selected images between NetBackup domains, allowing for a separate disaster recovery domain that may support several source domains.

With Auto Image Replication backups are written to deduplicating disk using a suitably configured SLP and then duplicated to a remote target NetBackup domain. An SLP in the target domain logs the duplicated copy in the catalog of the target domain and it is held there for a period specified by a retention level set in the source domain. This effectively ensures that a copy of the backup is available in the target domain for as long as required by the SLAs of the source domain. Where that copy is held is determined by the target SLP that catalogs the duplicated backup when the copy arrives at the target domain. This SLP can be configured to duplicate the automatically imported backup to other types of storage and expire it from the target storage, thus making space for other backups to be replicated across. Note that in NetBackup 7.1 the 'import' destination can only have a fixed retention. Support for the 'expire on duplication' and 'capacity managed' retention options will be included in a future release of NetBackup.

The primary purpose of Auto Image Replication is to create off-site copies of mission critical backups to protect against site loss, it is not intended to extend the storage capacity of a backup domain by allowing backups to be stored in a separate domain. NetBackup 7.1 is the 'phase 1' implementation of Auto Image Replication and does not support any direct communication between the source and target master servers. This means that the source domain has no record of the copy held in the target domain and that copy cannot be restored directly to the source domain (this capability will be available in a future release). The target domain copy can, however, be restored to a machine in the target location or duplicated to a tape which can then be imported back into the source domain.

Business Value

Auto Image Replication allows customers to automatically create copies of mission critical backups at a remote location that forms part of a separate NetBackup domain. This domain could be another data center belonging to the same organization, a dedicated disaster recovery site within the organization or even a third party disaster recovery facility.

Electronic off-siting in this way allows the backup to be duplicated to an off-site location as soon as the backup has completed and without the need for user intervention. It also means that the duplicate copy is available at the disaster recovery site as soon as the duplication has completed. This offers a significant advantage of traditional tape based off-siting in which backups must be duplicated to tapes that are then removed from the tape library and shipped to the disaster recovery site, a time consuming and costly activity with a risk that the tapes could be lost or stolen in transit.

Underlying Principles

In order to use Auto Image Replication you must have suitable disk storage devices configured in the source and target domains. In the NetBackup 7.1 release the Auto Image Replication feature only supports the Media Server Deduplication Option (MSDP) to duplicate backups between domains. Support for OpenStorage devices and PureDisk Deduplication Option are expected to follow in the first half of 2011.

Auto Image Replication works by duplicating backups from a disk pool in the source domain to a disk pool in the target domain. The replication operation requires two SLPs, one in the source domain and one in the target domain, both of which must have the same name. The SLP in the source domain is associated with the backup policy and controls the writing of backup and the subsequent duplication to the target domain.

The SLP in the target domain is not associated with a backup policy but is invoked by an alerting mechanism when a new image (duplicated from the source domain) is detected. This SLP runs the process to add the information about the backup to the target domain and can also be configured to duplicate the backup to other storage locations in the target domain.

One storage destination in this target SLP must have a retention setting of “replication retention”. This setting takes the retention period set in the duplication destination in source domain SLP and ensures that a copy of the backup is retained for the period requested by the source domain’s backup administrator. The target domain backup administrator can configure the SLP to create other copies with shorter or longer retention periods as well.

An Auto Image Replication backup image is different to a normal NetBackup backup image because once the backup has completed the image database information associated with the backup (the part of the NetBackup catalog that lists what the backup actually contains) is appended to the end of the backup image before it is duplicated to the target domain.

When a new backup is detected in the target domain this information is read from the backup and used to populate the target domain’s NetBackup catalog. This information only exists in the source domain copies and the initial copy in the target domain and is not included in any subsequent duplicate copies created in the target domain.

.

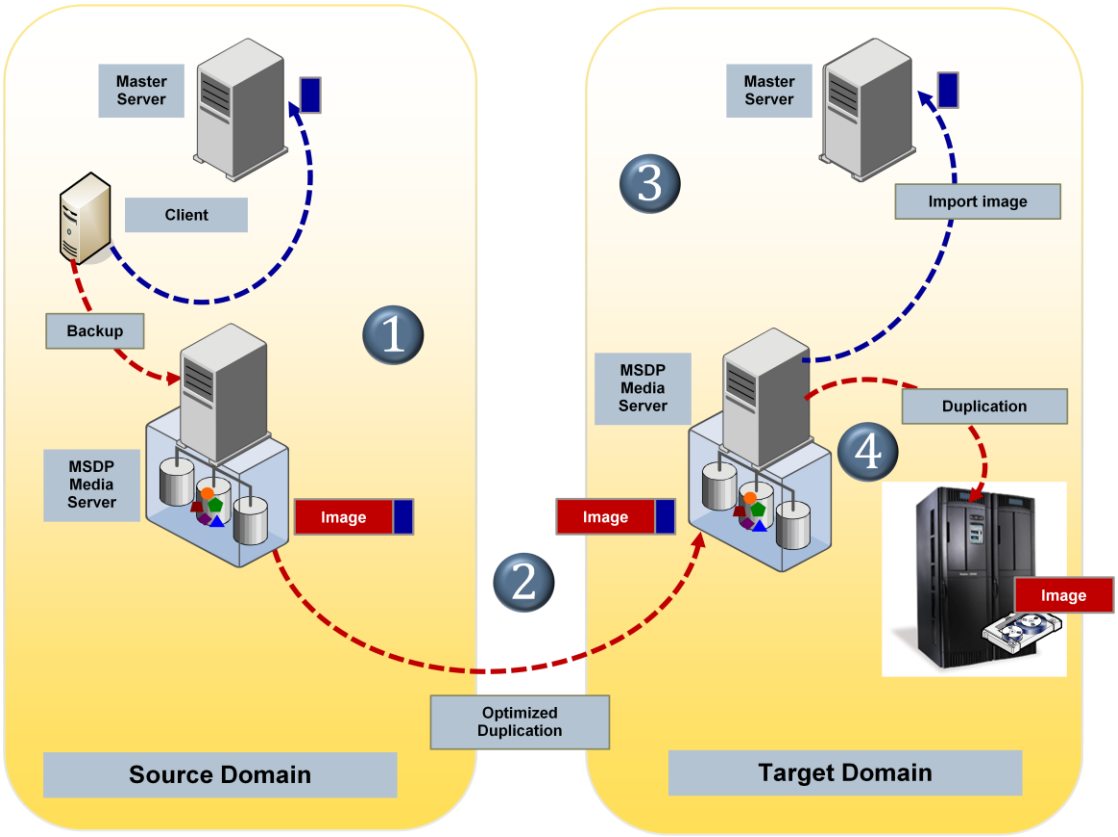

Figure 1 - The Auto Image Replication duplication sequence with MSDP

Figure 1 above shows the sequence of events in an Auto Image Replication backup and duplication operation broken down into 4 discrete steps:

Step 1 – The backup is written to disk storage in the source domain using a backup policy with an SLP configured for Auto Image Replication when the backup completes the catalog data it generates is appended to the end of the backup.

Step 2 – the backup is duplicated to the target domain

Step 3 – the storage device in the target domain alerts the target master server to the fact that a backup has been duplicated to it. This triggers the receiving SLP to run a 'fast import' operation in which the catalog data transferred from the source domain is added to the target domain's catalog.

Step 4 – the receiving SLP in the target domain can now duplicate the received backup to any desired location for storage.

Test Drive

As discussed in the previous sections, Auto Image Replication is based on the interaction between storage lifecycle policies in the source and target domains. The SLP in the source domain duplicates the backup images to the target domain and the SLP in the target domain imports them into the target domain's catalog.

The first step to setting up Auto Image Replication is to configure at least one MSDP media server in both the source and target domains. Ensure that the media server in the target domain is reachable over LAN connections from the media server in the source domain. Once the storage servers have been created in both domains the storage server in the target domain must be granted replication credentials in the

storage server in the source domain. This is set through the administration GUI on the source domain master server by adding the target domain storage server through the replication tab on the (source domain) storage server credentials as shown in Figure 2 below:

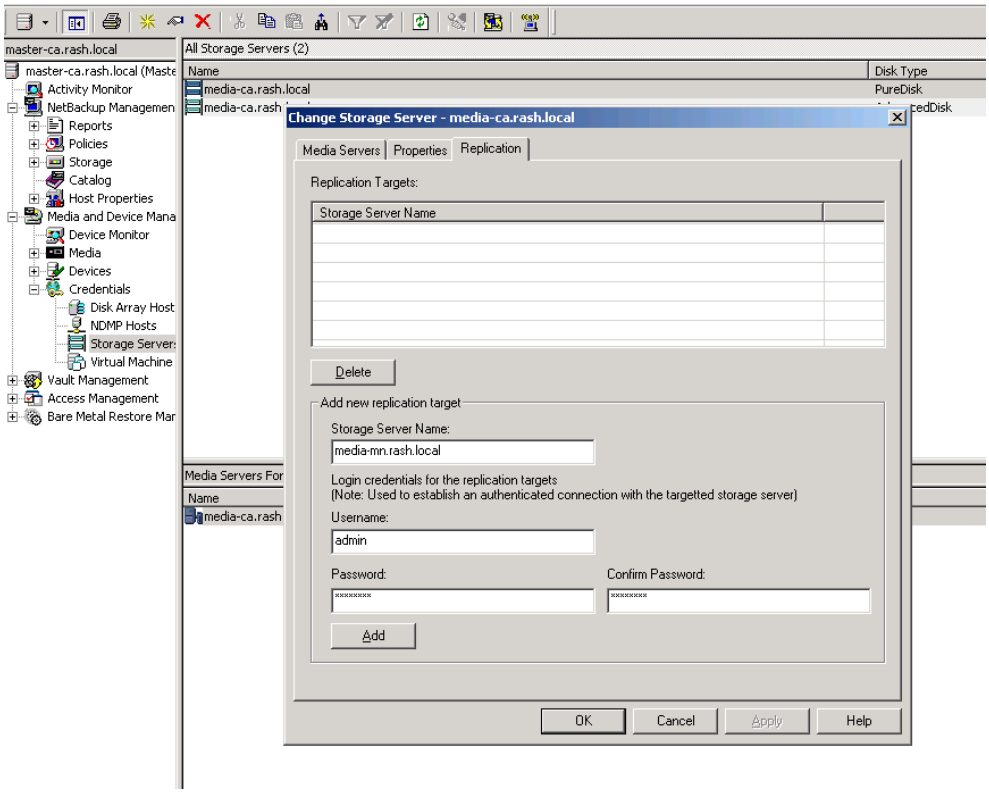

Figure 2 - Setting replication credentials on the source storage server.

The next step is to configure two storage lifecycle policies, one in each domain, with the same name. The SLP in the source domain must have at least one duplication storage destination that is configured as 'remote master' as shown in Figure 3 below:

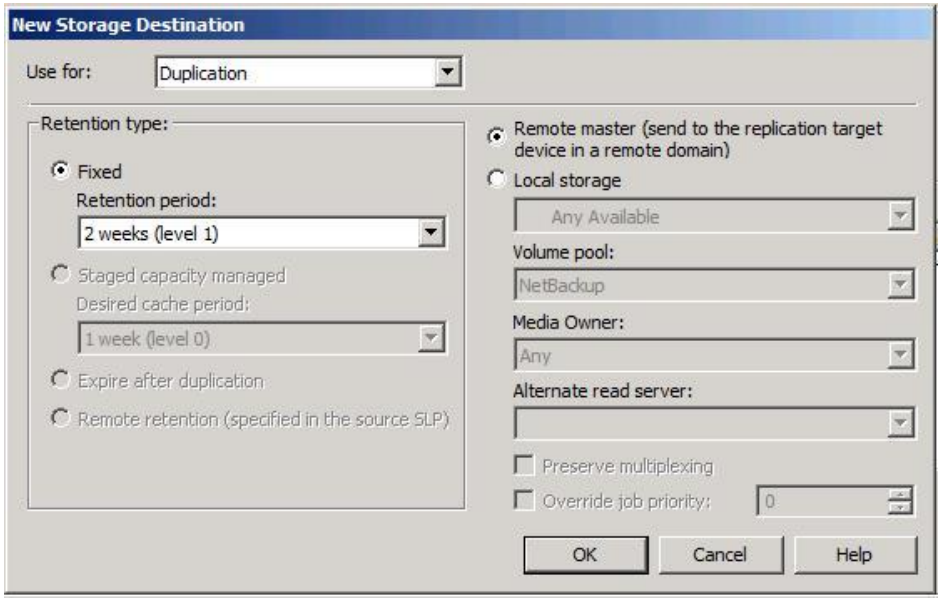

Figure 3 - Auto Image Replication remote duplication storage destination

The SLP in the target domain has no storage destination of type 'backup', instead the first storage destination has the type 'import' and is configured to use the storage unit associated with the remote media server. At least one storage destination in this SLP must have the retention type set to 'remote retention'. In Figure 4 below a single 'import' storage destination is configured with the 'remote retention' type selected.

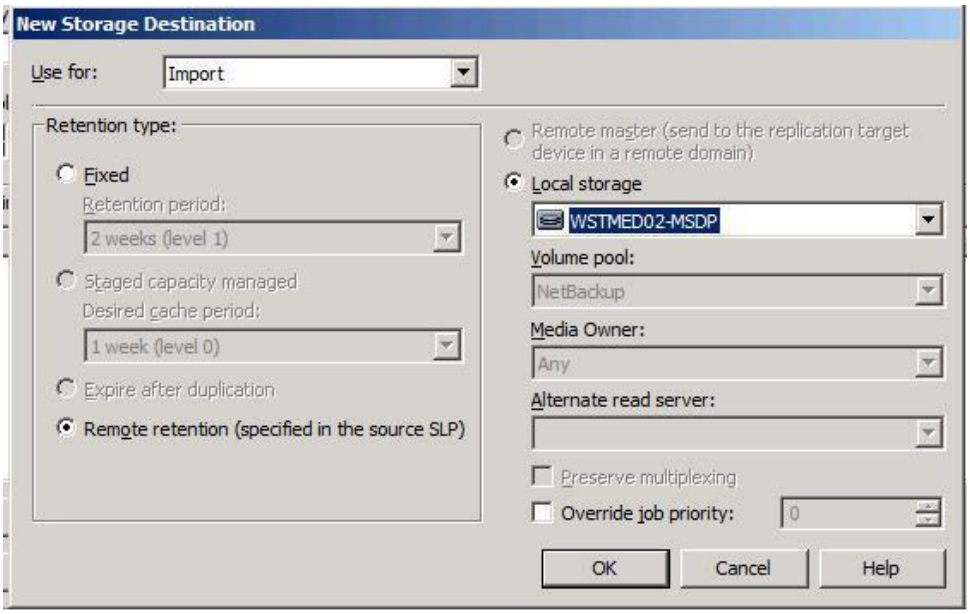

Figure 4 - Auto Image Replication import storage destination

Once the SLPs have been set up configure backup policies in the source domain that uses the Auto Image Replication storage lifecycle policy. When these policies run the backups will be written to the backup storage destination and then duplicated to the storage server in the target domain. Figure 5 below shows the job monitor view from the source domain, showing the initial backup and subsequent duplication job.

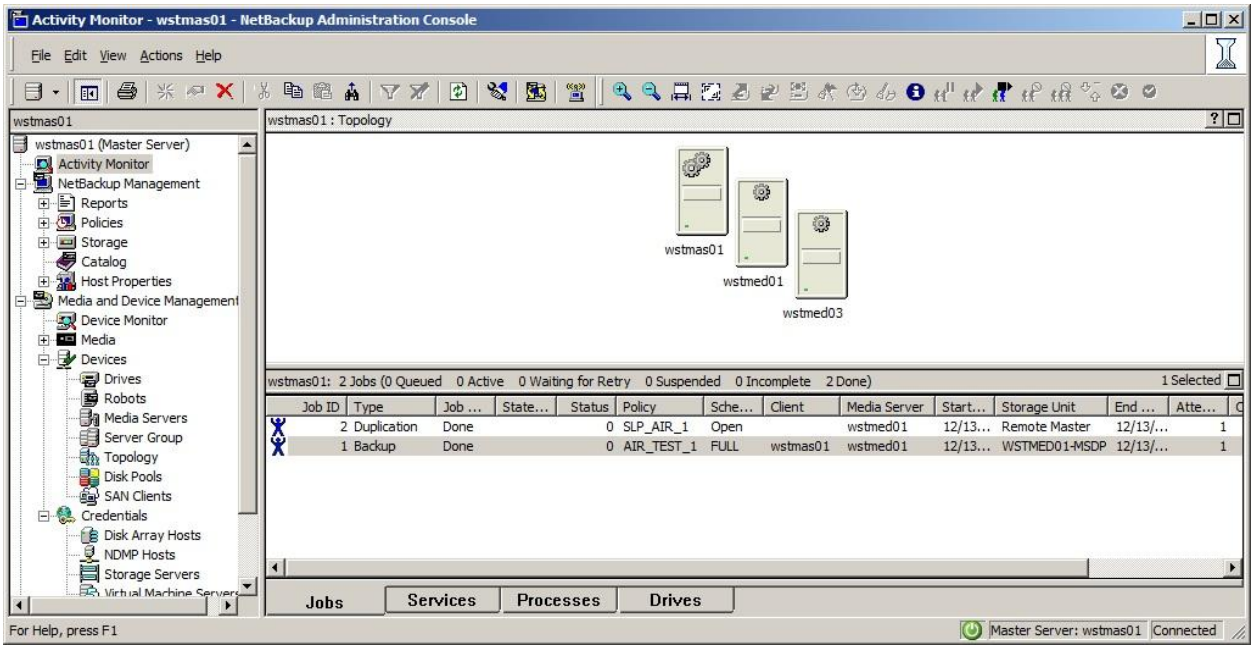

Figure 5 – Source domain, the backup is created and duplicated to the remote target domain

When the duplicated backup is detected in the target domain the target domain's SLP is trigged to import it and, if specified, duplicate it to other storage. Figure 6 below shows the job monitor view from the target domain showing the import of the image into the target domain and the subsequent duplication to alternative storage.

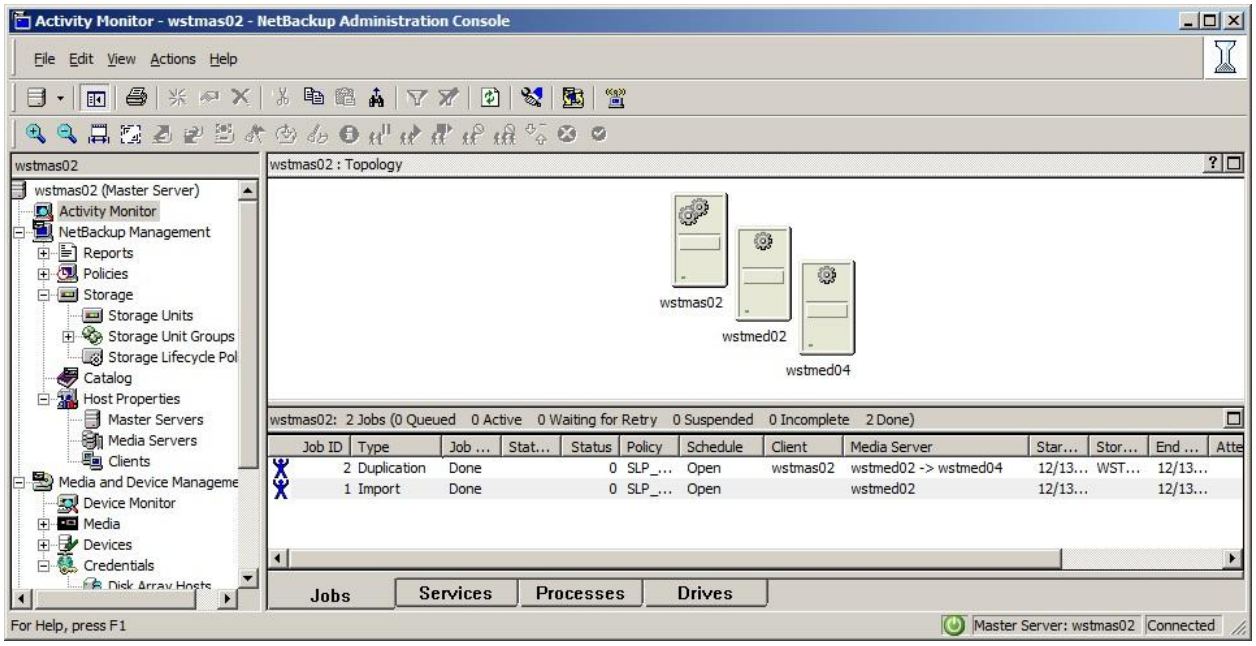

Figure 6 - Target domain, the backup from source domain is imported and duplicated to long term storage.

Once the backup has been imported into the target domain it is possible to browse the image and restore files from it in the target domain using the Backup, Archive and Restore (BAR) GUI as shown in Figure 7 below. (Remember to add the client from the source domain to the list of Source Clients for Restore in File -> Specify NetBackup Machines and Policy Types and select it):

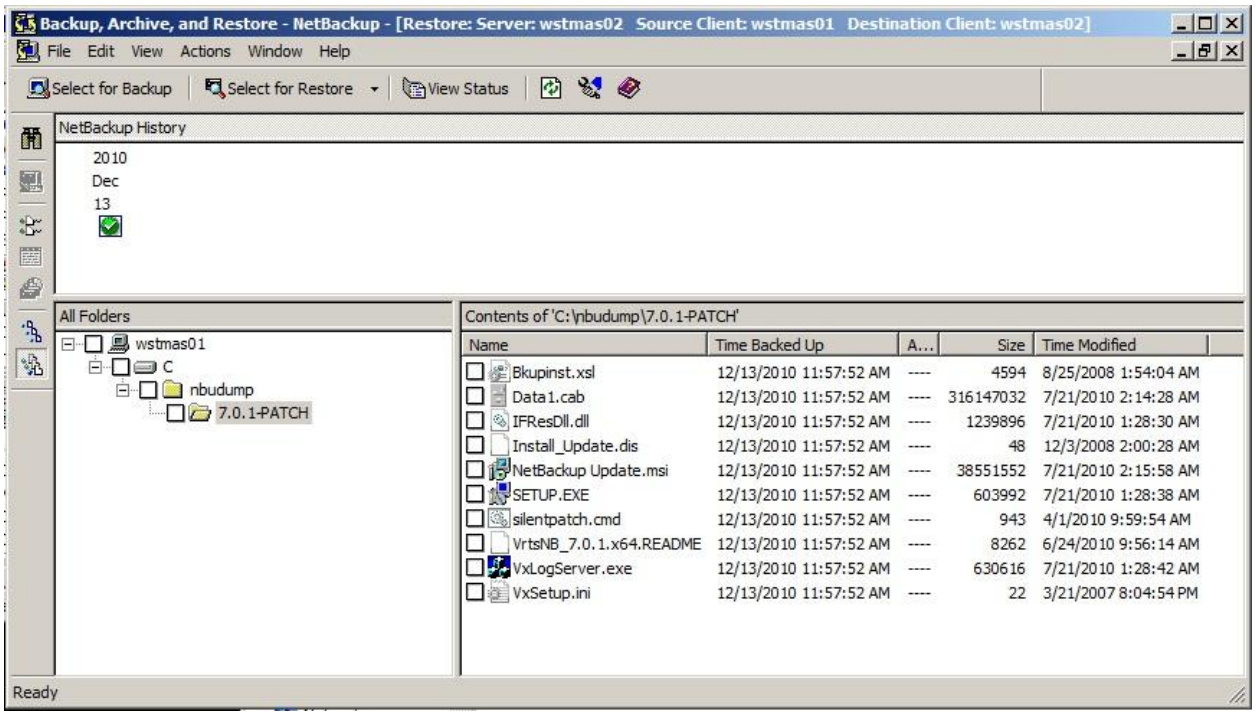

Figure 7 - Target domain BAR GUI browsing Auto Image Replication backup.

Licensing and support considerations

The Auto Image Replication feature forms part of the Storage Lifecycle Policies feature and requires the enterprise disk license and, where appropriate, the deduplication license, for both the source and target domains.

Customers using capacity based licensing are not required to purchase licenses for the target domain as capacity used by Auto Image Replication backups is only counted in the source domain and do not count towards the used capacity in the target domain.

Customers using traditional licensing are required to purchase infrastructure licenses for the master and media servers in the target domain when using AIR, even if that domain is a dedicated DR domain. This requirement is because the master and media servers in the target domain must be active at all times and are not in "idle" mode as mandated by the cold DR policy. Note that deduplication and enterprise disk licensing is based on front-end terabytes, so there is no additional cost for these licenses on the target domain. All other infrastructure licenses including servers, library tape drives and shared storage option are required in the target domain in all cases.

In the NetBackup 7.1 release the Auto Image Replication feature is only supported using the Media Server Deduplication Option to duplicate backups between domains. Support for Auto Image replication with OpenStorage devices and PureDisk Deduplication Option are expected to follow in the first half of 2011. Note that this is dependent of the plug-in developed by the OpenStorage vendors and will not require code updates to NetBackup 7.1 itself.

Backups replicated using Auto Image Replication will not be recoverable using Guided Application Recovery in the target domain (this capability will be available in a future release).

In the NetBackup 7.1 release the 'import' destination can only have a fixed retention and does not support the 'expire on duplication' or 'capacity managed' retention options. To minimize the amount of space occupied in the 'import' destination the target SLP should be configured to duplicate the images to other storage for the 'remote retention' period and the import storage destination should be given a short retention period. This limitation will be addressed in a future release.

In the NetBackup 7.1 release the source domain has no knowledge of copies of backups held in the target domain and thus, once duplicated, these copies cannot be tracked by the backup administrator in the source domain (this capability will be available in a future release).

In the NetBackup 7.1 release the Auto Image Replication feature does not support the restore of backups directly from the target domain to the source domain.