White paper
Use of local replication functions of ETERNUS DX with BS2000/OSD

Modern high-end memory subsystems offer functions for the uninterruptible mirroring and splitting of replicas. While the original data remains available for the main application, parallel backup procedures and evaluations can be executed on a copy. Parallelism leads to permanent data availability for strategic applications and a significant improvement in the online availability of data.
In the BS2000, one of the ways this is achieved is through the use of Fujitsu ETERNUS DX storage systems and their replication functions.
This white paper describes local mirroring with Equivalent Copy (EC) and SnapOPC+, possible scenarios in which this can be used and the benefits to the customer, while also explaining how the local replication functions are supported in the BS2000/OSD.

Content

Local replication functions of ETERNUS DX hardware 2
- Equivalent Copy (EC) 2
- Snapshots with SnapOPC+ 2

Application scenarios and customer benefits 2
Support for the replication functions in BS2000/OSD 3
Configuration 4
- Equivalent Copy clone units 4
- SnapOPC+ snap units 4

SHC-OSD 4
- Support of Equivalent Copy (EC) 4
- Support of Snapshots with SnapOPC+ 5

Pubset replication 6
- Using and addressing pubset copies 6
- Pubset replication with PVSREN 6
- Pubset replication with SHC-OSD 7

Data backup 7
- Physical data backup with FDDRL 7
- Data backup with HSMS 7
  - Data backup of clone units 7
  - Data backup of databases 7
- Disk-to-disk data backup using snapsets 8
  - Transferring backup data from snapsets to HSMS 9
- Use of disk copies with FDDRL, HSMS and Snapsets by comparison 9
- Volume-based reconstruction of pubsets 9
  - Reconstruction of pubsets from clone units 9
  - Reconstruction of pubsets from snap units 9

Data security based on standby pubsets 9
Exporting data (Migration) 10

Date of issue June 2013

Local replication functions of ETERNUS DX hardware

For ETERNUS DX storage systems the replication functions Equivalent Copy (EC) and SnapOPC+ are available. They permit additional copies of volumes to be created within a Fujitsu ETERNUS DX system so that they can be processed separately after being split off.

Equivalent Copy (EC)

Equivalent Copy (EC) creates a copy of the original unit (referred to as unit in the following) and maintains a local synchronous mirroring for the clone pair. At a certain point in time ("point-in-time copy") the copy, called clone unit can be activated for an independent use. It is available immediately after their activation. Activation can only be performed after the copy procedure has been completed between the original and the copy.

Together, the original and clone unit form the clone pair which is administered by Equivalent Copy (EC). Any ETERNUS DX units of the same size can be used as clone units. After activation, the unit and clone unit can be accessed, thus enabling both to be processed independently of each other, for example by various applications. After separate processing has been concluded, the copy can be updated with the original data in order to reconstitute the equality of both and to continue the replication.

Write accesses to the unit are also performed simultaneously to the clone unit in mirrored status.

During the initial synchronization period the performance of the application, which accesses the data on the original unit, can be impacted.

If the mirroring is interrupted or the clone pair dissolved, both units of the former pair can be used freely.

When mirroring is resumed, only the changed data are copied from the unit to the clone unit.

Recovering (restoring) the original unit from the clone unit must be carried out in several steps. At present, an SDF-P procedure can be provided for this purpose as part of a project solution. The general availability of the RESTORE function is planned for the follow-up version of SHC-OSD.

Snapshots with SnapOPC+

SnapOPC+ produces a “snapshot” of a logical unit (or possibly several units). The snapshot, known as the “snap unit” is a logical copy of the original unit at a specific point in time (“point-in-time copy”): while the data on the original unit is changed, the snap unit preserves the data as it existed at the time the snapshot was made.

The snapshot is available immediately after the snap pair has been created (and implicitly activated), in other words there is no copy process. The creation of the snap units is therefore highly efficient. Together the original unit and snap unit form a snap pair. SnapOPC+ manages it in a so-called snap session. An original unit can have up to 256 snap units.

SnapOPC+ uses the “copy-on-first-write strategy”: the original data is only saved in the ETERNUS DX if data on the original unit has been changed. Consequently, SnapOPC+ only needs a small disk capacity. Nonetheless, from a user perspective, a complete copy of the original data is available at all times. This copy is separate from the original, so that the original and copy can be processed separately, for example by different applications.

The various snap sessions of an original unit are **dependent** on one another. They can be ended individually, starting with the “oldest” snap session (/STOP-SNAP-SESSION FORCE=*NO). It is also possible to end a ”more recent” snap session; in this case, all the corresponding “older” snap sessions are also ended (/STOP-SNAP-SESSION FORCE=*YES). After the separate processing or the original and snap units has been ended, the data of the snap units can be reconstructed on the original units.

Application scenarios and customer benefits

General:

The additional copy enables critical applications to be productive consistently. While the original data is available for the main application, backup runs and reports (which otherwise require the application to be terminated/interrupted) can be performed in parallel on the copy. The parallelization leads to continuous data availability for strategic applications as well as significantly increasing the online availability of data.

A copy of the “live data” of the application can be created in order to test program changes without restricting the availability of the main application.

Loading and updating data resources of data warehouses can be done at short time intervals.

Local replication functions by comparison

	Equivalent Copy (EC)	Snapshots with SnapOPC+
Data availability	■ Clone units are highly available copies of the original unit ■ Unit and clone unit are independent copies ■ Clone units are available after termination of copy operation ■ Writing access on the clone unit has no influence on the original unit (separation against errors)	■ Snap units are logical copies of the original unit; a complete copy of the original data is available from user view. ■ Snap units are available directly at the beginning of the snap session ■ Snap units always assume the availability of the original volume

Disk storage demand	■ Requires 100% additional disk space per copy	■ Requires less than 30% additional disk space (mainly on the Snap Data Volume), depending on the extent of modifications and duration of snap session
Performance	■ Supports very high I/O load (independent access to unit and clone unit)	■ Supports moderate I/O load (parallel access to production I/O) ■ High-performance creation of snap units, as no copy process required
Utilization	■ Ideal for applications with high extend of modifications ■ Supports migration between different RAID levels	■ Ideal for applications with small extend of modifications

Comparison of the capacity requirement for periodic consistency statuses when Equivalent Copy (EC) and SnapOPC+ are used

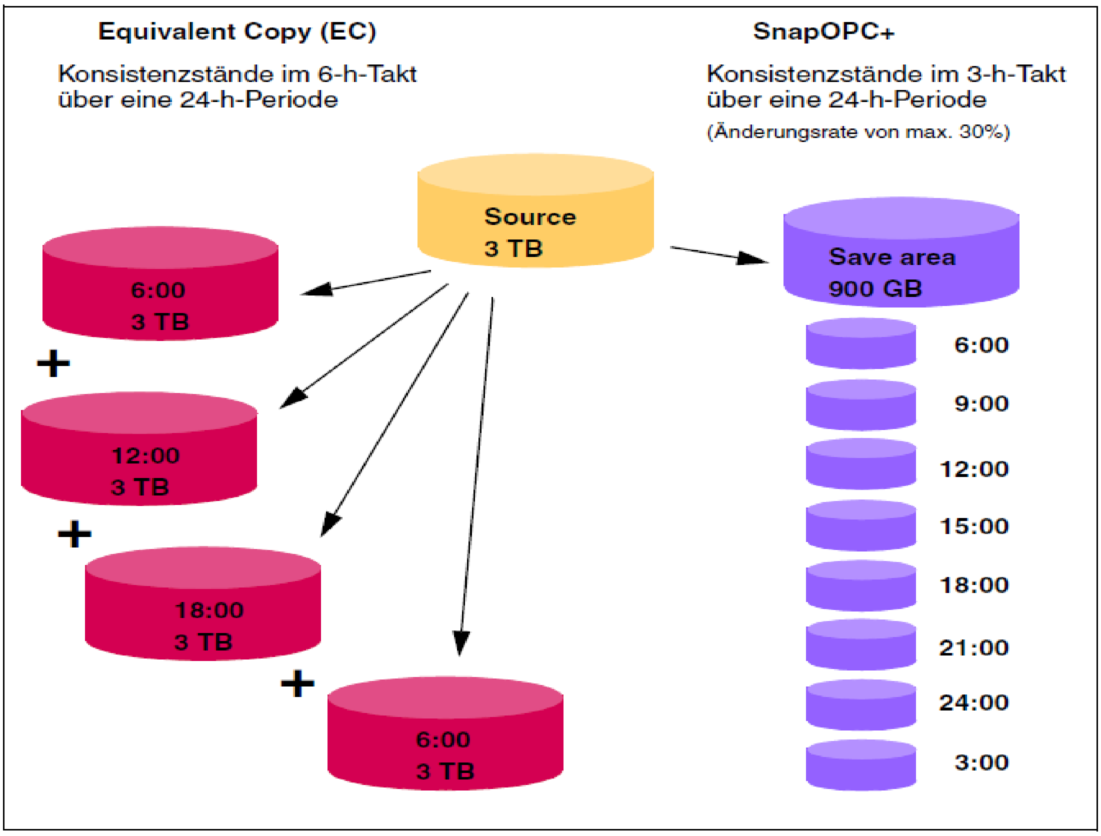

Four full copies with additional mirror units or clone units require 12 TB of additional capacity.
Up to 15 Point-In-Time copies with snap units require only 900 GB of additional capacity

The following chapters of this white paper describe how the operation scenarios of Equivalent Copy (EC) and SnapOPC+ mentioned above are supported by BS2000 in a comprehensive manner.

- In BS2000 the Equivalent Copy (EC) function is available at the volume level; it makes sense if one or perhaps two copies are to be used independently from each other. For clone pairs the synchronization times before split must be considered.
- The SnapOPC+ functionality is available in BS2000 on Pubset level. Snaps are good as backup copies because several copies can be kept as different backup versions, using only a modest amount of memory and always linked to the original pubset.

Support for the replication functions in BS2000/OSD

The information functions and replication function Equivalent Copy (EC) for the storage systems Fujitsu ETERNUS DX400, DX400 S2, DX8000 and DX8700 S2 are offered as of SHC-OSD V9.0 and BS2000/OSD-BC V7.0. Furthermore for ETERNUS DX S2 the function Snapshots with SnapOPC+ is

supported as of SHC-OSD V10.0 and BS2000/OSD V9.0. The additional function snapsets in BS2000 is available for ETERNUS DX S2 starting with BS2000/OSD-BC V9.0, KP1/2013.
The needed firmware statuses of the supported storage systems are to be retrieved from the release notes of SHC-OSD.

Configuration

The replication functions operate on the basis of "logical volume" granularity, referred to as the term „unit".
The physical disks of an ETERNUS DX are combined in so-called RAID groups and these are distributed over several logical volumes - the BS2000/OSD server views a logical volume as a disk (MN).

Equivalent Copy clone units

Any volume of the ETERNUS DX system can be used as clone units. The clone units must match the units to be copied in both capacity and emulated device type (D3435 and D3475-8F). The Raid type can be different.
The units, used as clone units must be generated like every other disk of the BS2000 server.
At present it is possible to configure a maximum of 32 clone units on a unit. This maximum number is defined by the Equivalent Copy (EC) function and relates to the total value for EC and REC (Remote Equivalent Copy).

SnapOPC+ snap units

As snap units SnapOPC+ uses devices of the ETERNUS DX specially configured and initialized, known as **Snap Data Volumes (SDV)**. Disk type D3435 is supported in BS2000/OSD for SDVs. SDVs must be available in a sufficient number and size – the configured capacity of an SDV must be greater than or equal to the capacity of the original unit.
SnapOPC+ operates with the "copy-on-first-write strategy", which means that the original data is only recorded on the snap unit if data has been previously changed on the original unit.
The data storage is initially on the SDV itself. If the capacity of an SDV is exhausted, then further changes are stored in a central storage area of the ETERNUS DX, the so-called **Snap Data Pool (SDP)**. This consists of so-called Snap Data Pool Volumes (SDPV), which provides temporary storage for the SDVs. The temporary storage provided for an SDV can be located on more than one SDPV. This storage is released again when the snap session is ended (i.e. the snap pair is dissolved).

The devices for the original unit and snap unit must be contained in the same storage system and be of the same type.
Up to 256 snap units can be configured for an original unit.
The original unit can also be the source unit of a remote copy pair. The clone unit can also be an original unit of a snap pair.

Details can be found in the "SHC-OSD/SCCA-BS2" manual as of V10.0.

SHC-OSD

SHC-OSD as of V9.0 provides information services via the global ETERNUS DX configuration, the ETERNUS DX device configuration and the remote replication functions via BS2000/OSD command interfaces.
Furthermore SHC-OSD also enables the replication functions EC, REC and SnapOPC+ of ETERNUS DX to be used and controlled via BS2000/OSD command interfaces. Control of these functions can be integrated into runtime procedures, thus achieving a high degree of automation and secure handling in critical operating situations.
The use of BS2000 software product SHC-CM-LR in combination with SHC-OSD is a prerequisite for the use of the local replication functions.
SHC-OSD permits units to be selected via their VSN, their mnemonic name, their serial number and the internal number of the logical volume within the ETERNUS DX storage system, or via the catalog ID (catid) of the pubset or volume set to which they belong. The most common application scenario is the processing of pubsets as a whole. The main features of this application scenario are described below.

Support of Equivalent Copy (EC)

For control of Equivalent Copy (EC) functions following commands are provided in SHC-OSD:

/START-CLONE-SESSION	Create clone pair

/START-CLONE-SESSION creates one or more clone pairs, with one volume being assigned as a clone unit to each original unit. Any suitable volume with the same capacity can be used as a clone unit. Clone units for all volumes of a pubset can be created with one /START-CLONE-SESSION command.
The execution of the command starts implicitly the synchronization of the clone pair (i.e. the copy process from the original unit onto the clone unit). The clone pair gets the SYNCHRONIZING status. After the first synchronization is finished, the pair gets the SYNCHRONIZED status. Afterwards, each data change is applied to both units, i.e. to the original data and to the copy. Through repeatedly entering the command /START-CLONE-SESSION it is possible to create several clone sessions (several clone pairs) for a unit.

/ACTIVATE-CLONE Activate clone

/ACTIVATE-CLONE activates one or more clone pairs, each consisting of an original unit and a clone unit. All clone pairs of a pubset can be activated with a single /ACTIVATE-CLONE command.

After activation the clone pair gets the SPLIT status.

As very few applications can work with inconsistent data (open files, etc.), it is strongly recommended to create a recovery point before the mirrors are split off. The coordinated separation of multiple mirrored volumes of a (shared) pubset is made possible by halting all I/Os to these volumes before the split (HOLD-IO), thus enabling the splitting operations to be performed more or less simultaneously. This enables parallel data resources to be made available for recovery-capable applications, even when a pubset consists of a number of volumes.

Following successful activation, the original unit and the clone unit are separated. Both can be accessed by applications from the host. From the viewpoint of the host applications, the clone unit contains all the data of the original unit at the time of the /ACTIVATE-CLONE command.

/SWAP-CLONE-SESSION Swap the features of the original and clone unit

The command /SWAP-CLONE-SESSION swaps the mirror features of a clone pair. This swaps the direction of the local mirroring: The first original units become clone units and the clone units become the new original units.

The swap can be performed when the clone pairs are in the SPLIT status and after the separate processing on the unit and clone unit has been ended.

After the swap the clone pairs remain in the SPLIT status. Original units and clone units continue to be accessible.

Swapping the mirror direction enables the changes that were made to the primary clone unit to be transferred to the original unit. In order to do so the command /RESTART-CLONE-SESSION must be called after the swap.

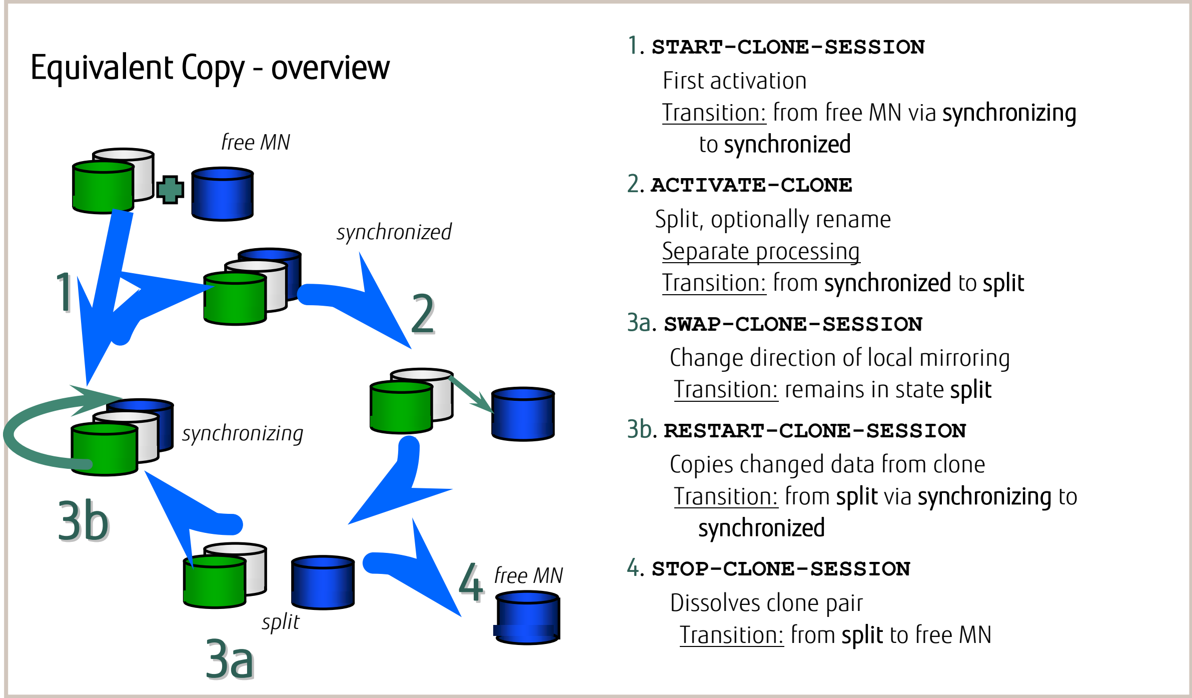

1. **START-CLONE-SESSION**
   First activation
   Transition: from free MN via **synchronizing** to **synchronized**
2. **ACTIVATE-CLONE**
   Split, optionally rename
   Separate processing
   Transition: from **synchronized** to **split**

3a. **SWAP-CLONE-SESSION**
   Change direction of local mirroring
   Transition: remains in state **split**

3b. **RESTART-CLONE-SESSION**
   Copies changed data from clone
   Transition: from **split** via **synchronizing** to **synchronized**

4. **STOP-CLONE-SESSION**
   Dissolves clone pair
   Transition: from **split** to free MN

/RESTART-CLONE-SESSION Recreate clone pair

/RESTART-CLONE-SESSION recreates a clone unit of a clone pair from the original unit, i.e. the previous copy is replaced by a copy taken at a later time. (Restart of the mirroring with the now updated version of the productive data.) The command is only executed if the clone pair has the SPLIT status. After the restart the clone pair changes via the SYNCHRONIZING status to SYNCHRONIZED.

/STOP-CLONE-SESSION Dissolve clone pair

/STOP-CLONE-SESSION terminates the assignment of one or more clone pairs. This command can be used to dissolve the first possible clone session, a specific prior clone session or all clone sessions for a unit.

Executing the command causes the clone unit to return to the status of a normal volume without clone function.

Support of Snapshots with SnapOPC+

SnapOPC+ is only integrated in SHC-OSD for ETERNUS DX S2 models from version SHC-OSD V10.0 and BS2000/OSD V9.0 onwards. The following specific functions are provided for the purpose of controlling the Snap functionality:

/START-SNAP-SESSION Create a snap pair and activate the snap

/START-SNAP-SESSION creates one or more snap pair(s) by assigning a SDV to an original unit as a snap unit. A /START-SNAP-SESSION command can be used to assign snap units for all volumes of a pubset. The snap pairs are immediately activated for ETERNUS DX. From the perspective of the application, the snap unit contains all the data of the original unit at the time when the command was processed. The execution of the command gives the snap session COPY-ON-WRITE status.
After successful activation, the normal unit and snap unit are separated: both are available for applications. From the perspective of the application, the snap unit contains all the data of the original unit at the time when the command was processed.
After the snap unit has been activated, all data pointers of the snap unit point to the data of the original unit. From this point onwards, the original data is backed up in the snap data volume or snap data pool before it will be changed on the original unit.

/RESTORE-FROM-SNAP Reconstruct the original from snap
/RESTORE-FROM-SNAP reconstructs the original unit of a snap pair from the snap unit. For the reconstruction, the snap units can be selected explicitly (e.g. via the catid of the renamed pubset) or implicitly according to their relative age. The reconstruction causes all modifications made to the original unit since the snap pair was created to be discarded. The snap session remains in COPY-ON-WRITE status. The reconstruction can be done multiple times for different snap units.

/STOP-SNAP-SESSION Dissolves the snap pair
/STOP-SNAP-SESSION terminates the assignment of one or more snap pairs. The snap session must be in COPY-ON-WRITE or FAILED status. If the snap pair is broken up, the data on the snap unit is discarded. Executing the command causes the snap unit to assume UNUSED status.
If there is more than one snap session (snap unit) for an original unit, then, by default, SnapOPC+ only allows the "oldest" snap session to be ended. The operand FORCE=*YES also allows "more recent" snap sessions to be ended. All other "older" snap sessions are also implicitly ended.

SHOW-STORAGE-CONFIGURATION enables information to be displayed about the snap data pool (SDP) of an ETERNUS DX.
Monitoring functions are available for checking the filling level of the snap data pool. When certain filling levels are reached or changed, warnings are issued to the console (limiting values set on the ETERNUS DX).
If the SDP is exhausted, the current snap sessions can only be continued until the capacity limit of the SDV has been reached and can then only be ended. The changed data on the snap units are then lost.

During the separate processing of the snap units, i.e. while data is recorded in the save pool starting at /START-SNAP-SESSION, actions with a large volume of changes, such as data reorganization with SPACEOPT are to be avoided on the original pubset, as this leads to a heavy load (I/O rate, memory requirement) on the snap data volume.

Pubset replication

Using and addressing pubset copies

A distinction should be made when using pubset copies between use by applications, e.g. use for test purposes or data mining, and use by the operating system and the backup products for physical or logical backup and restore. Pubset copies for use by applications must be assigned a new catalog ID. In this way they can be processed in any desired way in parallel with the output pubset.
With regard to using a pubset copy for backup and restore, renaming the copy to new catalog IDs is counterproductive because of the following reasons:

- A pubset copy can be imported (and it must be imported for file-oriented access that is necessary during HSMS/ARCHIVE backup). In doing so, there is no write protection with respect to other applications, as it would be necessary during backup (BS2000 has no read-only IMCAT).
- Renaming causes the association with the output pubset to be lost; in the case of HSMS backups, the data maintained in a directory for a renamed pubset cannot be assigned as a matter of course to the output pubset – not without intervention from outside.

For this reason, in the case of FDDRL and HSMS backups, as well as for online backups using snapsets, the replicas are addressed by "redirecting" the addressing of the volumes of the output pubset. For this purpose the volumes of the pubset copies that represent the "frozen" status of the pubset at a particular point in time are assigned a special VSN (SPECIAL-VSN) which is known only to the backup methods (FDDRL, HSMS/ARCHIVE, Snapset Management).

Pubset replication with PVSREN

The PVSREN utility routine can be used to generate a new, independent pubset from clone units of an SF or SM pubset. PVSREN provides the statement CREATE-PUBSET-FROM-MIRROR for this purpose.
Clone copies are suitable for creating pubsets which are to lead an independent existence.

Both SF and SM pubsets can be used as output pubsets.
Already separated pubset copies with special notation can also be made into independent pubsets.
The result is a new SF pubset with a new catalog ID or, analogously in the case of SM pubsets new volume sets with new catalog IDs.
Besides updating the information in the file catalog, PVSREN now also performs other changes on an optional basis:

- Renaming the default catalog ID in the user catalog of the HOME-PVS. This is an option if the newly created pubset is to replace the output pubset or if the replicated pubset is to be used on another system.
- Renaming the default catalog ID in the user catalog of the replicated pubset.
- Updating the IMON Software Configuration Inventory. This is an option, as above, if the replicated pubset is to replace the output pubset or if the replicated pubset is to be used on another system.

- Automated updating of the SMS storage classes. Optionally, the storage classes for the newly created pubset can be reset or taken over. If taken over, the volume lists are updated to match the new catalog IDs of the SM pubset.

Pubset replication with SHC-OSD

For SF pubsets (Single-Feature pubsets), SHC-OSD provides an integrated renaming function (operand NEW-PUBSET) as part of the commands for activating clone units (/ACTIVATE-CLONE) and starting snap units (/START-SNAP-SESSION). When a pubset copy is created, the catid of the pubset copy can be converted explicitly into a new catalog ID or automatically into special notation.
The PVSREN program can complete renaming performed by SHC-OSD in relation to IMON, SYSID and default pubset for user IDs. As a result, the pubset is identical to a pubset fully renamed by PVSREN.

Data backup

Physical data backup with FDDRL

The software product FDDRL (**F**ast **D**isk **D**ump and **R**e**L**oad) is used for physically backing up and restoring disks and pubsets in a BS2000 system.
FDDRL supports the backup of clone units explicitly by processing disks and pubsets with special notation.
FDDRL requests these disks renamed in special notation for processing using the FDDRL statements DUMP-DISK and DUMP-PUBSET and the operand SPECIAL-VSN=*YES.
The disks are backed up as though they had their original VSN. They can be restored under the original VSN.

Data backup with HSMS

Data backup of clone units

The software product HSMS can be used in combination with the CCOPY function for data backup using the ETERNUS DX function Equivalent Copy (EC) (HSMS as of V9.0 and BS2000/OSD-BC as of V7.0).
The system administrator sets up the clone session with /START-CLONE-SESSION. For HSMS backups using the Equivalent Copy (EC) function, the granularity “pubset”, i.e. for all volumes of a pubset, is required for mirroring. The operand
CONCURRENT-COPY=*YES(WORK-FILE-NAME=*BY-ADDITIONAL-UNIT) must be specified in the HSMS //BACKUP-FILES statement.
Embedding the function in this way enables Equivalent Copy (EC) to be used transparently for backup.
The clone pairs are split during the initialization of the backup job for all volumes of the pubset concerned. During the split, the inputs/outputs are held on the pubsets. This satisfies the following conditions:

- The data of the file set to be backed up is crash-consistent in itself. Even files opened for writing do not cause the backup to crash or can be backed up using the SAVE-ONLINE-FILES=*YES option. Only open files for which online backup was explicitly defined with the OPNBACK operand of the CATAL macro are backed up
- The meta data of the pubset on the split-off clone mirrors is consistent in the sense that it permits the reconstruction of the pubset at the time of the split.

This ensures that the consistent status of the data at the start of the backup is always available during the entire backup.
The data backup is performed using the split-off clone units.
When the clones are split off, they are renamed in special notation. The data sets are then backed up as though the pubset had its original catid. The backups are managed using the original catid and can be restored under the original catid.
If an error occurs, the backup can be repeated based on the status of the clone unit.

Data backup of databases

If databases are to be accessible 24 x 7, they cannot be shut down and closed for backups, so the backup is generally performed with the database open. The ETERNUS DX Equivalent Copy (EC) function can be used for this. With this variant of online backup, a database backup can be completed much faster, since the lock wait times for the database files are shorter and smaller LOG files for the reconstruction of a consistent status can sometimes result.

UDS/SQL

Using HSMS (SAVE-OPTIONS=*PARAMETERS(SAVE-ONLINE-FILES=YES) operand) it is possible to back up the UDS/SQL database files while the database is simultaneously being processed and modified by the DBH.
The following preconditions must be created for online backup of a database:

- The “online backup capability” of the database must be specified for the database using the BMEND utility routine (ENABLE-ONLINE-COPY statement) before the start of database operation. The online backup capability of a database is then noted both in the UDS management data and in the DVS catalog entries of the database files.
- AFIM logging (after-image logging) of the database must be activated, because only the online backup in combination with the ALOG files(s) (Archive Log files containing the after-images) generated during the backup form a consistent database status.

The corresponding HSMS statements for clone backup can be used for creating online backups of UDS databases or individual realms.

ORACLE

Open Oracle database files can also be backed up with HSMS. In order to back up individual tablespaces or DB files online, the identifier for online backup must have been set for the associated files using the INSTALL.C.OPNBACK utility (this causes the BS2000 CATAL macro to be invoked for setting the OPNBACK attribute for the file). This should be done before a file is added to a tablespace. For the online backup, the tablespaces must be set in backup mode. The online backup then takes place in the following steps:

- SQL> ALTER TABLESPACE name BEGIN BACKUP
- Backup of the files of the tablespace with HSMS for clone
- SQL> ALTER TABLESPACE name END BACKUP.
  Backup mode can be canceled as soon as the job JV (CONTROL-JV) of the HSMS //BACKUP-FILES statement contains the status value T in the CCS INIT STATUS field (monitoring of the Concurrent-Copy initialization).

SESAM/SQL

In contrast to UDS/SQL and ORACLE, SESAM/SQL provides use of HSMS backup of clones that is integrated into the DB backup functions.
Using the utility statement COPY, the database administrator can create SESAM backup sets of the entire SESAM/SQL database or of parts of the database such as catalog space and user spaces. The backup sets can optionally be generated using the software products HSMS or ARCHIVE. If database operation runs on an original unit mirrored with clone, COPY ... USING DIRECTORY *hsms_archive_name* BY_ADD_MIRROR_UNIT can be used to back up database files residing on a clone unit highly efficiently into an HSMS archive on disk or magnetic tape cartridge. When executing the COPY statement, SESAM/SQL uses the HSMS Concurrent Copy (CCOPY) function.
The backup is performed in three phases:

- First, the clone unit is separated from the original unit. Read DML accesses to the database are possible here.
- The specified database files are saved from the clone unit into the specified HSMS archive. The database on the original unit is processed by the DBH. During the backup interval, read accesses (with COPY … OFFLINE) or modify DML accesses (with COPY … ONLINE) to the database are possible.
- After the backup has been written into the HSMS archive, the clone unit and the original unit are resynchronized. Read or modify accesses are also possible during this time.
- Parallel to the online backup of files from clone into the HSMS archive, it is possible to run a formal check on the spaces to be backed up (CHECK FORMAL parameter of the COPY utility statement). The formal check on the spaces to be backed up is performed on the split-off clone unit. If an error is found during the formal check, the backup counts as unsuccessful. Thus, a copy of the database files that has been checked for formal consistency can be backed up to an HSMS archive without additional time overhead.

The advantage of the use of the HSMS backup of clones integrated into the SESAM backup methods is that the database administrator does not have to be concerned with the separation or the synchronization of the clone unit.

LEASY

In the current LEASY version V6.2 (Release in March 2007) a ReadOnly Mode for LEASY files has been introduced. In order to get consistent online backup copies the new function ROMS (ReadOnly Mode: Set) of the LEASY-MASTER utility sets the files of a LEASY catalog to ReadOnly mode.
Using HSMS (SAVE-OPTIONS=*PARAMETERS(SAVE-ONLINE-FILES=YES) operand) it is possible to back up the LEASY catalog files while the files are simultaneously being read. The corresponding HSMS statements for clone backup can be used for creating online backups.
The ReadOnly mode can be turned off again using the new function ROMR (ReadOnly Mode: Reset) of the LEASY-MASTER utility. ReadOnly mode can be canceled as soon as the job JV (CONTROL-JV) of the HSMS //BACKUP-FILES statement contains the status value T in the CCS INIT STATUS field (monitoring of the Concurrent-Copy initialization).

Disk-to-disk data backup using snapsets

BS2000/OSD-BC as of V9.0 supports snapshot-oriented backup/restore scenarios in ETERNUS DX S2 configurations. The virtual copy of a pubset which can be used for restore consists of the concurrently created snap units for all volumes of the pubset. A pubset copy consisting of snap units like this is referred to below as a “snapset”. Snapsets are not full-fledged pubsets, but pubset mirrors which can be accessed for reading, albeit only for restoring individual files. With SM pubsets, snapsets can be formed for the full pubset only, and not at the level of volume sets.
Support for snapsets in BS2000 for ETERNUS DX storage systems as of BS2000/OSD-BC V9.0 is based internally on the function SnapOPC+ in SHC-OSD as of V10.0.
For one pubset a maximum of 52 snapsets can be managed in BS2000. The advantage of snap-based backups is the lower space requirement compared with clones and is especially worthwhile for Pubsets with data from relatively minor change volume.
Snapsets are created by the administrator using the command CREATE-SNAPSET. They are deleted using the command DELETE-SNAPSET – perhaps only the eldest snapset (or the specified and all prior ones) can be deleted.
The maximum number of snapsets that is allowed for a single pubset can be specified per command (/SET-PUBSET-ATTRIBUTES ..., SNAPSET-LIMIT=).
The administrator can assign a special Snap Save Pool to the snapset also using the SET-SNAPSET-PARAMETERS command.
The functions SnapOPC+ (START-SNAP-SESSION) and Snapsets (CREATE-SNAPSET) are not at all allowed to be used at once on the same pubset, as otherwise data loss is threatening.
In shared pubset mode, the snapsets can be used by all the systems.
In a disaster recovery scenario using REC replication, snapsets can also be maintained in the target storage subsystem.

The end user can restore individual files and job variables from the available snapsets using new DMS functions. SHOW commands (/SHOW-SNAPSET-CONFIGURATION, /LIST-FILES-FROM-SNAPSET, /LIST-JVS-FROM-SNAPSET) and action commands (/RESTORE-FILE-FROM-SNAPSET, /RESTORE-JV-FROM-SNAPSET) are provided for this purpose.
The snapset end user functions of listing and restoring files and job variables from snapsets are also provided via program interfaces.

LMS also supports the restoration of library elements from snapsets on ETERNUS DX storage systems.

Starting with V10.0, **VM2000** supports the use of snapsets on guest systems running BS2000/OSD-BC V9.0 or higher.
Under VM2000, the AUTO-SNAP-ASSIGNMENT privilege permits the guest system on a VM to assign itself snap units of a snapset implicitly without VM and device being prepared for the implicit device assignment (i.e. no ASSIGN-BY-GUEST privilege or attribute for VM and device).
A VM is granted the AUTO-SNAP-ASSIGNMENT privilege by default at /CREATE-VM time.

Customer benefits of disk-to-disk backup using snapsets

- Unrivaled backup speeds "on-the-fly" and minimal storage requirements for backup data
- Fast restore of large data sets
- Enables new backup strategies: tighter sequence of backups in parallel with production, decoupled tape backup of snapsets at times of low I/O load
- Supports both logical and physical restore from the same backup (the snapset) for the first time.

Transferring backup data from snapsets to HSMS

Files and job variables backed up on snapsets can be transferred to a backup archive using the HSMS BACKUP-FILES statement. The catalog ID of the pubset and the identifier of the snapset to be backed up are specified in the statement:
//BACKUP-FILES ..., CONCURRENT-COPY=*YES(WORK-FILE-NAME= *FROM-SNAPSET(PUBSET-ID=<catid>,SNAPSET-ID=<snap-id>).
When backing up snapsets it is important to note that the files and job variables do not correspond to the current status at the time of BACKUP-FILES processing, but represent a backup status at the time of snapset creation. For that reason the backup version thus generated in the backup archive is assigned the snapset creation date. If more recent backup versions were already present in the backup archive, a new backup file must be created in order to maintain the monotony of the backup versions within a backup file. Continuation of the backup file is rejected.

Use of disk copies with FDDRL, HSMS and Snapsets by comparison

FDDRL backs up from disk copies in order to get a backup without interruption during the production running.
HSMS (and the data bases) back up from disk copies in order to get a backup with a very short interruption during the production running.
FDDRL and HSMS need the disk copies only while the backup run is performed, the backup result and the restore behaviour are the same as with a traditional backup without disk copies.
With the snapsets the disk copies themselves serve as backups, from which logical or physical restore can be performed.
Backup on Snapsets is not independent: if a disk controller falls through disaster (such as fire) or the original plates are damaged, the snapset backups are no longer available. Therefore snapset backups should either be doubled operated with remote mirroring (REC mode) or taken in a backup archive with HSMS and secured on an alternative medium, such as tape.

Volume-based reconstruction of pubsets

Reconstruction of pubsets from clone units

SHC-OSD can be used to reconstruct the original unit of a clone pair from the clone unit. Recovering (restoring) the original unit from the clone unit must be carried out in several steps. At present, an SDF-P procedure can be provided for this purpose as part of a project solution. The general availability of the RESTORE function is planned for the follow-up version of SHC-OSD. Alternatively the reconstruction is also possible by swapping the features of the original and the clone unit using the swap function (/SWAP-CLONE-SESSION) and the subsequent synchronization (/RESTART-CLONE-SESSION) of the clone pair.
Clone units that were created with a special notation by SHC-OSD must subsequently be renamed to their original VSN using PVSREN (statement RESTORE-LABELS-OF-PUBSET) in order to be able to work with the corresponding disks in normal operation. Renaming can be performed in units of SM/SF pubsets and volume sets.

Reconstruction of pubsets from snap units

SHC-OSD can be used to reconstruct the original unit of a snap pair multiple from the snap unit (RESTORE-FROM-SNAP).
Snap units that have been created by SHC-OSD with special notation must then be renamed with their original VSN with PVSREN (command RESTORE-LABELS-OF-PUBSET) in order to be able to work normally with the relevant disks. The renaming is possible in units of SM-/SF-Pubsets and volume sets.

Data security based on standby pubsets

Two (or even more) clone sessions associated with e.g. clone units C1 and C2 are started for each disk of the data/home pubset for the data/home pubset. At a defined point in time, the C1 clone units are split off by means of an ACTIVATE-CLONE command, thus creating a standby data pubset corresponding to the last day or production section or a standby home pubset following completion of the day's administration tasks.

At the next split time, the C2 clone units are separated by means of an ACTIVATE-CLONE command (thus making C2 the current standby pubset) and the disks of the now superfluous first standby pubset are then linked to the data/home pubset by means of a RESTART-CLONE-SESSION command to form a new clone pair.
If the original data pubset or the home pubset fails, operation can be resumed with the respective standby-pubset using the /SWAP-CLONE-SESSION and the subsequent /RESTART-CLONE-SESSION.
If the original data pubset or home pubset should fail, then operations can be resumed by reconstructing the pubsets from the relevant standby pubset – see Reconstruction of pubsets from clone units.

Exporting data (Migration)

The replication function Equivalent Copy (EC) can also be used for migrating data from one unit to another.
With Equivalent Copy (EC), a clone session is set up between the unit containing the data to be migrated and the target unit (= clone unit).
Following termination of the copy process and activation of the clone unit, the data is available on the clone unit. The clone session is now terminated and the clone unit can be used directly.
Alternatively, after migration has been completed, it is also possible to retain the mirroring and resume the replication function after the mirror direction has been reversed (swap function).

Contact
FUJITSU Technology Solutions
Barbara Stadler
Address: Mies-van-der-Rohe-Straße 8, 80807 München, DE
Phone: +49 (0)89-62060-1978
E-mail: Barbara.stadler@ts.fujitsu.com
Website: www.fujitsu.com/DE
2013-06-01 EM EN

Technical data subject to modification and delivery subject to availability. Any liability that the data and illustrations are complete, actual or correct is excluded. Designations may be trademarks and/or copyrights of the respective manufacturer, the use of which by third parties for their own purposes may infringe the rights of such owner. [Other disclaimers]