FUJITSU

WHITE PAPER

Use of Symmetrix TimeFinder Functions with BS2000/OSD

Issue December 2010

Pages 13

Ideally, mission-critical applications should be usable 24 hours a day, every day. Scheduled and unscheduled application downtimes, e.g. for backing up or restoring production data, should be kept to an absolute minimum or eliminated altogether. Current methods used for this are parallelization of processes by non-disruptive mirroring and splitting-off of replicas of the production data and accelerated backup and restore using disk-based backup/restore techniques.

In BS2000/OSD this is achieved through use of EMC^{2}Symmetrix disk systems and their TimeFinder functionality.

The present white paper describes the TimeFinder Clone, Snap and Mirror(BCV) options, their potential application scenarios and the customer benefits they offer, and deals with the support for the TimeFinder replication functions in BS2000/OSD.

Contents

TimeFinder replication functions of the Symmetrix hardware

The Symmetrix TimeFinder product family permits additional copies of volumes to be created within a Symmetrix system so that they can be processed separately after being split off.

TimeFinder/Clone

TimeFinder/Clone is a flexible and powerful method for creating copies of entire Symmetrix DMX or V-Max volumes. It enables one or more clone units to be produced from an original unit. TimeFinder/Clone is a copy at a specific point in time and usually is not a permanent mirror.
From the user viewpoint, a clone unit is an instantly available copy of an original unit. With Clone no preliminary run is necessary to create a copy, since the copy-on-access technique is used for clones:
TimeFinder/Clone in each case copies the tracks to the clone unit which are accessed for writing on a unit of the clone pair or for reading on the clone unit.
Two operating modes are possible:

- Clone can generally (only) operate in copy-on-access mode (copy as snapshot) or
- The storage subsystem copies the entire original unit to the clone unit in the background.

In both cases TimeFinder/Clone requires 100% additional disk space.
TimeFinder/Clone doesn't require special pre-configured devices. Any unit of the Symmetrix system can be used as clone unit. After the "clone pair" relationship has been dissolved, the two units of the (former) pair can be used freely. In particular, new clone pairs can also be formed. Up to 16 sequential clone units can be generated for one original unit.

TimeFinder product family

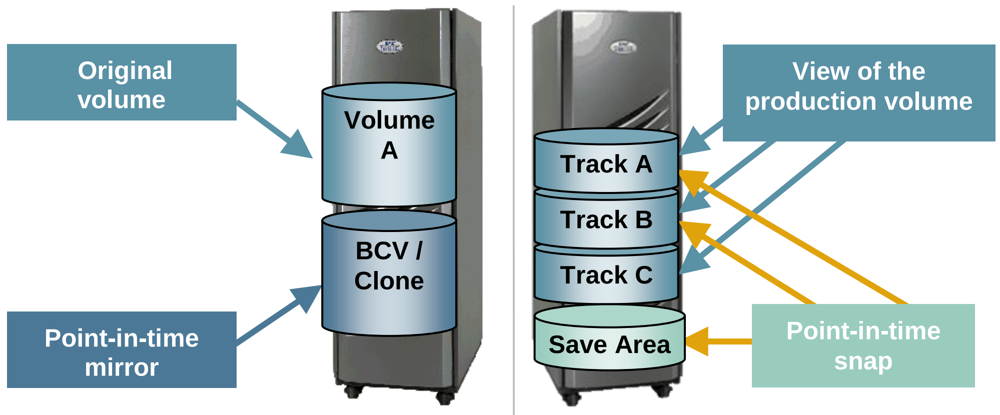

EMC TimeFinder/Mirror (BCV) or TimeFinder/Clone

- Full copies
- Requires 100% additional capacity
- Applications can access volumes independently

TimeFinder/Snap

- Pointer-based copies
- Requires ~30% additional capacity
- Applications have shared access to volumes

Source: EMC

TimeFinder/Snap

TimeFinder/Snap is a powerful, space-saving method for creating pointer-based copies of Symmetrix DMX or V-Max volumes. TimeFinder/Snap implements a virtual copy of an original unit at a random snap(shot) time, as follows:
At the first modification of a block, TimeFinder/Snap saves the contents of the original unit by writing the original contents of the block into a special Symmetrix area, called the Save Pool (copy-on-first-write technique).
At restore time, the original status is reinstated by writing back the differential blocks from the Save Pool to the original unit.
The Save Pool is present as standard as a single default pool. In Symmetrix DMX starting with Enginuity 5671 as well as in V-Max, further Save Pools can be set up in addition to the default pool. This can be used to partition the snap area and segregate different snaps from one another.
Considerably less disk space is required for snaps than for clones. Where there are several sequential snaps for an original unit, only the originals of the blocks modified since the next older status are stored. Up to 15 consecutive snaps can be generated for an original unit. As of Enginuity 5772 this number is elevated to a maximum of 127 consecutive snaps for FBA disks (the 127 snaps extension is not available for ESCON CKD Volumes). General release of this Enginuity 5772 function for operation in the BS2000 environment was effective in 2007, December.
No preliminary run for creating a copy is necessary with Snap.

In the current Symmetrix systems, up to 16 copy sessions are currently possible per original unit. This value comprises the total number of clone, snap and BCV sessions active at a given time. Clone sessions in which the storage subsystem makes a note of changed tracks (/START-CLONE-SESSION ... DIFFERENTIAL=YES), and the first snap session count double in this case.

As of Enginuity 5772 the Symmetrix systems can manage up to 127 snap units per original unit. These snap sessions count altogether as a single copy session.

TimeFinder/Mirror (BCV)

TimeFinder/Mirror (referred to previously and in the following as BCV) is the classic method for creating high-performance, highly available mirrors of Symmetrix volumes without disrupting production operation.
The volume containing the production data is referred to as the original unit and the volume containing the copy of the data as the BCV unit (**B**usiness **C**ontinuance **V**olume). The original unit and the BCV unit together form a BCV pair.
BCV pairs (original unit and BCV unit) consist of pre-defined unit types, which must be configured and hold available specially. In difference to TimeFinder/Clone and Snap, a forerun for creating a copy is necessary by TimeFinder/Mirror.
When working with TimeFinder/Mirror, the data of the original unit is initially copied to the BCV unit. On completion of the copy operation, both units continue to be maintained in synchronism until the separation of the BCV pair is initiated (split), e.g. at a recovery point or checkpoint for the data. The BCV unit is then available for separate processing. Once the separate processing has been terminated, the BCV pair can be restored.

Concurrent BCV

TimeFinder/Mirror provides the concurrent BCV function as standard. With concurrent BCV, up to two BCV units are maintained in synchronism with an original unit in the Symmetrix system.

Multi-BCV

TimeFinder provides the multi-BCV function as standard. This permits up to 16 different BCVs to be assigned to a single original unit. This means that copies from earlier points in time can be kept in addition to the synchronous BCV units.
All assignments and change information of an original unit relating to max. 16 BCV units are managed in the Symmetrix system, thereby optimizing e.g. the copying times of BCV mirrors used in alternation.

The product TimeFinder/Mirror is no more offered with Symmetrix V-Max. The compatibility is guaranteed furthermore with emulated TimeFinder/Clones.

Application scenarios and customer benefits

General: The additional copy enables critical applications to be productive for longer. 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 extending 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.
If the parallel data resources are not modified, the checkpoint is retained and can be used for fast restores.

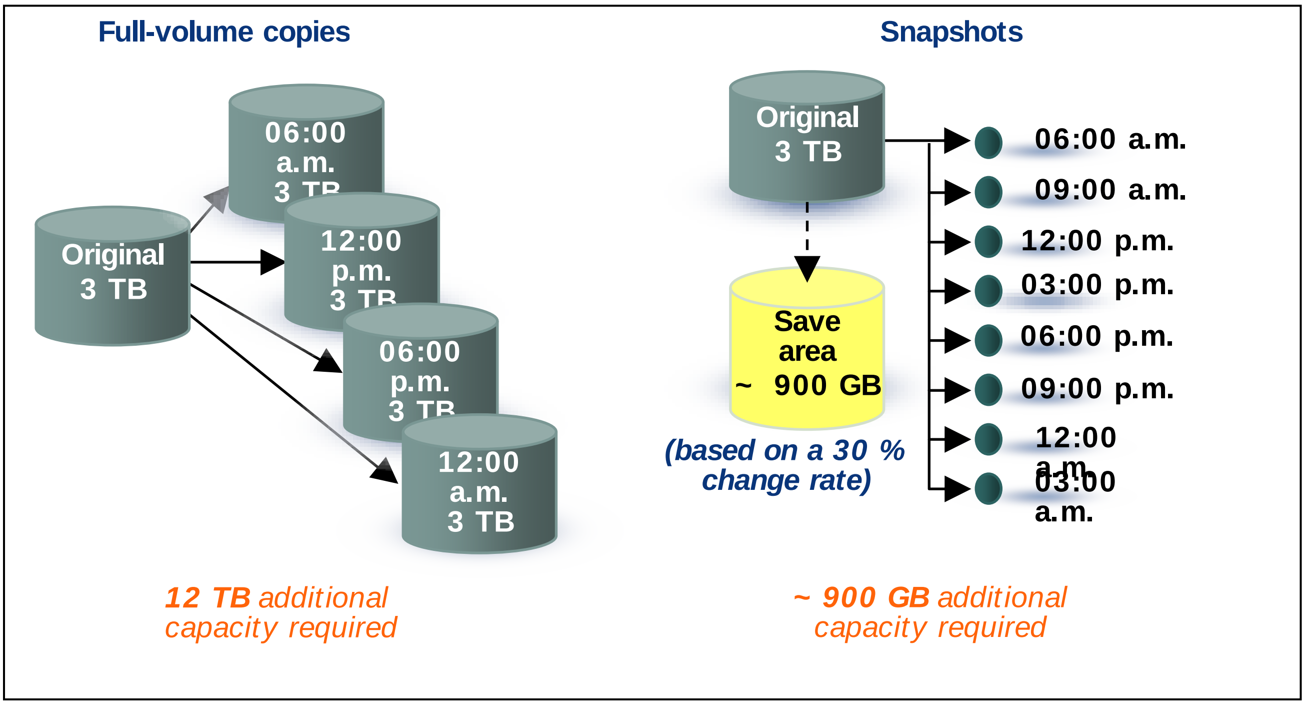

Positioning: Clone vs. BCV

TimeFinder/Clone provides everything a customer needs to create data replicas for parallel processing. Clones are used (like BCVs earlier) as input for tape backup and as a free disk copy. In some cases clone offers even more than BCV. The main advantage of clones is their flexibility: no specially preconfigured BCVs are necessary. A clone can become an original for a further clone session, whereas a BCV always remains a mirror.

Beginning with Symmetrix V-Max, the clone-mirroring replaces the BCV-mirroring. All BCV-scenarios can be mapped to clone-scenarios. Also in reference to the license BCV is an AddOn to Clone.

TimeFinder/Snap
The primary application scenario for TimeFinder/Snap is fast and efficient backup and restore from disk; it can also be used as input for tape backup and as a free disk copy (preferably for read accesses).
The copy is available immediately. There are no wait times for synchronization as with BCVs. Disk space is saved, since only space for pointers and changed blocks is required.

TimeFinder options compared

	TimeFinder/Clone	**TimeFinder/Snap**	**TimeFinder/Mirror**
Performance (impact on the host application with production I/O)	Supports high I/O load (copy runs in the background, parallel access to production I/O only as long as copying operation has not been completed)	Supports moderate I/O load (parallel access to production I/O)	Supports very high I/O load (independent access to original and BCV)
Data availability	Pointer-based copy, requires original volume to be available until copy operation has been completed	Pointer-based copy, requires original volume to be available at all times	Original and BCV are independent copies
Currency	Copy at Activate time	Copy at Activate time	Permanent mirror
Functionality	Clones are instantly accessible, no need to wait for end of copying operation Supports migration between different RAID levels	Snaps are instantly accessible	BCVs must be created as a BCV pair before the split
Costs	Requires 100% additional disk space	Requires only some additional disk space (mainly on the Save Pool), dependent on change volume and duration of the snap session	Requires 100% additional disk space

The following chapters of this white paper describe that the TimeFinder options Clone and Snap operation scenarios mentioned above are supported by BS2000 in a comprehensive manner.

- TimeFinder/Clone covers all scenarios that could be realized with BCV before. In BS2000 the Clone and BCV function is available at the pubset level. For BCVs the synchronization times before split must be considered, for clones they do not apply.
- TimeFinder Clones make sense if one or perhaps two copies are to be used independently from each other, because they provide pubset copies for different tasks. In BS2000 the TimeFinder/Snap function is available at the pubset level. Snaps are suitable as backups, because several copies can be hold as different backup instances requiring only some additional disk space, always bound to the original pubset.

Support for TimeFinder replication functions in BS2000/OSD

TimeFinder/Clone and TimeFinder/Snap are supported for Symmetrix DMX systems as of Enginuity 5670 and SHC-OSD V6.0. The release took place in December 2006 for BS2000/OSD-BC V5.0 or higher. Add-on functions such as e.g. snapsets are available starting with BS2000/OSD-BC V7.0, which was released in March 2007. TimeFinder /Clone and TimeFinder/Snap are supported on Symmetrix V-Max as of SHC-OSD V7.0A01, released in September 2009 for BS2000/OSD-BC V6.0 and higher. With BS2000/OSD-BC V8.0 and SHC-OSD as of V7.0, TimeFinder support is being extended to the BS2000/OSD SQ series Business Servers.
This White Paper describes the state SHC-OSD V8.0, released in January 2011 for BS2000/OSD-BC as of V6.0.

The TimeFinder/Mirror (BCV) is the classic TimeFinder function and has been supported in BS2000/OSD since the first release of SHC-OSD V1.0, the BS2000/OSD host component for Symmetrix systems. It is however no longer offered as of Symmetrix V-Max. Migration to the flexible and innovative TimeFinder/Clones is recommended. The compatibility is guaranteed furthermore by emulated TimeFinder/Clones.

Configuration

All TimeFinder functions operate on the basis of 'logical volume' granularity.

- One physical disk in the Symmetrix system can be distributed over multiple logical volumes - the BS2000/OSD server sees a logical volume as a disk (MN).

- With meta volumes (striped or concatenated), the Symmetrix system combines several logical volumes internally into a single logical volume and presents them externally (to the BS2000/OSD server) as one 'logical volume'.
- The disk type D3475-8F for Fibre Channel disks attached to SX and SQ servers permits a logical volume to be partitioned into up to 7 slices, thus enabling multiple MNs to be placed on one logical volume. When TimeFinder is deployed, however, disks with type D3475-8F may not be partitioned.

TimeFinder/Clone

TimeFinder/Clone requires no special preconfigured devices. Any units of the Symmetrix system can be used as clone units. The clone units must match the units to be copied in capacity, stripe size (if applicable) and emulated device type (FBA or CKD). As of Enginuity 5772 it is possible to create TimeFinder Clone units with a capacity larger than the original units.
The units used as clone units must be generated like any other disk during the hardware generation of the BS2000 servers with IOGEN (S servers) or configured in X2000 (SX/SQ servers).

TimeFinder/Snap

TimeFinder/Snap requires specially configured (by the service engineer) devices in the Symmetrix system, known as “save devices” and “virtual devices”.

- If, following activation of a snap unit, data is changed on the unit (or even on the snap unit), the original data is written to a **save device**. The **Save Pool** is a group of save devices.
- The **virtual device** consists of the pointers to the original data. For unchanged data, the pointers address the original unit. For changed data, the pointers address the corresponding save device. The virtual device of the Symmetrix system is the snap unit which can be accessed by the host (S or SX/SQ server).

Virtual devices must be generated using IOGEN for S servers, or configured in the X2000 device configuration of the SX/SQ servers.

TimeFinder/Mirror (BCV)

BCV units are set up in the Symmetrix system by the service engineer. The BCV units must be of the same disk type and have the same capacity as the original units to which they are assigned.
S servers
With S servers, the BCV units must be defined like any other disk in the file containing the I/O configuration data (IOCF) at hardware generation time using IOGEN.
SX and SQ servers
The BCV units belong to the FC peripherals of the SX and SQ servers; they are identified to the BS2000/OSD system by way of the configuration in X2000.

SHC-OSD

SHC-OSD provides information services via the global Symmetrix configuration, the Symmetrix device configuration and the SRDF and TimeFinder functions via BS2000/OSD command interfaces.
SHC-OSD also enables the SRDF and TimeFinder functions 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 commands for TimeFinder allow control of the entire TimeFinder product family, i.e. TimeFinder/Clone, TimeFinder/Snap and TimeFinder/Mirror (BCV).
SHC-OSD permits units to be selected via their VSN, their mnemonic name, their Symmetrix-internal number or via the catalog ID (catid) of the pubset or volume set to which they belong. Private disks are supported with restrictions. The most common application scenario is the processing of pubsets as a whole. The main features of this application scenario are described below.

Support for Clones

Clones can be integrated into SHC-OSD starting with SHC-OSD Version 6.0. Specifically, the following functions are provided for controlling the clone functionality:

TimeFinder/Clone Operating Sequence

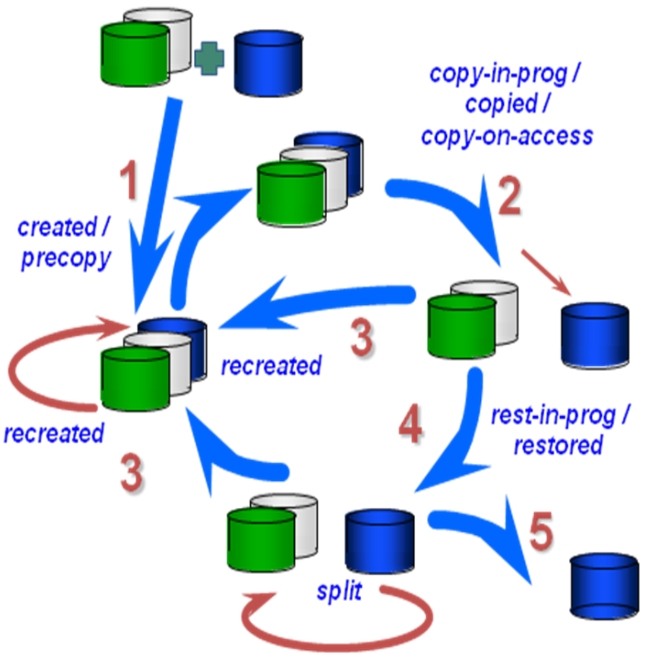

1 /START-CLONE-SESSION

Pairing of independent volumes

States: **created / precopy**

2 /ACTIVATE-CLONE

Split, optionally renaming for separate operation

States: **copy-in-prog / copied / copy-on-access**

3 /RESTART-CLONE-SESSION

copies changed data from original

Status: **recreated**

4 /RESTORE-FROM-CLONE

copies changed data onto original

States: **restore-in-prog / restored / split**

5 /STOP-CLONE-SESSION

finishes clone mirroring

Status: **-**

/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 and format (FBA/CKD) can be used as a clone unit. As of Enginuity 5772 it is possible to create TimeFinder Clones with a capacity larger than the original volumes. SHC-OSD V7.0 supports this function. Clone units for all volumes of a pubset can be created with one /START-CLONE-SESSION command.
The DIFFERENTIAL operand can be used to determine whether the storage subsystem makes a note of changed tracks. If it does not, no restore and restart will be possible later.
The COPY-COMPLETE-DATA operand can be used to determine whether the storage subsystem is to copy the entire original unit to the clone unit in the background or whether it is to copy only those tracks to the clone unit which are accessed for reading on a unit of the clone pair or for reading on the clone unit. If the entire unit is to be copied, it can additionally be determined whether the copy process will start after the creation of the clone pair or only after its activation.

/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.
Before the command is invoked, a recovery point (checkpoint) should be created for the data (see similar recommendation for /HOLD-MULTI-MIRRORING). In order to coordinate activation of all the clone units of a pubset, HOLD-IO = *UNTIL-ACTIVATED can be used to halt inputs/outputs to the clone pair until all clone pairs of the pubset are in the COPY-ON-ACCESS state. After this, the inputs/outputs on the units are resumed.
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 command processing.
After the clone pair has been activated, copy-on-access processing starts, i.e. (only) those tracks are copied to the clone unit which are accessed for writing on a unit of the clone pair or for reading on the clone unit. The pair is in the COPY-ON-ACCESS, COPY-IN-PROGRESS or COPIED state, depending on which options were specified when it was created.

/RESTORE-FROM-CLONE Restore clone original
/RESTORE-FROM-CLONE enables the original unit of a clone pair to be reconstructed from the clone unit. A restore is only possible for pairs in which the original unit was copied in full to the clone unit and in which the storage subsystem has kept a record of changed tracks.

/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. A restart is only possible for pairs in which the original unit was copied in full to the clone unit and in which the storage subsystem has kept a record of changed tracks.

/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.

TimeFinder/Clone can also be used in conjunction with SRDF for SRDF source units or SRDF target units.

As of SHC-OSD V8.0 the new function **Cascaded TimeFinder/Clone** is offered:

A clone unit can be used in parallel to its clone unit feature as an original for a further TimerFinder/Clone mirroring. A clone of the clone is basically created.
A consistency version already created on a clone unit can thus be provided or migrated again in an efficient manner and parallel to running applications without the original clone relationship being stopped; i.e. without any effect on the productive application and without a new complete copy of the data (only delta copies). For support of Cascaded TimeFinder/Clone SHC-OSD V8.0 was extended as follows:

- Action commands take cascaded TimeFinder/Clone into account,
- The command /SHOW-SYMMETRIX-DEVICE-CONFIG indicates in TimeFinder column for clone use a “B” (both) for TimeFinder cascaded clone volumes (volume, which are original of a TimeFinder/Clone-Mirroring as well as TimeFinder clone units)
- The command /SHOW-CLONE-SESSION displays 2 lines with the both clone pairs for cascaded clone pairs

Support for Snap

Snap is integrated in SHC-OSD starting with SHC-OSD Version 6.0. Specifically, the following functions are provided for controlling the snap functionality:

/START-SNAP-SESSION Create snap pair
/START-SNAP-SESSION creates one or more snap pairs, with one virtual device being assigned as a snap unit to each original unit. Snap units for all volumes of a pubset can be assigned with one /START-SNAP-SESSION command. A virtual device of the same type must be available as a snap unit for each volume of the pubset.

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

Before the command is invoked, a recovery point (checkpoint) should be created for the data (see similar recommendation for /START-MULTI-MIRRORING and /ACTIVATE-CLONE). When activating for the granularity “pubset”, the IOs can be halted (HOLD-IO), in a similar way to clones and BCVs, in order to ensure the pubset is in a crash-consistent state.
Following successful activation, original and snap unit are separated: Both are available to applications from the host. From the viewpoint of the applications, the snap unit contains all the data of the original unit at the time of command processing.
After the snap pair has been activated, all data pointers of the snap unit point to the original data of the original unit. From this time onward, the original data of all data changed on the original unit is saved on a device of the Save Pool and the data pointers of the snap unit are updated to point there.

/RESTORE-FROM-SNAP Restore original from snap
/RESTORE-FROM-SNAP enables an original unit of a snap pair to be reconstructed from the snap unit. The snap units can be selected for the restore either explicitly (e.g. via the catid of the renamed pubset) or implicitly. The restore causes all changes made on the original unit since the activation of the snap unit to be discarded.

/STOP-SNAP-SESSION Dissolve snap pair
/STOP-SNAP-SESSION terminates the assignment of one snap pair or multiple snap pairs. The snap is discarded, i.e. the data in the Save Pool and the snap unit can no longer be accessed.

The functions for partitioning the Save Pool (starting with Enginuity e5671) are supported as follows:

- A Save Pool can be selected at the start of the snap session.
- SHOW-SYMMETRIX-CONFIGURATION displays all the Save Pools in a Symmetrix system.
- Monitoring functions are available for monitoring the use of the Save Pool or individual Save Pools of the Symmetrix system. Warning messages are output to the console when certain fill levels are reached or changed (to be set via /MODIFY-SHC-PROCESSING ... SNAP-SAVE-POOL-SAT).

During the separate processing of the snap units i.e. when writing in the Save Pool after /ACTIVATE-SNAP, actions involving a large change volume, such as e.g. data reorganization with SPACEOPT on the original pubset, should be avoided, because this leads to a heavy load (I/O rate, space consumption) of the Save Pool.

Creation of a Standard Snap Save Pool
SHC-OSD V7.0 allows defining a Default Snap Save Pool for TimeFinder/Snap via command or via SHC-OSD parameter file.

If the SRDF function is used, the original unit can also simultaneously be the source unit of a remote-copy pair. The target unit can also be the original unit of a snap pair.

As of Enginuity 5772 a maximum of 127 consecutive snaps can be created per original unit for FBA disks. This extension is supported by SHC-OSD V6.1. Release of Enginuity 5772 with the 127 snap extension for operation in the BS2000 environment, together with the appropriate corrections for SHC-OSD V6.1, was effective in 12.2007.

Support for BCV

The following functions are provided for controlling the TimeFinder/Mirror (BCV) functions:

- Create Business Continuance Volumes (BCVs) which can be addressed independently of one another
- Output status information about BCV pairs and units
- Consistently separate BCV pairs for independent processing of the data
- Restore BCV pairs

Specifically, the following functions are provided for controlling the BCV functionality:
(Note on terminology: In the SHC-OSD manual, the terms “normal unit”, “additional mirror unit” and “multi-mirror pair” are mostly used instead of the terms “original unit”, “BCV unit” and “BCV pair”.)

/START-MULTI-MIRRORING Create BCV pair
/START-MULTI-MIRRORING creates one or more BCV pair(s), with one BCV unit being assigned to each original unit. With /START-MULTI-MIRRORING for a pubset, a BCV unit must also be available for each original unit of the pubset.
When /START-MULTI-MIRRORING is executed, a copy process is initiated in the Symmetrix system to copy the data of the original unit onto the BCV unit (or vice versa), i.e. the BCV pair is created. After being created, each data change is applied to both units, i.e. to the original data and to the copy.

/HOLD-MULTI-MIRRORING Separate BCV pair
/HOLD-MULTI-MIRRORING separates one or more BCV pairs, each consisting of one original unit and a BCV unit. /HOLD-MULTI-MIRRORING for a pubset separates all the BCV pairs of this pubset.
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 separation, the BCV unit is available for the host and contains all the data of the original unit at the time of command processing. Original unit and BCV unit can each be processed independently, by different applications, for example.

/RESUME-MULTI-MIRRORING Reconstruct BCV pair
/RESUME-MULTI-MIRRORING can be used to reconstruct a BCV pair that was separated previously with /HOLD-MULTI-MIRRORING. The data can either be copied from the original unit to the BCV unit or vice versa, i.e. mirroring can be resumed on the basis of the production data status updated in the interim or rolled back to the (saved) status at the time of the split.

/STOP-MULTI-MIRRORING Dissolve BCV pair
/STOP-MULTI-MIRRORING terminates the assignment of a BCV pair. Following this, the BCV unit can be used to mirror another unit.

Multi-BCV and concurrent BCV
The changes to an original unit compared to all associated, split-off BCV units are stored in the Symmetrix system. Consequently, when a BCV unit is reassigned to the original unit, only the changes made in the interim need to be copied. Concurrent BCV pairs can be created by two successive /START-MULTI-MIRRORING commands with the same original unit.

TimeFinder/Mirror can also be used in conjunction with SRDF. The BCV unit can be attached to the SRDF source unit or SRDF target unit.

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 or BCV units of an SF or SM pubset. PVSREN provides the statement CREATE-PUBSET-FROM-MIRROR for this purpose.
The creation of independent (“free”) pubsets based on snaps is also supported in PVSREN. Free snap-based pubsets are (only) useful for short-term use with predominantly read access. Only Clone and BCV 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.
As well as 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.
- Starting with BS2000/OSD-BC V7.0, PVSREN automates the 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 as part of the commands for separating the BCV units (/HOLD-MULTI-MIRRORING) and the commands for activating Clone and Snap (/ACTIVATE-CLONE, /ACTIVATE-SNAP). 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. If concurrent BCVs exist and these are separated together, both can be renamed directly to two different new catids.
The PVSREN program can complete renamings 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 clones and BCVs 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. After BCVs have been backed up, disk mirroring with SHC-OSD can be resumed (/RESUME-MULTI-MIRRORING).

Data backup with HSMS

Data backup of clone or BCV

The software product HSMS can be used in combination with the CCOPY function for data backup using the Symmetrix functions TimeFinder/Clone (Clone in HSMS as of V8.0B and BS2000/OSD-BC as of V7.0) and TimeFinder/Mirror (BCV).
The system administrator sets up the BCV mirror with /START-MULTI-MIRRORING and the clone session with /START-CLONE-SESSION. For HSMS backups using the TimeFinder function, the granularity “pubset”, i.e. for all volumes of a pubset, is required for mirroring. The HSMS user interface is identical for clones and BCVs. The operand CONCURRENT-COPY=*YES(WORK-FILE-NAME=*BY-ADDITIONAL-UNIT) must be specified in the HSMS //BACKUP-FILES statement.
(In order to refer to the neutral use of clone or BCV disk copies, the term “BY-ADDITIONAL-MIRROR-UNIT” in the HSMS //BACKUP-FILES statement is replaced by “BY-ADDITIONAL-UNIT” in HSMS V8.0B. However, the former term is also supported due to compatibility reasons.)
Embedding the function in this way enables TimeFinder to be used transparently for backup.
The TimeFinder 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/BCV 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 of BCV units.
When the Clone/BCV 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/BCV unit.
After the backup has been successfully terminated, mirroring of the BCVs is resumed automatically. This can be prevented by specifying RESUME-MIRRORING=*NO in the BACKUP-FILES statement.
Note:
Clone support in HSMS is based on the same HSMS call interface as for HSMS backup of BCVs:
In other words, the system administrator initializes the clone session with /START-CLONE-SESSION (similar to initializing the BCV mirror with /START-MULTI-MIRRORING), HSMS is called with the same user interface for clones and BCVs, the differentiation for clone or BCV is done within CCOPY. If a clone is present, the clone pairs are then activated by HSMS internally via a privileged program interface of SHC-OSD (similar to the splitting of the BCV pairs using this interface). This mode of use requires an Enginuity version 5671 or later.

If for the same pubset there were created as well Clones (START-CLONE-SESSION) as well as BCVs (START-MULTI-MIRRORING), a BCV mirror is created by BACKUP-FILES ….,CONCURRENT-COPY=*YES(WORK-FILE-NAME=*BY-ADDITIONAL-UNIT), by compatibility to BS2000/OSD-BC < V7.0. Then backup is performed from the split BCV units.

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 Symmetrix TimeFinder/Mirror or TimeFinder/Clone functions 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 or BCV 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 or BCV
- 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 BCVs 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 BCV, COPY ... USING DIRECTORY *hsms_archive_name* BY_ADD_MIRROR_UNIT can be used to back up database files residing on a BCV 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 BCV unit is separated from the original unit. Read DML accesses to the database are possible here.
- The specified database files are saved from the BCV 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 BCV 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 BCV into the HSMS archive, it is possible (starting with SESAM/SQL V4.0 and BS2000/OSD-BC V6.0) 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 BCV 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 BCVs 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 BCV unit.

In combination with HSMS as of V8.0 and BS2000/OSD-BC as of V7.0, SESAM/SQL also provides use of the HSMS backup of TimeFinder/Clone that is integrated into the DB backup functions, since clone support in HSMS is implemented using the same HSMS call interface as the HSMS backup of BCVs (see above, section titled “HSMS – Data backup of Clone or BCV).

When backing up BCVs, the BCV mirror of a remote data set managed via SRDF can also be used instead of a local BCV mirror (requires SESAM/SQL V4.0 or higher).

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 or BCV 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 in BS2000/OSD-BC as of V7.0

BS2000/OSD-BC as of V7.0 supports snapshot-oriented backup/restore scenarios in DMX 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 and entire pubsets. 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 is based internally on the functions for TimeFinder/Snap control in SHC-OSD as of V6.0.
With the Enginuity versions until 5771, a maximum of 15 snapsets can be managed. For the extension as of the Enginuity 5772 update, 26 snapsets were implemented in BS2000/OSD-BC V7.0 in advance. In the following version BS2000/OSD-BC V8.0 the maximum was doubled to 52 snapsets. So a customer taking backups on working days can cover more than one monthly period using only snapset backups and does not have to resort to conventional HSMS tape backups in the period preceding this.
The support of more than 15 snaps as of Enginuity 5772 is not available for ESCON CKD disks. That means more than 15 snapsets are only possible for pubsets consisting exclusively of type D3435 disks (FBA disks).
The maximum number of snapsets that is allowed for a single pubset can be specified per command (/SET-PUBSET-ATTRIBUTES ..., SNAPSET-LIMIT=).
Snapsets are created and deleted by the administrator using the new commands CREATE-SNAPSET and DELETE-SNAPSET.
The administrator can assign a special Snap Save Pool to the snapset with the SET-SNAPSET-PARAMETERS command.
In shared pubset mode, the snapsets can be used by all the systems.
In a disaster recovery scenario using SRDF mirroring, 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-FILES-FROM-SNAPSET, RESTORE-JVS-FROM-SNAPSET) are provided for this purpose.
In BS2000/OSD-BC V8.0, the snapset end user functions of listing and restoring files and job variables from snapsets are also provided via program interfaces.
In addition several information functions have been extended in BS2000/OSD-BC V8.0 to support the larger number of snapsets.

LMS also supports the restoration of library elements from snapsets (as of supplementary version LMS V3.3B).

Starting with V9.0, VM2000 supports the use of snapsets on guest systems running BS2000/OSD-BC V7.0 or higher.
Under VM2000 V9.0, 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>).
As the total number of snapsets for a pubset is limited by the disk subsystem, snapsets normally will be created and deleted in the same cycle. Before a snapset is deleted, its data can be transferred in a backup archive using HSMS, if the corresponding backup level has to be conserved for a longer time.
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 compared

FDDRL backs up from disk copies in order to get a backup without interruption during the production running (e.g. from the home pubset).
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.
A standby pubset concept is included with the snapsets.
If a disk controller fails because of a catastrophe (e.g. a fire), FDDRL or HSMS backups on tape are still available, in contrast to the snapsets that are in the same controller as the original disks. The latter can (should) be compensated by doubled snapsets with SRDF operation or by transferring backup data from snapsets to HSMS.

Volume-based restore of pubsets

When restoring pubsets from TimeFinder copies it should be borne in mind that the different TimeFinder functions, in particular TimeFinder/Clone and TimeFinder/Snap, mutually influence one another. Thus, for example, when a pubset is restored from a snapset, no clone session must be running, and vice versa. Restored snap and clone sessions should also be terminated as quickly as possible after completion of the restore. Detailed information on this can be found in the SHC-OSD manual.

Restoring pubsets from snapsets

Using the command /RESTORE PUBSET-FROM-SNAPSET, the administrator can restore an entire pubset from the latest snapset. In volume-based restore of pubsets, the most recent snapset is copied back onto the volumes of the pubset. It is taken into account here that the number of volumes of the original pubset at the snapset creation time may have been less than the current number of snap units, since, of course, volumes can be added dynamically to a pubset. Restore pubset is not possible if the pubset has been reduced in size since the creation of the snapset.

Restoring pubsets from clone units

SHC-OSD (/RESTORE-FROM-CLONE) enables the original unit of a clone pair to be restored from the clone unit if the clone unit was created as a full copy. In this case too, clone units with special notation must subsequently be renamed to their original VSN using PVSREN.

Restoring pubsets from BCV units

A pubset can be restored from BCV volumes by rolling back the mirroring with SHC-OSD (/RESUME-MULTI-MIRRORING) to the status of the BCV units. BCV volumes created by SHC-OSD with special notation must subsequently be renamed back to their original VSN using PVSREN (RESTORE-LABELS-OF-PUBSET statement) in order to be able to work with the corresponding disks in normal mode. The renaming is possible in units of SM/SF pubsets and volume sets.

Data security based on standby pubsets

Using two alternating standby data pubsets or standby home pubsets

This is a classic application scenario for multi-BCV and/or concurrent BCV. Two additional mirror disks per disk of the data/home pubset are held in reserve for the data/home pubset. One of the two is always an active mirror of the corresponding disk of the data/home pubset. The other is a mirror, split off at a defined point in time by means of a HOLD-MULTI-MIRRORING command, which, together with the other split-off mirrors, represents 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 currently active mirrors are separated by means of the HOLD-MULTI-MIRRORING command and the disks of the now superfluous first standby pubset are then resynchronized with the data/home pubset using the RESUME-MULTI-MIRRORING command. Only the data that has changed is synchronized in this process. This ensures that a permanently up-to-date standby data pubset (with the status of the last day or production section) is present and can replace the original data pubset if the latter should crash. Similarly, if the home pubset crashes due to logical errors and the resulting system crash, operation can be resumed with an up-to-date standby home pubset.

Standby pubset when using clones

Two clone sessions associated with 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 crashes, operation can be resumed with the relevant standby pubset by means of RESTORE-FROM-CLONE. An Enginuity version 5671 or later is a requirement for this usage mode.

Exporting data (migration)

The TimeFinder/Mirror and TimeFinder/Clone functions can also be used for migrating data from one unit to another.
With TimeFinder/Mirror, a BCV unit is assigned for this purpose to the unit containing the data to be migrated. Following synchronization, the BCV unit can be split off using /HOLD-MULTI-MIRRORING and assigned to a new original unit with /START-MULTI-MIRRORING. RESTORE=*TO-ORIGINAL causes the data to be copied from the BCV unit to the new original unit, which then also contains the migrated data.
Following termination of the copy process, the BCV unit can be separated again using /HOLD-MULTI-MIRRORING and used for other purposes.

With TimeFinder/Clone, a clone session with full data copy is set up between the unit containing the data to be migrated and the target unit (= clone unit). Following activation of the clone unit and termination of the copy process, the data is available on the clone unit. The clone session is now terminated and the clone unit can be used directly.
This means that the migration scenario is ideally suited to clones, since this allows the migration to be completed directly in a single step. Taking a detour via a special BCV is not necessary.

Migration via TimeFinder/Clone onto Clones > Original
As of Enginuity 5772 it is possible to create TimeFinder Clones with a capacity larger than the original volumes. SHC-OSD as of V7.0 supports this function. The extended capacity on the clone can be made usable for BS2000 by means of the SPACEOPT product.

Published by department:

Barbara Stadler
Phone: ++49 (0)89 62060 1978
Fax: ++49 (0)89 62060 329 1978
Barbara.Stadler@ts.fujitsu.com
ts.fujitsu.com

Extranet
extranet.ts.fujitsu.com