THE POSSIBILITIES ARE INFINITE FUJITSU

WHITE PAPER

BS2000/OSD V8.0

Issue Mai 2009

Pages 20

BS2000/OSD V8.0 is the new operating system version for BS2000/OSD.

This is the first BS2000/OSD version supporting BS2000/OSD on high-end Intel x86 servers. The new BS2000/OSD SQ-server line and BS2000/OSD V8.0 enable /390-object-compatible BS2000-applications to run on x86-servers. The x86-platform is widely regarded as convergence platform for many types of servers. It seems that the server-world follows the same path as the client world did before.

By the general release of BS2000/OSD-BC V8.0 the SQ100 models will be offered as the first servers of this new server line. The SQ100 is a powerful mainframe for the entry- and medium-sized performance range.
In the first stage BS2000 native mode will be supported on SQ servers with BS2000/OSD V8.0. Later on parallel operation of multiple BS2000 guest systems and mixed operation of BS2000, Linux and Windows will be added.

Further highlights in BS2000/OSD V8.0 include performance measures to optimize data backup to fast tape devices (LTO) as well as improved integration of the FibreCAT CX storage systems into BS2000/OSD.

Important functional enhancements in BS2000/OSD V8.0 relate to:

Support for new hardware

- **Support for the new business server line of the SQ series based on standard Intel processors**
- **Support for LTO-4 MTC devices and LTO-4 drive encryption support with MAREN V12.0**

Enhanced scalability/performance

- **Optimization of disk IOs**
- **Faster backup of small and medium-sized files with HSMS V9.0**
- **Multiplexing up to factor x4 in FDDRL V17.0**

Improved storage integration

- **BS2000/OSD integration for FibreCAT CX attached to SX and SQ servers with SHC-OSD V7.0**
- **Extended support for the Symmetrix DMX TimeFinder/Snap function**

More simple and effective BS2000 operation

- **Mail interface in BS2000/OSD**
- **BS2ZIP extensions**

Extended openness and integration ability

- **New POSIX A41 version providing transparent access to BS2000 files in POSIX and POSIX-Sockets extensions**
- **New version Apache V2.2 with integrated SSL support**

General release of BS2000/OSD-BC V8.0: May 2009

New versions of system-related software products will also be released together with BS2000/OSD-BC V8.0. The main functional enhancements to certain selected products (FDDRL, HSMS, MAREN, openNet Server, SECOS and SHC-OSD) are also presented in this preview.

Contents

Hardware support

Support for the BS2000/OSD business servers of the S series and SX series

BS2000/OSD-BC V8.0 supports all currently released servers of the S series and SX series. For the SX servers, the BS2000/OSD-BC V8.0 variant ported to the SPARC64 architecture is released as part of the OSD Extended Configuration package OSD/XC V4.0.

Support for the new SQ server line

BS2000/OSD-BC V8.0 supports the new SQ100 servers of the SQ series, with high-end Intel x86 servers as the system platform. The X2000 hardware abstraction layer based on Linux and Xen acts as the carrier system for running BS2000/OSD and its applications. Xen is a wide spread virtual machine monitor for the x86 platform. The use of the Xen virtualization software means that it will be possible in the medium term to run other operating systems such as Linux and Windows on the x86 hardware, alongside BS2000.
The release is slated to take place in stages:

- At the time of the general release of BS2000/OSD-BC V8.0, BS2000 native mode will be supported on SQ servers.
- Approx. one year later, an X2000 correction version will enable VM2000 to run on SQ100. A new VM2000 version (planned VM2000 V9.5) will be released for this purpose.
- As well as the parallel operation of multiple BS2000 guest systems using VM2000, it will also be possible at a further stage to run Linux and Windows as additional guest systems.

As well as for SX servers, the OSD Extended Configuration package OSD/XC V4.0 is available as an operating system package for the SQ servers. The release of OSD/XC V4.0 provides SQ servers with a derivative of BS2000/OSD-BC V8.0, ported to the x86 architecture.

HW/SW architecture of the SQ servers with BS2000/OSD and X2000

In order to run BS2000/OSD on the SQ servers, the architectural principle based on the X2000 HW abstraction layer used for SX servers will be transferred to the high-end Intel x86 servers. For this purpose X2000 uses Linux as base system.
On SQ servers, X2000 has almost the same functions as on SX servers, i.e.

- X2000 provides the hardware/software interface for running BS2000/OSD and its applications in /390-compatible mode by means of HAL (Hardware Abstraction Layer) and the CISC firmware,
- X2000 performs the I/Os and the server administration and operation.

On the SQ server Linux does not run in native mode but as a privileged Linux system controlled by the Xen hypervisor.
The main difference between the concept with Xen and the former SX concept is the fact that on the SQ server BS2000 runs as a Xen guest system.
The SQ servers' HW/SW architecture including the Xen function for running BS2000/OSD on SQ is described in more detail below.

What is Xen?
The **Xen** software is an open source virtual machine monitor (VMM) that is shipped as part of the Suse Linux distributions (SLES10).
Xen consists of a number of components. The Xen hypervisor runs directly on the hardware and handles CPU, memory and interrupt tasks. On top of the Xen hypervisor, the Linux Xen kernel is booted. This privileged Linux system is called Domain 0 and is the first guest operating system under Xen. In Xen terminology, the virtual machines are called domains.

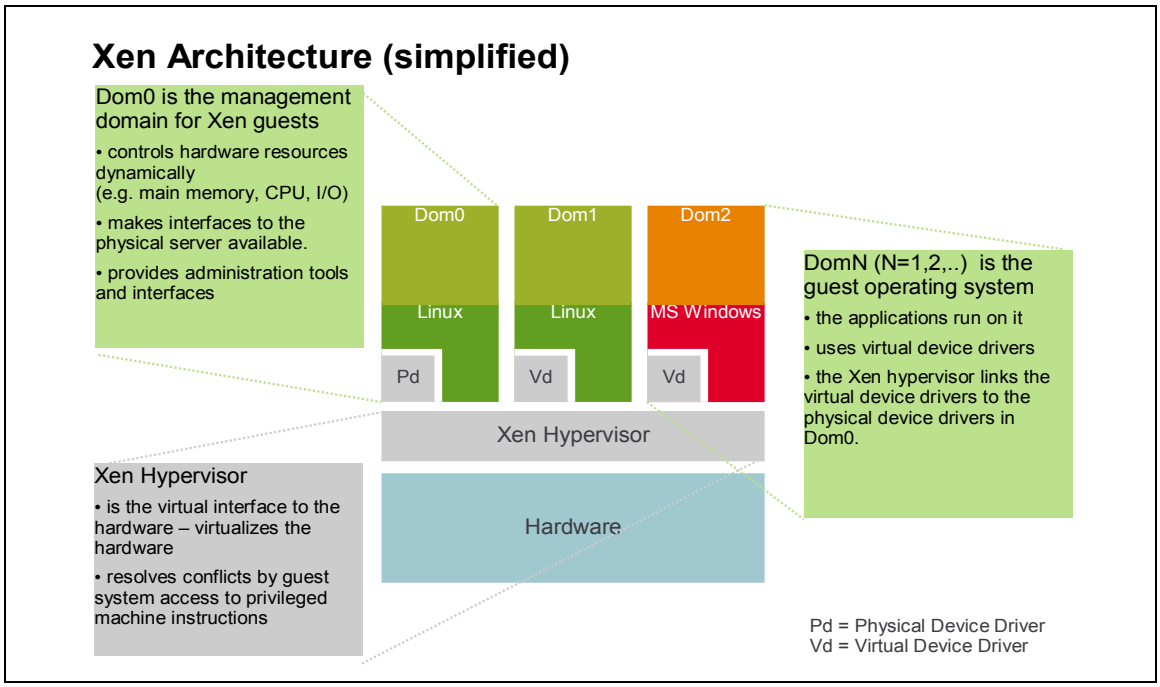

Domain 0 serves for managing the domains and hosts the drivers of the connected peripheral devices. The following domains invoke the backend drivers of Domain 0 via their frontend drivers in order to access hard disks or network devices of the underlying hardware. Privileged instructions from within the domains are forwarded via hypercalls directly to the Linux Xen kernel of Domain 0, the kernel then in turn accessing the underlying hardware.

BS2000 as a guest system on Xen
BS2000 runs on the SQ server as a Xen guest system in a domain called “DomBS2” under the control of the “X2000 µkernel”. The µkernel is a software layer that is needed so that BS2000 can run as a guest system under a Xen hypervisor. This µkernel corresponds to the part of the current SX server firmware X2000, which on SX servers runs on the BS2000 processors. The HAL (Hardware Abstraction Layer) and the CISC firmware (inter alia) are integrated into the µkernel. On SQ servers HAL is the firmware component that maps the privileged hardware/software interfaces of the /390 mode to equivalent interfaces of the x86 mode. On SQ servers the CISC firmware serves to map unprivileged /390 code to x86 code. The X2000 µkernel also handles communication with domain Dom0 for managing the BS2000 I/Os.
Running BS2000 as a Xen guest system serves to realize BS2000 native operation on the SQ server. The VM2000 product will enable multiple BS2000/OSD systems to run in parallel also on SQ servers (see section VM concept for SQ servers on p.5).

Running BS2000 and its applications on SQ servers
The privileged parts of the BS2000/OSD-BC V8.0 operating system and the system software running in privileged mode are ported to the x86 hardware with the aid of compiler technology (ASSTRAN) developed at Fujitsu. Existing (unprivileged) customer applications in /390 code run unchanged in object-compatible mode by means of the CISC firmware, while new or recompiled applications are likewise compiled by default in /390 code using the compilers provided for the SQ servers.

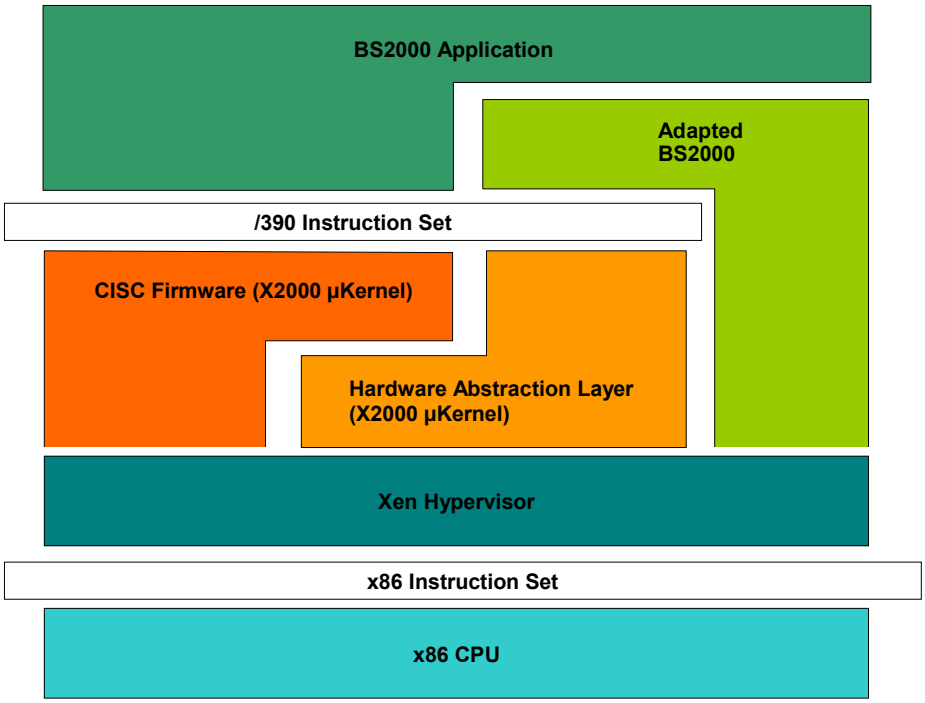

Handling inputs/outputs
The CPUs of the SQ servers are logically divided into CPUs for applications running under BS2000/OSD (BS2000 CPUs) and one CPU for performing inputs/outputs (I/O processor). The BS2000 CPUs are assigned to domain DomBS2, and the CPU for performing inputs/outputs is assigned to the privileged domain Dom0.

Dom0 is the privileged domain which provides the operating functionality and handles the physical I/Os. It contains X2000 with the /390 device emulation and administration functions.

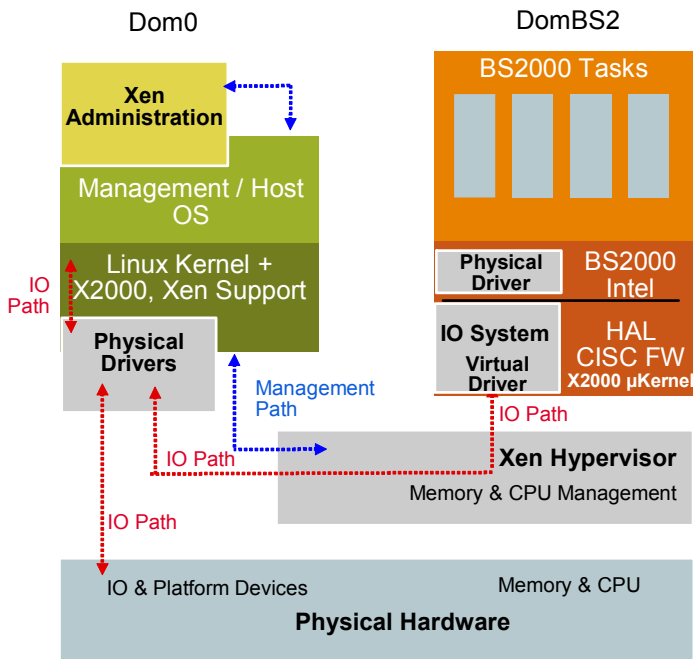

The X2000 functionality (BS2000 device definition and emulation, KVP, …) currently existing on the IOP side is ported to Linux and runs in the privileged domain Dom0.
BS2000 devices are emulated by the X2000 backend driver and mapped to devices and services that are supported in Linux.
The X2000 µkernel in DomBS2 maps the BS2000 device drivers to virtual drivers which communicate with Xen and X2000 in Dom0. I/Os are then handled via the physical device drivers in Dom0.

Preview: VM concept for SQ servers

VM2000 on Xen
VM2000 will enable multiple BS2000/OSD systems to run in parallel also on SQ servers. VM2000 is made available with a Xen-based architecture on SQ servers, while retaining the VM2000 user interfaces.
In this case the Xen hypervisor handles the operation and virtualization of the hardware.
On the S and SX servers, VM2000 essentially consists of the components VM2000 Monitor and VM2000 Hypervisor. The VM2000 Monitor component implements the user interface of VM2000, provides the internal accounting, logging and eventing functions, manages the VM administrators, and implements the $VMCONS functionality (i.e. network access for the VM administration and the virtual console).
On the SQ servers, the VM2000 Monitor continues to be available with a compatible command interface (initial exception: VM groups). The functions of the VM2000 Hypervisor are handled by the Xen hypervisor and a new VM2000 agent component in Dom0.

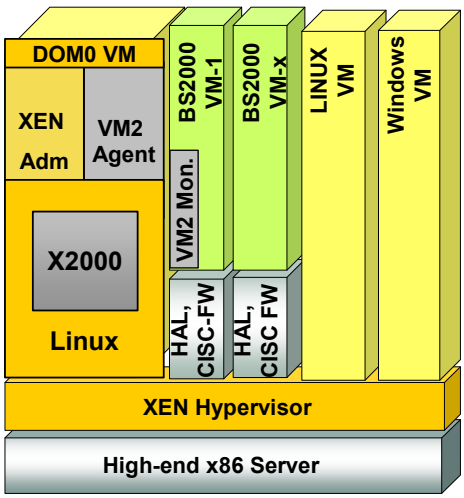

On SQ servers, the VM2000 command functions for operating VMs and managing the host are mapped to Xen administration functions.
The VM2000 Agent daemon running in Dom0 implements the mapping of the VM2000 Hypervisor functionality, with BS2000 guest systems also occurring as job submitters alongside the VM2000 Monitor. The agent maps these jobs to the XenAPI, to KVP calls (BS2-Start, BS2-Reset) and to jobs to BS2000 guest systems (SHUTDOWN-VM).

VM2000 functions in the guest systems
Currently, each BS2000 guest system has communication interfaces to the VM2000 Hypervisor in order to obtain information about VM2000 operation and initiate VM2000-specific actions in the guest system (e.g. implicit device assignment). On SQ servers, too, the BS2000 guest systems can continue using VM2000 functionality by means of the guest HPVCs (Hypervisor Calls) by mapping the HPVCs to Xen.
In VM2000 operation, each BS2000 domain (in contrast to the current global HAL and CISC firmware) has its own X2000 µkernel, in each case with HAL and CISC firmware.

VM2000 will be released for SQ servers approx. one year after the general release of BS2000/OSD-BC V8.0.

Running BS2000, Linux and Windows guest systems on SQ servers
As well as the parallel operation of multiple BS2000 guest systems supported by VM2000, it will also be possible at a later stage to run Linux and Windows as additional guest systems based on the Xen architecture.
A web-based (state-of-the-art) user interface is provided for the combined administration of BS2000, Linux and Windows guest systems.

Peripheral support: LTO-4 MTC devices

In addition to the existing LTO device types LTO-1, -2 and LTO-3, BS2000/OSD-BC V8.0 also supports the device type LTO-4. LTO-4 devices are designed for connection to the FC channel for operation with SQ and S servers in conjunction with a Quantum/ADIC Scalar 10K, i2000 or i500 library system. LTO-4 devices can also be used with SQ servers attached to an MTC autoloader via FC.
LTO-4 devices support even higher data rates than LTO-3 devices: The maximum data rate is 120 MB/sec (native), compared to 80 MB/sec with LTO-3. The minimum data rate required for streaming LTO-4 tapes remains nonetheless the same as for LTO-3, at 30 MB/sec.
When the high-performance LTO-4 devices are connected directly (i.e. without CentricStor) to a BS2000 host, a well-balanced configuration is a precondition, i.e. fast disk peripherals are also necessary, and the disks must be connected via FC channel, in order to reach the streaming rate.
In CentricStor, connection of LTO-4 devices was released starting with version V3.1D-SP02 P03 end of 2007.

LTO-4 tape encryption with MAREN V12.0

The LTO-4 drives are equipped with a 'tape encryption' hardware feature. With tape encryption enabled, tape access performance is reduced only by less than 1% in terms of data rate. This affords a highly effective means of implementing data protection at tape level, enabling tape contents to be protected against unauthorized reading, especially when in transit, when stored externally (e.g. in fireproof vaults) and when on loan.
Support for tape encryption is provided as an extension to the MAREN product for magnetic tape management in BS2000/OSD in MAREN V12.0. MAREN handles both the key management function and control of encryption and decryption. The encryption is performed in accordance with the AES standard using a symmetric 256-bit key. For details see section titled MAREN V12.0.
BS2000 control of LTO-4 encryption is only available when the LTO-4 devices are connected directly (i.e. without CentricStor) in conjunction with a library attached to a BS2000 host.

Scalability/performance

Optimizing disk inputs/outputs

Numerous fine-tuning tools are available for providing powerful support for the current online and nearline peripherals for BS2000/OSD servers. Optimal backup performance is predicated on a careful balance between online and nearline systems. The disk data rates are dependent on the degree of I/O parallelization, the transfer length and the alignment of the data on the disk. Multiple I/O jobs can be written in parallel to one logical volume using Parallel Access Volume (PAV) for disks attached to the FC or S channel type of the S servers and with Remote System Call (RSC) for emulated disks on SX and SQ servers. I/O times and I/O rates in TP operation together with I/O rates during sequential reading from a volume with RAID levels RAID5 and RAID1/0 (meta volumes) experience a significant increase with PAV/RSC.
To further optimize sequential reading from disk and consequently data backup to high-speed tapes (LTO), the internal transfer length will be increased in BS2000/OSD-BC V8.0.
The individual increases for the different disk formats are as follows:
- D3475-8F, from 32 to 80 PAM pages (= 160 KB)
- D3435,NK2 on Symmetrix, from 80 to 240 PAM pages (= 480 KB)
- D3435,NK2 on FibreCAT, from 80 to 128 PAM pages (= 256 KB)

The bigger internal transfer length is used in HSMS/ARCHIVE V9.0, in SPACEOPT V5.0 and in the COPY-FILE command in OSD V8.0. In OSD V8.0A the COPY-FILE I/O size will be augmented from 64 KB to 128 KB.
COPY-FILE measurements for files on Symmetrix D3435 disks were performed with an early OSD V8.0 version on an S190 server. The CPU demand for large files showed significant improvements up to 20%. The elapsed time was enhanced up to 15 %.

Performance improvements for small files during logical backups with HSMS/ARCHIVE V9.0

With the logical backup products ARCHIVE/HSMS, backup of small to medium-sized public files (size = several hundred PAM pages) runs 30-50% slower than in the case of large files due to the prorated OPEN/CLOSE overhead. Currently, out of the overall processing sequence “OPEN -> read useful data and transfer --> CLOSE”, which is executed per file in the backup product, the OPEN and CLOSE sections cannot be parallelized. Possibly, however, the “read useful data and write to tape” can be performed asynchronously (with the optimizations in HSMS V8.0B), thereby causing the proportion of time for OPEN/CLOSE to become [more] dominant.
Performance improvements by up to 30% for small files are likely as a result of the following measures:

- Disk I/Os for reading catalog entries minimized: The files are managed in order of their storage on the TSOSCAT files catalog. The overall result is fewer disk I/Os on the files catalog.
- The catalog entry write during the OPEN in the archive subtask eliminated: Previously, during OPEN in read mode (as with every OPEN), a write operation on the catalog entry was performed which updated Access-Date and Access-Count. In OSD V8.0, ARCHIVE performs a special OPEN without updating the access data, thereby dispensing with a write operation on the catalog entry.
- CLOSE overlapping with Wait-for-I/O: During backup and restore of a sequence of files, the CLOSE operation of a file A is delayed, i.e. a start is made on transferring the useful data of the next file B first, and only then is file A closed. In this way the closing operation for A coincides with a latency time that is unavoidable anyway on the disk side (i.e. due to waiting for the next bundle of B...).

Multiplexing up to factor x4 in FDDRL V17.0

Up to four disks are backed up in parallel to one tape; this allows operation at tape data rates >100MB/s in suitable configurations. Data is read alternately and asynchronously from multiple disks and in addition up to fourfold per disk if PAV is operated.
Multiplexing can then be combined with parallel tape devices, i.e. for example it is possible to back up simultaneously from 8 disks to 2 tapes.

Handling tape inputs/outputs on SQ servers via RSC

RSC (Remote System Call) for disks attached to SX servers was introduced in OSD V6.0B. This measure led to a considerable reduction in the load on the X2000 CPU (by approx. 20%) and an increase in throughput (particularly with multiple parallel I/Os). In a similar way to the RSC disk interface, RSC is also used on SQ servers for performance-related read/write I/Os to tape, and moreover, both for LTO devices and for Fibre Channel-connected CentricStor systems.
Handling tape I/Os via RSC for SQ servers produces the following improvements compared to SX servers:

- X2000 CPU load for tape I/O reduced by approx. 40-50%,
- Throughput increase: read 20%, write 5%.

Extended storage integration

BS2000/OSD integration for FibreCAT CX attached to SX and SQ servers with SHC-OSD V7.0

FibreCAT CX storage subsystems are mostly used with the SX servers and in future will be deployed with the SQ servers. Since the FibreCAT CX systems will increasingly meet high-end requirements, they are also subject to the requirement for BS2000 integration (as implemented for Symmetrix).
Functions for local replication in a FibreCAT system (SnapView software) and functions for remote replication between multiple FibreCAT systems (MirrorView software) are available for FibreCAT disk storage systems. SnapView is the equivalent of TimeFinder and comes in the editions SnapView Snapshot and SnapView Clone. MirrorView is the equivalent of SRDF.

Functionality provided

In SHC-OSD V7.0, information functions and BS2000 integration of local replication using SnapView are implemented. (FibreCAT MirrorView is supported in BS2000/OSD in the context of HIPLEX AF.)
The information functions for FibreCAT CX include global information relating to the storage system, information on the volume configuration and its mapping to BS2000 granulates (MN, VSN, catid) as well as information on local mirroring (pairs, statuses, type).
SHC-OSD V7.0 supports the control of local mirroring with SnapView clones and SnapView snapshots with the same functional scope as the control of local mirroring with TimeFinder clones and TimeFinder snaps, particularly at pubset level with consistent split and integrated renaming function.
Control of the mirroring functions can be built into runtime procedures, thereby achieving a high degree of automation and secure and reliable handling in critical operating situations.
The SHC-OSD interfaces for controlling local mirroring are being expanded for FibreCAT CX systems in a (mostly) compatible way. When they are called with the BS2000 granular pubset, MN or VSN, the SHC-OSD product checks internally which storage system is present, and calls the function via SYMAPI in case of Symmetrix, and via StorMan in case of FibreCAT CX (see solution architecture on p.8).
Interface compatibility for SnapView and TimeFinder snaps and SnapView and TimeFinder clones also enables SHC-OSD add-on products such as e.g. HSMS and FDDRL to support the backup of SnapView clones without modification; SESAM does not need to modify its HSMS calls. The BS2000/OSD-BC snapset function is being provided for FibreCAT CX also.

All in all, this means support for the following scenarios using SnapView (Clone and/or Snap)

- Pubset replication
  - Use and addressing of pubset copies
  - Pubset replication with PVSREN (for SM and SF pubsets)
  - Pubset replication with SHC-OSD (for SF pubsets)
- Data backup and recovery
  - Physical data backup with FDDRL
  - Disk-to-disk data backup using snapsets
  - Data backup of clones with HSMS, including data backup of databases
  - Transfer of backed-up data from snapsets to a backup archive with HSMS
- Volume-based recovery of pubsets
- Data security using standby pubsets
- Exporting of data (migration)

Solution architecture

The SHC-OSD product and the StorMan component of the DDC (Dynamic Data Center) product portfolio, available on Windows and Linux systems, are combined and extended as follows:

SHC-OSD

SHC-OSD acquires and manages configuration data of the storage system and maps it to the BS2000 configuration. SHC-OSD makes information and control functions available for the replication features at program and command level. All this was available for Symmetrix and was extended for FibreCAT CX. So SHC-OSD changes from the Symmetrix Host Component into the Storage Host Component.

StorMan

StorMan implements storage management functionality on the basis of the SMI-S standard (Storage Management Initiative Specification) and in this way encapsulates the storage specifics. In particular, StorMan controls local replication for SnapView snapshots and clones via SMI-S. The StorMan client and server components were ported to BS2000/POSIX.

SMI-S Provider

The SMI-S implementation for the actual FibreCAT hardware is carried out in EMC's SMI-S Provider. The SMI-S Provider runs on the Solaris platform in the case of SX servers and on the Linux platform in the case of SQ servers.
The SMI-S Provider is currently released both for Solaris and for EMC's SUSE SLES Linux.

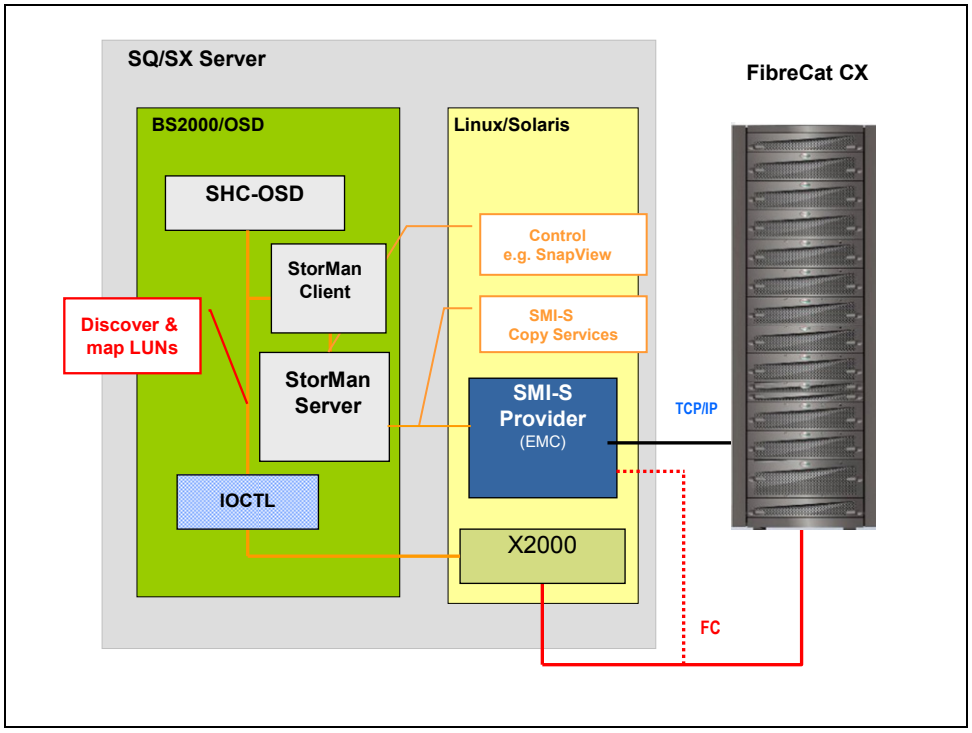

Hardware boundary conditions

The following main differences compared to Symmetrix systems apply to the FibreCAT CX systems and the associated SnapView software:

- Max. 256 devices of a FibreCAT CX can be connected to a BS2000 server. This limit applies when VM2000 is used for the entire VM system, i.e. for all guest systems in sum.
- Currently, SnapView can create up to 8 snapshots of an original and max. 2048 snapshots in one FibreCAT CX system. With SnapView, you can create up to 8 clones of an original and max. 1024 clones in one FibreCAT CX system.

Extended support for the Symmetrix DMX TimeFinder/Snap function

Save/restore to / from snapsets: basic functions and customer benefits

For Symmetrix DMX storage systems, EMC provides its Symmetrix-internal replication mechanism TimeFinder/Snap. Support for snapshot-oriented backup/restore scenarios in DMX configurations was implemented for the first time in BS2000/OSD-BC V7.0 in connection with SHC-OSD V6.0. The virtual copy of a pubset that can be used for restore purposes consists of the simultaneously created snap units for all volumes of the pubset, the “snapset”.
Snapsets are created and deleted by the administrator. The administrator can restore an entire pubset from the last snapset. New DMS functions allow the end user to restore individual files and job variables from the existing snapsets.

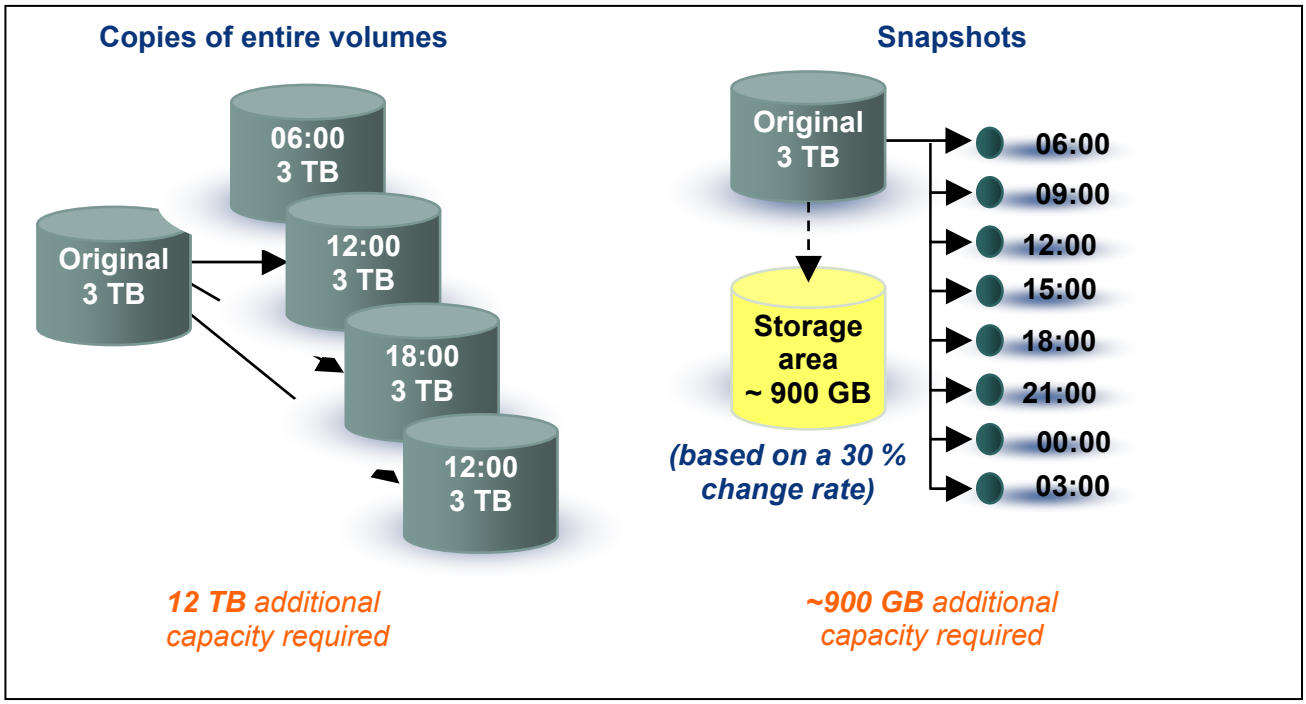

Backup using snap technology offers the following benefits compared with a normal tape backup:

- Ultrafast pubset backup as snapset – to snap units – in less than 1 minute, without I/Os, regardless of file quantity and size.
- Everything in the public volume set is saved, “in one go” and “simultaneously”, - open files, too, are thus backed up consistently in the same way as after a system crash.
- The disk subsystem only needs its own storage space for the snap units, provided the blocks in the original – i.e. in the current pubset – are changed.
- Files and job variables backed up to snapset can be listed and restored using simple, non-privileged commands; the process is disk-oriented, synchronous and independently parallel (as in the case of Copy File).
- The entire public volume set can be reset at lightning speed to the status of a snapset, based on the volumes and without I/Os.

Snapsets are implemented as special pubsets. Addressing takes place using a “pseudo notation” based on the original notation. DMS access by normal user applications is not possible (i.e. no standard IMPORT). In BS2000/OSD-BC V7.0 the VSNs of a snapset are derived from the VSNs of the associated pubset using a lowercase letter.

Extension to 52 snapsets

Starting with EMC microcode 5772, release for BS2000 effective in December 2007, the upper limit was increased from the former max. 15 snap units to max. 127 snap units for FBA disks. This extension is supported by SHC-OSD V6.1A01 (release 12.07) or higher. In preparation for the extension starting with the Enginuity 5772 update, 26 snapsets have already been implemented in BS2000/OSD-BC V7.0, with all lowercase letters of the alphabet being used for the snapset VSN.
Provision has been made in BS2000/OSD-BC V8.0 to increase the limit of the snapsets from 26 to 52 by using lowercase and uppercase letters for the snapset VSN, so that 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.

Extended information output

Several information functions have been extended to support the larger number of snapsets:

- SHOW-SNAPSET-CONFIGURATION also displays the relative age (-1 ... -52) of the snapset in addition to the snapset ID.
- Previously, with SHOW-SNAPSET-CONFIGURATION, the terminal output for all snapsets displayed one snapset per line. With the previous max.15 snapsets, the output always fitted on one screen. With the maximum number raised to 52, the terminal output is built using one “half-line” per snapset, with two succeeding snapsets per line, including with relative numbers -1 to -52.
- In the LIST and RESTORE functions, a snapset selection “*ALL” covering all snapsets was provided previously. Starting with BS2000/OSD-BC V8.0, to reduce the processing time, these functions additionally offer the option of specifying an interval using the relative numbers, i.e. e.g. (-1,-7) for the 7 most recent snapsets.

Program interfaces for snapsets

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.

SRDF switching including snapsets

In a disaster recovery scenario with SRDF mirroring, snapsets can be maintainedin the source as well as in the target storage subsystem. However, the snapsets are not mirrored by SRDF. Instead, the snap copies are created locally in each case on source and target. If the processing is switched over to the mirrored Symmetrix using Autoswap (or using the SHC-OSD command SET-REMOTE-COPY-ACCESS activated by it), access to the snapsets assigned to the pubset is not switched over together with this operation.
When a Symmetrix system is taken out of service as part of a scheduled downtime (e.g. for maintenance purposes), the HIPLEX-AF Autoswap function permits a switchover on command to an SRDF-mirrored Symmetrix, without interruption and with no need to restart the applications of the source server. After a short time, the applications continue running on the same server (where they already reside anyway) using the consistent data of the mirrored Symmetrix.
In BS2000/OSD-BC V8.0, the new ADAPT-SNAPSET-ACCESS command enables the switchover to be simulated for the snapsets assigned to a pubset. The command checks whether access to the snapsets takes place in the same Symmetrix box as for the pubset. If this is not the case, the currently attached snapsets are taken out of service and then the snapsets in the local Symmetrix box are attached. In this way it can be ensured that the snapsets continue to be available even after an Autoswap.

Manageability, ease of use

Mail interfaces in BS2000/OSD

BS2000/OSD V8.0 provides mail functions which permit system components and user programs to generate emails easily from within system processes.

Email address in the BS2000 user ID

For this purpose BS2000/OSD-BC V8.0 allows an email address to be assigned to every user. This email address is stored in the user entry of the BS2000 user ID under the new EMAIL-ADDRESS entry. The EMAIL-ADDRESS operand is supported in the ADD-USER, MODIFY-USER-ATTRIBUTES and SHOW-USER-ATTRIBUTES commands.
The email address is stored as a string. No syntax check takes place, though the syntax should correspond to the receiver address in the SEND-MAIL command of the interNet Services product. As with the SEND-MAIL command, the maximum possible length of the email address in BS2000 is 1800 bytes.

MAIL-FILE command and macro

BS2000/OSD-BC V8.0 provides a MAIL-FILE command and macro interface. Via MAIL-FILE a text file can be sent as an email appendix file to the mail address of the BS2000 user ID, as an alternative to print output. The actual sending of emails from within BS2000 is performed using the Mail-Sender (SEND-MAIL interface) of the inter*N*et Services product (new version V3.3). The BS2000 file to be sent (SAM or ISAM file) is handled as a text file. An automatic character set conversion based on the CCS file attribute is performed for it.
On MAIL-FILE, the email address is taken from the EMAIL-ADDRESS entry of the user ID. Similarly to PRINT-FILE, MAIL-FILE provides a Delete option and a Subject text (like the header text during printing). The Delete option can be used to control whether and how the BS2000 file will be deleted following successful sending.

Mailing of the system files *SYSLST and *SYSOUT

The MAIL-FILE command also allows to require the mailing of the system files *SYSLST und *SYSOUT.
In BS2000/OSD-BC V8.0, the EXIT-JOB/LOGOFF, CANCEL-JOB, ENTER-PROCEDURE commands, that (among other functions) control the print output of system files, include a new *MAIL operand enabling an email to be sent to the email address of the user ID of the relevant task, as an alternative for the print output.
So the print output of system files can be completely replaced by mail.
This is possible without changing existing procedures and jobs, because the new class-2-system parameter SSMOUT allows defaulting sending the system files *SYSLST and *SYSOUT output by mail.

MAIL-FILE usage in SW products running under OSD V8.0

The MAIL-FILE interface is being used in the following products:

- HSMS V9.0: Mailing of HSMS reports
- MAREN V12.0: Mailing of the OUTPUT-FILE for all statements that generate an OUTPUT-FILE
- HIPLEX-MSCF V6.0: Mailing of extremely important system messages: An email is sent when a critical situation occurs. It is sent to an email address specified in a user entry via a /SET-MSCF-ENVIRONMENT or /MODIFY-MSCF- ENVIRONMENT adminstrator command; this is the TSOS user ID as a standard.

Stemming the flood of console messages

The SET-MSG-SUPPRESSION command enables the output of messages with a required message number at the console to be suppressed. The messages concerned must originate from a message file and have been created using the MSG7 or MSG7X macro. All types of unanswerable messages can be suppressed.
The messages for up to 128 different message numbers can be suppressed during a system run.
In BS2000/OSD-BC V8.0, the previous limit of 128 message numbers will be increased to approx. 1,000 message numbers. To increase the limit, a new insert/scan strategy is being implemented in order to keep lock sections in the UCON task short (each message output to console must be synchronized with the exception list).

BS2ZIP enhancements

BS2ZIP is the WinZip-compliant compression tool of BS2000/OSD. The new BS2ZIP V1.2 version will be part of BS2000/OSD-BC V8.0. It will offer the following functional enhancements:

- The BS2ZIP SHOW-FILE-ATTRIBUTES statement displays a list of the ZIP archive's files. This information can be moved into S-variables to exploit in SDF-P procedures.
- The files within the ZIP container can be protected by a container-associated crypto password. The Winzip 2.0 standard encryption mechanism is being used for this purpose. It provides a relatively weak protection, but the main goal is to allow BS2000 users to extract encrypted zip files from a Zip archive generated on Windows, and vice-versa.
- When transferring ZIP files from/to Windows with openFT, a converter must be used, either before transfer (transfer to Windows) or after transfer (transfer from Windows), in order to convert the BS2ZIP PAM file in a SAM file that can be managed by openFT. In BS2ZIP V1.2 the converter function is being integrated into BS2ZIP via the CONVERT-ZIP-CONTAINER statement.
- BS2ZIP processes PLAM libraries as a whole file. Now the PLAM-LIB indicator is set when a PLAM library is extracted. It is no more necessary to open the file in output mode with LMS in order to set this indicator.

The new BS2ZIP version V1.2 is being released under BS2000/OSD-BC V6.0 or higher.

Common Memory Pools Status Information

The new SHOW-MEMORY-POOL-STATUS command provides information about common memory pools currently set up in the system. Output is to SYSOUT and supplies a list of memory pool identifiers (name, scope, user ID/group ID) together with a list of the TSNs of all tasks connected to the respective memory pool.
The output can be restricted to memory pools with a specified name range or sharer group, or to memory pools with specific attributes.
Information on privileged memory pools is also available to privileged callers (TSOS, SW-MONITOR-ADMINISTRATION).

Information on and modification of NK-ISAM caches

Since BS2000/OSD-BC V6.0B, all multi-task NK-ISAM pools have been set up in data spaces. The buffer areas for standard ISAM pools are created automatically by the system and basically set up for individual files.
The system parameter MAXDSBN specifies the maximum number of data spaces to be provided for multi-task NK-ISAM pools. A range from 1 - 127 can be specified, with 2 as the default value.
In BS2000/OSD-BC V8.0, new commands permit information on the NK-ISAM caches and the data space load to be output and data spaces to be added and detached dynamically. This allows the system administration to check which setting is most appropriate for the particular configuration, and to adjust this setting during live operation.
The SHOW-ISAM-CACHING command keeps the system administration updated about the data spaces that are currently being used as ISAM caches and that are created and managed for accommodating multi-task ISAM pools. The information output can include both global information about data spaces, i.e. essentially information on occupied and free areas, and detailed information about data spaces, ISAM pools and ISAM files buffered in these.
The MODIFY-ISAM-CACHING command enables the system administration to dynamically change (create and detach) the number of data spaces used by ISAM and created and managed for the purpose of accommodating multi-task ISAM pools.

Other change requests

JV V15.0: Performance improvement in the mass handling of JVs
The sorting routines in the JV product (in SHOW-JV-ATTRIBUTES, DEL-JV) have been optimized to improve sorting efficiency for large JV quantities, as supported since the extension to max. 6-digit numbers of JVs starting with JV V14.0C under BS2000/OSD-BC V6.0B or higher. As an additional option, the sorted output can also be switched off.

SPACEOPT V5.0: Reduction in the number of file extents for large files
In SPACEOPT V5.0 file reorganization is possible also in the case of less free disk space available than the file size requires. Reducing the number of file extents using REDUCE-FILE-EXTENT-NUMBER is currently only possible if a free work area equal to the size of the files is available. In SPACEOPT V5.0, the function will also allow local improvements based on the free disk space available.

DAB V9.2: New filter possibilities for AutoDAB Caching
Currently AutoDAB caching allows excluding selected catalogued files from caching explicitely. In DAB V9.2 partial filenames can be specified within the exception list. So for example, temporary work files can be excluded from being cached, and even before they are generated. In addition, the manual file selection can be reset, i.e. AutoDAB can be reactivated for files that formerly were excluded.

Recovery of pubset reconfiguration locks
Pubsets can be extended or reduced in size during online operation (dynamic pubset reconfiguration). In rare cases, the pubset reconfiguration locks can remain hanging in the MSCF cluster. Previously, if this happened, BS2000 had to be restarted in order to make the pubset fully operational again. Functions for displaying and resetting the locks have been implemented to avoid a system failure in such situations. The following new commands are provided:

- SHOW-PUBSET-LOCKS
- REMOVE-PUBSET-LOCK

Pubset renaming with transfer of all catalogue settings
When renaming a pubset it is now possible to take over all MRSCAT settings of the original pubsets into the MRSCAT of the renamed pubset (included e.g. the Speedcat settings).

SHOW-AUDIT-STATUS
The new SHOW-AUDIT-STATUS command informs about HARDWARE-AUDIT or LINKAGE-AUDIT being switched on for specific operating sequences (e.g. tasks). This enables a better diagnosis of performance problems.

Inserting pubsets with paging files into an SM pubset
Pubsets with paging files can be took over with SMPGEN into an SM pubset, if they are not marked as active paging pubsets in the MRSCAT.

COPY-FILE with transfer of change date
By means of the new PROTECTI-ON=*SAME-AND-CHANGE-DATE operand, the COPY-FILE command allows to transfer the protection attributes and the change date from the source file to the target file, thus e.g. supporting cloning of a userid or a pubset.

More parallel memory pools for shared code
The number of memory pools in which the user can store shared code by means of the ASHARE macro is up to now limited to 8 per scope (ENAMP operand SCOPE) for each user ID. With BLSSERV V2.7 the maximum number of memory pools per scope is being doubled from 8 to 16.

Starting SCA (Speedcat) automatically
If the SCA (Speedcat) product is installed, Speedcat will be started in *SPEEDCAT-TASK mode during a pubset import as of BS2000/OSD-BC V8.0 as a standard, in order to grant a performing catalogue access for SF pubsets.
If the start of Speedcat is not desired for a pubset, this must be set up explicitly within the catalogue entry using the START-SPEEDCAT=*NO operand of the ADD-MASTER-CATALOG-ENTRY resp. MODIFY-MASTER-CATALOG-ENTRY command. In the START-SPEEDCAT operand of the ADD-MASTER-CATALOG-ENTRY command the former default value *NO was changed to the new value *AUTOMATIC.

Openness and Integration ability

The years-long strategy of openness and integration ability is being systematically continued with BS2000/OSD-BC V8.0.

POSIX A41

POSIX as the technological basis of the BS2000/OSD open system strategy is being developed further in response with user requirements. Release of the new POSIX A41 version is planned in conjunction with BS2000/OSD-BC V8.0. Functional enhancements are:
bs2fs – transparent access to BS2000 files in POSIX
The BS2000 file system bs2fs allows to access BS2000 files transparently from POSIX and via existing POSIX interfaces (commands and program interfaces). For this purpose a new file system type bs2fs was introduced in addition to the existing file system types ufs or NFS. In order to work with bs2fs, a subset of a user's files is mounted with a bs2fs mount command to a position within the POSIX file system. This mount process makes the BS2000 files accessible to the user. Now he can work on these files with POSIX means, e.g. perform an open() on such a file or apply commands on such files.
In order to access BS2000 files with POSIX interfaces, the files are copied on demand - and not visible for the user - in a container file system especially foreseen for bs2fs (of file system type ufs). This copying task is performed by one or several daemons. A daemon opens this file in BS2000 and copies it to POSIX, at the same time the BS2000 file is locked for other users in BS2000, while the POSIX copy in the container is not locked for other POSIX users. Then operations can be performed on thefile, e.g. reads, positions and writes. When the file is closed, it is copied back to BS2000.
Usage scenarios are:
- Browse of BS2000 files resp. PLAM library elements with respect to certain patterns with the mighty POSIX grep command.
- Use of the make function for efficient generation of programs or program systems.
- Nested procedures (Call of the POSIX shell from the BS2000 command mode, execution of POSIX commands and then return to the BS2000 command mode) can be replaced by pure POSIX shell scripts. File manipulations in BS2000 are replaced by manipulation of these files from within POSIX, once having mounted these files via bs2fs. The amount of switches from BS2000 ito the shell and vice versa is declining.
- Offer BS2000 files in web by simple means.

Extensions for Posix sockets
- Transfer of a connected socket to an independent process (not a "child")
- Support for the sockets option TCP_NODELAY after connect() or accept()
- Support for a maximum buffer size of 32K for udp datagrams
- Support for the default length 16 for I/O vectors both for UDP and for TCP

The new POSIX A41 version is being released under BS2000/OSD-BC V7.0 or higher.

Apache

The latest version of Apache, the world's most popular web server, is also available for BS2000/OSD in addition to Windows, Linux and Solaris, and is included in the operating system basic configuration BS2000/OSD-BC.
BS2000/OSD-BC V8.0 includes the version APACHE (BS2000/OSD) V2.2.
The migration to the Apache http server 2.2 of the Apache Software Foundation takes place with the release of APACHE (BS2000/OSD) V2.2, with support for PHP V5.2, PERL V5.8, TOMCAT V5.5 instead of JSERV/JSP. APACHE (BS2000/OSD)

V2.2 includes support for the SSL (Secure Socket Layer) protocol for secure transfer of documents and data over the internet; the existing add-on product interNet Security ("Apache+SSL") is omitted.

Java

With BS2000/OSD Environment for Java (short name JENV), Java is provided as part of BS2000/OSD-BC. BS2000/OSD Environment for Java (JENV) enables any Java programs, written on any platforms, to be run on BS2000 systems. BS2000/OSD-BC V8.0 includes the version JENV (BS2000/OSD) V5.1. Version V5.1 of the Environment for Java is a rebasement on Java2 SDK Standard Edition 5.0 Update 15.

Enhancements in Software Products

The functional enhancements in new versions of certain key software products to be released in the same timeframe as BS2000/OSD V8.0 are presented below. These will also be available for BS2000/OSD V7.0 and V6.0.

FDDRL V17.0

Multiplexing with a factor of 4

With FDDRL V17.0, physical disk backup supports the ongoing technical development of MTC technology (increased throughput for LTO devices).
With the Disk-Set backup unit, multiple disks can be backed up to one tape/tapeset and restored from there.
The Disk-Set backup unit was introduced with FDDRL V16.0. Large tape blocks are created. The block size is approx. 160 KB, irrespective of the maximum I/O size of the individual disks. Multiple disks are saved to a tape series with (in V16.0) dual multiplexing between disk and tape. A multiplexing factor of 4 is supported for the Disk-Set backup unit in FDDRL V17.0. This also enables support for configurations in which the throughput rate of the tape device cannot be reached by backing up only 2 disks (e.g. LTO-4). With a multiplexing factor of 4, the overall throughput can therefore be doubled in the ideal case.
The FDDRL control parameter TAPE-FORMAT specifies the tape format and hence the compatibility of the backups. The higher the tape format, the better the performance that can be achieved. A new FDDRL tape format (V17-FORMAT) is being introduced with FDDRL V17.0. The multiplexing factor 4 is set via the default value TAPE-FORMAT= V17-FORMAT.
With FDDRL V17.0, the former tape formats are supported in compatible mode.

Disaster recovery: restoring pubsets following total disk failure

In disk configurations without SRDF, if the controller fails, pubsets can be restored by FDDRL using First-Tape and Restore. In this disaster scenario, the server is not affected, so the generation of the devices is still present.
FDDRL V17.0 provides functions for this scenario to simplify handling of the disks that are to be restored following a disaster. These include an information function relating to the properties of the disks to be provided, simplified mounting of disks, and initialization of the disks by VOLIN. This enables faster restoration of the disks by FDDRL.
The process of restoring pubsets from the FDDRL tape is improved by the following measures:

- With the new SHOW-FDDRL-TAPE-INFORMATION statement, FDDRL provides the system administrator with information on which disks are required in order to restore the pubset: VSN, MN, disk type, formatting, capacity. This helps avoid errors in disk provisioning and unnecessary correction cycles.
- Disk formatting is initiated by FDDRL. This enables the restore operation to be completed at maximum speed and free of error. The steps “format disk”, “configure pubset” and “restore pubset” can be completed in one FDDRL run.
- The FDDRL user interface is extended by the specification of the MN of the output disks (COPY-DISK/COPY-PUBSET and RELOAD-DISK/RELOAD-PUBSET).
  To restore the disks, the device type, format and capacity are the most important parameters. Using the same MN for restore as for backup makes sense particularly if all the disks have been destroyed. The user’s work is made easier by the operand *SAME-UNIT, which is provided as an alternative to a list of MNs. In this case the user should be advised beforehand of the original MNs, for which reason the MN should be output in addition with the information relating to the backup.
  If the output disks are specified using their MN, they can also be used if they have not been initialized, or if they have the wrong format. They are then initialized automatically and formatted in the required format.

FDDRL V17.0 is released for BS2000/OSD-BC V6.0 or higher.

HSMS/ARCHIVE V9.0

Performance improvements for small and medium-sized files during logical backups

To improve performance during logical backups for small and midsized files, the following measures will be implemented in HSMS V9.0:

- Minimization of disk I/Os for reading catalog entries
- The catalog entry write during the OPEN in the archive subtask eliminated
- CLOSE overlapping with Wait-for-IO

This is produces improvements of up to 30% for small files. For details, see section Scalability / performance.

Use of optimal transfer lengths for disk I/O (ARCHIVE in connection with BS2000/OSD-BC V8.0)

- Use of larger disk I/Os (with D3435NK, up to 240 PAM pages) during backup and restore of files. Only max. 128 PAM pages were possible previously.
- Standardization of the minimum transfer length for all disks.

Long-term fixing of I/O buffers for disk and tape I/Os (ARCHIVE in connection with BS2000/OSD-BC V8.0)

Previously, the I/O buffers for disk and tape I/Os were fixed for each I/O job by IOCTRL. In ARCHIVE V9.0, the I/O buffers are fixed once per ARCHIVE subtask.
Fixed CI.5 pages are accepted by the system in BS2000/OSD-BC V8.0 on termination of ARCHIVE tasks. Automatic defixing takes place at CI.5 RELEASE-ALL time.

Fast (emergency) restore of the simultaneously saved directory

Under HSMS V7.0 or higher, a directory saved using “Dirsave” at backup time can be conveniently restored by means of a special operand. For this, the tape is read up to end-of-tape – an operation which is very time-consuming, especially with large LTO capacities.
To perform a fast directory restore in HSMS V9.0, the tape position of the last directory saved to this tape is noted in the end-of-tape label. During the (emergency) restore, the tape is positioned to that point starting from the end of tape and the directory is therefore found quickly (even with LTO) and restored.

Fast re-migration of files

In file migration using HSMS, many files are migrated again even though they have not been modified in the meantime. To detect this, in HSMS V9.0 the Recall-Date is stored in the catalog entry. A fresh migration is then performed without tape operation using the already existing migration tapes (only the location information is noted).

Report output to library elements and via mail interface

For reports, HSMS V9.0 will also provide output via email in addition to output to library elements. The reports are controlled using the operands REPORT=*NONE/*SUMMARY/*FULL and OUTPUT=*PRINT / <filename> in the user statements (e.g. Save and Restore). The second operand is extended with the option to specify a library element or a *MAIL value. With relatively lengthy asynchronous processes such as, say, HSMS backups, it is now clear with the email reception of the report that the backup job has been completed.
In HSMS, the email is sent in the HSMS server task under TSOS to the address of the HSMS caller with the latter’s other user ID.
The mail output of the reports is supported only in connection with BS2000/OSD-BC V8.0.

HSMS V9.0 is released for BS2000/OSD-BC V6.0 or higher.

MAREN V12.0

Tape encryption with LTO-4

Starting with BS2000/OSD-BC V8.0, BS2000/OSD will provide support for the first time for LTO-4 drives equipped with a 'tape encryption' hardware feature. With tape encryption enabled, tape access performance is reduced by only less than 1% in terms of data rate.
Support for tape encryption has been implemented as an enhancement to the MAREN magnetic tape management product in BS2000/OSD in MAREN V12.0. MAREN not only handles the key management function but also takes control of encryption and decryption. The encryption conforms to the AES standard and uses a symmetric 256-bit key.

Key management

Key management is performed in two stages:

- The keys used by the tape device to encrypt the volume contents are called data keys. They are created and managed on a volume-specific basis. The keys are stored both in the MAREN catalog and on the tapes, in encrypted form in each case. On the tape, the tape device stores the keys on a block-specific basis.
- The keys used to encrypt the data keys
  are derived internally in MAREN from the encryption password assigned by the MAREN administrator and managed by MAREN within a key box.

New statements in MAREN have been provided and existing ones extended for the purpose of creating and managing the encryption password.

Control of encryption and decryption

The new volume type TAPE-U4E is being introduced for LTO-4 tapes with encryption. This will be implicitly supported in addition at all DMS and HSMS interfaces; e.g. during processing of tape files, the device is requested via the volume type specification in the DEVICE-TYPE operand of the respective DMS/NDM command. This means that encryption of LTO-4 tapes is supported in OSD V8.0 via all user interfaces at which tape processing is controlled by means of the volume type, assuming encryption passwords have been stored in MAREN in preparation for this.
Encrypted backup tapes can be generated with HSMS, ARCHIVE and FDDRL without having to change the control statements for the backup products.

Hardware requirements
LTO-4 tapes with volume type TAPE-U4E are only supported in LTO-4 devices connected directly (i.e. without CentricStor) to the FC channel of an S or SQ server via a Quantum/ADIC Scalar 10K, i2000 or i500 library system. LTO-4 devices can also be operated on SQ servers on an MTC autoloader connected via FC. With this solution, encryption is not possible for CentricStor tapes, since CentricStor supports BS2000 tapes via volume type TAPE-C4. Encryption appliances are recommended for CentricStor.
The tape encryption feature allows data to be stored on a tape in encrypted form without significant performance loss. This affords a highly effective means of implementing data protection at tape level, enabling tape contents to be protected against unauthorized reading, especially when in transit, when stored externally (e.g. in fireproof vaults) and when on loan.

Output of volume information as CSV file and sending as email

The MAREN and MARENADM statements used to create an OUTPUT-FILE can also generate the output in csv (comma separated value) format. Files of this type lend themselves very well to post-processing, e.g. for table generation in EXCEL, because they can be imported in column-true format.
In OSD V8.0, in addition, the SEND-BY-MAIL parameter can be specified in these statements to define whether a created file is to be sent by email. The email address assigned to the caller's user ID is used as the receiving address.

Easier and faster recovery of the MAREN catalog

To ensure MAREN operation can be resumed in the event of the loss of the catalog file, a backup of the catalog must be taken. A copy can be created for this purpose using the //COPY-VOLUME-CATALOG statement.
Currently the logging files of all the systems must be merged with the SORT utility into a single file before //UPDATE-MAREN-CATALOG is invoked.
In MAREN V12.0 the administrator can specify all logging files in the //UPDATE-MAREN-CATALOG statement. All that the administrator is now required to do is ensure that all the logging files are accessible. The LOGGING-FILE parameter has been extended with "list-poss(99)". This enables up to 99 logging files per statement to be specified, in any order.

Easier transfer of HSMS/ARCHIVE data

The MARENADM.ARCHIVE program helps the administrator analyze ARCHIVE directories and import the data residing in them into the MAREN catalog. The function is integrated into the MARENADM statement //UPDATE-MAREN-CATALOG (new operand values SELECT = *BY-HSMS (..) | *BY-ARCHIVE (..)), thus rendering the MARENADM.ARCHIVE program superfluous.
The interface to HSMS is supported for HSMS V7.0 or higher.

Changes to UCP parameters effective immediately

Since the release of MAREN V11.0, almost all parameter changes on the local system take effect immediately. The only remaining gap relates to the parameters of MARENUCP:
Changes to the parameters defined in the MARENUCP Enter procedure do not take effect until after a restart.
In MAREN V12.0, the parameters that now have to be defined in the MARENUCP Enter procedure will be included in the computer-specific parameter set. //MOD-MAR-PAR and //SHOW-MAR-PAR are being extended accordingly. The first time MARENUCP is started in MAREN V12.0, the parameter values contained in the ENTER procedure are copied over. After that, they are ignored.

Other MAREN enhancements

Logging of cross-computer administrator actions
Starting with MAREN V11.0, a MAREN administrator can carry out status changes of the MAREN system (stopping MARENCP and MARENUCP, switching the logging file, etc.) for other systems in the MAREN cluster as well. When this takes place, it is not transparent on the system in question, by whom and from which system the action was initiated.
In MAREN V12.0, corresponding information is output to the console on the system concerned.

Automatic recovery after temporary MARENCP failure
Batch tasks waiting for the MAR0085 console message automatically continue running in MAREN V12.0 once the cause (e.g. temporary failure of MARENCP due to loss of the RFA connection) has been removed.
In future, batch tasks will output the MAR0085 as an asynchronous query to the console. After a wait time has elapsed, the task checks whether the cause of the problem is still present. If it is, the waiting period is extended. If it is not, the message is withdrawn and processing resumed. The operator can respond to the message as previously.
The solution for batch tasks has also been implemented retroactively for MAREN V11.0 as a source correction. It was released with correction package 1/07.
Runtime performance has also been improved for MARENADM or MAREN statements in interactive dialog mode. Currently, if MARENCP fails, the statement is aborted. In future the user will be prompted to decide whether to wait or abort. The old response is retained in procedure mode.

Inheritance of administrator rights by batch tasks
The current administrator or all-domain administrator authorization also applies to batch tasks with MARENADM calls started by the administrator. For this purpose an //ENTER-MAREN-PROCEDURE statement is being introduced in MARENADM. This includes a subset of the parameters of the /ENTER-PROCEDURE command.
The currently necessary administrator password and authentication as all-domain administrator, among other things, can then be omitted from the ENTER procedures.

Free naming of operator role
A new operator role was introduced in MAREN V11.0 and its name was defined as “SYSMAREN”. In MAREN V12.0, this name can be freely assigned by the MAREN administrator:
A system-specific parameter is being introduced for the operator role (default: SYSMAREN). This can be changed using //MOD-MAR-PAR and displayed using //SHOW-MAR-PAR. The parameter is interpreted by MARENUCP and MAREN-INIT.

“All-Domain-Admin” as default state
In MAREN V12.0, the all-domain administrator (ADA) can specify that he/she is to be the all-domain administrator again immediately and automatically every time MARENADM is started. Previously, the administrator had to enter //MODIFY-ADMINISTRATION-SCOPE each time for this.
A system-specific parameter is being introduced for the feature “ADA as basic setting” (preset default: Domain-Administrator). This can be changed using //MOD-MAR-PAR and displayed using //SHOW-MAR-PAR. It is read at MARENADM start time. The subsystem call is then initiated as for //MOD-ADM-SCOPE and the authorization checked – provided domain protection is enabled.

Easier initialization of new volumes
The initialization of new volumes up to now required the following steps: //AD-FREE-VOLUMES; //FREE-VOLUMES; //INIT-VOLUMES. Via the new operand INIT-FILE in //ADD-FREE-VOLUMES the second step can now be omitted.

Improved interface with ARCHIVE/HSMS
- IMPORT-FILE and COPY-EXPORT-FILE without explicit directory-specification
  As of ARCHIVE/HSMS V9.0 a new interface has been introduced for IMPORT-FILE and COPY-EXPORT-SAVE-FILE for specifying a Save-File-ID (SFID). You don’t have to specify a directory or the corresponding volumes any longer. MAREN delivers all tapes belonging to a Save-File to HSMS.
- Directory-Indicator in MAREN
  For disaster recovery, i.e. the tape-based reconstruction of the MAREN-catalog in emergency cases, ARCHIVE offers the backup of the current MAREN-directory at the end of the tape. So it can be restored from tape – in case it has been lost on disk – via a simple ARCHIVE/HSMS-interface.
  As of ARCHIVE/HSMS V9.0 and MAREN V12.0 a volume with MAREN-CAT-backup gets a new identification informing on the presence of the MAREN-CAT on tape. This identification is shown for a tape with //SHOW-VOLUME-ATTR and can be used as selection parameter.
- Implicit purge in HSMS
  As of HSMS V8.0 it is possible to reuse ‘obsolete’ tapes faster. Instead of an explicit purge with deleting the Save-File and tape-release an implicit purge is executed while saving, so that before the actual save-procedure obsolete Save-Files are deleted. In ideal cases the released tape can be used again as scratch-tape for the next backup. Before MAREN V12.0 tapes had to be released with //FREE-VOLUMES prior to their usage. As of MAREN V12.0 the explicit tape-release is only necessary in the following cases:
- The MAREN-administrator or the user has changed the release-date, which has been fixed when reserving the tape or writing on the tape, with //MOD-VOL-ATTR FREE-DATE = or EXPIRATION-DATE. MAREN interprets this case as a tape that doesn’t follow normal cycles.
- The MAREN-administrator or the user has decided - by issuing the command //MOD-VOL-ATTR INIT=ERASE - that the tape has to be erased after its release. This can only be done by the MARENADM-command //INIT-VOL.

MAREN V12.0 is released under BS2000/OSD-BC V6.0 or higher.
MAREN V12.0 can be used in combination with MAREN V11.0 and V10.0.

openNet Server V3.3

IPv6 extensions in BCAM and SOCKETS

IP multicast enables data to be sent efficiently to many recipients simultaneously in TCP/IP networks. This takes place by means of a special multicast address. Every address beginning with FF00::/8 is reserved for this in IPv6. The ICMPv6 protocol is used in addition for control purposes in IPv6.
- IPv6 Multicast (RFC3493)
  The SOCKETS interface in version 2.0 permits communication in accordance with IPv6 conventions (RFC 2553). SOCKETS is the porting platform for Fujitsu and customer products.
  Support for multicast ensures that SOCKETS can continue being used as a porting platform for BS2000.
- ICMPv6 (RFC4443 supersedes RFC2463)
  IPv6 multicast and compatibility – incoming IP segments must not cause any damage – require extension.
  (ICMP functionality in IPv4 is included in ICMPv6.)

Classless Interdomain Routing (CIDR)

Classless Inter-Domain Routing (CIDR) is intended to enable more efficient use of the existing IPv4 32-bit IP address space. CIDR removes the need for fixed assignment of an IP address to a network class, subdivision into further subnetworks and aggregation of multiple networks of one class. There is now only one network mask, which subdivides the IP address into network part and host part.
- CDIR mitigates the consequences of IP address shortage; more than 90% of IP routers use CIDR.
  This means that BCAM must support CDIR in order to be able to integrate BS2000 into the data center network. Without CIDR, BS2000 would make using CDIR more difficult for all network partners.

- The IPv4 Class A,B,C network types are being discontinued. The subnetwork addresses are noted by specifying the prefix length.
- By “monitoring” OSPF (router) packets and evaluating them, BCAM can establish which CIDR blocks exist in the local network.

Extended OSPF Handling

OSPF (Open Shortest Path First) is used for the exchange of routing information between IP router systems that belong to an autonomous system, for example an enterprise network. Within an autonomous system, all router systems involved have identical routing tables, which can be maintained consistently with the help of OSPF.
In openNet Server V3.3, besides the automatic discovery of IPv4 routers the subnet information contained in the OSPF protocol is evaluated.
So the necessary amount of administrative net definition can be reduced.
When OSPF is available, on the side of BCAM only

- own cable connections and
- own addresses

must be defined.
All other objects can be automatically discovered or generated.

Security (IPSec Level 4)

IPSec adds functions to the TCP/IP protocol stack that are missing from the standard protocols, i.e. encryption, integrity and authentication of security-related data.
IPSec uses openCRYPT™ to implement a broad range of security mechanisms. The IPSec implementation provides flexible control techniques enabling messages to be transmitted in encrypted form, and reliable authentication of communication partners, without any need to intervene in existing communication applications. IPSec supports the use of cryptography in Layer 3 (Network Layer) of the OSI reference model.
An additional “secure channel” or a key exchange protocol is required for the key exchange between sender and recipient. This can be accomplished by means of a manual exchange or automatically by a key management system. IPSec in BS2000/OSD will support automatic key exchange via ‘Internet Key Exchange’ protocol version 2 (IKEv2) in Stage 4 of the IPSec implementation.
Note: IKEv1 is already included in openNet Server V3.2B.

- IKEv1 and IKEv2 (not compatible) are supported in parallel. IKEv2 offers the customer functional advantages (NAT traversal (“passing encrypted through the firewall”)) and better performance (simplified protocols).
  Porting of an IKE daemon
  RFC4301ff (incl. IKEv2), 3DES is a MUST and default for the required cryptographic algorithms.
- Support for IPCOMP (IP Payload Compression Protocol) in openCrypt V1.3A. In IPSec, IPCOMP is considered the “default transformation”.
- Improvements in operation of IPSec and in keeping it up-to-date.

LWRESD (Light Weight Resolver Daemon) – rebasing on BIND 9.4

To provide simple and consistent management of names and addresses of the partner systems connected to BS2000/OSD, these can be stored in external DNS servers (DNS stands for Domain Name Service). BCAM can access this DNS server, which handles the translation of names to addresses and vice versa, via the LWRESD product also supplied.
The Domain Name Service is a distributed, replicated database containing DNS servers and DNS clients (resolvers).
The resolver functionality in BS2000/OSD is a porting of the BIND coding, which is considered the standard DNS implementation.

- With openNet Server V3.3, BIND 9.2.x is rebased on BIND 9.4.x., because BIND V9.2 was withdrawn.
- So functional extensions are available such as cache dump, lifetime query for cache entries and cache deletion.

Online backup and restore of the BCAM configuration

With the present length of BS2000 sessions, it is usual to issue a large number of BCAM commands in some cases during the course of the sessions (e.g. to dynamically generate new partner systems, redefine existing network partners, etc.). In future, these configuration changes will be permanently available and retrievable.
The BCAM configuration can be backed up as a command file during online operation. A SOF file is generated from a complete BCAM configuration backup with the suitable BCAM SDF commands, which recovers the backed up configuration when referred in the DCSTART commando.
The migration from an existing configuration, for example RDF generation or „old“ BCIN commands to a network definition via SDF commands is supported.
In addition it is possible to backup selected BCAM object groups. They are stored in a file that can be managed with CALL-BCAM-COMMANDS.
Add-on display functions are also being implemented:

- Display the BCAM-internal gateway table,
- Display the assignment of a route to a network access.

openNet Server V3.3 is released under BS2000/OSD-BC V6.0 or higher.

SECOS V5.2

Security audits with inclusion of BS2000/OSD business servers

More and more customers want their data center to be subjected to a security audit. For this, an organization drafts its own security policy and commissions an internal or external auditor to evaluate the policy itself and compliance with it. The motivation to carry out a security audit generally comes from outside the organization, with the terms of reference being set by statutory requirements such as the Control and Transparency in Business Act (KonTraG) in Germany or the Sarbanes-Oxley Act in the USA.
As with quality management, security policies can use ISO standards as a basis, most notably ISO/IEC 27001, which bears the title “Information technology - Security techniques - Information security management systems - Requirements”. In Germany specifically, the Federal Office for Information Security (BSI) has established comprehensive implementation requirements for this in the form of the IT-Grundschutzhandbuch (*“IT basic protection manual”*) and offers its own security certification process. With their security functions, the BS2000/OSD business servers can be successfully included in security audits. The BS2000/OSD-BC operating system includes extensive security mechanisms by default. SECOS enables extended, fine-grained control of the security mechanisms; in the ISO27001 environment, SECOS has the key functions:

- Audit trail,
- Personal login also under system administrator Ids,
- Ensuring the integrity of system files (also implemented via SAT logging).

Functional enhancements in SECOS V5.2

SECOS V5.2 implements functions in response to requirements raised in customer workshops as well as requirements resulting from the monitoring of security audits conducted in accordance with the BSI IT basic protection manual (and which previously had to be bypassed by organizational measures). Specifically, these are:

- User Ids can be barred automatically because of false password input or if they have not been used.
  Previously, failed attempts during password input were sanctioned with time penalties or connection cleardown; even automated intrusion attempts could be prevented in this way, however. In SECOS V5.2 user IDs / terminals can be barred after n failed attempts. For User IDs that cannot be barred or may not be barred (e.g. TSOS or common user IDs), it is also possible to bar the initiator. The initiator can be e.g. the personal user ID for dialogue with personal LOGON, the Kerberos Principal , the caller’s user ID or the caller’s Principal for batchim Batch die Benutzerkennung des Aufrufers oder der Principal des Aufrufers sein.
  User IDs that have not been used for n days can also be barred. Currently, the passwords – not the user IDs – are time-monitored by SECOS.
  The SET-LOGON-PROTECTION command is being extended fort he barring specifications (new SUSPEND-ATTRIBUTES and INACTIVITY-LIMIT operands). The new SHOW-SUSPEND-USER command displays all suspensions of a user ID, with UNLOCK-SUSPEND-USER they can be cancelled.
- Default protection for user IDs (“policy enforcement”). User management can specify a protection setting to be used by default for all user IDs. Currently, the protection settings must be applied for each user ID individually using /SET-LOGON-PROTECTION. With this concept a user manager has the possibility to enforce a global security policy, but to accept exceptions in certain circumstances.
  The Logon Protection’s default attributes are managed with the new MODIFY-LOGON-DEFAULTS command. They are diplayed with SHOW-LOGON-DEFAULTS.
- Centralization of SECOS administration: A system administrator performing the 3 roles system, security and security audit file administrator on his own can concentrate his tasks under the single user ID TSOS. The additional overhead due to the previous rigid separation into different user IDs (SYSAUDIT, SYSPRIV) has been removed.
  The central SECOS administration is specified with the new SECMAN=*DECENTRAL / *CENTRAL statement for SRPM in the startup parameter file.
- Extensions to SAT logging:
- The SAT outputs can optionally be output also in XML format for post-processing.
- With the LTO-4 tape encryption feature, new security-related events which can be logged and analyzed will occur in MAREN.

- SECOS V5.2 is released under BS2000/OSD-BC V6.0 or higher.

SHC-OSD V7.0

BS2000/OSD integration for FibreCAT CX
A priority topic in SHC-OSD V7.0 is the BS2000/OSD integration of FibreCAT CX for SX and SQ servers with support for the local replication functions provided with SnapView. SnapView for FibreCAT CX is the counterpart to TimeFinder for Symmetrix and comes in the variants SnapView/Snap and SnapView/Clone. For details, see the chapter ‘Extended storage integration’.
As for Symmetrix, program interfaces for integration in backup procedures using HSMS/CCOPY for BS2000/OSD V8.0 and for snap usage from DMS functions are provided here in addition to command interfaces.
SHC-OSD V7.0 also supports the following enhancements provided with Enginuity 5771 and 5772 as well as SYMAPI V6.0 or higher:
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.
Migration via TimeFinder/Clone onto Clones > Original
As of e5772 it is possible to create TimeFinder Clones with a capacity larger than the original volumes. SHC-OSD V7.0 supports this function. The extended capacity on the clone can be made usable for BS2000 by means of the SPACEOPT product.

Performance measures
SHC-OSD manages and observes the entire configuration of Symmetrix arrays and for this purpose performs periodic updates of the configuration data. The trend toward ever larger configurations based on storage consolidation and new Symmetrix models (e.g. DMX-3 and DMX-4 with > 64K volumes) makes it necessary to optimize response times and the CPU requirement of the SHC-OSD user task. The following measures are provided in SHC-OSD V7.0:
- Optimization of the internal configuration data update,
- Introduction of selective monitoring of the Symmetrix arrays with respect to global status or the entire device configuration.

Performance improvements were already achieved in SHC-OSD V6.1 as a result of making event monitoring more flexible:
- Controllable SHC user task priority,
- Fast failure detection for Symmetrix.

SHC-OSD V7.0 is released under BS2000/OSD-BC V6.0 or higher.

Software product overview

The following table contains all software products for which a new version is being released to provide support for BS2000/OSD V8.0, along with a summary of the relevant new functions.

Product	Version	■ New function with OSD V8.0
AID	3.4	■ Updated in line with BS2000/OSD-BC V8.0, support for SQ servers
ARCHIVE	9.0	■ Use of optimal transfer length for disk I/O ■ Long-term fixing of I/O buffers for disk and tape I/Os
COSMOS	17.0	■ Updated in line with BS2000/OSD-BC V8.0, support for SQ servers (technically coupled product)
CRTE	2.8	■ Updated in line with BS2000/OSD-BC V8.0, support for SQ servers (technically coupled product)
DAB	9.2	■ New filter possibilies for AutoDAB caching ■ Updated in line with BS2000/OSD-BC V8.0, support for SQ servers (technically coupled product)
DPRINT	1.2	■ Updated in line with BS2000/OSD-BC V8.0, support for SQ servers
DRV	3.2	■ Updated in line with BS2000/OSD-BC V8.0, support for SQ servers
FDDRL	17.0	■ Multiplexing factor 4 (performance) ■ Extension for disaster recovery
HIPLEX-AF	3.3	■ Updated in line with BS2000/OSD-BC V8.0, support for SQ servers
HIPLEX-MSCF	6.0	■ Updated in line with BS2000/OSD-BC V8.0, support for SQ servers (technically coupled product)
HSMS	9.0	■ Performance improvements for small files ■ Fast migration ■ Fast directory restore
interNet Services	3.3	■ Support for the OSD V8.0 MAIL-FILE command ■ Support for SQ servers
JV	15.0	■ Performance improvement in the mass handling of JVs ■ Updated in line with BS2000/OSD-BC V8.0, support for SQ servers
LMS	3.4	■ Updated in line with BS2000/OSD-BC V8.0, support for SQ servers
MAREN	12.0	■ Support for LTO-4 encryption (with OSD V8.0) ■ Volume information output as csv file ■ Easier MAREN catalog recovery ■ Easier transfer of HSMS / ARCHIVE data
openCRYPT-SERV	1.3	■ Updated in line with BS2000/OSD-BC V8.0, support for SQ servers
openNet Server	3.3	■ IPv6 extensions ■ Classless Interdomain Routing ■ IPSec - IKE V2 ■ Online backup and restore of BCAM configuration ■ LWRESD rebasing
openSM2	8.0	■ Updated in line with BS2000/OSD-BC V8.0, support for SQ servers (technically coupled product)
PCS	2.9	■ Updated in line with BS2000/OSD-BC V8.0, support for SQ servers (technically coupled product)
PROP-XT	1.3	■ Updated in line with BS2000/OSD-BC V8.0, support for SQ servers
RFA	17.0	■ Updated in line with BS2000/OSD-BC V8.0, support for SQ servers (technically coupled product)

Product	Version	■ New function with OSD V8.0
RSO	3.6	■ Secure printing with IPP protocol via encryption ■ Updated in line with BS2000/OSD-BC V8.0, support for SQ servers
SBA-BS2	6.2	■ Updated in line with BS2000/OSD-BC V8.0, support for SQ servers
SCA	17.0	■ Updated in line with BS2000/OSD-BC V8.0, support for SQ servers (technically coupled product)
SCCA-BS2	2.0	■ Support for Pubset Support + FibreCAT CX extensions with EMC ControlCenter V6.1
SDF-P	2.5	■ Updated in line with BS2000/OSD-BC V8.0, support for SQ servers
SECOS	5.2	■ Barring of user IDs / terminals ■ Default protection for user IDs (“policy enforcement”) ■ Centralization of SECOS administration ■ Extensions to SAT logging (XML output, additional MAREN events with LTO-4 encryption)
SHC-OSD	7.0	■ Updated in line with BS2000/OSD-BC V8.0, support for SQ servers ■ New DMX, Enginuity, SYMAPI version ■ Performance Part 2 ■ SNAP save device activation ■ FibreCAT CX integration for SX/SQ servers, support for SnapView
SPACEOPT	5.0	■ Reduction of file extents for large files ■ Updated in line with BS2000/OSD-BC V8.0, support for SQ servers (technically coupled product)
TASKDATE	17.0	■ Updated in line with BS2000/OSD-BC V8.0, support for SQ servers (technically coupled product)
TIAM	13.2	■ Updated in line with BS2000/OSD-BC V8.0, support for SQ servers
VTSU-X.29	1.5	■ Updated in line with BS2000/OSD-BC V8.0, support for SQ servers

OSD/XC package for OSD V8.0: OSD/XC V4.0 for SX and SQ servers

Functional area	Product name	Version
Operating system	BS2000/OSD-BC	V8.0
Communications	openNet Server	V3.3
	TIAM	V13.2
Job scheduling	JV	V15.0
Performance management	SCA	V17.0
Print management	RSO	V3.6
Programming systems	CRTE	V2.8
	EDT	V17.0
Storage management	ARCHIVE	V9.0
	HSMS	V9.0
Utilities	LMS	V3.4
	PERCON	V2.9
	SORT	V7.9

Published by department:

Margret Germann
Phone: ++49 (0)89 3222 2623
Fax: ++49 (0)89 3222 329 2623
Margret.Germann@ts.fujitsu.com
ts.fujitsu.com

Extranet
extranet.ts.fujitsu.com