

DAB (BS2000/OSD) Version 9.0

Issue October 2004

Pages 3

Product Characteristics

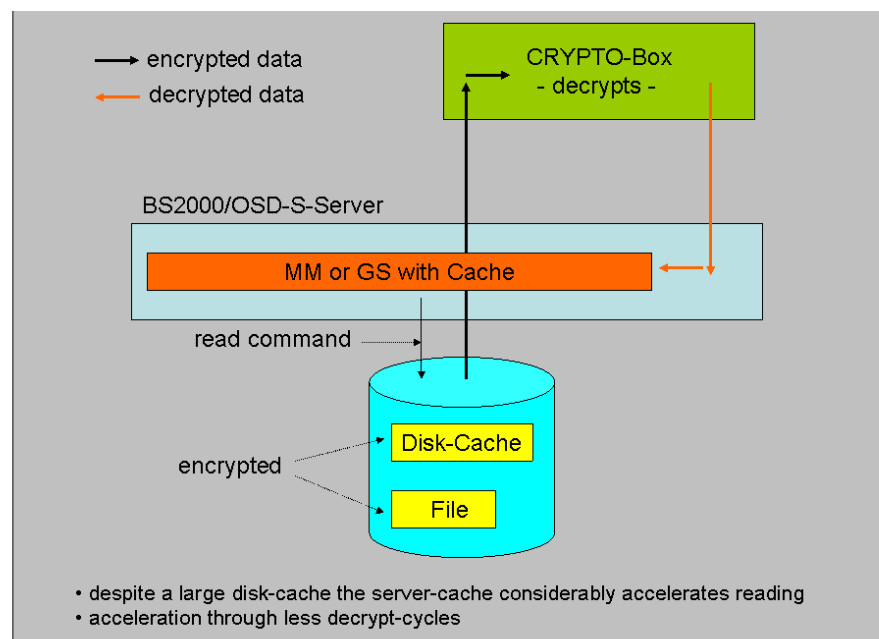
Sophisticated, scalable and self-adaptive caching technologies are fundamental for performant, self-optimising E-Business-applications with guaranteed response-times (Business Critical Computing). Caches accelerate the reading and writing of data by factors. They are applied in networks, servers and disk subsystems. DAB is the BS2000/OSD-software for disk data caching in main memories or/and in Global Storages of BS2000-Servers.

Disk Access Buffer (DAB)

- n manages the caching of peripheral data on high-speed semiconductor memories. It caches data according to its temporary access profile. This results in significantly enhanced I/O performance and much quicker applications.
- n buffers the data in intermediate storage (caching) with far shorter access times than those possible for external data storage (disk storage), even if the disk storage has its own cache. This considerably reduces the average I/O response time. I/O-intensive applications benefit in particular as their throughput and response times are significantly improved by DAB.
- n can use various storage media (such as main memory (MM), expansion storage (ES) and global storage (GS)) as buffers.
- n supports main memories up to 2 TBytes and files up to 4 TBytes..
- n supports data spaces on SX-servers and file encryption.

The advantages of DAB are apparent in both sequential processing and in very localized random-access processing. DAB also provides write caching for significant performance gains in applications with high write rates.

Sketch of DAB-Operation with encrypted data: Both encryption/decryption and DAB-operation are transparent for the applications



Functional Description

To achieve appropriate performance gains, caching can be tuned to the data access behavior of the programs. For this purpose, DAB offers three different caching modes and two different caching techniques, with considerable differences as to optimum usability, possible performance increases and data security. In addition to the read cache, a write cache and a combined read/write cache are also supported.

With the read cache only read accesses to the data areas served by DAB are buffered. On every read access, the data (including any relevant adjoining areas) is stored in the cache, if it has not already been stored. Subsequent read accesses to the same data can then be satisfied from the cache much more efficiently. This data stored in the cache is transferred directly to the user input area, i.e. without disk access. In the case of writing, a record is always transferred to the disk. If this record also exists in the cache area, it is updated there in parallel. This ensures that the current status is always stored both on the disk and in the cache, to ensure an efficient response to subsequent read accesses.

With the write cache only write accesses to the data areas served by DAB are buffered. On such write accesses, the data to be written is first stored in a cache area. This data is saved at a later stage by DAB. If all available cache segments are occupied, the cache segment which has not been accessed for the longest time and whose data has been saved on the data medium is overwritten. In the case of read accesses, only the read hits are served by the DAB buffer, i.e. in the event of a read miss, no data is stored in the cache.

With the read/write cache both the read and the write accesses to the data areas served by DAB are buffered. This function therefore represents a combination of the two functions detailed above.

DAB supports all caching modes in all storage media. However, it should be noted that in volatile cache media such as MM/ES, the write cache data is lost in the event of a system crash. With the non-volatile GS cache medium, DAB allows failsafe, rapid writing. In this case, DAB rebuilds the data in the GS following a server failure due to a software crash, and the battery backup provides GS protection in the event of system failures due to power interrupts. DAB V9.0 operates with global storages up to 128 terabytes, main memories up to 2 terabytes and files up to 4 terabytes.

Displacement algorithms

DAB has two displacement algorithms which can be set via parameter control; these are displacement according to LRU and resident buffering.

Displacement according to LRU

Here a DAB cache area is set up with a freely selectable, fixed size. If the size of the cache area is smaller than the sum of the disk areas to be supported, the cache area is managed according to the LRU (least recently used) algorithm. If there is insufficient space for storing new data, a check is done to see which data has not been used for the longest time. This data is then overwritten by the new data to be stored.

Resident buffering

With resident buffering, the data of the disk areas supported resides permanently in the cache medium and is not displaced. After initial cache tuning, this cache mode produces only hits, all I/Os are served from the cache. Resident buffering helps speed up critical user processes by avoiding physical input/output, regardless of the access location.

Recommendations for successful use

In DAB V9.0, the system administrator simply selects the disks for caching, where the cacheable files are stored – DAB does the rest.

DAB dynamically calculates the current access profile for all the files on selected disks. A distinction is made between sequential processing, where large data areas are read in by DAB in advance for short periods (requiring only a small amount of cache storage), random access processing which is largely localized (e.g. transaction processing or database accesses), where only the data required by the application is read into the cache, and random access processing which is not localized, where caching has no positive effect on performance but uses cache storage. For this reason, DAB excludes such files from caching.

AutoDAB simplifies cache administration by allowing DAB to make the selection dynamically, automatically and self-optimising rather than the system administrator (or the user).

Benefits:

- n best possible hit rate and therefore optimum I/O throughput
- n optimum cache utilization through automation,
- n reduced system administration effort.

The new "Automatic Caching" functionality is offered for all files: Files on SF and SM pubsets as well as on private disks - on all disks in the form of ADM-PFA caching and also on public disks in the form of USER-PFA caching.

DAB extends the System Managed Storage concept (SMS). SMS is a concept for data and storage management, designed to make more efficient use of external storage resources. With other products and components of BS2000/OSD-BC V6.0, DAB meets the demand for high-performance data access.

Supplementary Data on DAB (BS2000/OSD) Version 9.0

Technical requirements

Hardware

BS2000/OSD Business Server

Memory requirements:

In addition to the main memory required by BS2000, memory is also required for administration tables in the cache areas. The size of the MM/ES/GS is dependent on the memory upgrade options of the relevant server and on the number and size of the disks and disk areas to be buffered.

Software

DAB V9.0 runs with BS2000/OSD V6.0 or OSD/XC V2.0.

Also recommended:

SM2 (for analyzing system utilization)

Operating mode

Batch, interactive operation

Implementation language

Assembler/SPL

User interface

Command interface for operator and system administrator

Installation

DAB can be used for all servers and disk storage devices. No modifications to application programs are necessary. DSSM is used for installation.

Documentation

DAB User Guide

DAB Flyer

Demands on the user

Knowledge of system administration

DAB can be managed using five system administrator commands (START, STOP, MODIFY, FORCE-STOP, SHOW-DAB CACHING commands).

Conditions

This software product is supplied to the customer under our conditions against a single payment or installments.

Warranty

Class: A

Delivery format: Machine language

Ordering and Delivery

This software product may be obtained from your local Fujitsu Siemens Computers regional office.