

VM2000 (BS2000/OSD) Version 8.0

Issue December 2006

Pages 3

Virtual Machine System

VM2000 supports the concurrent operation of several BS2000/OSD systems on one server. The server operating resources of a real server can currently be allocated to up to 15 virtual systems (guest systems). The guest systems are totally segregated from one another, i.e. problems in one guest system (up to total crash) have no impact on the rest. No guest system can access main memory areas, global storage areas or devices assigned to other guest systems. The sharing and assignment of peripherals can be modified and expanded online.

VM2000 automatically performs dynamic CPU scheduling across all guest systems (CPU sharing), thereby substantially increasing utilization of CPU capacity. The ability to modify VM2000's CPU scheduling strategy online according to the measured load and predefinable rules means that VM2000 can be used to provide an efficient service level management tool.

The guest systems running under VM2000 have the same functionality as in native mode in terms of command set, networking capability, and testing and debugging tools. A central VM accounting function is available for global accounting of the guest systems.



Functional description

In order to optimize the handling of diverse IT tasks, there is a growing requirement for the BS2000/OSD operating system to support different system environments concurrently. VM2000 supports concurrent operation of different system environments on one server, each totally insulated from the others. This allows the deployment options outlined below.

Version migration

Problem-free production operation has top priority, even during version upgrades. With VM2000, the individual production applications can be tested in a second guest system until all errors have been eliminated. Only then are the applications deployed in the production system. Migration with VM2000 thus proceeds free of time pressure. Precautions that were necessary in the past, such as lights-out test runs and weekend work, can be forgotten with VM2000. A further advantage is that operating personnel can get to know the new version at their own pace during the test operations. Using VM2000 thus also helps minimize subsequent operator errors during operation of the production system.

Concurrent operation of several production systems

Sometimes it is not possible to run production applications in the same system environment. If same names (for Userids, Pubsets and so on) are defined, they have to run in different systems. If an application should no longer be migrated to the new BS2000 version, it needs an older system environment to run.

Concurrent production, development and testing

Test systems always harbor an inherent risk of generating errors. For this reason it is not usual to perform tests in parallel with live operation on production systems. VM2000 allows this. Test and production systems can run on the same installation because system errors in one system have no impact on the other concurrently active systems.

Differentiated systems for e.g. service providers

Data centers that provide different customers with computing services (for production or backup) must have a number of computers available to support the different system environments necessary for this. VM2000 gives data center operators the option of installing a single high-performance computer, thus saving on floor space as well as on administrative overhead.

Balancing of load peaks, use for backup concepts

VM2000 also includes a function enabling the main memory domain to be increased or reduced in size dynamically, allowing further optimization of deployment scenarios. With this feature, there is greater scope for balancing out load peaks on guest systems; temporary virtual machines can be set up "on the fly"; and backup guest systems can be provided with minimal use of resources – all without the need to terminate live guest systems.

VM2000-Linux add-on product

VM2000-Linux extends the VM2000 virtual machine system of BS2000/OSD so that Linux systems can also run on the servers. Up to fourteen Linux systems can run simultaneously on selected BS2000/OSD business servers of the S series. Alternatively, several BS2000 systems and several Linux systems can run concurrently in parallel with and independently of one another (max. 15 systems in total).

Program description

Hypervisor

This is the central execution scheduler and monitoring component in VM2000. The Hypervisor runs in a special privileged state, which allows the Firmware settings for insulating the VM resources. Its most important tasks are:

- Resource management: If a resource is assigned to a VM, Firmware setting have to be done by the Hypervisor.
- Switching the CPUs between the VMs. The context of their virtual CPUs have to be switched and stored.

Monitor

The monitor performs the communication of the Hypervisor with the outside world. The administration of the VM2000 system is done within the monitor system. The monitor is implemented as a DSSM subsystem which runs in guest system 1 in BS2000 using standard BS2000 functions. This ensures a standardized user interface for VM2000 and BS2000. The following functions are implemented in the monitor:

- Command server for the VM2000 commands
- Message output
- VM2000 accounting
- Management and operation of virtual consoles (\$VMCONS application)
- Management of administration dialogs for the entire VM2000 system as well as (optionally) for the individual VMs
- VM2000 hardware error logging

VM1 – VMn

The individual VMs encapsulate the virtualized HW-environment for the guest systems. An own administration environment can be established for each VM.

Guest systems

The BS2000 systems running in the VMs are known as guest systems. Guest systems essentially use the hardware resources assigned to their VM (CPU/main memory/devices) directly, i.e. without the hypervisor as intermediary. The SPD and SPVS disks operated by the hypervisor via "indirect I/O" for BS2000 guest systems are an exception.

The operating of a single system is done by independent KVP-consoles. The function set of the guest systems in VM2000 operation is identical to that of the BS2000 systems in native mode. A connection can be set up to any guest system from the LAN network. At least one HNC 9185x interface adapter is required in order to connect a guest system to the network. At SX servers the VMs share the integrated LAN-Controllers. Additionally a FW-based Inter-VM-communication (VMnet) is implemented there.

New functions of version V8.0

Sophisticated possibilities of CPU-assignment

1. VM-Groups
VMs can be combined to groups. The VM2000-administrator can assign CPU-share in two steps: In a first step Quota and limitation for the group and then in a second step he portions this to the VMs of the group. If some members do not utilize their share, then other VMs of the group can prior use that CPU power. The administrator can offer better service level agreements to his customers who own several VMs.
2. CPU-Pools
The administrator can partition the CPUs of the server into CPU-Pools and can assign the VMs to these pools. New business models can be defined, better performance at big configurations is possible. CPU-Pools and VM-Groups can be established and changed "on the fly".
3. Dedicated CPUs
Is a CPU-Pool configured in a way that each active virtual CPU can be assigned to an own CPU, then a binding of the VMs to the CPUs is done. Best performance for guest systems is possible.
4. Affinity-based VM-Scheduling
The follow-up scheduling of a virtual CPU will be done as possible on the last used CPU. A reuse of still existing cache- and TLB-lines is made possible and so the processor utilization is improved. The Hypervisor pays attention that the CPUs will be utilized equally and with virtual CPUs of similar priority.

Dynamic Reconfiguration at SX150 servers

Within specified scenarios concerning maintenance and Capacity on Demand system boards (XSBs) can be switched on or off.

Capacity on Demand for SX150 servers

On the SX150 servers, the system operator can attach dynamically preconfigured, designated CPUs to cover temporary performance peaks.

Technical requirements

Hardware

BS2000/OSD S series and SX series Business Servers. Business servers S120, S140, S145, S155, S165, S170, S180, S190 and S200 enable the operation of Linux systems. On SX servers VM2000 V8.0 requires the version 3.0 of the X2000 carrier system.

Operating the guest systems using operator terminals connected to the S server SCP (except SCP 3970-40) resp. the SX server console requires a special firmware supplied with the VM2000 function disk resp. VM2000 function CD.

Software

S-Server:

Monitor system:

BS2000/OSD-BC V5.0, V6.0 or V7.0

Guest systems:

BS2000/OSD-BC V4.0 or higher

On S165/S200: BS2000/OSD-BC V5.0 or higher

SX-Server:

Monitor system:

OSD/XC V1.x*, V2.x or V3.0

Guest systems:

OSD/XC V1.x or higher

*V1.x only with SX100-B, SX130, SX140.

OMNIS (as an option)

OMNIS is needed for guest system operating over virtual consoles or VM administration over VMCONS.

Max. number of guest systems supported:

On S Servers: 15 (7 on S110, S115, S120)

On SX Servers: dependent on model, up to 15

Operating mode

Interactive dialog and batch mode

Implementation language

Assembler

User interface

SDF command language

Installation

In accordance with release notice

Documentation

User guide

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