FUJITSU

WHITE PAPER

HIPLEX
The BS2000/OSD-Cluster

Issue April 2010

Pages 6

To stay ahead of the game in the era of digital communication, enterprises make high demands on the information technology (IT).

For instance, IT platforms have to be more open than ever before, which makes for extremely flexible functionality and easy integration into all relevant infrastructures. Above all, IT systems must attain a maximum of availability and scalability.

As a result of global business transactions the demand for permanent IT availability is growing dramatically. Permanent availability means just 24 x 7 x 365 operation and fast system response times.

Contents

Business Critical Computing

Today's travel and seat reservation systems, for example, must be available round the clock. Electronic mail order houses must be able to serve their customers day and night without problems and on a worldwide scale. It is equally unthinkable for operators of information systems and service facilities (such as credit card organizations) to shut down the host, even for only a few minutes.

Customers expect permanent availability

Interruptions to business processes generate a high financial demage. Quite apart from financial losses due to lost sales or compensation claims, stricken organizations also have to contend with loss of image and customer satisfaction, if not even loss of creditworthiness.

Operators of IT infrastructures therefore have to find ways of keeping IT applications permanently available. While increasing the availability of IT subsystems is a necessary precondition, it is not enough on its own. Maximum application availability can only be achieved by building redundancies into the configuration, with multiple instances of the runtime components supporting these applications.

Key success factor in Business Critical Computing

Permanent availability of business processes and of the IT processes supporting them has become a key success factor for companies. In other words, network access points should never fail, servers never stand idle, and data never is inaccessible. The IT response to these high availability requirements is network, server and data redundancy, and mechanisms that use this redundancy to bypass local failures.

Solutions like these are often centered round mainframe systems such as BS2000/OSD, which support Business Critical Computing. And when computing is critical to business, this in essence means it is imperative for all the data and applications to be available permanently to the parties in the business process, since productive working is not possible without the data and applications in question.

Business Critical Computing thus includes:

- protection against data loss
- hardware and software reliability
- ability to perform all maintenance and administration tasks during online operation
- disaster protection
- scalability of performance and management of unplanned load situations
- data security

All these features are traditional strengths of BS2000/OSD.

High availability

Today's BS2000/OSD servers achieve a system availability of 99.9 % thanks to the fine granularity of the recovery features that are standard on every server. They carry the main load in Business Critical Computing in many organizations.

This top-level reliability derives from decades- long development and refinement of the fault detection and reconfiguration capabilities in the server hardware and operating system. These features make the BS2000/OSD servers extremely reliable and highly available enterprise servers.

However, business critical applications and the need for mobile connection of users generate demands for more availability than an individual server is able to deliver. Automatic disaster recovery, for instance, is a tough challenge for a single server.

What is meant by high availability?

A business process is highly available if its resources are always ready when needed and the overall process flow can be maintained or quickly restored in the event of a fault. This means, among other things, that the supporting IT process must also be made highly available.

In the IT configuration, the availability behavior of the subsystems such as

- networks with their operating software,
- data and its backup,
- servers with operating systems,
- application software and middleware, and
- IT operation must be optimized to deal with scheduled and unscheduled outages.

High availability through failover

High system availability is basically achieved using two techniques:

- Fault tolerance:
  Enough redundancy and fault detection logic are built into an existing system to ensure that faults are almost always detected immediately and a switchover to redundant components can be made without delay.
  The application is then never out of action – but this high availability comes at a price: In normal operation, the redundant components perform no productive work and are therefore very expensive.
- Failover:
  An essentially identical second system is connected to an existing system. The two systems monitor each other and each acts as the backup for the other. Both systems are productive during normal operation.
  If a problem occurs, business-critical applications on the failing system, including the associated resources, are moved to the live system. This process, called failover, is not based on unused redundancies, which makes it a very cost-effective solution.

Clusters provide failover capability in the mainframe environment

HIPLEX, or **H**ighly **I**ntegrated System Com**plex**, is the Fujitsu Technology Solutions clustering concept for supporting an availability and performance cluster consisting of local or regional BS2000/OSD production servers.

BS2000/OSD failover clusters

In failover clusters, maximum application availability is achieved by moving an application, its resources and its network access point from one server to another (standby system) for troubleshooting. This eliminates the impact of downtimes, repair times and system startup time on application availability.

Application availability attains values of up to 99.999%, which is considerably above the availability achievable with only one server.
In a non-fault situation, the standby system is used for other purposes, i.e. it is productive too. Moreover, the failover architecture is symmetrical: Each system can be a production system and a standby system at the same time. The system on which a failover-capable application normally runs is the production system.

Failover clustering with BS2000/OSD offers lots of advantages and benefits:

- High availability: Very high application availability at reasonable cost
- Universality: It can be implemented for virtually all applications
- Ease of use: No interventions required in existing applications
- No entry threshold: Easy and affordable transition from a single server to a failover cluster
- Cost-effective: No unproductive redundancy
- Full automation: Failover is a fully automated process
- Workload management: The failover function can also be used when exporting applications (e.g. for server maintenance, upgrade or replacement)
- Maximum performance: Free scalability of the failover cluster up to 16 servers
- Design granulate equal to system: Works for systems, and therefore can also be used by VM2000 systems (guests)
- Failover granulate: Individual applications or application groups in a system
- Failover diversity: Provided their resources are independent, different applications of a system can be moved to different standby systems
- Disaster recovery: High availability also in the incalculable event of a total failure of a data center by automatic failover to a distant data center
- Concurrent maintenance, upgrade: Individual servers can be split off from and reattached to the HIPLEX during cluster operation for maintenance, upgrading or replacement.

High-end clustering with BS2000/OSD

Servers consisting of mono- or multiprocessors, each with autonomous BS2000/OSD or VM2000 guest systems, are the objects between which applications can be moved in the HIPLEX. A cluster implemented to increase availability therefore consists of at least two business servers and, optionally, two or more guest systems.

Failover operates between pairs of servers, or between guest systems on pairs of servers. These are linked by way of a shared pubset volume set or Global Storage. The software products HIPLEX MSCF and HIPLEX AF are required on all systems in the cluster. This creates enormous flexibility. It also saves on resources, notably in VM2000, as the resources of the back-up system can be kept to a minimum and dynamically adjusted in the event of a failover.
In the failover cluster, the HIPLEX operates as a hot standby, i.e. as a running operating system corresponding to the production system. If a fault occurs, only the application is restarted, with no need for intervention.

HIPLEX is the solution for implementing an availability and load balancing cluster with BS2000/OSD. The core products HIPLEX MSCF and HIPLEX AF provide the central, fully automatic functions in this cluster.

Disaster recovery

The HIPLEX can also cope with failure of the entire local IT infrastructure. A local failure is a failure restricted to a building, company site or the immediate vicinity of a company site, and is typically caused by flooding or fire or a regional power failure. It also covers events such as the collapse of a building due to earthquake. A regional or global IT process is disaster tolerant if recovery from a complete local failure of its supporting IT infrastructure can be accomplished without significant time delay. Applications on a BS2000/OSD server are disaster-tolerant because they continue running, possibly after a short interruption, on a remote standby server. Various mechanisms (e.g. mirroring techniques) are available in order to generate a suitable HIPLEX infrastructure – extensive connection of network and peripherals of the local server to a standby server.

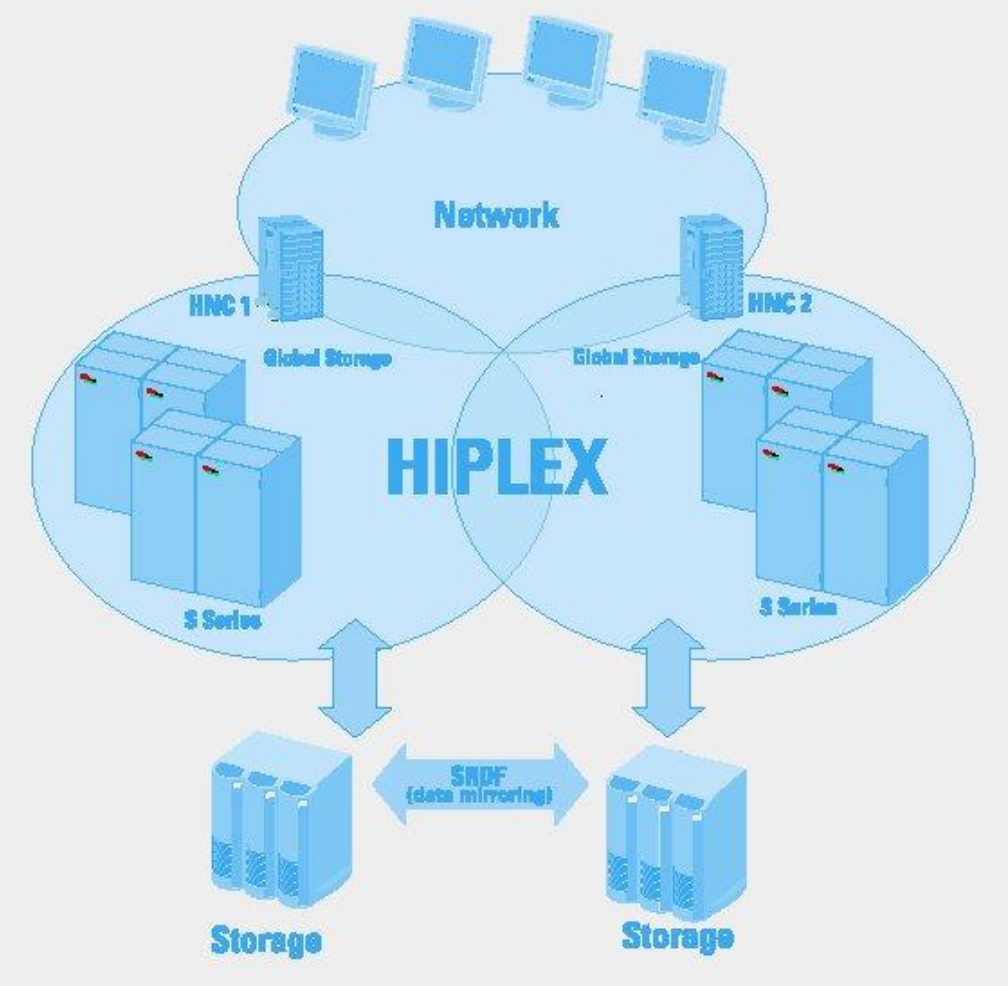

HIPLEX offers more

The central software products in the HIPLEX are HIPLEX MSCF, HIPLEX AF. These are optimized to work with existing administration, automation and application distribution products in BS2000/OSD cluster systems, and play the following roles in the different types of cluster:

Availability cluster:

The availability cluster provides all the advantages of the load balancing cluster and, in addition, guarantees very high application availability because it can respond to fault situations by exporting applications and their resources fully automatically to systems on intact business servers within the cluster.

HIPLEX MSCF detects failures in the availability cluster and **HIPLEX AF** is then able to switch over affected applications and their data selectively and fully automatically to a standby system.

Load balancing cluster:

Since the HIPLEX comprises multiple BS2000/OSD servers and systems, the overall processing workload can be distributed between them either statically or dynamically.

This load balancing cluster increases scalability (with global store more than 10.000 RPF) and also makes for more flexible deployment: applications can be distributed to the guest systems and servers in the HIPLEX either statically or dynamically. In this scheme, HIPLEX MSCF implements the basic cluster mechanisms (shared pubset, lock management, shared library, distributed catalogue management) and static task distribution.

Static methods:

- Application distribution according to the VM2000 model
- Job allocation, administration and monitoring in multiprocessor systems with HIPLEX MSCF and AVAS
- Automatic application export with the aid of HIPLEX AF

Dynamic methods:

- Distributed applications with *open*UTM, UDS/SQL, SESAM/SQL and ORACLE Parallel Server
- Dynamic job distribution with HIPLEX, JMS and/or AVAS

Highlights of HIPLEX

High availability of applications and data

- Automatic application failover
- Disaster recovery

Productive use of new and existing servers

- Cost-efficient, because all the servers and systems are productive
- Investment protection, integration of existing servers and their systems without modifying applications
- Coexistence of different operating system versions

Extremely scalable performance

- Servers and systems can be added during online operation
- Server upgrade without interruption to HIPLEX operation
- Mix of model series possible

Published by department:

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

Partner login
partners.ts.fujitsu.com