

Integration of PRIMECLUSTER and Mission-Critical IA Server PRIMEQUEST

● Masaru Sakai

(Manuscript received May 20, 2005)

Information Technology (IT) systems for today's ubiquitous computing age must be able to flexibly accommodate large-scale changes in workload as well as 24/7 continuous operation. PRIMEQUEST is a mission-critical IA server developed with the key concepts of "open," "mission-critical," and "global." PRIMECLUSTER is a foundation software designed to maximize continuous operation time by increasing system availability via a redundant server, storage, and network configuration. We have combined PRIMECLUSTER with our advanced middleware, servers, and storage to further strengthen our competitiveness in this field. The fusion of PRIMEQUEST, PRIMECLUSTER, and our vast experience in high-reliability technology for UNIX servers has enabled Fujitsu to provide the highest reliability and maximum continuous operation time in open cluster systems. This paper describes how maximum availability is achieved through the integration of PRIMEQUEST and PRIMECLUSTER.

1. Introduction

The use of mission-critical systems has typically been limited to financial accounting systems on mainframes. However, with the arrival of the ubiquitous computing age, mission-critical systems have become more popular and there are increasing needs for them to operate 24 hours a day, 7 days a week. Fujitsu has a lot of expertise in mission-critical systems, and PRIMECLUSTER¹⁾ is designed specifically for such systems. It was developed as a high-availability product to be combined with middleware, UNIX servers, and storage devices. Cluster systems providing redundancy through the use of several servers to keep downtime to the bare minimum are effective for increasing availability, but in such systems it is necessary to rapidly and accurately check the server status and also rapidly switch between servers. The combination of Fujitsu's mission-critical PRIMEQUEST IA servers and PRIMECLUSTER software is proof of our technologies in this area

and realizes the maximum degree of continuous operation in open systems. This paper describes how maximum availability is achieved through the integration of PRIMEQUEST and PRIMECLUSTER.

2. Issues with continuous operation

The most important requirement for continuous 24/7 operation is uninterrupted service, and in the rare cases when service is interrupted, service must be reestablished very quickly. To provide against unplanned system stops due to hardware failures and other failures, cluster systems need redundant configurations so that failover between servers and services can be done rapidly and securely (**Figure 1**).

The main requirements for cluster systems are as follows:

- 1) Rapid detection of faults
- 2) Dependable failover (prevention of double

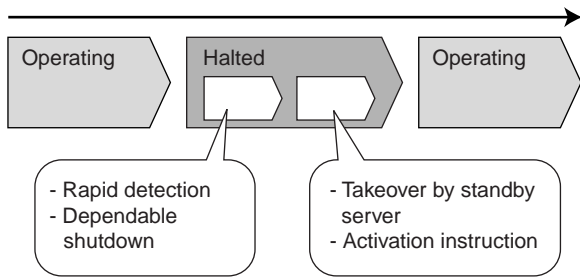


Figure 1
Main requirements for rapid failover.

activation)

3) A high-reliability failover mechanism

Having systems suspended for as little time as possible during failover greatly increases the availability. One way to detect and report OS hangups is to make these tasks event-driven. However, it is difficult to implement them solely in software, and there is therefore a need for collaboration among software, hardware, and firmware to resolve this issue. For instance, in the detection of server hangups, it is common to use a heartbeat that periodically communicates with software. However, a very short heartbeat interval can lead to an undesirably high system load, so heartbeat intervals are usually several 10s of seconds. As a result, it takes several 10s of seconds to detect an OS panic, by which time the server has already halted.

PRIMEQUEST solves the conventional heartbeat-monitoring problems of conventional IA servers by providing special hardware for detecting and reporting an OS panic.

3. Integration of PRIMEQUEST and PRIMECLUSTER

To raise the completeness level of cluster systems, efforts were made from the design stage to solve problems connected with continuous operation in PRIMEQUEST. As a result, PRIMECLUSTER achieves high availability and reliability by establishing a close collaboration among hardware, firmware, and the Linux kernel.

3.1 High-speed failover achieved through redundant server management board (MMB)

PRIMEQUEST has a function by which system boards (SBs) and independent server management units called management boards (MMBs) monitor the OS status and report it to the cluster software. It also has a function by which the cluster software issues a forced panic request to a faulty server (**Figure 2**).

The use of these functions and the integration of PRIMEQUEST and PRIMECLUSTER enables rapid, secure failover between servers.

Normally, the cluster software of ordinary open systems uses heartbeats for monitoring. If there is no response to a heartbeat, failover is initiated because this indicates that the OS has hung up. However, failover is usually not initiated until several 10s of seconds because that is the usual interval between heartbeats. In contrast, it takes the MMB only about a second to report an OS panic.

The combined use of MMBs and heartbeats enables individual faults to be pinpointed, which improves accuracy in fault determination and increases reliability (**Figure 3**) as follows.

- 1) If the MMB determines that the OS is operating but the heartbeat is not getting through, this may indicate that a fault has occurred in the heartbeat communication route or one or more communication processes are proceeding too slowly.
- 2) If an MMB communication error occurs, the status of the heartbeat communication route is checked. If it is operating normally, a fault may have occurred in the MMB.

To increase the reliability of PRIMEQUEST, the MMB and its interface with cluster software are completely redundant.

3.2 I/O fencing incorporated for rapid failover

In a server failover, all access to shared resources must be completed and further access

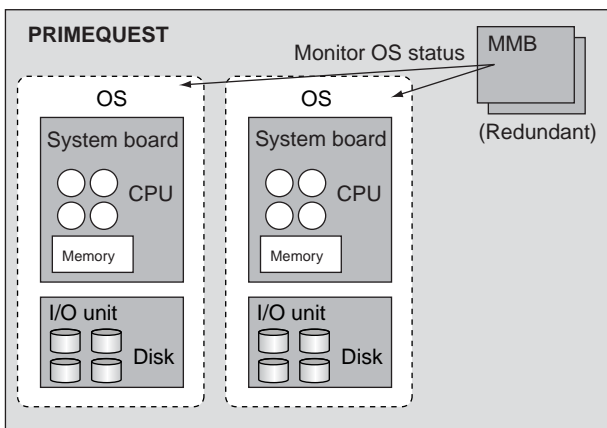


Figure 2
Configuration of MMBs and system boards.

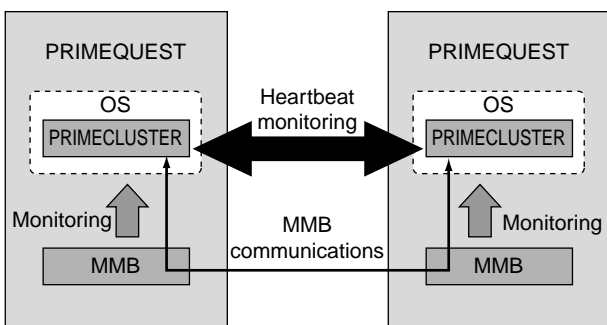


Figure 3
How PRIMECLUSTER uses MMB.

must be blocked before activating the standby server. If the shared disks, IP failover, and other resources are still running, then both servers might try to access the resources, which could lead to the very serious problem of data damage.

For this reason, the PRIMEQUEST OS, drivers, and firmware implement an I/O fencing function that shuts off the I/O immediately at OS panic and other system-down events. When an MMB reports a failure to the cluster software of the standby server, the standby server can be rapidly activated because the I/O fencing secures the I/O.

3.3 Reason for redundant configuration

If a cluster failover mechanism is unreliable, failover between units may fail and service may

become unavailable. This would be very serious and could stop operation of the entire system.

From this viewpoint, as shown in the configuration in **Figure 4**, PRIMEQUEST's MMBs and communications channels are completely redundant, as are its internode communications channels. Moreover, communications are conducted between the kernel and process layers, with the process layer using the redundant communication channels of the kernel layer.

4. Overview of PRIMECLUSTER

PRIMECLUSTER is a high-reliability core software that maximizes operation time in a redundant configuration of enhanced-availability servers, storage devices, and networks.

4.1 Redundancy in servers

Monitoring is conducted between the servers of a cluster using heartbeats and a LAN. If there is no response to a heartbeat, failover to the standby server is initiated and the standby server takes over the jobs that were running. This method in conjunction with hardware monitoring achieves rapid failover and high-reliability (**Figure 5**).

PRIMECLUSTER enables heartbeats and the hardware server monitoring mechanism to be used together to monitor the server status. This combination ensures high availability by reducing the time needed to detect server faults. It also ensures high reliability in terms of secure and dependable server failover.

Next, we describe the main features of PRIMECLUSTER in detail.

1) Synchronous and asynchronous monitoring

By combining monitoring using several channels with the monitoring mechanism incorporated in the hardware, PRIMECLUSTER enables rapid, high-accuracy detection of server failures and secure and dependable failover to another server.

- Synchronous monitoring

Synchronous monitoring reliably detects server downs and hangups using fixed-cycle

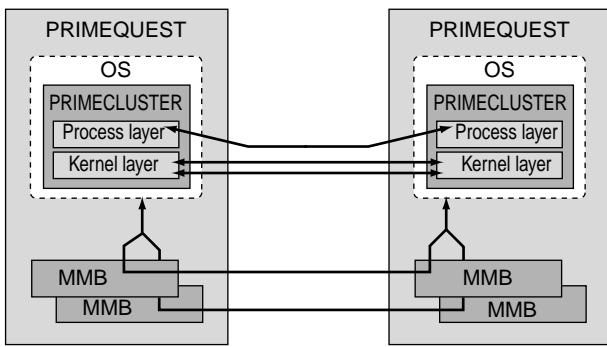


Figure 4 Redundant configuration of PRIMEQUEST and PRIMECLUSTER.

heartbeats transmitted over a dedicated LAN (intercluster LAN) and the redundant routes of the kernel space and user space.

- Asynchronous monitoring

In asynchronous monitoring, when MMBs in PRIMEQUEST detect a server panic, PRIMECLUSTER is immediately notified using a dedicated channel.

2) Disconnection of faulty servers

If there is no response to the heartbeat, the hardware server monitoring mechanism forces the faulty server to shut down. This ensures secure and dependable server failover.

3) Hot standby

The operations for taking over data and the activation of service applications on the standby server when a failure occurs in the operating server are not the same as those in a conventional standby system. With PRIMECLUSTER, these operations are supported by a hot standby function that makes preparations to resume service on the standby server in advance of a failure. Fujitsu's "Symfoware Server" database management software and "Interstage" core application software are designed for use with hot standby. To prepare for a failure, Symfoware and Interstage are activated in the standby server, shared disk devices are opened, and applications are activated to create a standby situation in which processing can be resumed without delay. This substantially reduces the time needed to restart

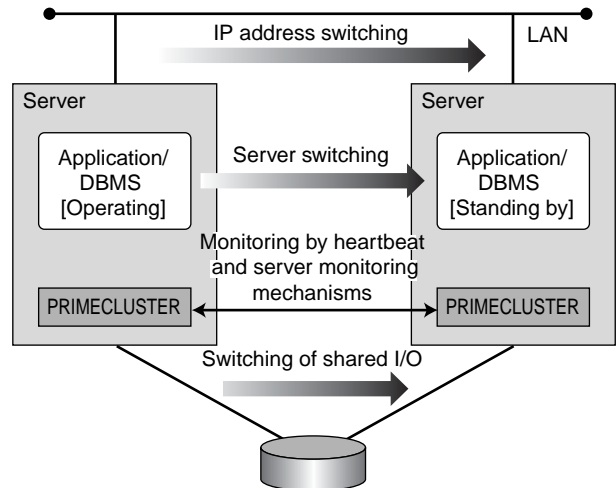


Figure 5 Outline of PRIMECLUSTER.

service.

4) Patrol diagnosis (standby patrol)

In ordinary failover of cluster systems, the server, network, storage devices, and applications of the operating unit of the cluster are monitored, but they are not monitored in standby units. Therefore, service might become completely unavailable if a failure occurs on a standby unit after failover. To guard against this worst-case scenario, PRIMECLUSTER has a standby patrol function by which the server, network, storage devices, and applications are also monitored on the standby units to prevent failure when service is transferred to them. If a fault is detected on a standby unit, the affected unit is disconnected and an alarm is sounded on the PRIMECLUSTER console to alert the system administrator.

4.2 Redundancy of storage and system disks

PRIMECLUSTER GDS provides software volume manager features that enable high-availability logical volumes to be created on each physical volume.

1) Mirroring of system volumes

If a failure occurs on a system volume disk, the system might be unavailable for a long time because it will be necessary to change the faulty

disk and perform recovery processing. However, if system volumes are mirrored and the faulty disk can be disconnected, operations can be taken over by a normally functioning disk. PRIMECLUSTER also has a hot spare function that automatically mirrors a faulty disk to a spare disk that contains no data and then replaces the faulty disk with the spare so system operation can continue.

2) Mirroring among RAID devices

RAID devices are widely used to increase the availability of storage; however, for mission-critical systems, a higher degree of data availability than normal is required. A higher level of continuity in data access is achieved by configuring RAID devices with an even greater degree of redundancy and mirroring among them.

4.3 Redundancy of networks

PRIMECLUSTER GLS provides several IP addresses to high-level software as a single virtual IP address.

1) NIC switching (various machine/multivendor environments)

With Network Interface Card (NIC) switching, switching of transmission routes is controlled through the exclusive use of duplicated NICs connected to the same network. Because no restrictions are placed on opposite-party devices, it is possible to communicate with multiple servers and vendor network equipment.

2) Rapid failover

During normal operation, bandwidth is expanded through the use of parallel transmission routes. Also, by immediately detecting faults using monitoring profiles and rapidly disconnecting faulty transmission routes, service can be continued without applications having to be concerned about faulty or disconnected networks.

4.4 High availability achieved by minimizing scheduled downtime

To maximize the operating time of information systems, it is of course necessary to minimize unplanned system downtime (from the occurrence

of faults to resumption of service). Moreover, planned downtime, such as that for maintenance and system upgrades, must also be minimized.

With conventional IT systems, down time could be scheduled for system maintenance, but this is unacceptable for today's IT systems that must operate 24/7.

The following functions of PRIMECLUSTER systems enable planned downtime to be reduced to the bare minimum.

1) Hot system replacement

The redundancy in the networks and storage devices of this system enables faulty units to be autonomously disconnected from service so they can be changed and redundancy can be reestablished without affecting the service.

2) Rolling update

With this feature, the servers of a cluster system can be stopped one at a time so their hardware and software can be maintained while their jobs are taken over by other servers. This enables the downtime needed for such maintenance to be minimized.

3) Hot system expansion

When business is rapidly expanding, and there are shortfalls in processing power and/or file system capacity, processing power can be increased by adding servers or the file system can be expanded online. By using the hot-system expansion functions of PRIMEQUEST and ETERNUS²⁾ CPUs, the memory and disk capacity can also be increased.

5. Issues for the future

By continuing to target mission-critical systems and high-reliability operation, Fujitsu will apply PRIMECLUSTER and PRIMEQUEST to the TRIOLE IT infrastructure³⁾ in addition to cluster systems. With TRIOLE, Fujitsu has been developing resource visualization (centralized appraisal of configuration, fault locations, performance, etc.) and autonomous control (automatic reallocation of server resources, etc. as required) functions for complex systems. In the future, these

functions will be applied as key technologies for the fault detection and job takeover functions that have been implemented using cluster technology.

6. Conclusion

The integration of PRIMECLUSTER and PRIMEQUEST has enabled high availability and reliability in cluster systems consisting of mission-critical open servers. In our present ubiquitous computing age, there are increasing needs for high availability, optimization, and convenience in the operation of IT systems. Fujitsu will continue to improve the reliability of PRIMECLUSTER and PRIMEQUEST in view of its vital importance to

our business. We will also center the future development of PRIMECLUSTER and PRIMEQUEST on the TRIOLE infrastructure in order to fulfill the expectations of our customers and maintain their confidence in Fujitsu.

References

- 1) PRIMECLUSTER homepage.
http://www.fujitsu.com/global/services/computing/server/unix/optionalsw/PRMPWR_pcl.html
- 2) Disk Storage Systems (ETERNUS) homepage.
<http://www.fujitsu.com/global/services/computing/storage/system/>
- 3) TRIOLE homepage.
<http://www.fujitsu.com/global/services/solutions/triole/>



Masaru Sakai received the B.E. degree in Mathematical Informatics from Tokyo University of Agriculture and Technology, Tokyo, Japan in 1981. He joined Fujitsu Ltd., Kawasaki, Japan in 1981, where he developed OSs for small computers. Since 1994, he has been developing cluster software for open systems.

E-mail: sakai@soft.fujitsu.com