High Availability Server Supports Dependable Infrastructures

MIZUTANI Fumitoshi, ABE Shinji

Abstract

Since IT devices permeate the social infrastructure as life lines that support the ubiquitous society, they are required to be dependable systems with a higher reliability than ever before.

The ft server announced in February this year was developed for the purpose of achieving a robustness and availability in response to such tall orders on open platforms of industry standards. Features and critical technologies of the newly developed “GeminiEngine” LSI, which is the core of such systems, are introduced in this paper.

Keywords

fault tolerant, availability, high-speed resynchronization, ft server

1. Introduction

The phenomenal progress of IT technologies has dramatically changed the foundation of our lives, as well as businesses. These days, as the realization of the ubiquitous society becomes tangibly visible in various shapes and forms, in the background of such development, the permeation of IT devices and network infrastructures is evident in our living environment. The concept of dependability is increasingly important with regards to usability, safety and certainty, in order for us to use such technologies as life lines. In reality, however, it is not possible to say that satisfactory technologies have been realized. It is, therefore, necessary for IT devices that will be supporting the ubiquitous society in the future to have an availability provided by existing backbone systems and a usability that makes it possible for anyone to use them.

2. Commercialization of ft Server Product

Mission critical systems, such as accounting systems for banks and operational management systems for railroads, have always realized a high degree of availability by using mainframe computers or cluster systems. The use of such systems will surely continue in the future, however, in the ubiquitous society it will also be necessary to offer a high degree of availability that is not confined to such a limited scope, but which is available even from IT devices found much closer to home. Since expensive mainframes and cluster systems, which are complicated to operate, are not necessarily suitable for such purposes, there has been a need for low cost products, easily used by anyone while offering a high degree of availability. In response to such anticipation, NEC in June 2001 succeeded in commercializing a fault tolerant server (hereinafter referred to as “ft server”) product based on the IA server, with a dramatically raised level of availability, in cooperation with Stratus of the United States.

(1) Fault Tolerant

The redundancy of hardware makes it possible to continue operations even if one of the two systems fails.

(2) Uninterrupted Maintenance

Replacement of failed parts is possible while the system continues to operate.

(3) Possible Use of General Software

General purpose operating systems, such as Windows or Linux are loaded to achieve an operability similar to that of general purpose servers, in order to make it possible for anyone to easily use the system.

3. New ft Server

As the market recognition of the ft server continues to increase, expectations are raised even further for the ft server to become the platform for the ubiquitous society. In order to further accelerate this trend NEC has been involved in research of the ft server using proprietary technologies since 2003. Up to that point NEC was engaged in product commercialization under a cooperative agreement with Stratus, based on technologies developed by Stratus. In order to pursue the latest techno-

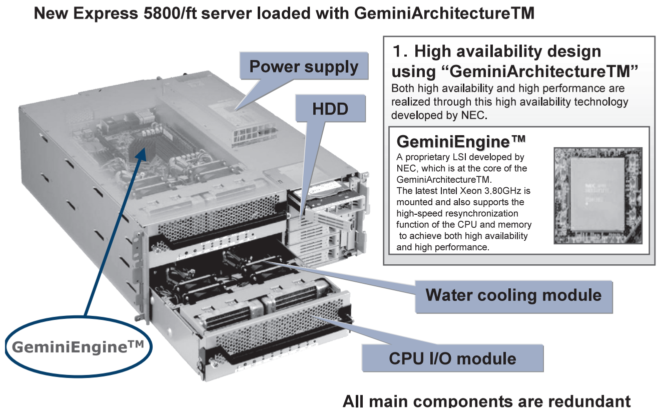
New Express 5800/ft server loaded with GeminiArchitectureTM
Power supply
HDD
1. High availability design using "GeminiArchitectureTM"
Both high availability and high performance are realized through this high availability technology developed by NEC.
GeminiEngine™
A proprietary LSI developed by NEC, which is at the core of the GeminiArchitectureTM. The latest Intel Xeon 3.80GHz is mounted and also supports the high-speed resynchronization function of the CPU and memory to achieve both high availability and high performance.
Water cooling module
GeminiEngine™
CPU I/O module
All main components are redundant

Photo New ft server Express5800/320Fa and GeminiEngine^{TM}.

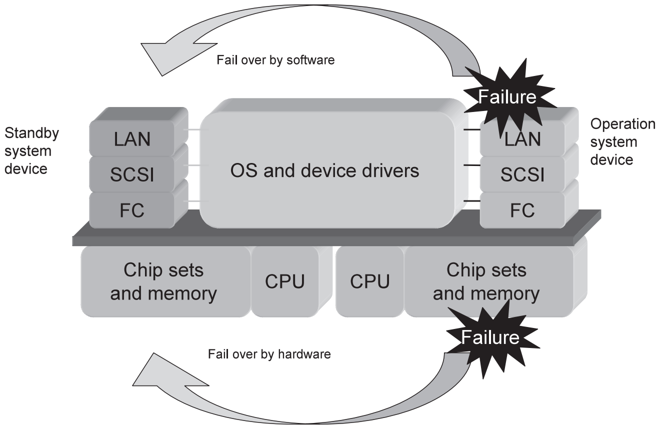

Fig.1 Hardware and software of ft server.

logical trends, reduce costs and make the system available to as broad a range of customers as possible, while responding to a diverse range of requests from customers, NEC decided that it would be necessary to conduct proprietary developments that utilize hardware development technological capabilities, which are a strength of NEC. The activities coincided with the objectives of the "Semiconductor Application Chip Project" of the New Energy and Industrial Technology Development Organization (NEDO) and as a result, product commercialization proceeded with the assistance of NEDO. The product was announced on January 23, 2006 and shipping of the product has already commenced **(Photo)**.

4. Development of New ft Server

The mission assigned to the development team for developing the new ft server was to get the ft server infinitely as close to the general purpose servers as possible in order to make the system easy for anyone to use. Conventional ft servers, due to their peculiar technologies, were not necessarily able to pursue the latest CPUs and technologies that have been commercialized one after another. Furthermore, although general purpose operating systems were used to operate them, it was necessary to make partial adjustments. These were issues that got in the way of speedy product commercialization. The new ft server development project was started with the objective of conquering such issues and making the ft server more widely accepted.

Although a redundancy is already being implemented with many IA servers, storage devices, LANs and power supplies, the redundancy of CPUs and chip sets, which are the main components of systems, is still considered to be difficult to accomplish, due to the technical and cost aspects.

Fig. 1 shows a conceptual diagram of the relationship between the software and hardware in ft server systems. The main components, such as CPUs and chip sets, which are the platform of the system, provide support for the entire system. The operating system and other software run on this platform. Furthermore, I/O components, such as LAN and SCSI, are mounted on this platform as well and operate by receiving instructions from the operating system and drivers.

All components are redundant in the ft server and if an I/O component fails, the software switches the device in use (fail over). On the other hand if CPUs or chip sets, which constitute the platform, fail a hardware fail over occurs, since normal operations of the software are no longer possible. In order to realize this fail over by the hardware, a technology to operate both components in synchrony, as well as with switching, becomes necessary.

5. Synchronization of CPU and Determinism

First of all the basic idea of the ft server is exceedingly simple. The idea is to have two sets of hardware and to keep them synchronized at all times, so that in the event of a failure occurring anywhere, the set that is operating normally can continue operating, which makes it possible to guarantee the continuous operation for the system as a whole. More specifically, since CPUs operate in synchrony with clock signals, they should always be in synchrony as long as the same clock signal is fed to both sets of CPUs from an external clock. This condition is called a lock step synchronization and the feature by which this condition is sustained is called determinism.

In the past, ft servers used such features. These days, however, since every component and interface has been accelerated to faster speeds, the environment surrounding ft servers that utilize determinism is changing dramatically. For example, the acceleration of CPU operating frequencies to high speeds or shifting of technologies, such as from the PCI bus of "parallel bus/low-speed clock synchronous" types to the PCI-Express of "serial link/high-speed clock asynchronous" types, resulting in an amplification of the causes that inhibit the synchronization of a lock step synchronization, such as asynchronous operating elements and clock discrepancies.

These were the most difficult issues that needed to be resolved in order to adopt the latest CPUs and technologies.

6. GeminiEngine

Synchronization of the two sets of hardware was determined by comparing the operations of the PCI bus in both sets with conventional ft servers. Sustaining synchronization of the latest servers that use an asynchronous PCI-Express interface is extremely difficult, as mentioned earlier. Furthermore, it would also be necessary to implement some kind of strategy to counter desynchronization due to the acceleration of CPU operating speeds. Since complete synchronization is difficult to achieve it was decided for the purpose of developing the new ft server, to resolve this issue by establishing a mechanism that can tolerate a certain degree of desynchronization and instantaneously compensate for any desynchronization, even if the gap becomes substantial.

The result of such efforts is the newly developed LSI, "GeminiEngine". The GeminiEngine incorporates functions of North Bridge, including memories and FSB controls, as shown in **Fig. 2** and is connected to both redundant systems through asynchronous high-speed cross links. It is also equipped with a mechanism for operating the system bus of CPUs (FSB), as well as monitoring and controlling both I/O I/F.

A comparison of both systems was conducted using FSB, as well as both I/O I/F, while a high-speed resynchronization process, described later, is made possible by detecting the precursor to desynchronization of CPUs with FSB.

Although the cross links that connect both systems are asynchronous interfaces, the problem is solved by synchronizing them internally in GeminiEngine. Furthermore, the functions that require two chips in conventional ft servers, have been integrated into a single chip with GeminiEngine, contributing to the miniaturization of the equipment, a reduction in costs and improvement to reliability.

7. High-Speed CPU Resynchronization Function

A comparison of both systems is conducted with FSB, as well as both I/O I/F with the GeminiEngine. Desynchronization of CPUs in many cases starts at FSB, including CPU failures or the simple desynchronization of CPUs. However, an immediate resynchronization is not conducted when desynchronization is detected with the FSB.

Causes for the desynchronization of CPUs may arise from the fluctuation of the operating timing within the normal oper-

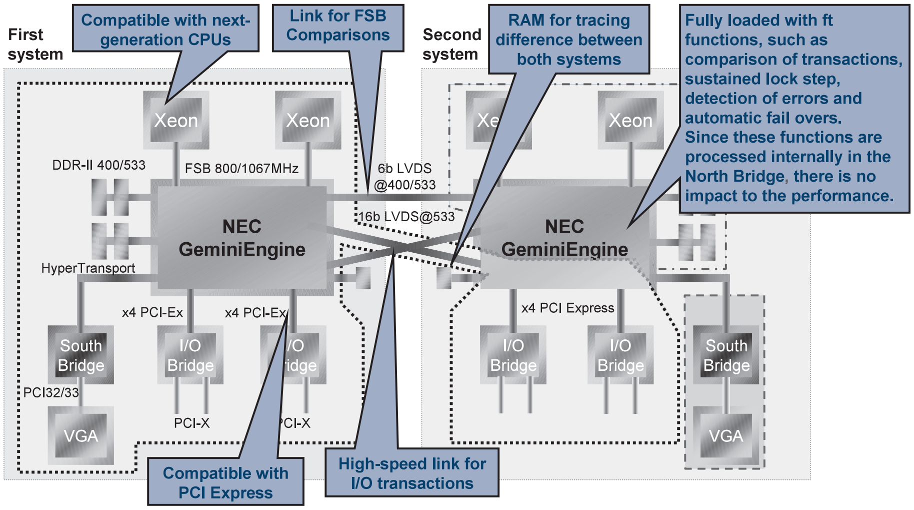

Compatible with next-generation CPUs

Link for FSB Comparisons

RAM for tracing difference between both systems

Fully loaded with ft functions, such as comparison of transactions, sustained lock step, detection of errors and automatic fail overs. Since these functions are processed internally in the North Bridge, there is no impact to the performance.

Compatible with PCI Express

High-speed link for I/O transactions

Fig. 2 Configuration of system using GeminiEngine^{TM}.

1. FSB operation comparison is performed at all times. Retention of write address information at tracer memory (TM) starts when desynchronization is detected.

Operation signal information on FSB is exchanged and compared to detect the desynchronization of CPUs. Since the content of the memory in both systems can potentially become mismatched, due to operations of CPUs that vary because of desynchronization, the write address of the memory is recorded in the tracer memory.

2. Once the detected non-error desynchronization is verified the varying portion only is copied.

The DMA engine of both systems loads information from the tracer memory to copy the applicable address portions.
Once copying has been completed, the CPUs of both systems are reset to synchronize and the context, immediately prior to the termination of the CPUs of the both systems, is restored and the high-speed resynchronization process is completed.

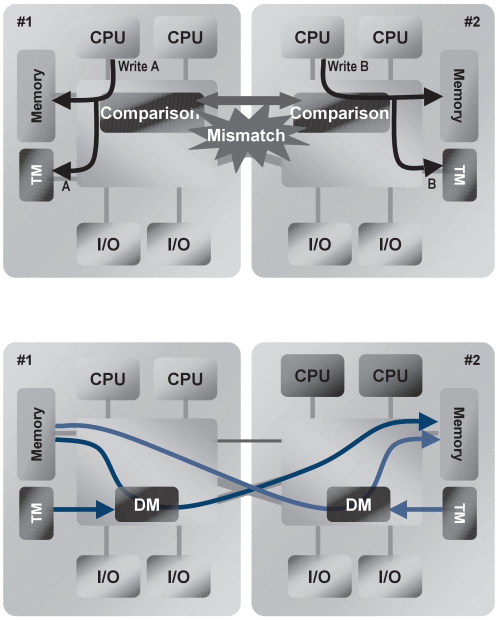

Fig. 3 High-speed resynchronization process flow.

ating range due to fluctuations in the asynchronous circuits inside CPUs, while the possibility of such desynchronization arising from fatal failures inside CPUs also exists. It is not possible to identify the cause and determine the failed CPU at the time desynchronization is detected with FSB.

A redundancy is sustained and retention of the updating information of the main memory starts in the trace memory under GeminiEngine for this reason, when a certain degree of "fluctuation" takes place between the CPUs of both systems. Operation with a certain degree of fluctuation is tolerated while verifying that no errors have been detected. It is only then that the high-speed resynchronization mechanism of GeminiEngine is set in motion. Since this high-speed resynchronization process involves using trace memory to copy only the updated regions in the main memory after desynchronization has been detected, resynchronization is completed within a time interval of about 200 milliseconds. For this reason, neither the software nor the user will ever become aware of this process **(Fig. 3)**.

The GeminiEngine CPU is continuously sustained in a redundant condition during the period starting from when desynchronization is detected until the high-speed resynchronization process starts. It is, therefore, possible to accelerate the high-speed processing of the resynchronization process while disengaging the module with which a failure was detected and performing a degeneration to the CPU that is operating normally, even when a true failure, which caused the desynchronization, is identified during this period.

8. Synchronization Function for Both Systems Prevents Interruption of Services

Dynamic redundancy and the degeneration of CPUs are necessary to realize the hot swapping of CPUs and I/O modules. In other words it is necessary to sustain the condition wherein the software is running on one of the CPU modules, while completely copying the conditions of the operating CPU module to the other CPU module.

Furthermore, since the time required for copying information becomes longer as the capacity of the main memory becomes larger, a mechanism to copy such information in the background without stopping the system is incorporated in GeminiEngine in order to minimize such effects. While the system is in operation the contents of the main memory, which is the source of the copy, continues to change with time. A process is performed for controlling the sequential transfers of data relating to such changes to the target of the copying.

9. Virtualization of I/O Devices

A method of operating one of the redundant I/O devices is adopted while the other set of devices are kept in standby and to which a switch over takes place when a device failure occurs. This is realized by building a virtually singular device as each of two devices, through the coordination of hardware and device drivers. As for conventional servers, when fatal failures occur with an I/O device or routing to I/O, such as PCI-X or PCI-Express, notifications of such incidents are sent to the operating system and often develop into an immediate downing of the system. In order to avoid such developments, all fatal failures of I/O systems are concealed from the operating system with GeminiEngine to prevent downing of the system. This is achieved based on the scheme of the hot plug supported by the operating system, known as surprized removal. In other words, even if a fatal failure occurs, the operating system is not notified immediately but instead, it appears as if an I/O device had suddenly been pulled out. The operating system and drivers detect the disappearance of an I/O device and transfers control to a standby device that will be operating in the place of the missing device, thereby performing a fail over.

Since a series of such procedures are performed within the scope of device drivers, application software will not be aware that a failure occurred or that a fail over has taken place. For this reason, it is possible to use existing applications without any changes.

10. Detection of Error and Detachment

The types of errors detected by GeminiEngine exceed several hundred. A vast number of error detection circuitry is packaged, in comparison with ordinary chip sets. The LSI is internally divided into finely sectioned blocks and when an error is detected, an applicable block is logically detached in order to ensure that erroneous operations are not transmitted. Furthermore, air-tight failure strategies are in place that provide parity protection on all I/F between blocks, as well as the buffers inside. Also, different reset signals are distributed to each individual block to make it possible to restore individual blocks by triggering individual resets.

For example, when an intermittent failure occurs with a CPU or memory and if the control logic also becomes impossible to operate, only the applicable block inside the LSI is detached. The I/O control section continues to be available by taking over the operation from the other set. It is possible to resynchronize by restoring the overall functions of the LSI by individually resetting each block.

11. Conclusion

This paper introduced features of GeminiEngine, the LSI at the core of the new ft server, which was developed as a high availability platform to support dependable infrastructures.

NEC is one of only a few companies offering products of both cluster technologies that dramatically raise the availability of conventional servers, as well as an ft server, which features a high availability technology. By taking advantage of the features of these systems and by putting the right system in the right place, NEC is able to provide solutions that respond to various customer requirements to meet the demands for creating a safe and secure society.

* Windows is a trademark or registered trademark of Microsoft Corporation in the United States and other countries.

* Linux is an appellation that is a registered trademark or trademark of Linus Torvalds in the United States and other countries.

Authors' Profiles

MIZUTANI Fumitoshi
Assistant Manager,
Client and Server Division,
2nd Computers Operations Unit,
NEC Corporation

ABE Shinji
Manager,
NEC Solutions (America), Inc.

●The details about this paper can be seen at the following.
Related URL: http://www.express.nec.co.jp/