Server Platform Optimized for Data Centers

● Franz-Josef Bathe ● Toshio Sugimoto ● Hideaki Maeda
● Teruhisa Taji

Fujitsu began developing its industry-standard server series in the early 1990s under the name FM server and has been marketing these servers globally since 2000 as PRIMERGY servers. Using mainly the Intel architecture in its hardware and supporting standard operating systems having a high industry share like Windows and Linux, the PRIMERGY series continues to provide an extensive lineup of servers to meet diverse customer needs. This paper outlines the features of the PRIMERGY series of servers and describes the energy-saving and cooling technologies of the PRIMERGY CX1000 platform optimized for data centers.

1. Introduction

To meet a wide array of customer needs, Fujitsu is using the latest technologies for its PRIMERGY industry-standard servers and providing diverse types of servers, enabling the most optimal ones to be selected for target applications. The outsourcing business has been attracting attention for some years now, and data centers are coming to deploy a huge number of servers whose performance and mounting density have been improving dramatically. As a result, total power consumption in data centers has been increasing, cooling costs to dissipate the heat emitted by servers have been rising, and the space needed for cooling (cold aisles, etc.) is increasing. At the same time, the high-density complex design of server interiors is resulting in longer maintenance times to deal with failures. In response to these problems, Fujitsu has developed and been shipping the PRIMERGY CX1000 (hereinafter, CX1000) industry-standard server platform for data-center use featuring a new architecture. This paper outlines the entire PRIMERGY product group and describes how Fujitsu is solving these data-center problems with the CX1000 server platform.

2. Features of PRIMERGY series

The PRIMERGY series of servers has been developed in accordance with customer needs. It features an extensive lineup of servers in three types—tower, rack, and blade—that enables the customer to use different servers for different usage scenarios in an optimal manner.

1) Tower servers

PRIMERGY tower servers are usually set up in or close to offices. Featuring whisper-quiet operation, extendibility, versatility, and a compact configuration, they target business operations of various scales. To achieve quiet operation, special care has been paid to the layout of cooling fans and internal components inside the servers, and a structure that promotes the flow of cooling air straight through the server housing (straight cooling architecture) has been used. This is because air flows that bend generate cooling losses and noise.

2) Rack servers

Rack servers have been developed for industry-standard 19-inch racks with the expectation that large numbers of them will be installed in data centers. The idea is to provide high performance, extendibility, and versatility along with energy-saving features. In this regard, achieving high performance and high functionality by high-density mounting within a limited volume of space has become an important objective, but the amount of energy needed for cooling purposes increases as the mounting density rises, which means that more electric power is consumed by the equipment. Fujitsu has therefore designed recent rack-type servers to minimize power consumption by increasing the area of intake vents, using power supplies having good power-conversion efficiency and voltage regulator modules (VRMs), and using low-power-consuming central processing units (CPUs) and memory modules. These design measures make up Fujitsu's Cool Safe energy-saving cooling concept, which supports the entire family of PRIMERGY servers.

3) Blade servers

While blade servers are developed for installation in data centers like rack servers, they aim, in particular, to achieve high server and firmware manageability to provide flexible support when equipment failures occur or business load fluctuates. Blade servers have therefore been designed to provide high levels of performance, availability, and operability as demanded by large-scale system operations. Server blades, storage blades, and other types of blades can be densely mounted, which improves extendibility, and chassis-mounted components can be modularized and deployed in redundant configurations. In short, blade servers aim for high availability and ease-of-use despite high-density mounting. From an operation perspective, blade servers take environmental aspects into account by providing energy-saving controls and visualization of energy usage.

3. Server market trends

As server performance continues to advance, consolidation by server virtualization and integration has been accelerating. According to a survey conducted by Gartner, Inc.,[1] Cloud-oriented servers are expected to account for about 28% of all servers shipped in 2012. At the same time, infrastructures such as data centers that support server consolidation are coming to suffer from two main problems as servers increase dramatically in number and come to be arranged in high-density layouts.

1) Rising operating costs

Cost of energy for running servers and air conditioning, cost of space for housing a huge number of servers, and cost of labor for maintaining and managing servers

2) Facility limitations

Limits on available power, floor load capacity, and cooling capacity

To solve these problems, Fujitsu has been developing its servers on the basis of the concept of simple designs specialized for data centers.

4. Development objectives for data-center-dedicated servers

This section describes Fujitsu's development objectives for servers to be used in data centers aiming for scale-out computing and its approach to meeting those objectives. There are three main development objectives as described below.

1) Reduced operating costs

Operating costs will be reduced in three ways: by reducing power consumption by using server nodes with fanless specifications based on a comprehensive energy-saving design; by enabling failed server nodes to be replaced on-site on a node-by-node basis; and by grouping cable connections at the front of the rack to simplify maintenance, shorten system maintenance time, and improve operating efficiency.

2) Reduced facility investment

Facility deployment expenses will be reduced by increasing the server-mounting density and saving space by reducing the rack footprint and increasing the number of servers that can be accommodated by a rack and by enabling a back-to-back rack setup to eliminate the need for hot aisles behind racks, which are used in conventional layouts. The idea is to provide a simple structural design and the ability to replace main boards on-site so that subcomponents can be used more effectively and server installation (in terms of cooling and cable connections) can be optimized for data-center infrastructures. These measures will help to keep down initial deployment costs and reduce facility costs.

3) Simplified maintenance and operations

Maintenance and operations have already been simplified by enabling customers themselves to replace nodes, and flexibility and customization capabilities have been provided. Since supporting inherited existing technologies and facilities is an important element of system deployment, support for a front-to-back cooling infrastructure can be provided to maintain compatibility with existing data-center layouts. Moreover, a plug-and-play design is also provided for server nodes to enable on-site replacement and general-purpose properties can be built in to make it easy to reuse standard main boards. Design specifications are also adaptable to future delivery models.

The following specific techniques help to reduce costs and power consumption by optimizing the rack structure and adopting simple specifications.

1) Reduced use of cooling fans

While 304 compact fans were needed in the past (RX200: 8 fans × 38 nodes), the development of a dedicated rack with optimized cooling efficiency enables cooling requirements to be met by two large fans.

2) High-efficiency power supply

To support a high-conversion-efficiency 200-V power-supply environment, Fujitsu uses a power supply with a power-conversion efficiency of 92% (conforming to CSCI gold and 80 PLUS Gold at 50% load) achieved by using high-efficiency standard parts and materials.

3) New cooling system

Fujitsu has achieved a high-density rack setup by changing the conventional rear-exhaust design to a top-exhaust one. As a result, there is no need to allocate space for hot aisles, which enables a 40% reduction in rack setup space.

5. CX1000 features and performance

This section introduces development concepts and product features of the CX1000 platform, which is optimized for data centers.

5.1 Concepts and features (improvements over existing servers)

5.1.1 Concepts

There are three key concepts: the server design is optimized for large scale-out in Cloud environments; the same housing size as for standard 19-inch racks is used, enabling installation alongside existing servers in a data center; and the need to allocate space behind racks for expelling heat and performing maintenance work, as in existing data centers, is eliminated.

5.1.2 Features

There are three key features: a server-node design with a power supply unit on each node to reduce material costs and weight; an innovative cooling system featuring central exhaust fans and chimney-type air flow and a simplified rack structure; and a small rack footprint (850 mm × 700 mm) and a rack structure conducive to a back-to-back arrangement, which doubles the server-node density per unit area in the data center (**Figure 1**).

5.2 CX1000 deployment effects

Compared with other rack servers, the CX1000 has a 40% smaller footprint and uses 13% less energy, which produces the following effects. Its independent nodes are unaffected by faults in other nodes (shorter mean time to repair). No space is needed in the data center for hot aisles. The maximum number of racks per unit area is higher. There are no extra costs for unused functions, which reduces facility and operating costs. The CX1000 has the lowest power consumption and highest energy efficiency among scale-out-type data-center architectures.

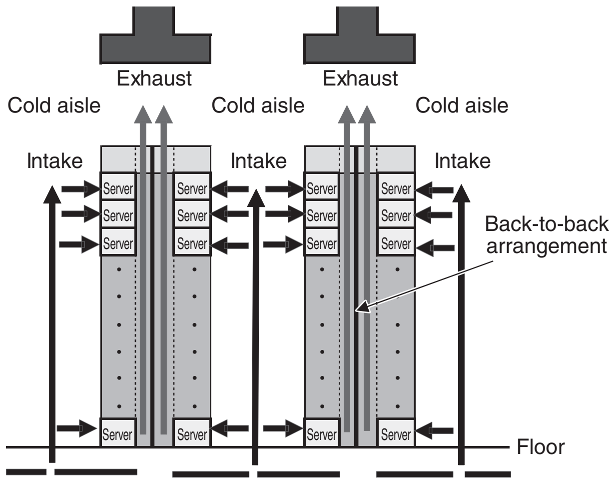

Figure 1
Overview of back-to-back arrangement and cooling system.

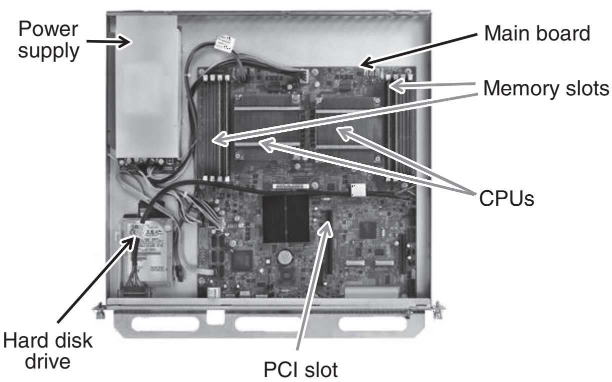

PCI: Peripheral Component Interconnect

Figure 2
Server node.

5.3 CX1000 configuration

The CX1000 rack features a simple and highly cost-efficient design that significantly improves energy efficiency per rack compared with existing rack designs. Up to 38 nodes can be mounted in one rack. Its simple and highly cost-efficient design makes the CX1000 server node a highly flexible, energy-efficient component (**Figure 2**). This server node uses a standard 13-inch × 12-inch main board and a high-efficiency VRM and has a fanless structure. It can also accommodate different types of CPUs and smaller main boards, reflecting its general-purpose properties.

5.4 Cooling technology for energy-saving effects

The CX1000 has a centralized cooling system for the entire rack that holds down power consumption (**Figure 3**). In contrast to existing rear- and front-exhaust systems, this system has fans for cooling the entire rack positioned on top of the rack to direct exhaust air upwards. In addition, the aperture ratio of the exhaust vents differs between the top and bottom of the rack to

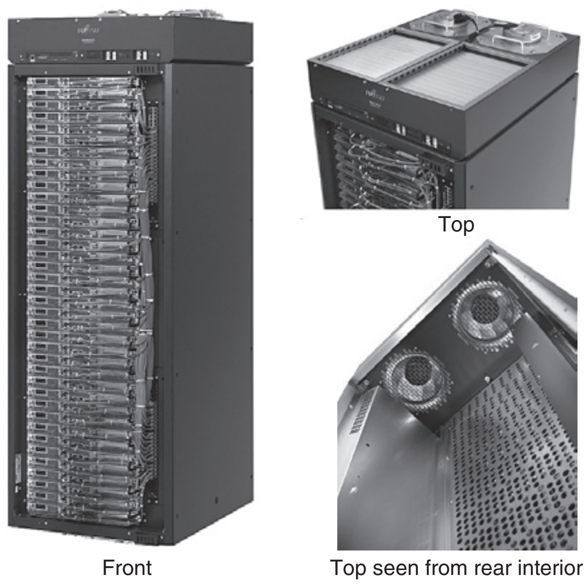

Figure 3
Appearance of CX1000 and its large central cooling fans.

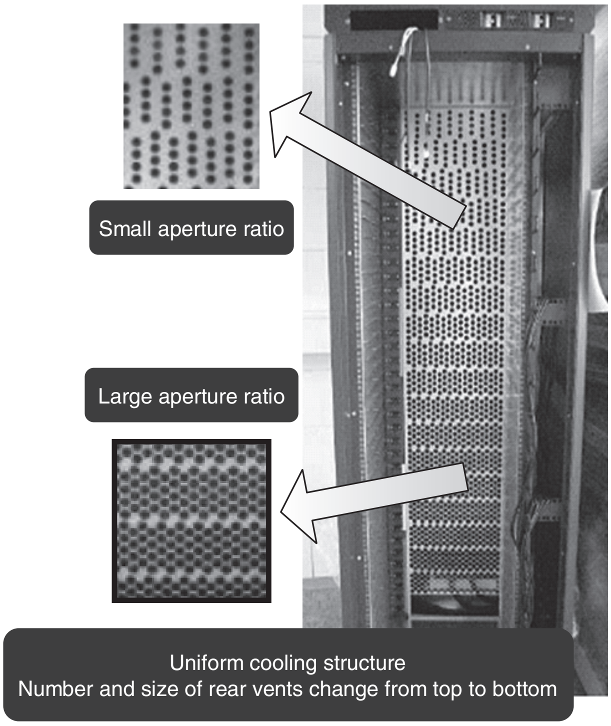

Figure 4
Rear exhaust vents aperture ratios.

adjust the airflow and provide uniform cooling across all server nodes (**Figure 4**). This rack structure was designed and tested by heating/cooling computer-simulation technology.

5.5 Performance

The CX1000 set a world record at 2320 overall ssj_ops/watt[note 1)] in tests performed using the benchmark test program SPECpower_ssj2008[note 2)] on April 7, 2010. This achievement demonstrates that the CX1000 has the highest energy efficiency and performance in the two-socket Xeon 5500 series server class.[2)]

note 1) A unit of computer-system throughput (performance) per watt.

note 2) A benchmark test program developed by the Standard Performance Evaluation Corporation (SPEC) for evaluating energy efficiency by assessing both the power consumption and performance characteristics of a general-purpose server. It indicates server performance per watt where a high value means high efficiency.

5.6 Energy-cost reduction effect

The CX1000 has achieved an energy-cost reduction effect of 1.23 billion yen in three years and the average amount saved annually is about 400 million yen (about 20%) for a large-scale data center (10 000 server nodes). This is composed of 210 million yen representing the annual saving calculated as a proportion of the one-off infrastructure creation saving for hardware procurement (about 15% less than with existing models), annual savings of 120 million yen (about 40% less than with existing models) for floor-space charges (facilities), and 80 million yen for reduced power consumption (up to 7.5 million kWh less [about 25% less annually] than for existing models).

6. Future issues

As a server, the CX1000 is targeted at data centers that aim for scale-out computing, and as a platform, it provides an optimal solution for constructing Cloud computing environments and data centers. From here on, Fujitsu will pursue the development of three key techniques to deal with the increasing diversity of data-center applications: multi-node techniques, control techniques for decreasing power consumption in racks, and localized cooling techniques for precise equipment cooling.

7. Conclusion

This paper introduced the PRIMERGY CX1000 server platform developed to meet data-center needs. It provided some background to product development, described new technologies developed to meet design targets, and discussed the effects achieved. Since the product announcement in spring 2010, Fujitsu has received questions and inquiries from many data centers, and new issues related to product deployment have become clear. Looking forward, Fujitsu intends to resolve these issues to deal with a higher level of customer needs and to perform surveys and research and development to respond flexibly to future trends such as the modularization of data centers.

References

1) Gartner: Dataquest Insight: The Impact of Cloud Computing on Server Sales, 2009 Update. Adrian O'Connell, 25 November 2009.
http://www.gartner.com/DisplayDocument?id=1236313

2) Fujitsu Press Release: Fujitsu's PRIMERGY CX1000 Sets Ground-Breaking SPECpower Benchmark in 2-Socket Xeon 5500 series server class.
http://ts.fujitsu.com/ps2/press/read/news_details.aspx?id=4157

Franz-Josef Bathe
Fujitsu Technology Solutions
Mr. Bathe is engaged in drafting specifications for next-generation data center products.

Hideaki Maeda
Fujitsu Ltd.
Mr. Maeda is engaged in the development of industry-standard server equipment.

Toshio Sugimoto
Fujitsu Laboratories Ltd.
Mr. Sugimoto is engaged in the research and development of elemental power-saving technologies.

Teruhisa Taji
Fujitsu Ltd.
Mr. Taji is engaged in drafting specifications for PRIMERGY products.