

Fujitsu develops prototype of world's first server that simultaneously delivers high performance, flexibility

September 29 2011




Figure 1: Resource Pool Architecture

Fujitsu Laboratories Limited today announced the construction of a next-generation server that, using resource pool architecture, is the world's first to succeed in the simultaneous delivery of high performance and flexibility.

In datacenters used to provide cloud services, ICT systems are configured to enable the efficient delivery of web services and other

tasks. As cloud services become increasingly diverse, however, the ability to configure systems to upgrade performance and [flexibility](#) has become a pressing issue.

By pooling CPUs, [hard disk](#) drives (HDDs) and other hardware components that comprise the ICT [infrastructure](#), and connecting these resources together using high-speed interconnects, Fujitsu Laboratories has enabled the delivery of high-performance server and storage capabilities as needed without losing any of the hardware's original [functionality](#) or performance.

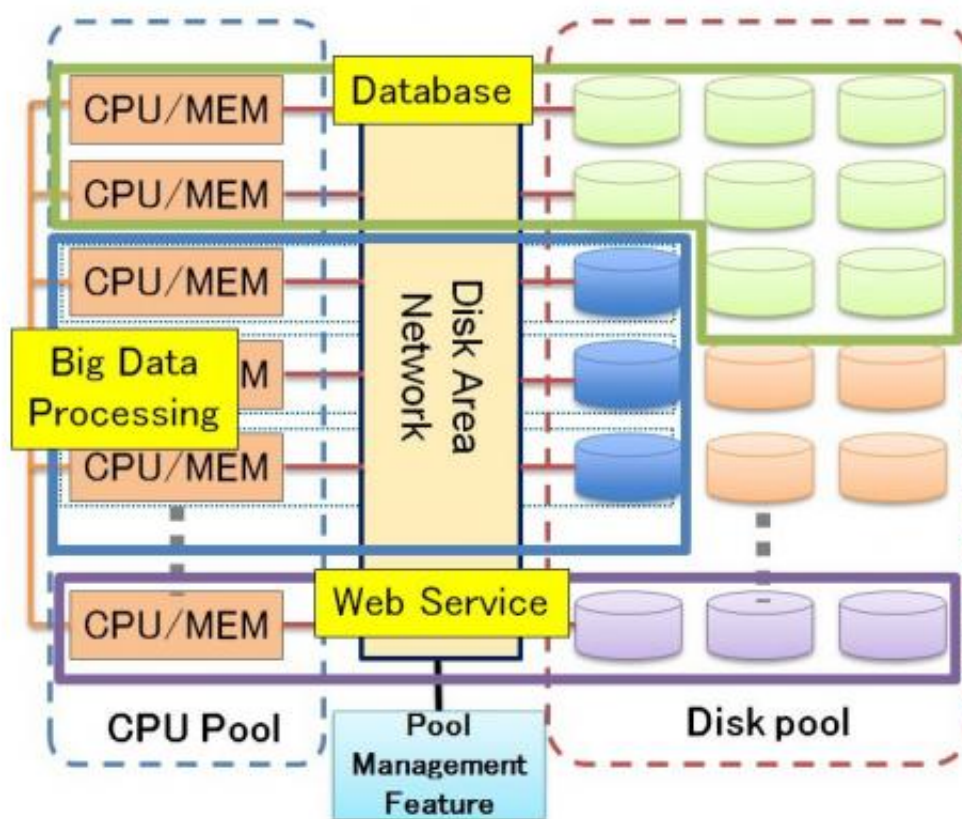


Figure 2: Variety of server/storage resource configurations in one system

This technology makes it possible to configure high-performance ICT infrastructure that can flexibly handle a variety of services. By consistently being able to rapidly configure optimal server and storage resources to fit the unique requirements of the task at hand, whether it is delivering the type of web services currently provided via the cloud, or new services yet to emerge, the new architecture makes possible higher value-added services. As a result, a variety of services are anticipated, including new ICT services, such as those involving the processing of huge sets of data, known as "big data," an essential component of human-centric computing.

With the spread of cloud computing, the role expected of datacenters delivering cloud services is significantly changing. For example, in addition to traditional web services, new services being offered are making backend technologies ever-more important. Furthermore, big data, transmitted by large volumes of [sensors](#) in such fields as lifelogging, medicine, and agriculture, is collected in datacenters and put to use. Those changes drive the datacenters to meet the needs of efficient and flexible processing for an increasing diversity of services.

In datacenters that deliver cloud services, the ICT infrastructure is built by connecting multiple servers and storage machines through a network. In most cases, the configuration of these servers and storage machines is tailored to the delivery of cloud-based services, such as web services. Now, however, in line with the diversification of [cloud services](#), the need to offer services from [datacenters](#) is rising. This includes database services requiring high I/O performance or large-scale data processing tasks that use the local disks of servers, as well as for other services requiring a level of performance that had been difficult to satisfy with configurations geared toward traditional cloud systems.

[Fujitsu](#) Laboratories has developed a resource pool architecture in which the hardware components, such as CPUs and HDDs, are linked together

with high-speed interconnects (Fig. 1). This marked success in building the world's first next-generation server prototype that can be configured either as a server equipped with local disks or as a system with built-in storage capabilities.




Figure 3: Next-generation server prototype

In comparison with typical system configurations tailored to traditional web services, the new prototype delivered a benchmark performance of approximately four times higher I/O throughput as well as increased performance of about 40% when running actual applications.

Features of the newly developed technology are as follows:

1. Pool management feature

In accordance with user requirements for CPUs, HDDs and other needs, the pool management feature allows for necessary resources to be allocated from the pool, the deployment of OS and middleware resources, and the on-demand provision of servers in a required configuration.

2. Middleware that offers storage function using servers apportioned from pool

Using server resources from the pool, storage capabilities are delivered by configuring the middleware, which controls HDD management and data management functions. Whether it is a server with multiple local disks tailored for large-scale data processing tasks, or RAID functions for improved data reliability, the system can be flexibly configured to meet performance and power consumption requirements.

3. High-speed interconnect technology that connects the disk pool

The disk pool comprised of multiple HDDs is connected to the CPU pool via a high-speed interconnect disk area network. The HDDs linked to the CPUs through the disk area network have the same disk access capabilities as the local disks in a typical server, and their performance is not affected by other CPUs. A disk area network was created using prototype interconnects that connect a CPU to a given HDD at a speed of 6 Gbps without any mutual interference.

This technology enables the creation of datacenter ICT systems that deliver both high performance and flexibility. For the delivery of existing web services via the cloud, or even for new services yet to emerge, the optimal server and storage units can be consistently configured to fit the unique requirements of the task at hand, swiftly enabling the delivery of ever-higher value-added services. By leveraging the technology's unique flexibility to adapt a system's configuration, capacity planning requirements for the budgeting of essential resources

are eliminated when designing a system for a new service for which it is difficult to accurately predict swings in load requirements.

Similarly, for cases in which the proportion of the workload accounted for by web services, database functions, or large-scale data processing tasks varies from day to day or year by year, by simply changing the server/storage configuration, optimal services can be delivered at all times with the same ICT infrastructure, resulting in improved datacenter utilization rates.

Moreover, even if any of the CPUs, HDDs or other hardware fails, simply switching the connections to the failed parts will reduce replacement frequency and lower maintenance costs.

The technology's functionality and performance will be verified with a goal of commercially launching the technology in fiscal 2013. Providing flexible and high performance IT infrastructure will contribute to processing big data and the realization of new ICT services.

Source: Fujitsu

Citation: Fujitsu develops prototype of world's first server that simultaneously delivers high performance, flexibility (2011, September 29) retrieved 23 April 2024 from <https://phys.org/news/2011-09-fujitsu-prototype-world-server-simultaneously.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.