Service Oriented Platform

● Hiroshi Yoshida ● Riichiro Take ● Hiro Kishimoto ● Yoshinobu Hibi

Fujitsu is developing a service oriented platform (SOP) as an infrastructure for providing emerging Cloud services. As a combination of servers, storage systems, networks, and software, SOP is an optimal platform for Cloud-era services merging virtualization and operation technologies. SOP design objectives and the architecture for achieving them are discussed with a focus on server-centric virtualization for automatically developing application systems from the viewpoint of servers and an evolution-oriented architecture for extending and upgrading the service platform without interrupting services. Virtual system packaging, dynamic resource management, and operations management are also presented as key elemental technologies for SOP design.

1. Introduction

Fujitsu is developing a service oriented platform (SOP) as an information and communications technology (ICT) infrastructure for data centers to support the provision of emerging Cloud services.[1)] Based on virtualization and all-IP technologies, SOP is a Cloud-service platform combining servers, storage systems, networks, and software (IP: Internet protocol). Two policies will be followed in the development of this platform.

1) Exploit the computing and network technologies that Fujitsu has developed especially for servers, storage, and software and utilize operations and management technologies and know-how built up in the process of providing hosting and other data-center services.
2) Use ICT infrastructure products currently being developed and provided for private Cloud services as core components. This approach will, first and foremost, make development more efficient, but for the customer, it will also facilitate the mixed use of public Clouds, virtual private Clouds, and private Clouds so that customers can use hybrid Clouds, depending on the type of work being performed.

In this paper, we first present the design objectives established for SOP development and discuss the architecture chosen to meet them. We then describe the elemental technologies that are the key to achieving SOP.

2. SOP design objectives

Fujitsu has established five design objectives in the development of SOP as an infrastructure for Cloud services.[2)]

1) Application-centric

In most cases, the development of a business system requires many tasks of various types in addition to application logic development. These might include the procurement, installation, configuration, and operation of servers, storage devices, and network equipment as well as the preparation of databases and user authentication

FUJITSU Sci. Tech. J., Vol. 46, No. 4, pp. 410–419 (October 2010)

mechanisms. The idea behind SOP is to provide non-application-logic parts of the development process as services to ease the burden on application-program developers.

2) Hands-on

One of concern among current and potential Cloud service users is the inability to observe the status of one's system or business operations within the Cloud. SOP addresses this concern by providing a dashboard that can visualize performance so that the user can monitor the response and throughput of a business system in real time. If the monitoring results should indicate that the business system is not performing as desired, the user can add more servers as needed by making use of the Cloud-computing characteristic of quick ICT resource procurement at optimal cost. This design objective of providing visualization and control aims to give the user a "hands-on" role in operations.

3) End-to-end security

Another user concern is security, but this is alleviated by equipping SOP with security technologies that cover the entire system from the network, operating system (OS), and storage systems to connections with outside services. For example, SOP is configured to provide not only public services over the Internet, but also virtual private Cloud services that connect to the user's corporate intranet. SOP also applies the network protection technologies and authentication technologies that Fujitsu has developed and accumulated through its network services. Furthermore, as multiple user systems share physical resources, Fujitsu has introduced a secure platform mechanism that clearly separates guests from each other and guests from the host through the use of virtual machines (VMs), which are the key infrastructure component of SOP.

4) Green

Reducing power consumption is an important issue in the provision of Cloud services considering that they are achieved by concentrating a large number of ICT resources in a data center. In SOP, this reduction will be achieved by reducing the power consumed by hardware itself as well as by raising the hardware utilization rate through total virtualization.

The central processing unit (CPU) utilization rate of a business server is generally no more than 20%. To improve upon this, SOP consolidates multiple VMs on a single physical server. It also raises the utilization rate of physical servers by increasing the consolidation factor through a VM live migration function that can move a VM on the fly whenever space on a physical server becomes free. The storage utilization rate is also being improved by using various techniques including deduplication and thin-provisioning in which actual disks are not allocated to unused areas.

5) Open

We can expect future enterprise ICT systems to be achieved by optimally combining a private Cloud with Cloud services provided by multiple providers. In such a scenario, interconnectivity between Clouds is important from the viewpoint of protecting a company's investments in applications. At present, standardization of an application programming interface (API) for Cloud platforms is progressing at a number of standardization bodies such as the Distributed Management Task Force and Open Grid Forum. As part of its SOP development efforts, Fujitsu is participating in and making proposals at such standardization forums with the intention of implementing standard API specifications and improving application portability and inter-Cloud connectivity. Through these activities, Fujitsu hopes to dispel user concerns about being locked into a specific Cloud.

3. SOP architecture

A basic Cloud service achieved by SOP is infrastructure as a service (IaaS), which lets a user deploy servers, storage devices, and network equipment from an SOP operated by a service provider (Fujitsu) inside a data center to build an application system in an on-demand, self-service manner. The user can then install on that system the middleware, specific applications, and packages needed to run a business system. These functions are shown in **Figure 1**.

IaaS requires the installation of a group of servers, a group of storages devices, and a group of network equipment inside a data center and the creation of a virtual resource pool. The operations management of these physical resources must also be automated and it must be possible to dynamically deploy the application systems of many users (i.e., dynamically allocate necessary resources from the virtual resource pool). Also required is a mechanism for presenting an operations management environment to both data-center operators and users and, since IaaS is a service, mechanisms for managing and billing users. The SOP structure chosen for meeting these requirements is outlined in **Figure 2**. SOP is configured with hardware resources as an underlying platform and management software above it for achieving IaaS. The hardware resources consist of large numbers of servers and storage devices uniformly connected by IP networks. The management software consists of operations- and service-management components centered on mechanisms for virtualizing hardware resources and interconnecting and deploying virtual resources according to system models specified by templates.

The basic concepts behind this SOP architecture are server-centric virtualization and evolution-oriented architecture.

3.1 Server-centric virtualization

To provide application-centric, on-demand, and self-service features, all servers, storage devices, and network equipment on SOP are virtualized and all resources visible to users are separated from the actual hardware (total virtualization). These resources can be used in a server-oriented self-service manner. Here, the user only has to be concerned with the number of servers required, the logical connections between them, and the amount of storage needed. There is no need for the user to worry about the network and storage settings that have to be made to configure the specified system. This is the server-centric virtualization concept.

Each physical server in SOP is equipped with Xen as a hypervisor (VM control program) above which guest VMs operate. To achieve

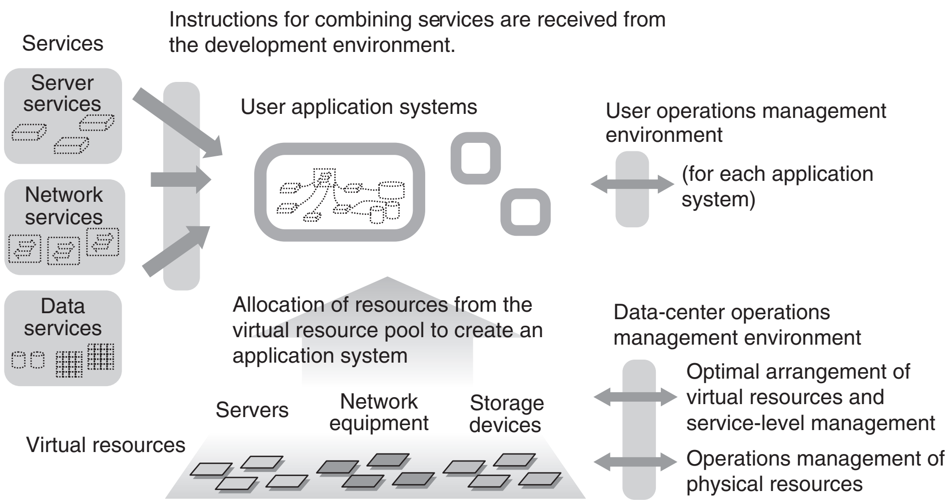

Figure 1
IaaS provided by SOP.

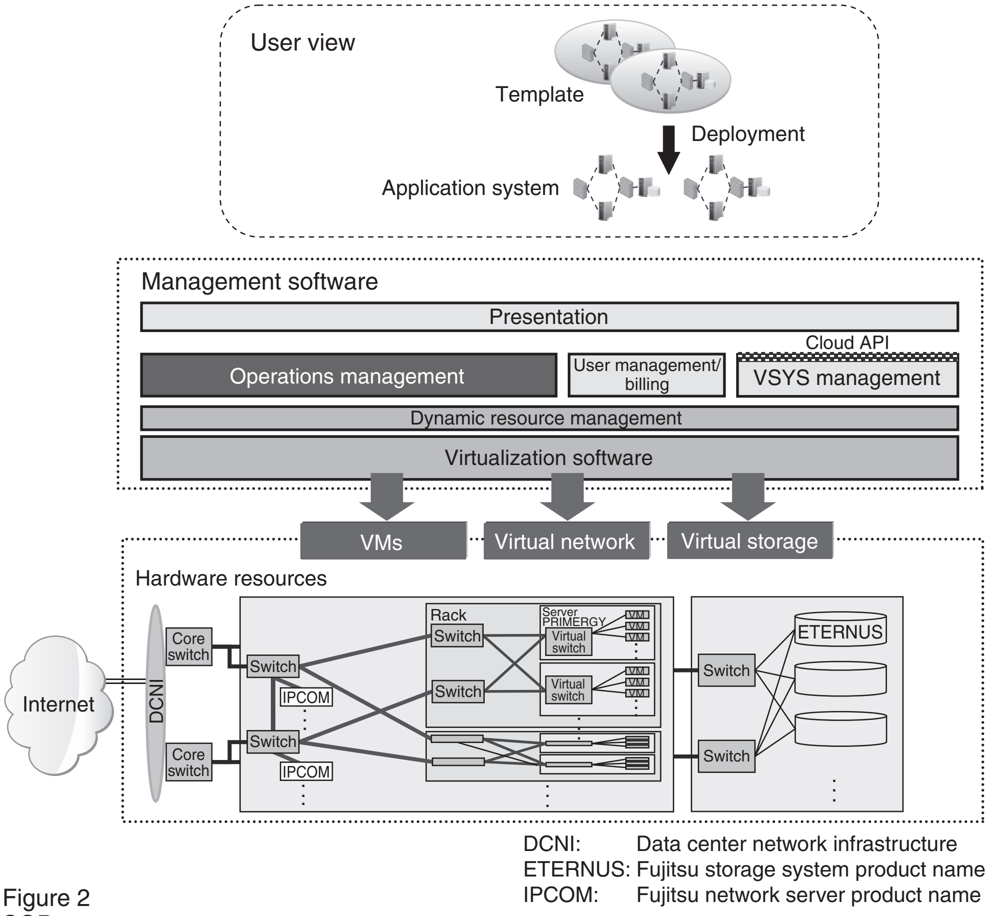

Figure 2
SOP structure.

server-centric virtualization, the control of virtual resources is integrated with virtualization software on the physical servers. Here, we note that while conventional virtualization software is applied to only physical server virtualization, the virtualization software in SOP also performs network and storage virtualization. That is to say, SOP does not provide separate virtualization functions for server, storage, and network resources. In this case, users must combine all these functions to achieve required application systems by themselves. SOP uniformly manages all server, storage device, and network equipment virtualization. For example, the storage virtualization is accomplished by loading a volume manager on top of the Xen management OS, and network migration is achieved by linking with server migration.

Network virtualization (**Figure 3**) provides the on-demand creation of a subnet configuration, which is commonly used to ensure security in an enterprise system. This function is achieved by network server technologies like firewalls and server load balancers that run as software on VMs (technologies accumulated via Fujitsu's IPCOM series of network servers) and management technologies that control the network dynamically in collaboration with the above technologies. In addition, the application of technologies like a tag-based virtual local area network (VLAN) enables the configuration of overlay networks that can allocate a private IP-address space to each service user. This approach enables a virtual private Cloud having a high affinity with the existing enterprise systems to be built quickly and flexibly.

Finally, any network that interconnects physical resources on SOP including storage

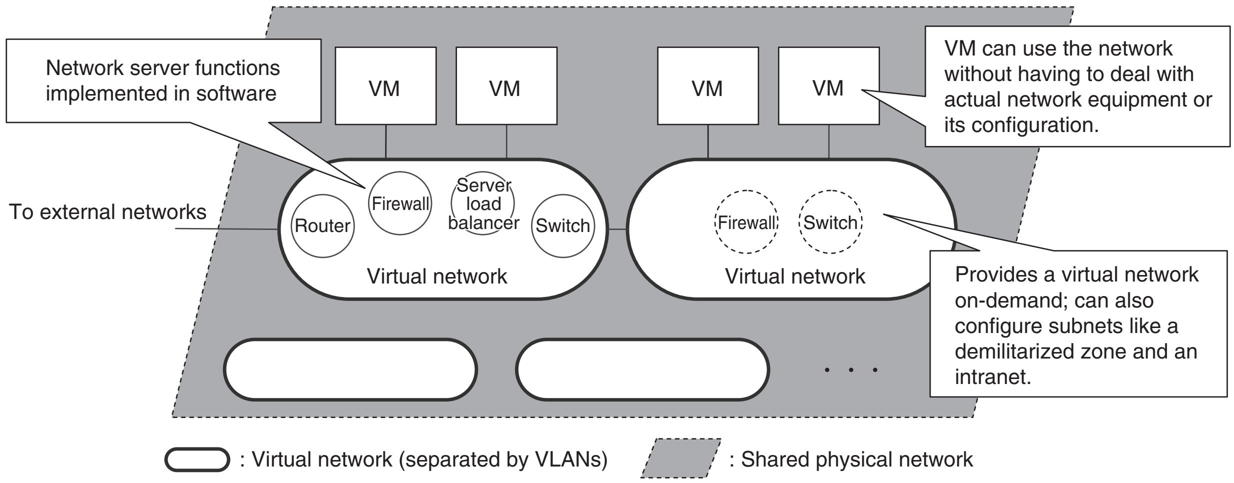

Figure 3
Network virtualization in server-centric virtualization.

devices is achieved by IP networking (Ethernet). This has the effect of reducing the number of physical network equipment types and cables, so it achieves significant cost reductions in network construction, configuration, and management.

3.2 Evolution-oriented architecture

A platform oriented to Cloud services must be able to provide a large number of users with uninterrupted services amidst updates, extensions, and upgrades, which can occur at any time. A mechanism enabling platform scalability and service-menu expansion while ensuring service continuity is called an evolution-oriented architecture. It features growth potential based on an island structure and a framework for accumulating operations know-how.

1) Island structure

Lowering the cost of operations by expanding a highly uniform configuration in a scalable manner is a basic characteristic of a Cloud computing platform. However, while this objective can be achieved in the short-term, the addition of new hardware and software and the ad hoc revision of existing hardware and software over several years could break the uniformity and drive up operating costs.

To prevent an entire system from becoming increasingly complicated over the long-term, SOP is configured in units of “islands” (each consisting of from 100 to 1000 physical servers). The hardware and software within any particular island are uniform (that is, of the same generation) and have a low degree of complexity. On the other hand, different islands may have different generations of hardware and/or software, which allows the deployment of new hardware and software (**Figure 4**).

An island can be thought of as a unit of metabolism since the addition of new resources or the abandoning of old resources is carried out in units of islands. For example, the revision of island elements (as in software patching) can be performed in a batch throughout islands of the same generation. This has the effect of keeping the diversity in element combinations to less than the number of island generations. SOP also features an internal interface for observing and controlling any island. This interface lets operation managers perform integrated management of a multi-generation group of islands.

2) Framework for accumulating operations know-how

The accumulation of operations know-how is a valuable asset for Cloud service providers.

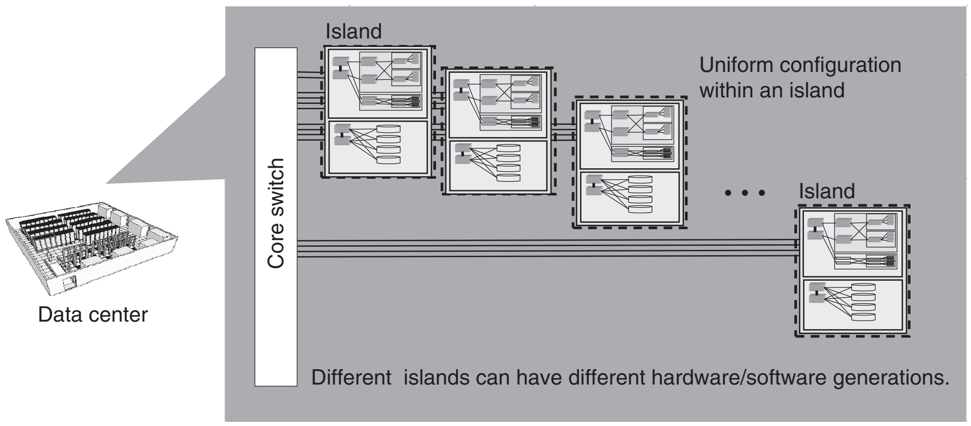

Figure 4
Island structure.

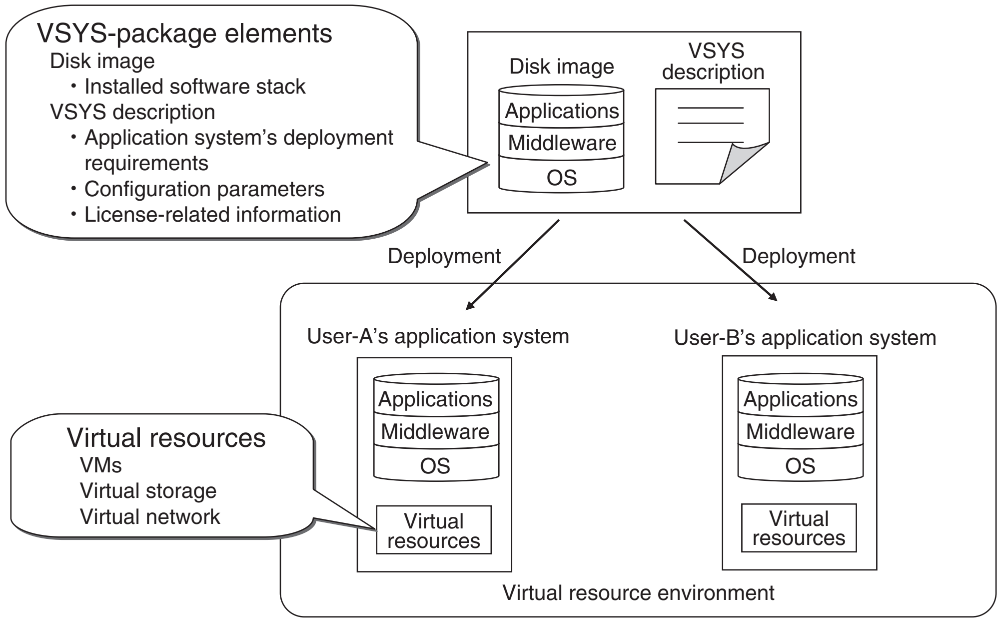

Figure 5
VSYS package.

The progressive automation of operations as know-how is accumulated can reduce the cost of providing these services.

To accumulate and exploit this know-how, SOP will feature a language for transforming experience into formalized knowledge and an engine for executing operation policies written in that language. Here, it is important that the island structure described above helps to narrow down the configuration of management targets and hence accelerate operations automation and know-how accumulation.

4. Virtual system packaging

With SOP, the provision of typical, commonly used application systems to users supports the application-centric design objective. Such an application system, which is called a VSYS package, is slated to become the unit of software distribution in the virtualization era (**Figure 5**). A VSYS package eliminates the need to install individual software programs on a virtual machine since an application system can be provided once the software stack has been checked for consistency and application tuning has been completed.

A VSYS package consists of an extensible markup language (XML) document (VSYS description file) and disk-image files. The VSYS description file describes platform requirements for deployment of the application system (more than one virtual machine can be described) as well as configuration parameters and license-related information. The disk-image file, which includes a pre-installed software stack, conforms to the Open Virtualization Format (OVF)[3] established by the Distributed Management Task Force with the aim of maintaining interconnectivity between different Clouds. The user (service provider) can deploy and launch multiple application systems from a single VSYS package.

Specifications that aim to make software deployment efficient and consistent on a system composed of multiple servers were developed and standardized even prior to OVF. These include ACS (Application Contents Service)[4] and CDDLM (Configuration Description, Deployment and Lifecycle Management)[5] established by the OGF (Open Grid Forum) and SDD (Solution Deployment Descriptor)[6] established by the OASIS (Organization for the Advancement of Structured Information Standards). The maturation of virtualization technologies in recent years and their widespread diffusion indicate that these technologies have now reached a stage of practical use.

Here, a whole application system built by the user (developer) is designed using a service builder. Although the user may build a new system entirely from scratch if so desired, using a VSYS package as a basis to work from makes the development process more efficient. In combination with server-centric virtualization, the service building process comes down to determining how servers are connected using a service builder. In this way, a complex system like a three-tier Web system can be easily built even without getting into details such as firewall settings, network configurations between application servers and databases, etc.

In addition, the use of an existing VSYS package means that operating procedures such as for backing up data and responding to failures of components as well as security policies related to the operation log maintenance and data encryption are all predefined. This feature can help reduce the amount of labor involved in system operation and monitoring.

5. Dynamic resource management

With SOP, the multiple virtual machines making up an application system are deployed automatically, and the connections between those virtual machines are made via a virtual network. The design map of a user's application system is described in a VSYS description file. When the user instructs the Cloud management system from a dashboard to deploy or withdraw an application system instance, SOP's dynamic resource management will automatically execute the necessary processing in accordance with that VSYS description. This scheme enables the user to deploy an application system in a self-service manner.

The dynamic resource management must satisfy the following requirements.

1) Immediacy: Must provide virtual resources whenever the user desires without complicated initial settings
2) Flexibility: Must deal with rapid changes in workload and business environment
3) Efficiency: Must automate resource rearrangement and reconfiguration in the event of hardware failures or upgrades

In addition to server virtualization management software called ServerView Resource Coordinator VE,[7] which has already been successfully used in the enterprise market, Fujitsu has developed dynamic resource management software called ServerView Resource Orchestrator.[8] Both software products are used as SOP management software while also adding functions targeted at large-scale Cloud services.

The following functions are being provided for dynamic resource management.

1) Virtual resource pool function

Manages virtual resources as shared resources and rapidly provides necessary resources in accordance with deployment requests from users

2) Resource orchestration function

Flexibly deploys application systems according to business characteristics and changes by maintaining and managing the relationships among servers, storage systems, and networks

3) Automatic resource optimization function

Deploys alternative virtual platforms and performs failover or migration processing automatically in response to the failures and planned upgrades of a server or other equipment

Various types of policies can be considered when arranging virtual machines on physical servers. In SOP, virtual machines are dynamically arranged on appropriate physical servers according to the policy specified by the service provider (data center manager). Two typical arrangement policies are the energy-saving policy and high-availability policy. The energy-saving policy concentrates VMs on as few physical servers as possible to free up a certain number of physical servers whose power can then be turned off. This has the effect of reducing the total amount of power consumed by the data center. On the other hand, the high-availability policy distributes across as many physical servers as possible to minimize the number of VMs affected by a physical server failure.

6. Operations management environments

SOP provides two operations management environments.

1) Operations management environment for users

This environment provides users with a dashboard as a collection of monitoring and control functions for self-deployed application systems. Control functions enable the user to boot, terminate, and back up an application system and individual servers composing the system. Monitoring functions enable the user to observe the state of resource usage including CPU loads plus the response of each virtual machine and the responses of the services being used.

2) Operations management environment for service operators (data center operators)

This environment provides functions for monitoring and controlling physical and virtual resource pools. These functions are executed through the use of a cockpit console and display oriented to platform managers. The control functions support the management of physical resource pools and the optimal arrangement of virtual resources on physical resources. The monitoring functions use visualization techniques to let platform managers observe the operational status of physical resources and the usage and response of virtual resources on the cockpit display.

In this environment, the same information provided on the dashboard of each user application system can also be accessed via the same mechanism so that service-level information sharing can be achieved between users (service providers) and data center managers.

In SOP, in which the operation of about 10 000 physical servers and 100 000 VMs must be managed, the ability to implement a large amount of physical equipment (servers, network switches, etc.) all together and make monitoring and maintenance more efficient is important. Efficient and straightforward operations management technology is needed not only to reduce labor expenses and other operational costs, but also to maintain a uniform level of maintenance quality across a large amount of equipment. Fujitsu has developed platform construction technology that makes the operations management of a huge amount of physical equipment significantly more efficient.

The implementation of equipment in a data center has traditionally required various types of manually performed tasks such as mounting equipment on racks, wiring it, and visually checking connections, configuring the server BIOS (basic input/output system), and checking and registering media access control addresses (MAC addresses). These tasks have put a large workload on on-site engineers and have made equipment implementation inefficient. For SOP, Fujitsu is developing functions for automating the execution of whole series of on-site tasks at the time of implementation while also developing a system for grouping physical equipment into islands in units of racks. The aim is to minimize human error and reduce costs when equipment is added to a data center.

7. Conclusion

Cloud computing represents an expanded and converged aspect of technologies like grid computing, utility computing, and the service oriented architecture. This concept represents a rethinking of those technologies in terms of services from the user's viewpoint, and, as such, it has come to fascinate many people. The essence of Cloud computing is "services provided over the network," which means that no borders exist when the Cloud is being used. For this reason, a Cloud platform that provides extensive interconnectivity using standard technologies must be achieved through open innovation.

As with electric power, ICT consolidation can result in significant cost savings.[9] Thus, the construction of Cloud data centers on a scale from several thousand to several tens of thousands of servers and the automation of operations should make possible the provision of ICT services at much lower prices than conventional hosting services and the exploration of new ICT applications including ones in fields to which ICT could not be applied in the past. Thus, raising the infrastructure scale and automation level can been seen as technical requirements for a Cloud service platform.

Looking forward, Fujitsu plans to intensify its efforts to develop Cloud-platform technologies in this direction.

References

1) C. Sagawa et al.: Cloud Computing Based on Service Oriented Platform. *Fujitsu Sci. Tech. J.*, Vol. 45, No. 3, pp. 283–289 (2009).
2) R. Take: Virtualized IT Platform for Cloud Computing Era. (in Japanese), *FUJITSU*, Vol. 60, No. 3, pp. 266–273 (2009).
3) Distributed Management Task Force: Open Virtualization Format Specification Version 1.1.0. DSP0243, 2010.
4) Open Grid Forum: Application Contents Service Specification 1.0. GFD-R-P.073, 2006.
5) Open Grid Forum: CDDLM Configuration Description Language Specification Version 1.0. GFD-R-P.085, 2006.
6) Organization for the Advancement of Structured Information Standards: Solution Deployment Descriptor Specification 1.0. 2008.
7) Fujitsu: ServerView Resource Coordinator VE. (in Japanese). *http://software.fujitsu.com/jp/rcve/*
8) Fujitsu: ServerView Resource Orchestrator. (in Japanese). *http://software.fujitsu.com/jp/ror/*
9) N. Carr: The Big Switch: Rewiring the World, From Edison to Google. W. W. Norton & Company, Inc., 2008.

Hiroshi Yoshida
Fujitsu Ltd.
Mr. Yoshida is a senior vice president of the Service Oriented Platform Strategy and Development Office and its principal architect. He is responsible for the development of the Service Oriented Platform.

Hiro Kishimoto
Fujitsu Ltd. and Fujitsu Laboratories Ltd.
Dr. Kishimoto is leading Cloud computing research in Fujitsu and SOP global service delivery. He is a visiting professor at the National Institute of Informatics, Vice President of the Grid Consortium in Japan, and a member of the board of directors of DMTF. He received the IEEE Gordon Bell Award and the inaugural GGF/OGF Leadership Award.

Riichiro Take
Fujitsu Ltd. and Fujitsu Laboratories Ltd.
Mr. Take has been working on the SOP project as the principal architect since 2008.

Yoshinobu Hibi
Fujitsu Ltd.
Mr. Hibi has been working on the SOP project since 2008.