Infrastructure Technology for Cloud Services

● Toru Kino

The application of a new information and communications technology (ICT) model called cloud computing is finally moving into a full-scale deployment phase in the corporate world as a new business infrastructure and in diverse fields such as primary industry, medical care, and nursing to improve productivity. To expand use of the "cloud," that is, to dispel worries about using cloud computing, Fujitsu came to reevaluate, redesign, and redevelop its cloud infrastructure technology from the architecture level to achieve a higher dimension in both convenience and reliability. These efforts led to the launching in October 2010 of a commercialized cloud service called "On-demand Virtual System Service" (now known as the "Fujitsu Global Cloud Platform 'FGCP/S5'" service in Japan and as the "Global Cloud Platform [GCP]" service globally). This "Made in Fujitsu" infrastructure technology is being expanded to data centers in six countries including Japan to promote the provision of cloud services on a global scale. This paper introduces Fujitsu's cloud infrastructure technology and the global expansion of Fujitsu cloud services.

1. Introduction

As reflected by the papers in this special issue, the application of cloud computing (the "cloud") as a new information and communications technology (ICT) model is finally entering a full-scale deployment phase. For example, companies whose core business consists of *monozukuri* (innovative manufacturing) are using the cloud to provide network services directly to end users, and primary industries and fields like medical care and nursing that have traditionally relied on experience and human resources are using the cloud to improve productivity in the face of an aging society. There are many examples of this new way of using ICT, but what supports them all is the unique economic characteristic of the cloud: instead of making an up-front investment in facilities, customers use only as many computer resources as they need when they need them and pay for only what they use.

With the aim of expanding the application of cloud-based ICT to more areas in society, Fujitsu began to reevaluate its cloud infrastructure technology for data centers from the architecture level in 2009 and proceeded to redesign and redevelop this technology. In this process, prototype systems were developed to gain operational experience through actual in-house use and to obtain feedback from the departments using those systems. The result was a gradual maturation of cloud infrastructure technology, leading to the commercialization of a full-scale infrastructure as a service (IaaS) type of public service in October 2010 under the name "On-demand Virtual System Service." In June 2011, it was renamed "Fujitsu Global Cloud Platform 'FGCP/S5'" service. Moreover, Fujitsu now offers a "Global Cloud Platform (GCP)" service in Australia, Singapore, the U.K., the U.S., and Germany as well. Many of the application examples referred to above are being

 FUJITSU Sci. Tech. J., Vol. 47, No. 4, pp. 434–442 (October 2011)

constructed on this IaaS platform in a speedier, easier, more scalable, and more cost efficient manner than on existing ICT infrastructures.

This paper introduces Fujitsu's infrastructure technology for cloud services and the global expansion of these services.

2. Requirements of cloud infrastructure

Services provided to corporate customers in the form of a cloud include software as a service (SaaS), which provides users with business application functions, platform as a service (PaaS), which provides users with application development and execution environments, and IaaS, which provides users with basic ICT resources such as servers and storage as a service. Fujitsu has been involved in the development of cloud services from the very beginning of this field, devoting particular efforts to SaaS for all sorts of business tasks and operations and to IaaS for on-demand hosting of physical resources.

Through these experiences, Fujitsu came to realize that cloud infrastructure technology for data centers had to be reevaluated and optimized from the architecture level to enable customers to enjoy a broader range of cloud benefits in an ongoing manner. To this end, Fujitsu performed a company-wide consolidation and integration of all essential technologies related to servers, storage, networks, software, and service management and redeveloped and reconstructed the cloud infrastructure. The basic concept behind these efforts was to achieve both convenience and reliability on a higher dimension, enabling anyone to use cloud services in a safe and secure manner.

2.1 Convenience

Cloud services have expanded from their beginnings in the United States as free or low-cost convenient services such as Web search and E-mail provided by vendors to the general consumer. Since then, an ever increasing number of people have unconsciously stepped inside the world of cloud computing as smartphones and multi-function mobile devices with network connections have rapidly penetrated society and users of social network services (SNS) have multiplied. These "public services" from American vendors have two key features:

- Support of many users attracted by a high level of convenience
- Low costs achieved by high-volume activity and uniformity in services

These features differ greatly from the conventional ICT model in which each company constructs and operates its own specialized system.

Against this background, Fujitsu has been making technical innovations in intensive automation and power-savings techniques specifically for the high-volume operation of uniform services at mega data centers that accommodate several hundred thousand servers.

2.2 Reliability

At the same time, public services such as these, while providing a high level of convenience, feature a shared environment used by many people. As a consequence, companies tend to have the following concerns with regard to reliability when considering the use of public cloud services:

- How can users be adequately separated when multiple user systems share the same physical resources?
- How could a company with data stored on its own intranet feel comfortable placing that data on an Internet system, and wouldn't it be difficult to link the cloud with the company's intranet system?
- Since the cloud appears as a black box to users, how could a company deal with any problems that occur?

These concerns were justified during the expansion phase of the Internet. In the 1990s, the Internet diffused rapidly through society as a means of providing and exchanging information among general consumers. From the viewpoint of corporate use, however, a shared network raised concerns about security and response time, and, as a result, many companies decided to use dedicated networks and robust but complex security technologies. In time, though, these concerns began to disappear as Internet security technologies progressed and successful use of the Internet in the corporate world increased: electronic commerce, Internet banking, and other types of business began to expand rapidly on the Internet.

3. Service oriented platform

Highly convenient services are bound to find widespread use if concerns about their safety can be eliminated, that is, if users can feel comfortable in using them. With the aim of achieving both a high level of convenience typical of public services and a high level of reliability that dispels concerns about safety, Fujitsu has revamped its cloud infrastructure from the architecture level, resulting in the development of a service oriented platform (SOP).[1),2)] An SOP is an optimal platform for services in the cloud era, integrating the server, storage, network, and software technologies that Fujitsu has accumulated over many years as a domestic computer vendor in Japan and merging virtualization and operation technologies. Furthermore, by using technologies common to Fujitsu software products, an SOP serves to improve interconnectivity when configuring a hybrid cloud.

To achieve cloud services, it is necessary to deploy groups of servers, storage devices, and network devices inside a data center and to construct a pool of virtual resources. It is also necessary to automate the operations management of these physical resources and to be able to dynamically deploy the individual virtual systems of many users (i.e., to allocate needed resources from the virtual resource pool). Also required is a mechanism that presents an operations management environment for the virtual systems to cloud service users and one for the entire cloud infrastructure to data center operators.

The SOP architecture is outlined in **Figure 1**. An SOP consists of hardware resources as infrastructure and of management software for achieving services on that infrastructure. The hardware resources consist of a large number of servers and storage devices, all interconnected by an IP network. The management software consists of operations-management and service-management components centered about a mechanism for virtualizing hardware resources and interconnecting and deploying virtual resources in accordance with system models specified by templates. Underlying this SOP architecture are two key concepts—server-centric virtualization (ScV) and evolution-oriented architecture (EoA)—as described below.

3.1 Server-centric virtualization

Fujitsu's SOP virtualizes all resources, that is, servers, storage, and networks. This approach separates the ICT resources provided to the user from the actual hardware and enables resources to be reconfigured and reused from a server-centric point of view. As a result, the user need only be concerned with the number of logical servers needed, the network connections involved, and storage capacity. In other words, the user does not have to worry about the physical network settings, storage settings, etc. that have to be made to achieve the above virtual system. This is the ScV concept.

While the objective of standard virtualization software is to virtualize only physical servers, SOP virtualization software also has the role of virtualizing storage and networks. That is to say, an SOP does not simply combine separate virtualization functions for servers, storage, and networks; rather, it uniformly manages total

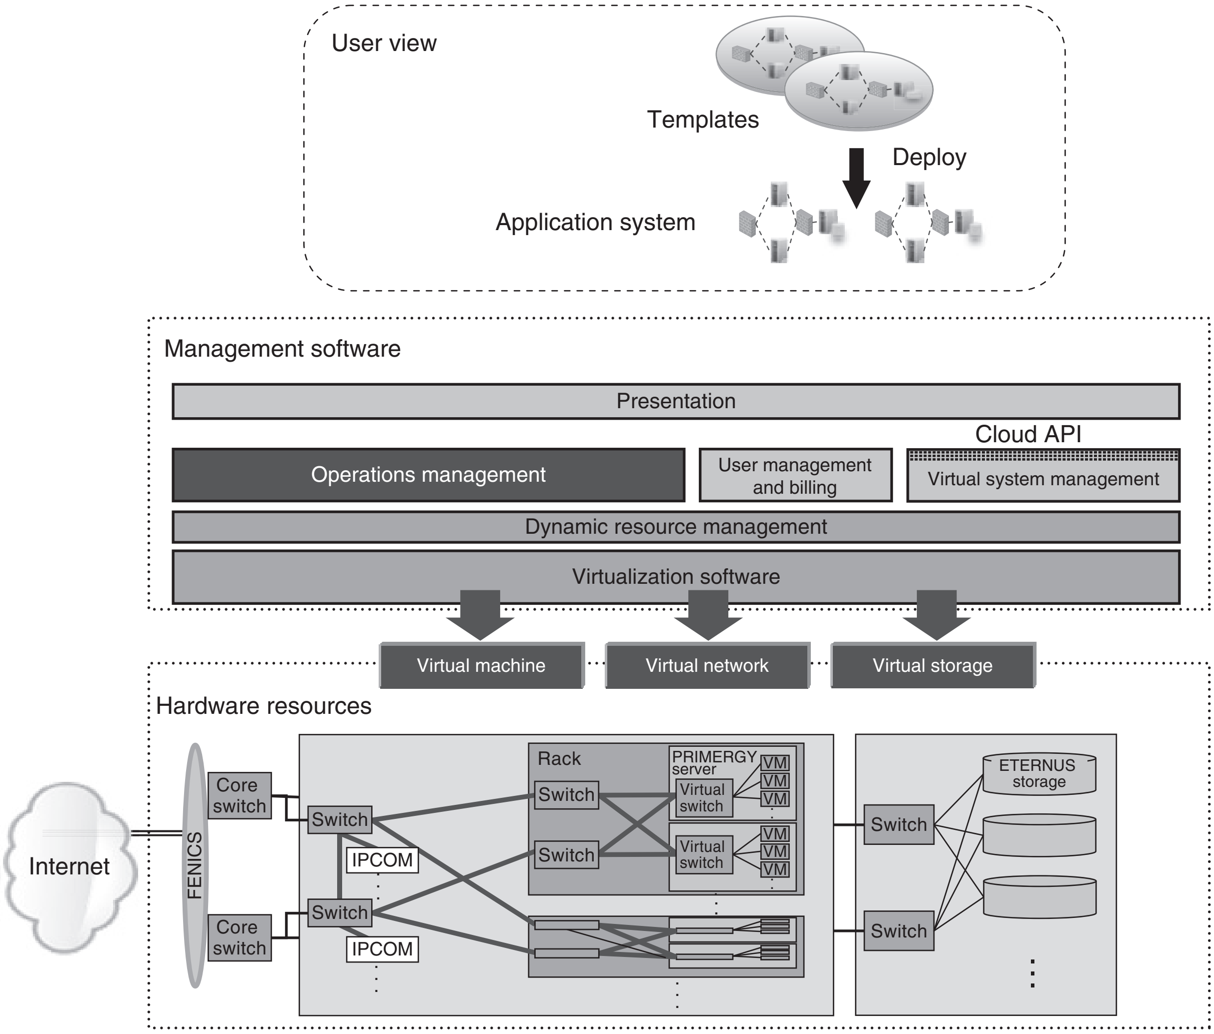

Figure 1
Structure of SOP.

virtualization including that of servers, storage, and networks.

Furthermore, for network virtualization, an SOP provides a function for creating a subnet-partitioned configuration in an on-demand manner. This function is achieved by taking technologies like firewalls and server load balancers developed as part of Fujitsu's IPCOM series of network servers and combining them with management technologies to dynamically control the network.

3.2 Evolution-oriented architecture

A platform designed for services that continue to expand and develop as in the case of cloud services must still be able to provide a large number of users with uninterrupted service. A mechanism that enables a platform to evolve over time to ensure scalability and expansion of the service menu while providing service continuity is EoA. This type of architecture is achieved by using an island structure that promotes growth and a framework for accumulating operations know-how.

1) Island structure

Lowering operational costs by expanding a highly uniform physical configuration in a scalable manner is basic to a platform oriented to cloud computing. Over time, however, the addition of new hardware and software and the revision of existing hardware and software have the potential of breaking down this uniformity and increasing the complexity of operations.

To prevent a system from becoming overly complicated over the long term, Fujitsu configures an SOP as a set of unit structures called "islands" (each consisting of from 100 to 1000 physical servers). The hardware and software within any one island is made to be of the same generation (to be as uniform as possible), but between islands, the generation of hardware/software may differ. This approach enables new generations of hardware and software with higher cost performance to be introduced into the system. An island in an SOP can be thought of as a unit of metabolism (a cell) because the addition of new resources and the discarding of old resources are carried out in units of islands.

2) Framework for accumulating operations know-how

The accumulation of operations know-how is of great value to providers of cloud services as it can facilitate the automation of operations, thereby improving service quality and operation efficiency. To accumulate and exploit operations know-how, an SOP features a language for transforming experience into formalized knowledge and an engine for executing operation policies written in that language. An important element of this process is that the uniform island structure described above narrows down the configuration pattern targeted for management, which helps to accelerate automation and the accumulation of more know-how.

4. FGCP/S5 service

Using an SOP as an engine, Fujitsu launched an "On-demand Virtual System Service" (now called "FGCP/S5" service) in October 2010 as an IaaS-type of public cloud service for allocating and providing virtual system environments for the exclusive use of customers on the large-scale resources installed in Fujitsu data centers.

FGCP/S5 service is similar to no other cloud service in the world. In addition to being highly convenient like a public service, it is also highly reliable, enabling a highly secure three-tier Web system to be easily constructed.[3] The following summarizes the features of FGCP/S5 service.

4.1 Convenience

1) Design-free

When using IaaS, all design and operation tasks related to physical resources such as the positioning of device racks, the connection of LAN cables, and the design of redundant configurations are handled by the cloud provider. This frees the user from having to perform design work related to physical devices and wiring. On the other hand, the user must perform system-configuration design and logical-connections design using virtual ICT resources such as virtual machines (VMs). This design work involves commands and operations that differ from one IaaS provider to another and consequently requires cloud-oriented system configuration skills that differ from those used in conventional physical-type ICT infrastructure design.

Fujitsu has much experience in system integration and outsourcing, and it has accumulated much know-how by configuring servers and networks and interconnecting them in a variety of in-house systems. This know-how has been used to extract a number of prototypical configurations differing in business scale, business operations, and security level. The FGCP/S5 service provides system templates based on these configurations. Thus, to create a virtual platform for a business system, the user only has to select a system template and input required information because the interconnections and performance levels of the virtual ICT resources will be automatically set. The user's efforts can therefore be focused on the development of business applications. A system created in this way can be saved as a user-defined system template, and a system similar to a desired one can be copied. This approach has been found to be effective in creating a development environment or test bed when upgrading an application.

2) On-demand

Fujitsu had previously provided an on-demand hosting service to enable customers to use ICT resources in accordance with their needs. This service progressed to enable customers to use these resources on an “as-needed, when-needed basis.” In the past, several days would be needed to complete the procedures required for adding ICT resources to a system, and the contract period for using these resources was normally one year. Today, FGCP/S5 service can deploy the resources that the customer needs in about 30 minutes and allows these resources to be used in one-hour units. The following dynamic-resource-management functions are used to achieve this agility in resource allocation.

- Virtual resource pool

This function manages virtual resources as shared resources and rapidly provides required resources in response to deployment requests from users.

- Resource orchestration

This function maintains and manages the relationships among servers, storage, and networks so that application systems can be flexibly deployed in accordance with the nature of the customer’s business, changes to that business, etc.

- Automatic resource optimization

This function prepares alternative equipment for responding to a failure in a server or other physical device and provides for automatic failover.

3) Self-service

A self-service format and automation are essential to achieving a full-function on-demand service. In past services, the user would have to specify what ICT resources were needed and from when they would be needed by submitting an application by E-mail. This information-transfer process would generate much manual work on the part of both user and service provider by the time that requested resources were deployed. Fujitsu recognized that minimizing the amount of manual work in the provision and use of ICT resources could enable users to start using the resources that they need in a relatively short time.

The service portal in **Figure 2** acts as a self-service window. A user can use this service portal, for example, to apply the system templates described above and to specify the start and end times of resource usage. All subsequent processing is automatically controlled, enabling ICT resources to be provided with no time lag.

This service portal provides an icon-filled, graphical user interface designed with the user in mind. It is also possible to use a service application programming interface (API) that enables the same operations performed by the user on the service portal to be specified from a program. An API can therefore be used to cope with increases and decreases in processing load by automating the booting and termination of virtual machines in accordance with a predetermined schedule or the actual number of server accesses or to carry out a periodic

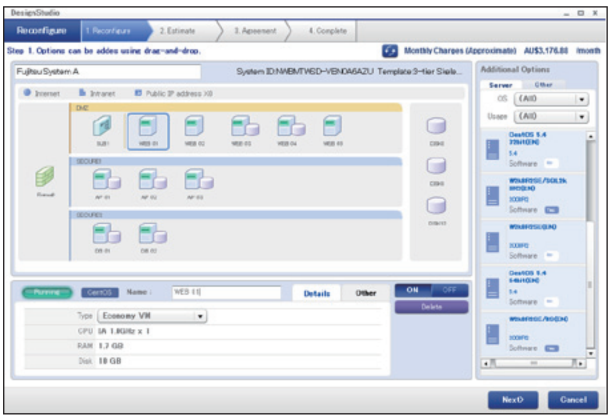

Figure 2
Service portal.

backup. Fujitsu is participating in service-API standardization activities toward an open cloud. Much work is currently being performed on the standardization of APIs for cloud platforms at a number of standardization bodies including the Distributed Management Task Force (DMTF) and Open Grid Forum (OGF). In addition to participating in these activities and making proposals, Fujitsu is also implementing such standard specifications to improve application portability and interconnectivity between clouds.

4.2 Reliability

1) Secure three-tier virtual system

On an SOP, servers, storage, and networks can be integrated and virtualized in units of systems. This makes it possible to separate virtual servers within a demilitarized zone[note)] having global IP addresses from virtual servers within a secure zone having private IP addresses through the use of virtual firewalls.

This sort of configuration, while being commonplace in physical environments, has not previously been available in services provided in a cloud virtual environment. In ordinary cloud services, all virtual servers have global IP addresses and are therefore exposed to the Internet. Some type of measure for blocking access from the Internet such as closing off ports by virtual-server settings is therefore necessary, but this in itself cannot completely achieve a private zone.

A secure three-tier virtual system enables even a full-scale business system to enjoy the key feature of immediate use of public clouds.

2) Virtual private systems and hybrid configurations

Corporate customers express a great desire to enjoy the advantages of a public cloud even if their business systems are difficult to place on the Internet.

With this in mind, Fujitsu provides an environment that enables the provision of virtual private cloud services that connect to the corporate intranet by linking the FENICS corporate network service with cloud services. By combining an integrated authentication service with an Internet virtual private network (VPN) as well as with an IP-VPN, wide-area Ethernet, and even a mobile network, secure access to business systems on public clouds can be achieved.

Additionally, a customer's business system may consist of many and varied interconnected subsystems. There may be subsystems that can enjoy the great advantage of using a standardized cloud environment as well as subsystems having separate configurations and operating conditions. Thus, on the whole, not all types of business conducted by a customer can be constructed using only a cloud.

In light of the above, there is a need for a hybrid environment in which systems on the cloud and systems in dedicated environments can coexist and connect to each other. In response to this need, Fujitsu's cloud services provide as standard a service for connecting to dedicated customer hosting zones within the same data center. In this hybrid environment, optimal integration and service management of the customer's business systems are made possible by incorporating data center services that Fujitsu has accumulated over many years.

3) "Made in FUJITSU" global cloud

Fujitsu developed its SOP by consolidating a wide variety of its technologies, and no third-party products are used in core technologies like server virtualization, networking, and integration. This policy enables rapid and flexible problem solving and security upgrading, which helps to put customers at ease about using the cloud.

Taking advantage of this "Made in FUJITSU" cloud, Fujitsu is confidently moving forward on expanding SOP-based cloud services on a worldwide basis. These services were

note) A zone isolated from both internal and external networks by firewalls.

launched in June 2011 under the name "Global Cloud Platform (GCP)" starting with five countries (Australia, Singapore, the United Kingdom, the United States, and Germany). The plan is to install uniformly configured SOP islands in the cloud data centers of each country so that services providing the same functions can be accessed from anywhere.

All of these SOP islands will be remotely monitored in a uniform manner from a consolidated control center, which will be staffed by full-time operators having specialized knowledge of cloud environments. The control center will also maintain close ties with the SOP development department to make operations even more efficient and accelerate problem solving.

5. Toward global cloud services

The use of FGCP/S5 service has been expanding rapidly, exceeding all expectations, and efforts to expand this service to global data centers have been intensifying. Against this background, the following issues must be addressed to provide cloud services on a global scale.

5.1 Capacity management

An SOP has an architecture that enables expansion in units of islands. Within an island, servers, storage, and network devices are configured on the basis of a fixed capacity ratio. It has become clear through actual use, however, that, as the number of users increase, storage consumption comes to exceed original expectations, creating an imbalance. Thus, in actual island design, the problem arises as to how to optimally balance capacity between the number of users and the number of VMs, amount of storage, and network bandwidth. Fujitsu intends to tackle this problem by applying the valuable experience it is gaining in operating this service.

5.2 Change and configuration management on a global scale

Fujitsu control centers have been performing service management based on Information Technology Infrastructure Library (ITIL) v3, but to accommodate large-scale sharing and a distributed environment, it has become clear that ITIL processes and tools will have to be optimized for the cloud.

For example, if SOP islands scattered around the world are to be remotely managed in a uniform manner, it would be preferable that device configuration, rack configuration, software-version control, etc. be the same among all islands. In this regard, it is relatively easy to match-up configurations in the initial phase of data center construction, but upgrading and expansion at these data centers will occur at different times in accordance with local business conditions. As a result, hardware and software in a variety of versions will come to coexist. To therefore achieve change and configuration management as well as security auditing in such an environment, a management process must be established on the basis of a virtual organization made up of departments located throughout the world.

5.3 Global business infrastructure

In addition to using an expanded SOP as a common platform for providing uniform cloud services on a global basis, it is equally important to provide local services in tune with the market needs of that region.

These services will include, for example, multilingual support for portals and multicurrency support for billing, as well as multi-site coordinated billing and ID management for global enterprises that use services from multiple data centers. In light of these needs, Fujitsu intends to refine and optimize the SOP architecture to achieve more flexibility in the way it handles the global-sharing section and localized section of each data center.

6. Conclusion

Cloud computing represents the ongoing fusion of a variety of technologies. These technologies have come to be supported by many users due to providing them in the form of services that are secure and easy to use.[4] The fees for using these services can be optimized by enabling cloud service users to use ICT resources efficiently on an as-needed basis. Since no large initial investments are required, the use of cloud services is beginning to spread even among business fields and industries that had not previously been utilizing ICT.

As the cloud has no national borders, the need is felt for a cloud infrastructure that achieves broad interconnectivity using standard technologies through open innovation. In tune with this trend, Fujitsu is committed to developing increasingly robust cloud infrastructure technology.

References

1) H. Yoshida: Evolution of Virtualization Technology and Development in Fujitsu. (in Japanese), *FUJITSU*, Vol. 60, No. 3, pp. 216–220 (2009).
2) H. Yoshida et al.: Service Oriented Platform. *Fujitsu Sci. Tech. J.*, Vol. 46, No. 4, pp. 410–419 (2010).
3) Y. Taniuchi: On-Demand Virtual System Service. *Fujitsu Sci. Tech. J.*, Vol. 46, No. 4, pp. 420–426 (2010).
4) N. Carr: The Big Switch: Rewiring the World, From Edison to Google. W. W. Norton & Company, Inc., 2008.

Toru Kino
Fujitsu Ltd.
Mr. Kino is engaged in the planning, development, and delivery of cloud services.