FUJITSU Cloud Service K5 IaaS for Growth and Innovation of Digital Business

● Yuji Wada ● Satoshi Iyoda ● Hiroshi Takamure ● Masashi Kameta
● Yasutaka Taniuchi

FUJITSU Cloud Service K5 IaaS (Infrastructure as a Service) is a new public cloud platform that supports cloud-native systems leveraging cloud-unique functions to dynamically allocate resources. It also supports promising new systems such as Systems of Record (SoR), Systems of Engagement (SoE), and the Internet of Things (IoT). K5 IaaS can respond quickly to user needs due to the use of open technologies and the use of OpenStack software for platform implementation. Additionally, the platform itself and services on the platform can be continuously enhanced through ongoing collaboration with the OpenStack community and its ecosystem. The K5 IaaS architecture is one common to private clouds and on-premise systems, enabling implementation of hybrid systems that seamlessly link existing on-premise business systems with virtual systems on K5 IaaS. This paper describes the functions of K5 IaaS and presents typical use cases.

1. Introduction

The cloud-first concept in which a company first creates a business system in the cloud and then expands it to its own premises as the need arises has found widespread acceptance. Indeed, an era has arrived in which diverse enterprises are using cloud services for all sorts of systems. In addition, new usage formats are appearing thanks to an environment that makes it easy and inexpensive to use many virtual resources. These include Systems of Record (SoR) for accumulating and analyzing business data and Systems of Engagement (SoE) / Internet of Things (IoT) for creating new value by processing massive volumes of data generated by people and things.

Fujitsu originally came to provide FUJITSU Cloud Service S5[1)] (hereafter, S5), a public cloud service that enables a conventional-type system typified by a Web 3-tier service (a Web/applications/database server configuration) to be easily constructed by the users themselves in a cloud using a graphical user interface (GUI). However, as usage formats such as SoR and SoE/IoT began to spread, together with cloud-native systems that use cloud-unique functions capable of dynamically allocating resources, demand arose for a public cloud that placed greater emphasis on scalability. Fujitsu met this demand on the basis the experience gained with S5 along with the knowledge and know-how of its system engineers (SEs) by providing FUJITSU Cloud Service K5[2)] as a common platform supporting both conventional and cloud-native usage forms.

This paper provides an overview of the Infrastructure as a Service (IaaS) provided by the K5 (hereafter, K5 IaaS) and describes its key features and typical use cases.

2. K5 IaaS overview

Many companies are beginning to adopt a cloud-first approach as the use of cloud services has become a viable option not only for implementing new systems but also for upgrading existing ones. For new systems, the reason often given for choosing cloud services is a desire for high scalability and availability through the use of a cloud-native system configuration. Meanwhile, many companies state that they would like to directly migrate existing systems to the cloud without changing their current configuration.

K5 IaaS provides a variety of functions for dealing with these needs. For example, it enables the construction of a conventional network configuration in the cloud by combining virtual servers, virtual routers, etc.

as needed. At the same time, it enables the development of systems in accordance with design patterns on a cloud-native system ("cloud design patterns").

A cloud-native system is attractive for several reasons.

1) High availability and load balancing are achieved by simultaneously running multiple virtual servers having the same functions.
2) Good balance of high performance and low cost is maintained by increasing or decreasing the number of virtual servers depending on the loads for user access and data processing.
3) Business continuity and data robustness are improved by deploying systems and data in multiple environments.

Fujitsu adopts open technologies in its public cloud services and applies OpenStack[3)] as the basic element for implementing platforms. Furthermore, by using application programming interfaces (APIs) compliant with open standards and interacting with the software community to form an industry-wide ecosystem, Fujitsu strives to improve the speed and sustainability of functional extensions. Fujitsu is contributing to the community by participating in the OpenStack Foundation as a Gold Member.

In the case of cloud services, it must be possible to perform continuous maintenance in a service environment so that new functions can be added and revisions and fixes can be applied promptly. K5 IaaS provides a mechanism for executing maintenance operations without having to shut down virtual servers or API access. This benefits both users and providers of cloud services: continuous operation of business systems and preservation of convenience for the former and ongoing maintenance and improvement of cloud services for the latter.

The increasing sophistication of cyber attacks in recent years has made advanced defensive systems and know-how a necessity even for in-house data centers. From the viewpoint of security, the advantages of cloud services that are monitored 24/7 by security specialists have been reaffirmed. K5 IaaS, which is expanding services on a global basis, provides users with a level of security equivalent to or better than that of on-premise systems through security designs and real-time monitoring in line with global standards.

3. K5 IaaS functions

In this section the most basic "distribution over multiple zone" functions for constructing a cloud-native system are described.

In the construction of a virtual system, K5 IaaS enables resources to be distributed and deployed in multiple zones with facilities in separate locations. This improves availability and fault tolerance against facility problems at a particular data center. Resilience to bursty loads can be improved by using a scale-out configuration spanning multiple zones. The following functions are provided to achieve a zone-spanning configuration.

3.1 Load balancer

The load balancer function treats a group of servers as a single server. On receiving a request, the load balancer selects a server among the group and forwards the request to it. K5 IaaS provides a load balancer as a service spanning multiple zones. This enables the user to specify the load-distribution target as virtual servers deployed across multiple zones.

3.2 Auto-scaling

The auto-scaling function automatically creates or deletes virtual servers depending on a user-defined threshold (for example, the load on virtual servers), as shown in **Figure 1**. This function is often used in combination with the load balancer function. For example, when virtual servers are created by auto-scaling (scale-out), the added servers are automatically added to the load-distribution target in coordination with the load balancer function. Conversely, when virtual servers are deleted by auto-scaling (scale-in), the deleted servers are automatically removed from the load-distribution target.

K5 IaaS provides an auto-scaling function as a service spanning multiple zones. The user can specify multiple zones as the target for increasing or decreasing the number of virtual servers. In this way, the user can have scale-out or scale-in of virtual servers automatically performed across multiple zones depending on load conditions defined by the user.

3.3 Security group

The security group function serves as a virtual firewall for virtual servers. As multiple security settings

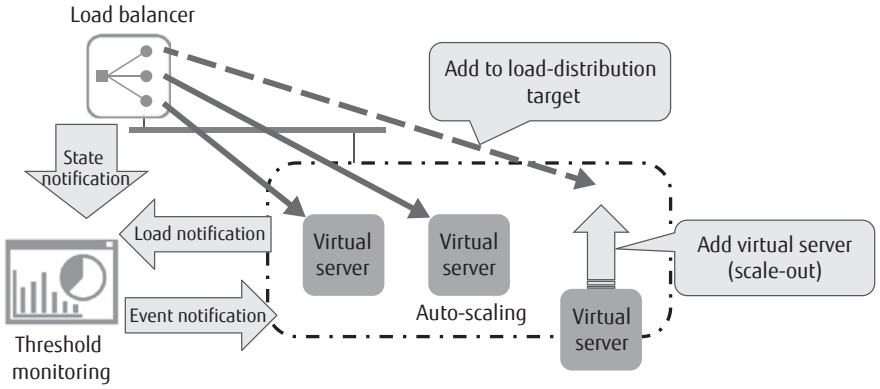

Figure 1
Auto-scaling.

are grouped into a security group, the user can easily manage them by including multiple virtual servers with the same security requirements in the same security group. Furthermore, if communications are allowed between different server groups, the user can define rules between the corresponding security groups. In other words, the security group function makes it unnecessary to define rules for each virtual server.

In K5 IaaS, security groups can be defined for virtual servers deployed in different zones. In addition, the auto-scaling and security group functions collaborate to enable specific security requirements to be automatically set for virtual servers added by scale-out.

3.4 DNS service

The Domain Name Service (DNS) enables K5 IaaS virtual servers, the load balancing service, and the database service to be accessed using fully qualified domain names (FQDNs). This service can be used to group the IP addresses of resources with FQDNs deployed in multiple zones and to return a randomly ordered list of IP addresses (DNS round robin). In this way, connection requests to resources can be distributed among different zones to disperse the load. Additionally, since the user will attempt to connect in the order given by the IP address list and will therefore connect to responsive IP addresses, this service can automatically degrade by detaching any faulty zones, thereby ensuring business continuity.

4. Use cases

This section describes specific systems that a user can construct using the functions and services provided by K5 IaaS.

4.1 Migration of existing Web 3-tier system

With K5 IaaS, a user can freely combine virtual networks, virtual routers, virtual firewalls, etc. to configure the target network. Furthermore, since the components (networks, routers, firewalls) necessary for configuring a network system in a conventional physical environment are also provided in a virtual environment, a network configuration almost the same as an existing business system in an on-premise environment can be achieved using K5 IaaS. In short, an on-premise business system can be migrated to K5 IaaS without changing its existing configuration.

An example of achieving a Web 3-tier system typical in on-premise environments using K5 IaaS is shown in **Figure 2**. Since virtual routers and virtual firewalls can be freely added using K5 IaaS, the user can easily create a multi-tier firewall configuration for minimizing security risks. There is therefore no need to change the role, function, and configuration of Web, application, and database servers from that of the on-premise environment. That is, a business application used in an on-premise environment can be migrated to K5 IaaS with practically no modification.

Once a system has been directly migrated to K5 IaaS from an on-premise environment in the above way, the use of system resources can be made even more efficient, and operation and management can

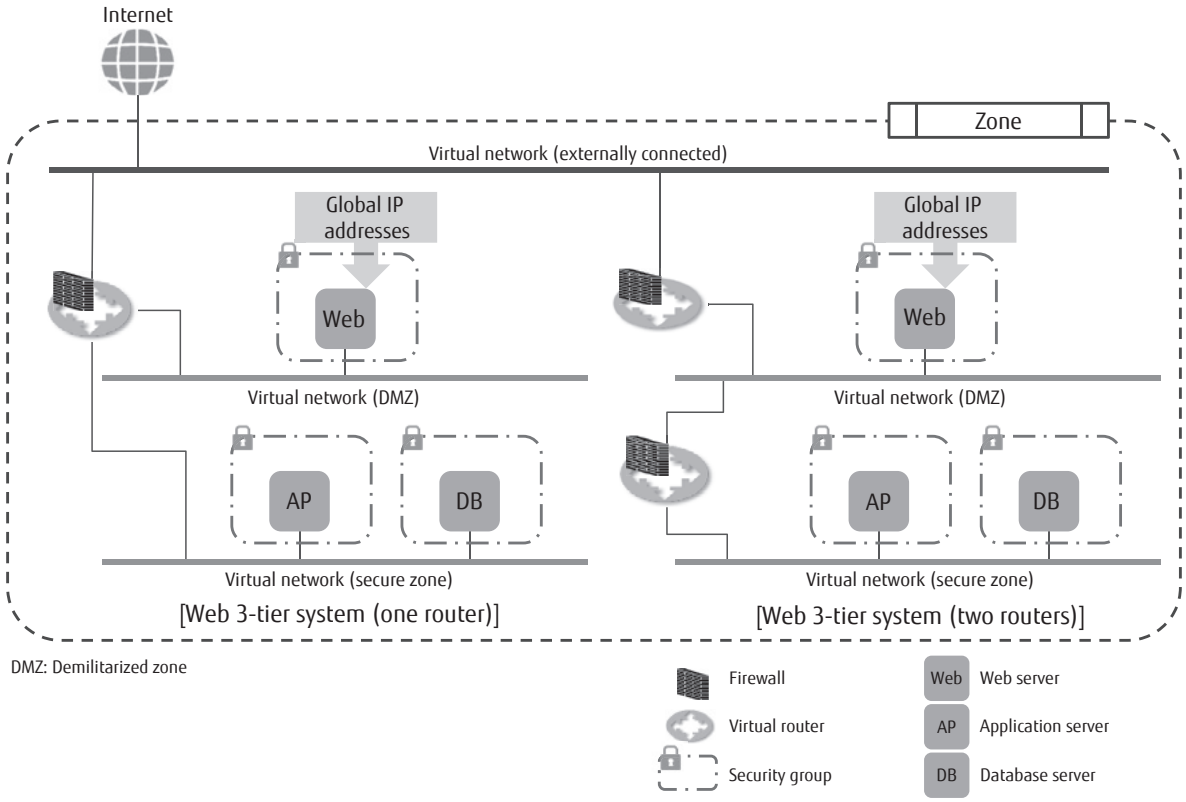

Figure 2
Web 3-tier system.

be simplified by incorporating the functions of a cloud-native system in a step-by-step process.

4.2 Cloud-native system

Developers and operators of user systems can entrust K5 IaaS with typical functions like load balancing and data management by utilizing the various services offered by K5 IaaS. They can then focus on developing and implementing the business logic of their systems. It is also possible to achieve high availability robustness against facility failures in data centers by configuring a system that spans multiple zones. High resilience to bursty loads can also be achieved by applying a scale-out configuration. An example of a system using the following services is shown in **Figure 3**.

1) Load balancing service

The load balancer forwards requests to virtual servers across multiple zones.

2) Auto-scaling function

This function automatically creates and deletes virtual servers in accordance with increases and decreases in load and failures of virtual servers. It can be used to deploy virtual servers evenly across multiple zones and to increase the scale of a system after a small start as business expands.

3) Database service

This service automatically deploys, operates, and maintains (through patches, etc.) a relational database (RDB) with a redundant configuration for high availability. A redundant configuration spanning multiple zones can also be specified.

4) Monitoring service

This service monitors the loads of the virtual servers and the activity of the load balancer. In the event that a predefined threshold is exceeded, it can send an alarm to the system operators and simultaneously execute auto-scaling.

4.3 Connection with on-premise system

It is possible to connect K5 IaaS to other environments. Examples of such environments include Fujitsu Enhanced Information and Communication

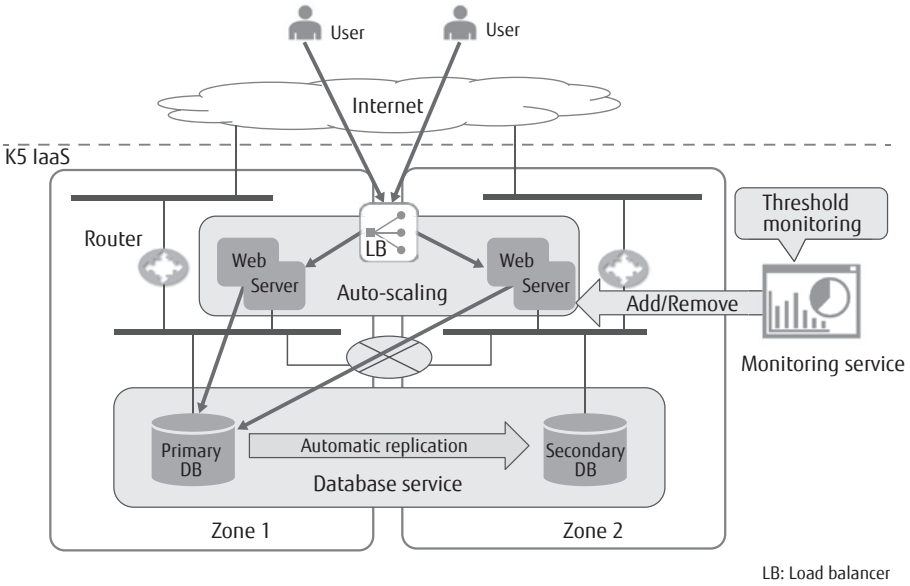

Figure3
Cloud-native system.

Services (FENICS), which provides mutual connections among company hubs and mobile systems in a multi-carrier environment, such as Fujitsu data centers in Tatebayashi and Akashi and communication carriers used by customers. A variety of environments can be consolidated by a private-connection option and connected to K5 IaaS (**Figure 4**).

5. Conclusion

This paper described FUJITSU Cloud Service K5 IaaS, a new public cloud platform provided by Fujitsu. In addition to providing systems based on the cloud-native concept for achieving SoE, K5 IaaS also supports the migration of existing SoR systems and the development of IoT-based businesses, which are expected to grow from here on. Additionally, to facilitate a hybrid type of usage connecting K5 IaaS with on-premise systems, Fujitsu is working to enhance platform commonality between on-premise products[4] and K5 IaaS.

Fujitsu is going to examine the functions that should be provided by K5 IaaS as a public cloud including its original functions with the aim of responding promptly to social needs. Fujitsu is also planning to consider new services such as Docker[5] containers and serverless architectures such as AWS Lambda.[6]

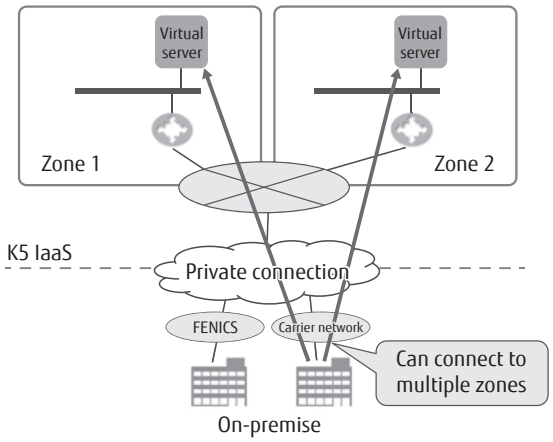

Figure 4
Connection with on-premise system.

References

1) Fujitsu: FUJITSU Cloud Service S5. *http://welcome.globalcloud.global.fujitsu.com*
2) Fujitsu: FUJITSU Enterprise Cloud Service K5. *http://www.fujitsu.com/global/solutions/cloud/k5/*
3) OpenStack: OpenStack. *http://www.openstack.com*
4) Fujitsu: FUJITSU Software ServerView Resource Orchestrator. *http://www.fujitsu.com/global/products/software/infrastructure-software/resource-management/serverview-resource-orchestrator/*

5) Docker: Docker.
https://www.docker.com/
6) AWS: Lambda
https://aws.amazon.com/lambda/

Yuji Wada
Fujitsu Ltd.
Mr. Wada is currently engaged in the planning and design of cloud services.

Satoshi Iyoda
Fujitsu Ltd.
Mr. Iyoda is currently engaged in the planning and design of cloud services.

Hiroshi Takamure
Fujitsu Ltd.
Mr. Takamure is currently engaged in the planning and design of cloud services.

Masashi Kameta
Fujitsu FSAS Inc.
Mr. Kameta is currently engaged in the development of services using K5 IaaS.

Yasutaka Taniuchi
Fujitsu Ltd.
Mr. Taniuchi is currently engaged in the planning, development, and operation of global cloud services.