Introducing the IIJ GIO Cloud Service

We began providing the IIJ GIO cloud service in 2009, and it is now in its fifth year.

We are continuing to expand our lineup of services to meet a variety of needs.

In this report we will once again give an overview of the service from a technological perspective, and in future reports we will continue to discuss the underlying technology that supports the service.

2.1 Introduction

Corporate adoption of cloud technology is proceeding at a faster pace than specialist institutions initially expected. One reason that cloud computing is growing so quickly is changes in the IT environment, such as faster and lower-cost networks, a significant reduction in the cost of deploying IT devices, unpredictable shifts in the business environment, and a wider range of applications for virtualization technology. Riding these waves of change, the market is beginning to understand the benefits to be gained from the use of cloud technology, including freeing companies from physical and geographical constraints, and alleviating the existing issue of IT infrastructure operating costs. Other factors include progress in moves toward asset-less environments and on-demand services, as well as the economic effect gained through the improved efficiency of resource utilization.

2.2 Service Overview

IIJ began providing the IIJ GIO cloud service in December of 2009, taking advantage of our accomplishments in the development and operation of IT resource services over more than 10 years, as well as the fact we have one of the largest Internet backbones in Japan. We began expanding our service lineup from June the following year, starting with the IIJ GIO Hosting Package Service. To this day we continue to provide high-quality cloud services that support our customers' business infrastructure.

The IIJ GIO Hosting Package Service and IIJ GIO Component Service in the IaaS space constitute the core of our varied service lineup.

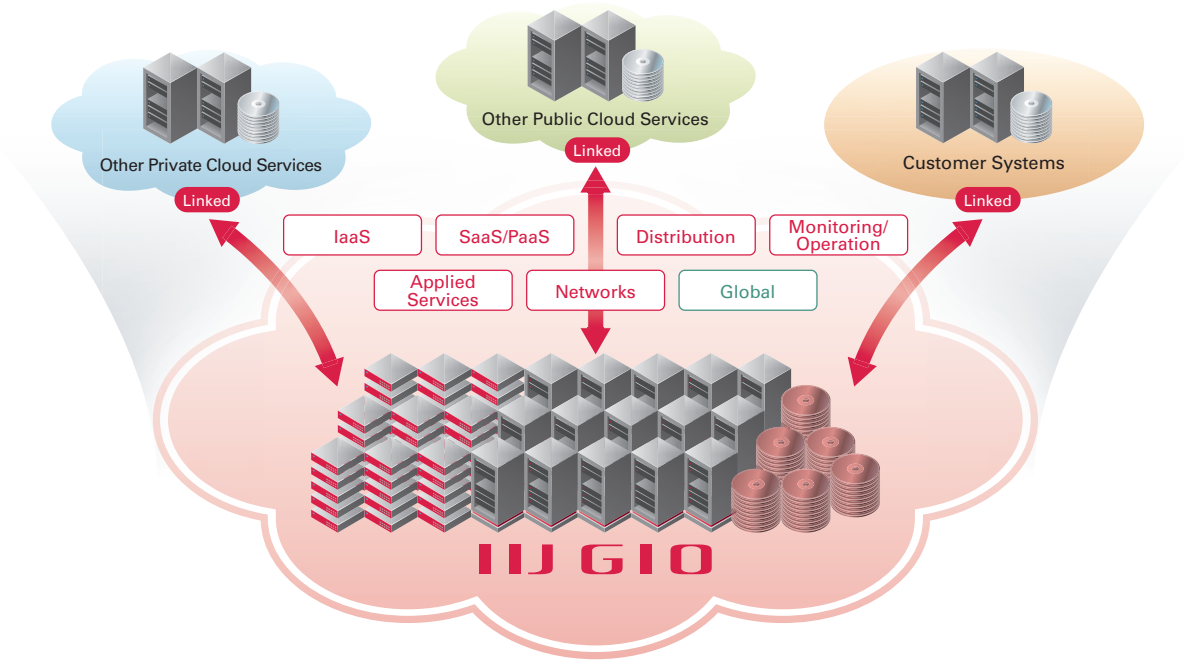

Figure 1: The Extensive IIJ GIO Cloud Service Lineup

2.2.1 Public Cloud-Based Services

The IIJ GIO Hosting Package Service*1 is an IaaS-based public cloud service available to use immediately through an online application. We made plans packaging system configurations for purposes in which the standard server software is fixed to a certain extent, such as Web mail and mail servers. Based on a Basic package that featured only CentOS, we have prepared a number of plans that combine software and hardware to meet various intended uses.

In September 2013 the API*2 was made public, answering the needs of those who wanted to automate routine operations by creating their own programs.

2.2.2 Virtual Private Cloud-Based Services

The IIJ GIO Component Service provides system configuration elements such as servers, storage, and networks as components. This is a flexible IaaS-based cloud service for enterprise customers that enables users to select the optimal configuration for their needs from a varied menu. Through direct connection with on-premise environments, it can also be used as a private cloud.

Its main components are the base server and the virtualization platform.

Two base servers are provided: the V Series virtual server in which the physical server is shared with other customers, and the X Series physical server that can be used exclusively. In January 2014 we updated the hypervisor, and released a V Series G2 option with enhanced management functions and virtual server specs.

For virtualization platforms we have the VW Series, which provides a physical server environment utilizing the VMware vSphere ESXi virtualization software from U.S. company VMware, along with VMware vCenter Server. Because VMware vSphere ESXi environments are provided with administrator privileges, they offer a degree of system configuration freedom comparable to on-premise environments. We think this provides an attractive option to users considering new server integration or cloud construction, as well as users already constructing or operating virtual infrastructure using VMware. We plan to offer a new version of vSphere ESXi and physical servers with more memory from spring 2014. We are also looking at adding packages using different virtualization software to our lineup.

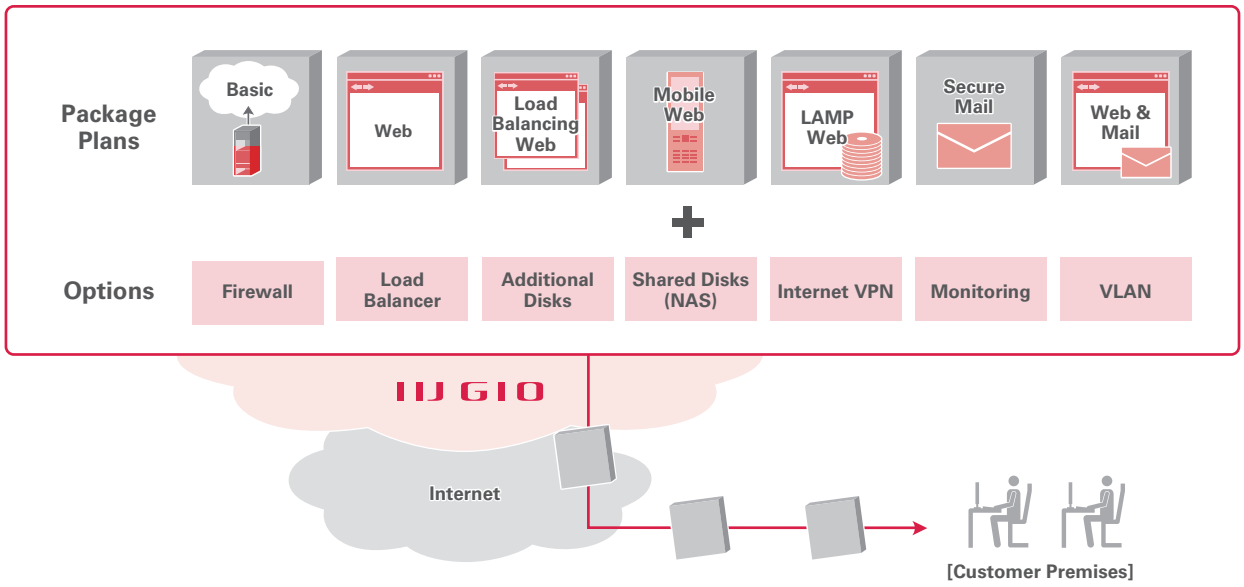

Figure 2: Overview of Public Cloud-Based Services

*1 See the following site for more information about the IIJ GIO Hosting Package Service. (http://www.iij.ad.jp/GIO/service/hosting/).
*2 See the IIJ GIO Hosting Package Service API Reference. (http://manual.iij.jp/gp/gpapi/).

Resources other than servers can also be selected as IaaS-based services.

The first of these is the Network Add-On*3, which enhances the network functionality provided as standard on base servers and virtualization platforms. By changing shared Internet access lines to dedicated lines (private connections), Internet VPN or walled garden networks (wide area networks) can be used to connect safely to on-premise environments. Multi-carrier configurations separated by carrier are also possible.

The second is the Storage Add-On*4. The two main options are Standard, which provides the same high-end storage used by financial institutions, and Basic, which is a mid-range option suitable for data management of general Web systems, etc. We offer network-based NAS storage, Fiber Channel (FC) SAN storage, and iSCSI SAN storage.

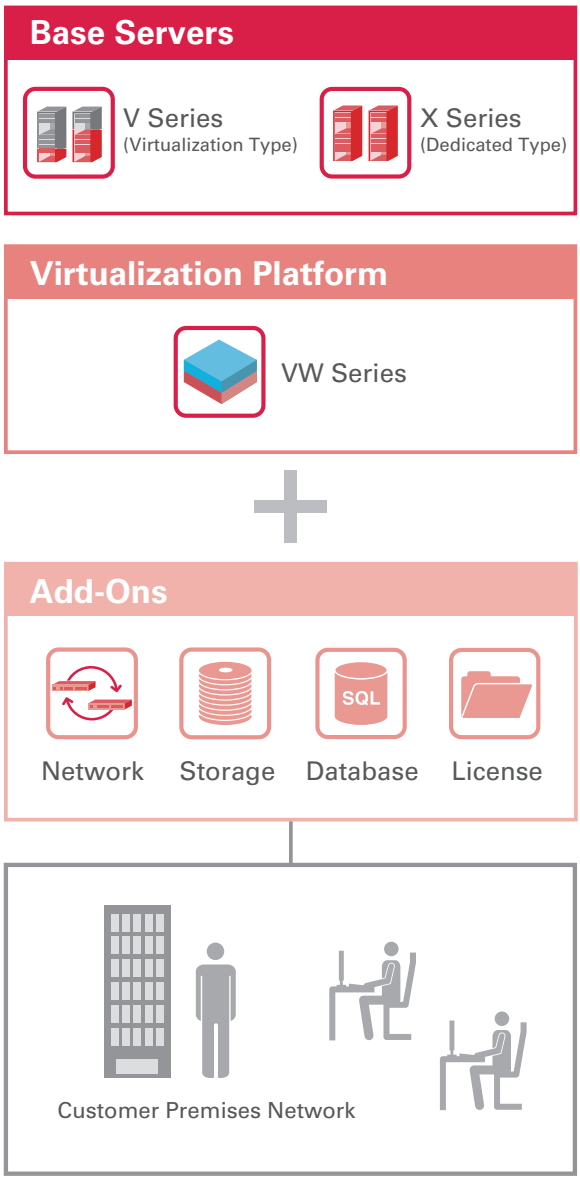

Figure 3: Overview of Virtual Private Cloud-Based Services

The third is the Database Add-On*5, which provides Oracle Database and MySQL as DBaaS (Database-as-a-Service). Oracle Database licenses are available for a monthly fee, enabling customers to alleviate the burden of investing in database licenses and avoid the risks associated with investment. This service provides database instances designed and operated by IIJ, based on our experience with implementing and operating large numbers of relational databases.

The fourth is the License Add-On, which provides software licenses for a monthly fee. Currently we provide software products that are in high demand for a monthly fee, including Microsoft SPLA products, as well as products from Citrix, VMware, and Symantec.

2.2.3 Cloud Service Usage Trends

IIJ GIO usage trends for the past few years show that, compared with April 2012, while the number of physical servers has almost doubled, storage capacity has almost tripled. One reason for this is that use of our cloud services has expanded from general Web systems to enterprise-oriented systems that handle large amounts of data. An indicator supporting this is the fact that trends in the number of servers for each of our services match the timing of increases in the number of VW Series corporate private cloud service servers used.

We will continue to monitor the usage status of equipment, and report information on usage trends, etc.

*3 See the following page for more information about network add-ons. (http://www.iij.ad.jp/GIO/service/component/network/).
*4 See the following page for more information about storage add-ons. (http://www.iij.ad.jp/GIO/service/component/storage/).
*5 See the following page for more information about database add-ons. (http://www.iij.ad.jp/GIO/service/component/db/).

2.3 Points for Building and Operation of Large-Scale Infrastructure

Large-scale infrastructure such as that used to provide cloud services can be boiled down to two key points: physical construction and infrastructure technology and operation. We plan to examine infrastructure technology and operation over several installments from our next report. In this first article, we will focus on physical construction.

Cloud services require infrastructure to be constructed on a massive scale in a short period of time. In light of these requirements, we evaluated a variety of implementation methods, and in the end we were able to shorten the lead time to about a third of what it was initially. Here we will discuss some of the issues we faced along the way, and how we overcame them.

2.3.1 On-Site Construction

On-site construction is still common practice for the physical construction of IT devices used on-premise or in colocation services involving data center rental.

Hardware such as the servers, storage, and network devices to be used are sent to the place of installation, and kitting, racking, and wiring of LAN cable and fiber optic cable is carried out on-site. We initially adopted the same approach for IIJ GIO. However, we faced issues unique to the construction of large-scale environments, such as data centers requiring a lot of space for unpacking and setup work due to the large amount of equipment involved. There were also issues with processing the huge quantity of waste material left over after construction, and the repeated need to deal with initial failures on-site. (Figure 4)

2.3.2 Container Transportation

In April 2011, we began container transportation following the establishment of Matsue Data Center Park in Matsue, Shimane Prefecture. This was the first commercial container-based data center with outside-air cooling in Japan. (Figure 5)

Previously servers were carefully packed, increasing their bulk about eight-fold, which meant that more vehicles were required to transport them. In contrast, transporting containers with racks to the server manufacturer's factory, and carrying out all the kitting, racking, and rack wiring there, realized the concept of hauling containers with servers already installed and ready to use straight away. At the same time, it enabled construction and transport costs to be reduced, and also solved the issue of needing to deal with a large amount of packaging material waste after installation work. (Figure 6)

Figure 4: Large Amounts of Packaging Material (Waste) Left After Physical Construction

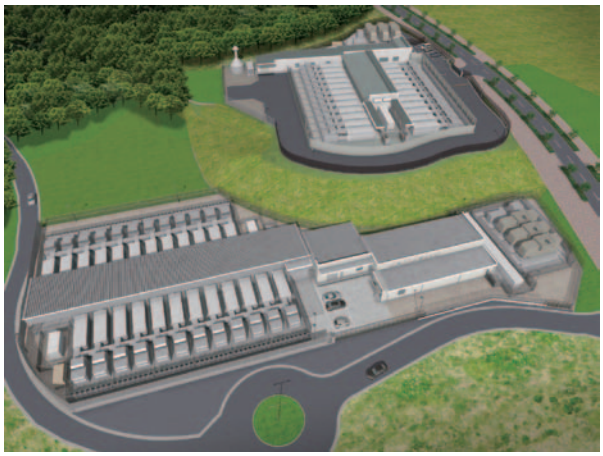

Figure 5: Illustration of Matsue Data Center Park Expansion After Completion

Figure 6: Installation of Transported Container

We also enabled low-cost, flexible transportation using standard heavy trucks*6 by developing slim containers that reduced the initial width to less than 2.5 meters. (Figure 7)

By implementing container transportation, we were able to provide services using large numbers of IT devices in a short space of time, and reduce the work carried out at data centers as much as possible. However, this still left issues such as the facility adjustment period for cooling units after on-site installation, and the handling of initial failures on-site. (Figure 8)

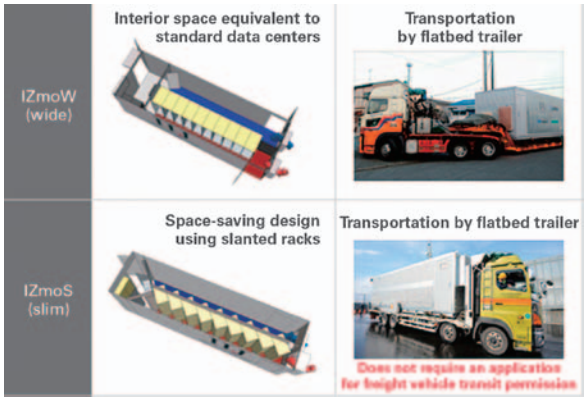

Figure 7: Container Transportation Using Heavy Trucks

Figure 8: Work Inside Containers and Initial Failure Handling

Figure 9: Racking Work at Server Factory

2.3.3 Rack Transportation

We adopted rack transportation at our urban data centers.

For container transportation, we bring containers fitted with racks to the server manufacturer's factory, but for rack transportation, we only bring the racks themselves. This enabled us to carry out the same kitting, racking, and rack wiring process as container transportation, but we were also able to achieve further reductions in the lead time after transportation by implementing electrical and initial failure handling at the factory as well. Of course, this would make no sense if device failures occurred during transportation, so we also incorporated efforts to limit the impact of transportation on installed devices to a minimum. (Figure 9)

Figure 10: Rack Transportation

*6 When a vehicle's width is over 2.5 meters, it becomes a "special-purpose vehicle" that must make a special application to travel over roads. Lead cars are required in front of and behind the truck, and the roads it can travel on are restricted.

Next, we need to consider the route for transporting racks to the data center. Because racks remain upright with servers installed, the transportation route must have sufficient load bearing capacity for the weight of the racks, and enough height to haul them upright on the truck. (Figure 10)

Weather is an even more important factor. Waterproofing measures are essential to proceed with transportation work without delay, as schedules may overlap with the rainy season or typhoon season. In addition to standard cushioning, racks are given three layers of waterproofing for good measure.

Here we have focused on urban data centers, but rack transportation is also a valid option for container-based data centers. With the containers themselves left connected to the on-site cooling units, the racks inside are transported to the server manufacturer's factory. This makes it possible to reduce the facility adjustment period with cooling units that occurs after installing containers on-site.

By transporting racks using the above approach, we greatly reduced the lead time before starting server operation.

2.4 Conclusion

In this report we gave on overview of the IIJ GIO service using IaaS as an example. In the future we will look at the underlying technology that supports the service.

We will also continue to provide you all with services and fundamental technology currently under development at IIJ that meet the needs of the market.

Author:

Shinri Kimura
Vice Manager, Platform Service Department, Platform Division, IIJ As service manager of the IBPS service that preceded the IIJ GIO cloud service, Mr. Kimura was responsible for its development and operation. He later participated in IIJ GIO from the planning and concept stage, and he is now engaged in the development and operation of IIJ GIO system infrastructure services.