Introducing the IIJ GIO Cloud Service - Part 2

Last year we presented the first part of our introduction to the IIJ GIO Cloud Service.

In this second installment, we examine the status of service utilization in fiscal 2014, and discuss our activities with regard to the development, operation, and maintenance of large-scale infrastructure.

2.1 Introduction

2.1.1 Cloud Market Trends

While the Internet business-oriented usage by parties such as social application providers that has supported the growth and expansion of the cloud market to date is slowing and leveling off, it is evident that enterprise usage at major corporations is growing, and this growth is expected to continue going forward. With these changes in usage trends, more focus is being placed on connecting to (or coordinating with) a variety of cloud platforms as part of a multi-cloud service, rather than a single company providing cloud services. IIJ is also paying close attention to these developments, and we are preparing to provide multi-cloud services in addition to our existing multi-carrier connectivity.

2.1.2 Service Overview

■ **Topics**

We will start offering the IIJ GIO Component Service Database Add-on In-memory Platform for SAP HANA® in April 2015. This service enables customers to run a production environment incorporating SAP's SAP HANA® in-memory database in the cloud. We previously launched the IIJ GIO for SAP Solution PoC for SAP HANA® (Figure 1) cloud environment for prior validation in June 2014, and this new service provides the same features for production environments. IIJ has obtained SAP

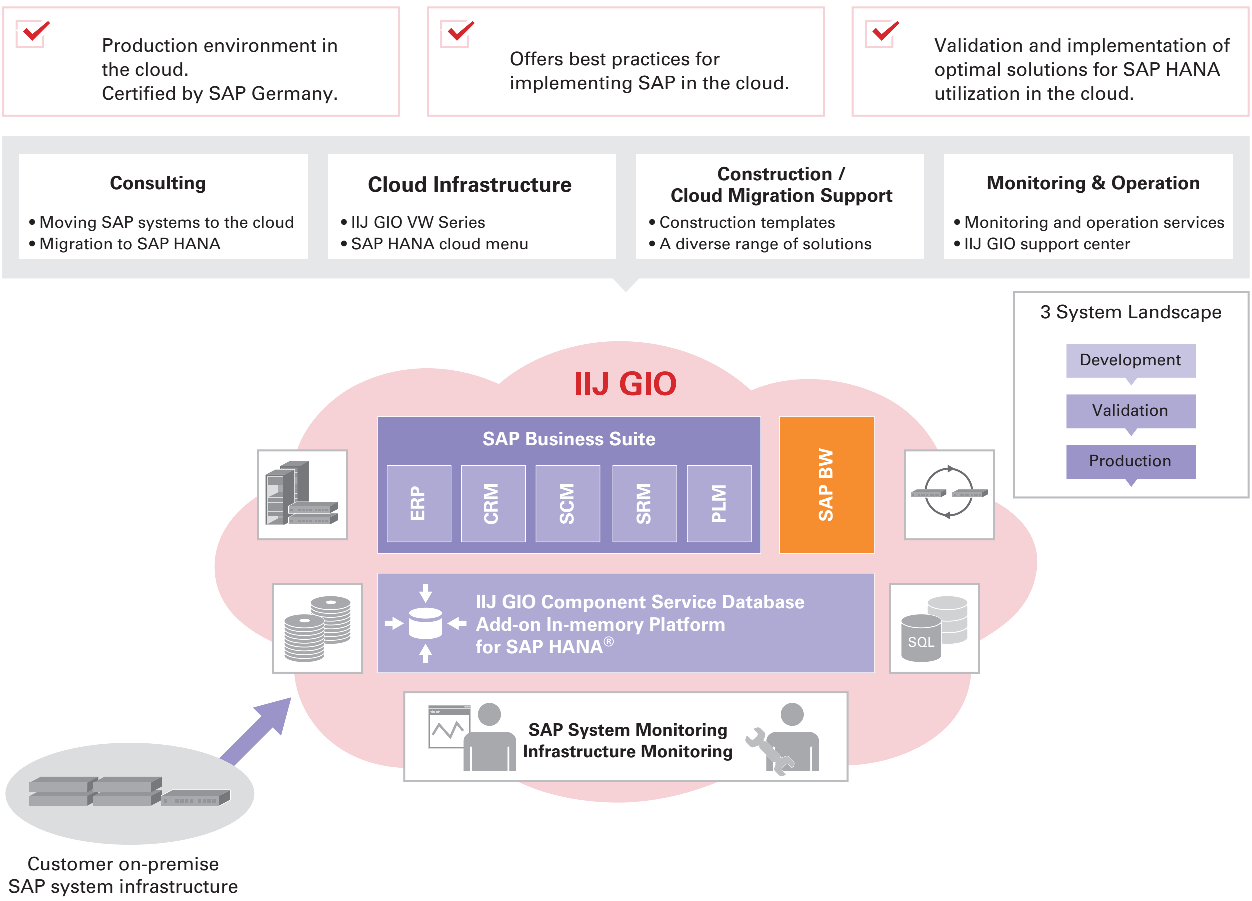

Figure 1: IIJ GIO for SAP Solution

Certified in Hosting Services and SAP Certified in Cloud Services*1 certification, and provides worry-free access to a high level of services and security in SAP production environments. This keeps initial investment and operation costs down, and enables swift development and deployment of systems to quickly respond to business decisions.

In January 2015, we began offering the IIJ Cloud Exchange Service for Microsoft Azure (Figure 2). This uses the Microsoft Azure ExpressRoute walled garden network service, which IIJ is the first partner in Japan to offer. Until now connections to Microsoft Azure were made only over the Internet, but ExpressRoute enables connection via a walled garden network. Starting with this service, IIJ plans to further expand its multi-cloud services.

■ Service Facility Development Status and Usage Trends

Regarding the status of facility development for IIJ GIO in fiscal 2014 (as of March 2015), we forecast around a 20 percent increase in physical server numbers and a 40 percent increase in physical storage space compared to the end of fiscal 2013. As far as usage trends are concerned, there continues to be fast-moving fluctuations each month in the use of virtual servers used mainly as front-end Web/AP servers for social application providers. Meanwhile, regarding physical servers chiefly used as database servers, we have seen further progress in the consolidation and shift from servers using basic local hard drives or FC storage to physical servers equipped with semiconductor storage that offer high I/O performance. In the area of enterprise-oriented equipment, this was a year that saw remarkable growth in storage over and above server numbers. The dramatic growth in storage usage suggests that major corporations requiring large amounts of storage are migrating mission-critical systems to the cloud, and real utilization of systems migrated to the cloud is progressing. We expect these trends to continue in fiscal 2015.

2.2 Large-Scale Infrastructure Technology Initiatives

2.2.1 Infrastructure Development

■ Cloud Infrastructure Considerations

As a cloud service provider, we are always thinking of ways to operate large numbers of hardware devices efficiently. If we were simply focused on optimization, then the homogenization of hardware would be ideal, but in reality this is still difficult

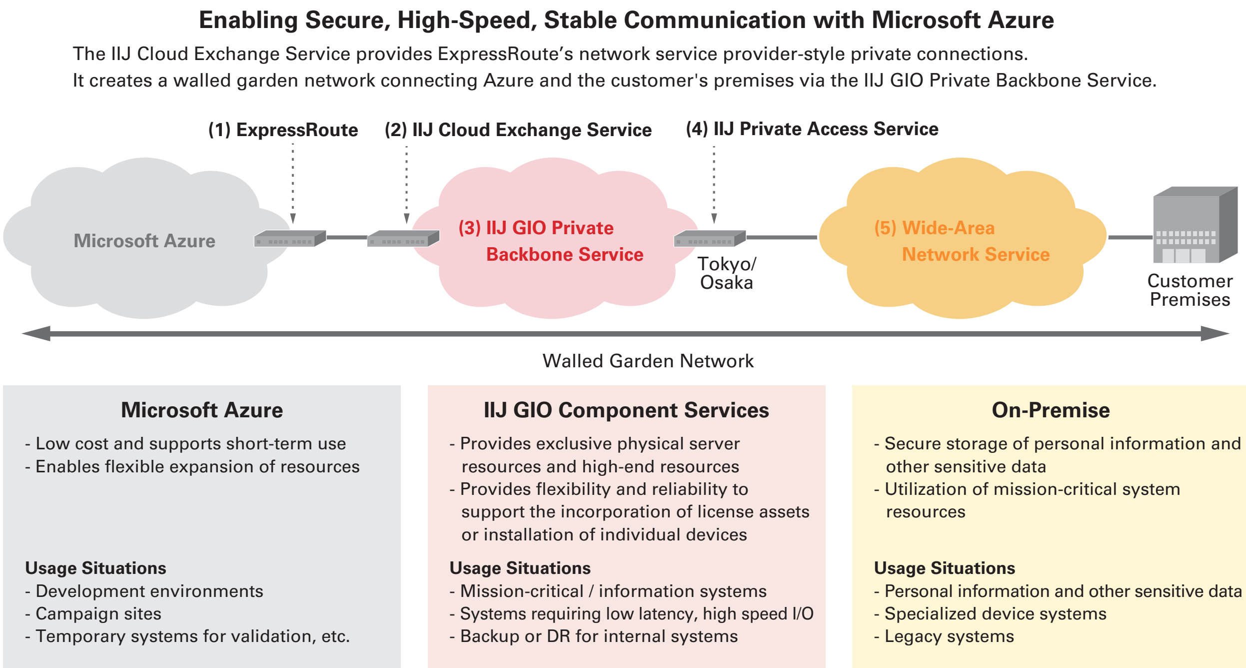

Figure 2: IIJ Cloud Exchange Service for Microsoft Azure

*1 Meets the service quality criteria established by SAP Germany, and is recognized as a SAP-certified cloud service provider and SAP-certified hosting service provider (certified April 18, 2013).

at present considering IIJ GIO's wide-ranging lineup of services, as well as its cost-effectiveness. As a result, we are currently aiming to implement only as much homogenization as possible.

For example, infrastructure for the public cloud-oriented IIJ GIO Hosting Package Service targeted at Web-based companies and social application providers must compete in the public cloud field where price competition is fierce, so we provide practical-minded hardware configuration and OSS combinations that eliminate excess redundancy. Meanwhile, infrastructure for the private cloud-oriented IIJ GIO Component Services targeted at the enterprise sector incorporates hardware redundancy at least equivalent to on-premise solutions, as on-premise expansion resources or a receptacle for migration from on-premise equipment. We also provide combinations incorporating commercial products to meet requirements for cloud infrastructure that guarantees operation with existing commercial products that a customer coordinates with or brings in. Above all else, it is essential to diversify technological risks to provide stable service infrastructure. Because it is not possible to completely eliminate the risk of issues with firmware or drivers, it is best to avoid using only products from a specific vendor even for the same component, to localize the impact to a certain extent should an issue occur. In other words, multi-vendor systems have benefits from the perspective of diversifying technological risks. Adopting multiple vendors also enables you to select from a range of vendor products, making choices based on commodity technology instead of being dependent on specific vendor product technology. As a result, the principle of competition applies when procuring equipment, contributing to reduced procurement costs. Adopting commodity technology also leads to optimizations in operational areas such as delivery and maintenance, making integrated management using common management tools possible.

In addition to diversifying technological risks, we are pushing forward with standardizing the base equipment for different service infrastructure with an eye toward homogenization. This makes it easy to reuse equipment in different service infrastructure by simply adding or changing components. In our current service development, this standardized approach is applied to the physical design of service infrastructure as long as there are no special requirements. Standardization enables the surplus inventory for each service infrastructure to be treated as a shared inventory pool, also contributing to reduced inventory risk. With regard to the hardware adopted, we have seen an increase in the range of options available centered around overseas cloud service providers over the past few years. In the server field, in addition to the purchase of general server vendor products as before, there have also been cases in which requests for manufacturing were made to ODM vendors, with a central focus on Taiwan. Servers produced through knowledge gained from the Open Compute Project (OCP) and certified by it have started appearing on the market as well.

Except for a few services, IIJ mainly uses general-purpose IA servers supplied by server vendors for the IIJ GIO cloud service infrastructure developed in-house. Of course, we have also evaluated and validated server products from ODM vendors. However, as a result of overall comparison of our service infrastructure requirements and the economies of scale gained from procurement, we came to the conclusion that server vendor products are more suitable for our current cloud service infrastructure. That said, this may not always be optimal in the future, so we will continue to keep tabs on trends in the aforementioned server vendor products, ODM vendor products, and OCP servers to respond to changes in the market swiftly.

With regard to software, we believe it is inefficient to create it all from scratch ourselves, and this would have few benefits for users. From our perspective, there is value in being able to provide services by adopting good products on the market based on customer needs, whether commercial products or open source software (OSS). However, there are also many difficulties associated with market products, such as the need to refine them when they don't support the scale required by cloud service providers. There is a lot of talk in public circles about how to handle Windows Server 2003, for which extended support will end on July 15, 2015, but the end of support is of course always an issue with commercial products, and even OSS. In light of this issue, for IIJ GIO we share information with and issue alerts to users of services that incorporate products for which support is scheduled to end. In particular, the private cloud-oriented IIJ GIO Component Services targeted at the enterprise segment provide a range of services that combine a large number of commercial products, and support lifecycles are defined for each vendor. To ensure that operation of these is always guaranteed and support can be provided, we list the support lifecycles of products for each service in the form of a service lifecycle, and periodically update information to help users continue using services safely and securely. The needs of users are not met by merely providing services at low cost, and IIJ takes this into account.

The Accumulation of Ongoing Technological Development

Infrastructure development is triggered by the dual action of service development requirements and accumulated technological development. In particular, developing the optimal infrastructure to meet service development requirements hinges upon how much day-to-day knowledge and experience you have in a variety of technological development fields (Figure 3). As an organization, IIJ continually repeats a technological survey cycle that involves gathering information on product trends and roadmaps for each manufacturer, comparing and evaluating technology and trends, and producing periodic reports. We also refer to feedback from the development site and management based on these reports, and perform internal product validation to assess whether products that show promise would stand up to service operation (Figure 4).

Our service development team plans and develops services in keeping with the technical knowledge we have accumulated and our business strategy. In the design phase infrastructure is designed based on the validation results produced by the product validation team, with the aim of guaranteeing service quality and reducing design time. Surveys are of course targeted at servers themselves, but also cover areas as diverse as hardware such as storage and network equipment, software such as OSes, hypervisors and SDNs, and control systems such as OpenStack and CloudStack.

Examples of our ongoing initiatives include keeping pace with server-side flash and Intel's next-generation CPUs. Each time a new product is released, IIJ obtains evaluation units from the vendor and works on identifying changes in performance characteristics over generations. We also clarify the impact on service and points of concern stemming from switching to successor models, and develop internal guidelines to facilitate their smooth introduction to the service environment. By continually tracking the above technological points even for comparatively new products and models with a short product

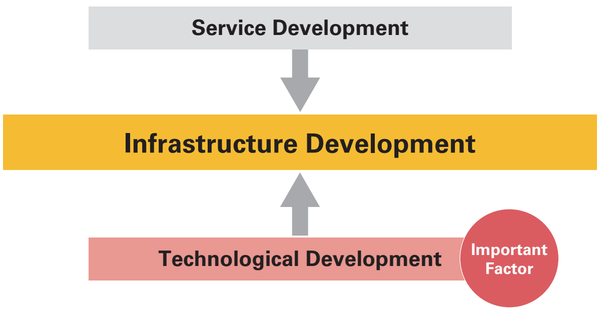

Figure 3: Infrastructure Development Approach

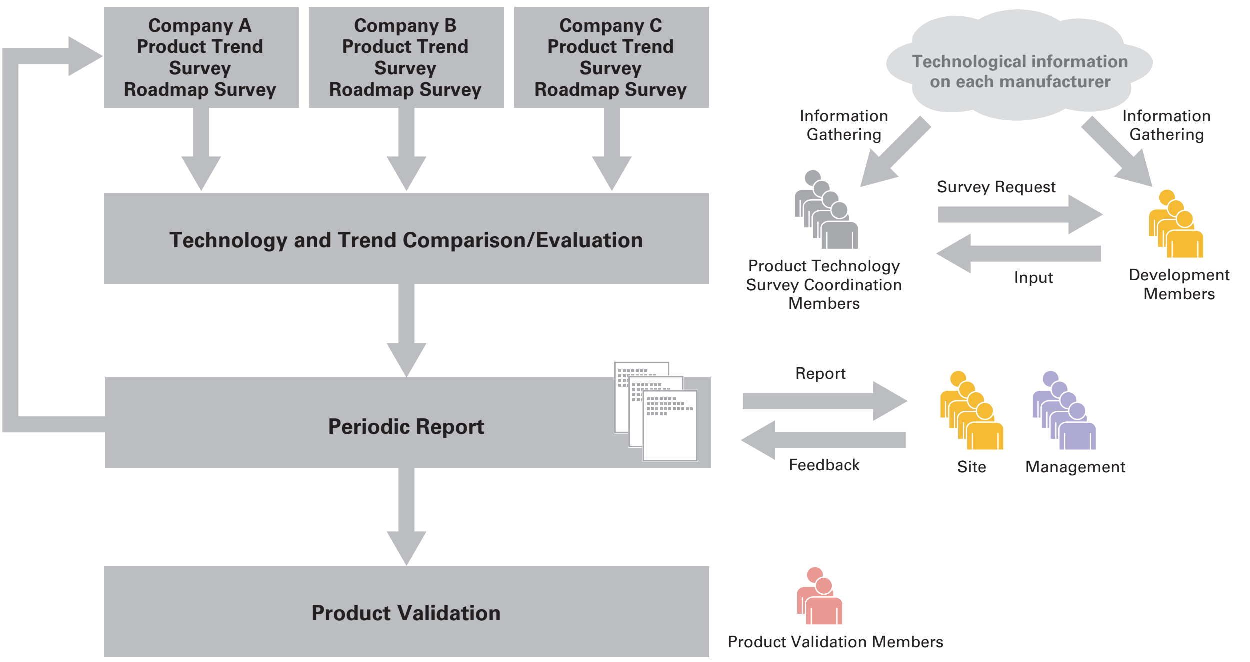

Figure 4: Technology Trend Survey Process

cycle, we can rapidly incorporate them into our infrastructure in response to service requirements and provide them. One of our initiatives in the area of software is participation in Microsoft's Technology Adoption Program (TAP). By participating in TAP, we can cultivate a better understanding of the design objectives, new features, and changes from the current version for the next versions of the Windows Server operating system and System Center management products. In addition, we can further our knowledge of implementing and upgrading these products. IIJ also submits feedback in the form of feature requests and notes on specification improvements to meet our service development requirements. Through these efforts, we are able to add new OS and hypervisor versions to our infrastructure and provide services immediately after their release. Having our requests as a cloud service provider incorporated makes it possible to migrate or run services with minimal downtime each time a new version is released.

2.2.2 Infrastructure Operation

Infrastructure operation is just as important as infrastructure and technology development. Here we discuss infrastructure maintenance activities such as failure and fault handling that are the mainstay of infrastructure operation. As with on-premise equipment, system failures also occur on the equipment of cloud service providers. However, even though we are not using new products straight after their release, extremely rare system failures with no precedent may be encountered due to the nature of multi-tenant environments, which feature large numbers of users applying a range of configurations and uses. Instead of simply dealing with rare failures on the spot, we check whether the failure affects the infrastructure as a whole, and monitor the situation for a set period to guard against recurrence. Although identifying the cause of a failure requires money and time, inaction during the initial response can lead to large-scale failure, so our stance is to aim for establishing surefire countermeasures.

■ Infrastructure Maintenance Activities Such as Failure and Fault Handling

Examples of infrastructure maintenance activities include proactive maintenance, measures for dealing with underlying causes, and replacement.

Proactive maintenance includes preventive maintenance for potential and fatal faults discovered after a service is released, as well as the application of patches to keep pace with the support lifecycle of firmware or software. It is implemented as part of the maintenance for receiving appropriate vendor support in case of unanticipated events. For information gathering, general security information is issued from an internal information sharing site daily, and we have a framework in place for checking failure and vulnerability information for the products and technology used in service infrastructure on an ongoing basis. The same applies to support lifecycles. However, we don't merely take vendor information on faith. We first carefully validate whether there is any impact on the existing environment, including parts we have built ourselves, using validation/staging environments practically identical to the production service environment. Then we select the parts we consider important for the stable operation of the service environment. It is also rare for the systems or tools provided by vendors to take deployment in large-scale infrastructure into consideration. In many cases, firmware and patches are provided on the assumption that they will be applied individually to equipment step-by-step, so it is necessary to devise a method such as automation to apply them reliably to large quantities of equipment. This demonstrates that proactive maintenance strongly resembles the technological development process detailed above. Dealing with underlying causes involves taking measures to tackle failures occurring after the release of a service that impact the infrastructure as a whole. As mentioned earlier, we sometimes encounter extremely rare system failures without precedent. Even for unprecedented failures, it is necessary to recreate the failure to resolve it. There is no way for vendor support to take measures without recreating the failure, leading to a dead end. However, due to vendor support performing reproduction tests using only a handful of validation units, in some cases a failure cannot be recreated even after we convey the details of the phenomenon. This is because some failures have an extremely low chance of occurring, so reproducing them requires many repeated attempts. Consequently, we occasionally have vendor support engineers visit us to confirm the failures actually occurring, and have ongoing reproduction tests carried out. We are able to reproduce even failures that normally only reoccur after a hundred attempts on ten machines comparatively easily because we can execute a range of patterns at once utilizing large amounts of equipment. Reproduction tests are performed by IIJ engineers who have accumulated reproduction knowledge through encountering many failures in past operational experience. Taking advantage of this superior reproduction capability, we constantly endeavor to encourage vendor support to have fixes made promptly to create more stable service infrastructure. For example, when introducing a large quantity of server equipment in the past, there were frequent NIC communication failure incidents right after initial installation. Because vendor support insisted that there had been no other cases in the past, we analyzed the failure events, established steps to reproduce them, and provided vendor support with information on environments the failures occurred in as well as the steps taken. This enabled vendor support to confirm the failure events, and a few months later the cause was identified. Vendor support came up with every workaround for the failures they could think of, such as updating the BIOS, firmware and drivers, as well as changing the settings, but ultimately the failures did not stop occurring. In the end, the vendor determined it would be difficult to avoid the failures with the existing NIC, so the radical measure of switching to a different manufacturer's NIC was taken.

We work on discovering and eliminating the risk of widespread failures on a daily basis, based on the reproducibility of faults like this. Due to an accumulation of frustrating experiences in infrastructure operation such as the one above, the diversification of technological risks is an indispensable requirement for service infrastructure today.

Replacement involves the renewal of equipment when devices you are using reach the end of their support lifecycle. For corporate IT systems, hardware is generally replaced in cycles of three to five years when the maintenance or lease expires. Service providers like us that offer cloud services also procure equipment (with some exceptions) from general IT hardware vendors, and replacement will be necessary after a set time (although this can be extended a little) based on the product lifecycle determined by each hardware vendor when equipment is installed. For example, the IIJ GIO Hosting Package Service infrastructure that provides Linux-based virtual servers faced the problem of maintenance support expiring for the storage devices it used in expansion disks. The existing storage devices required renewal due to end of manufacturer support and lease expiration. Cloud service infrastructure always calls for zero downtime when replacing equipment, but the existing models for these storage devices presented a catch-22 situation, as they did not support storage migration to different models, and the new models had no migration function to begin with. To solve this problem we collaborated with the hardware vendor. By meeting with the vendor many times over a number of months, we convinced them to add a function to existing models that enabled migration to different models. Going one step further, we also gained official support for this migration function by having the function added through the modification of publicly available firmware, instead of asking for the migration function to be provided to us individually. However, the problem was not solved by this. Another issue that came to light was that migration work for each chassis required the execution of several hundred commands. Overall, several tens of thousands of commands would need to be executed. Because there was no way to execute these manually in an error-free manner, we prepared a system for executing the migration function reliably by creating a new semiautomated script.

As the case above shows, collaborating with hardware vendors and taking user requirements into account are areas that require experience and knowledge. These kinds of responses, including the semiautomation of processes, are likely to be missed and thus fail to draw much attention from users, but they are hard to obtain for individual customers. I believe that aspects such as this are one of the benefits of using a cloud service provider like us.

2.3 Conclusion

In this report we have given an overview of the infrastructure technology that supports the IIJ GIO service. Hopefully this gives customers that use this service a glimpse of what goes on behind the scenes. We will continue to strive for day-to-day stable operation and improved quality in the IIJ GIO service, to provide safe and secure cloud services to everyone.

Author:

Shinri Kimura
Manager, System Infrastructure Technology 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.

Contributor:
Tsuyoshi Inaba, System Infrastructure Technology Department, Platform Division, IIJ.