

## IaaS Platform Using OpenStack and OpenFlow Overlay Technology

*Shintaro Mizuno, Hiroshi Sakai, Daigoro Yokozeki, Koji Iida, and Takaaki Koyama*

### Abstract

This article describes how OpenStack and a virtual network controller that uses OpenFlow overlay technology can be orchestrated to deploy virtual environments with complex network configurations, as in on-premise systems, in an on-demand manner.

### 1. Introduction

There is increasing use of infrastructure as a service (IaaS), in which virtual servers, virtual storage, and other such virtual environments are provided, as exemplified by Amazon Web Services. This situation has given rise to an open source software (OSS) community whose objectives are reduced cost and open development in the construction of IaaS platforms.

### 2. OpenStack

OpenStack [1] is an OSS IaaS platform development project established by NASA and Rackspace in 2010. The NTT Group is among the 183 companies and 3386 individuals that currently make up that growing community. OpenStack version 5 was released in April 2012 under the code name “Essex”. OpenStack is one of the most prominent of the OSS IaaS platforms and is now being used in commercial services as well as for other purposes. In April 2012, 19 companies including AT&T, HP, IBM, Rackspace, and Red Hat established the OpenStack Foundation to promote the development and use of OpenStack, raising expectations for even more widespread adoption of the platform.

OpenStack comprises multiple services (functional modules), including a virtual server management function, object storage, and network management. Users can select the services they need to flexibly construct an IaaS platform (Fig. 1). This article focuses on the platform’s virtual server management

and network control functions. We describe a method of constructing a flexible virtual environment that combines virtual servers and virtual networks through cooperation with those OpenStack functions.

### 3. Virtual networks

As in hybrid clouds, there are use cases for cloud systems that include virtual networks as well as virtual servers to constitute virtual systems in a cloud environment while maintaining the system configuration and integrity with conventional enterprise systems. The virtual network is one kind of technology for implementing such use scenarios [2].

### 4. Technology for implementing virtual networks

A number of ways of connecting virtual servers to isolated customer virtual networks in the cloud have been proposed. One is to allocate a virtual local area network (VLAN) to each customer to form virtual networks in units of a VLAN. In the approach using VMware\*1 or other such products, a particular VLAN is assigned to the virtual server. The virtual servers on the same VLAN engage in closed communication. This approach requires the prior setting up of a VLAN with switches and other communication devices in the datacenter.

\*1 VMware: A PC/AT compatible emulator that runs on Windows and Linux.

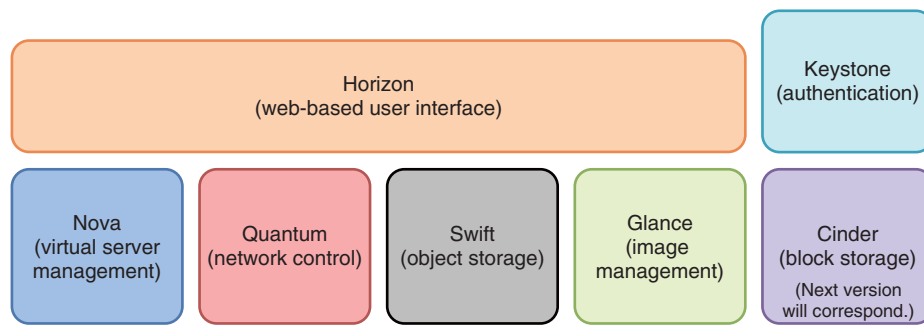


Fig. 1. OpenStack services.

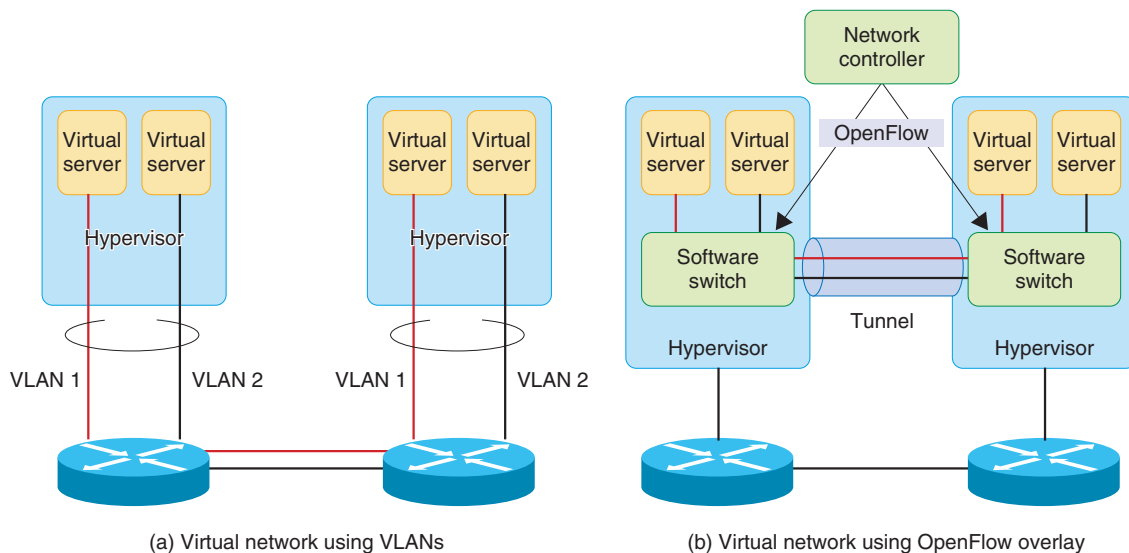


Fig. 2. Examples of virtual networks.

One method for connecting virtual servers that does not require control of the settings of communication devices is the *overlay* method, which involves extending tunnels between software switches on a hypervisor to establish communication paths between virtual servers running on the hypervisors. OpenFlow or other such tools are used to control the software switches (Fig. 2). Another method is the *hop-by-hop* method.

### 5. Network controller

In OpenFlow, a server referred to as the network controller manages the creation, setup, and deletion of virtual networks. The network controller manages the network devices, software switches, and other

functions on the hypervisor to construct virtual networks. Network controllers are provided by several companies (including NEC, Cisco, Nicira, and NTT), and each company has its own approach to network virtualization, including the methods described above.

### 6. OpenStack network control function

In OpenStack, the network control function is provided by the Quantum service, which provides an application programming interface (API) for creating, setting up, and deleting virtual networks. One feature of Quantum is a plugin architecture. Quantum itself provides only common functions such as authentication and an API for higher-level appli-

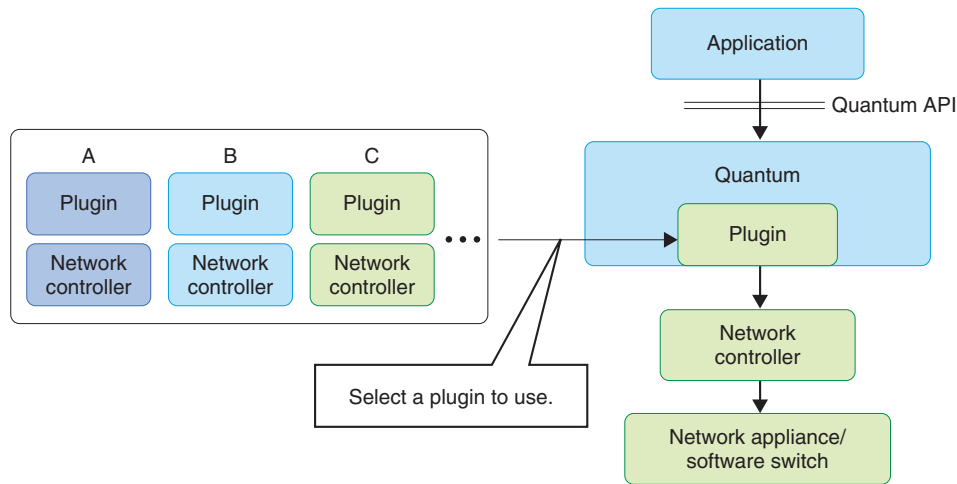


Fig. 3. Architecture of Quantum plugins.

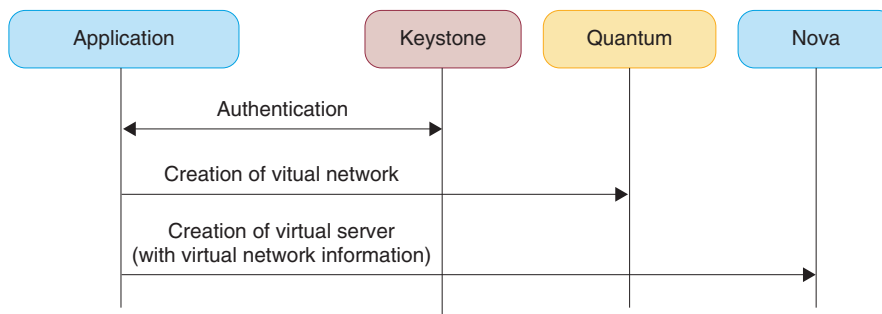


Fig. 4. Example of service call procedure.

cations. The virtual network functions are incorporated into the existing network controller. As described above, network controllers are already available from several third-party vendors, and flexible connection to the controllers is possible with Quantum plugins for the various controllers. That makes it possible to use any network function flexibly via a unified Quantum API (Fig. 3).

### 7. Cooperation between virtual server management and network control functions

In OpenStack, management functions for virtual servers and networks are individual services that are called by their respective APIs to construct a virtual environment from virtual servers and virtual networks.

A procedure for constructing a virtual environment

from virtual servers and virtual networks is illustrated in Fig. 4. When OpenStack is used to provide an IaaS cloud service, these procedures for calling the various services are performed according to the use case and customer requirements to create the virtual environment.

### 8. Work being done at NTT

NTT is using OpenStack to construct an environment for verifying IaaS platforms in various use cases (Fig. 5). That environment uses a Nicira network controller as a network controller attached to a Quantum network controller to implement OpenFlow-based overlay virtual networks. Three typical use cases that were verified in our laboratory are described below.

(1) Construction of a flexible virtual environment

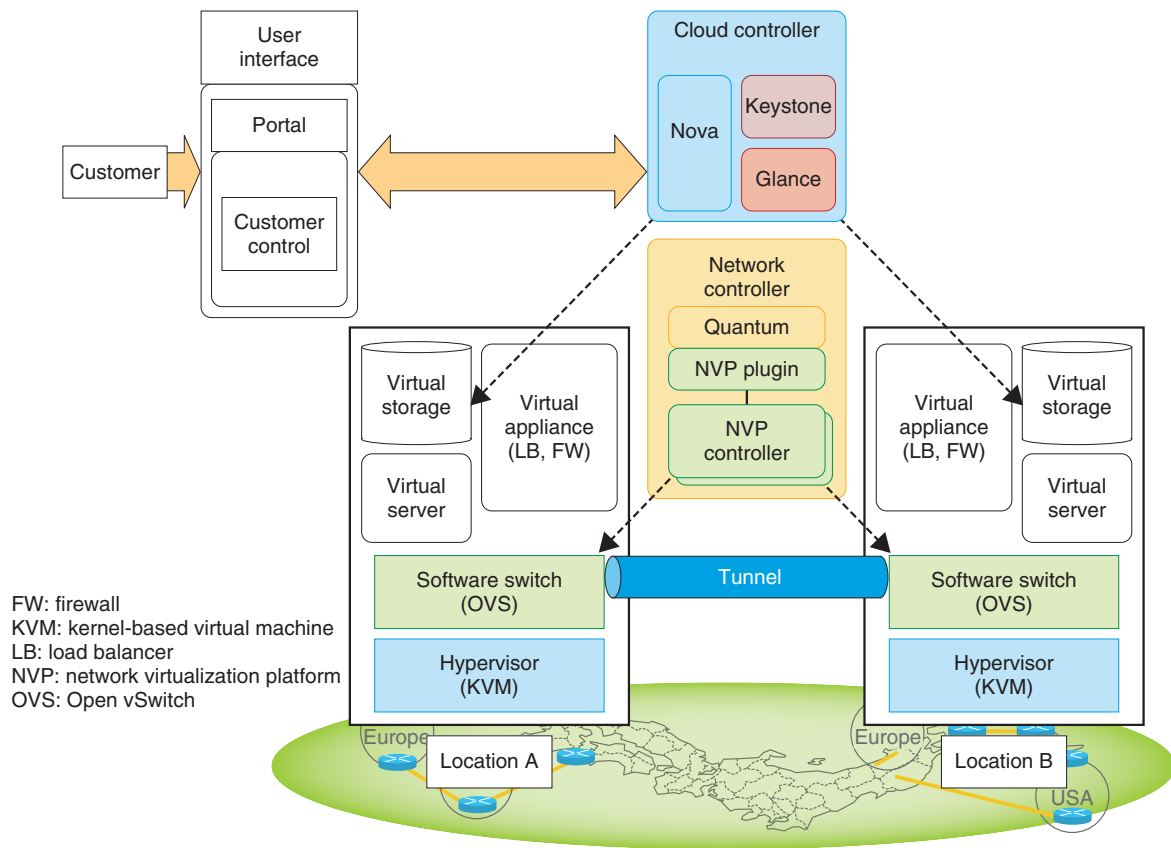


Fig. 5. NTT's test environment. Software switches distributed across Japan simulate a global virtual network.

Any number of virtual networks can be created for a customer. Combining them with virtual servers enables the same configuration flexibility as provided by on-premises systems, including cascaded multi-layer networks and demilitarized zones\*2.

Flexible system configuration enables effective use of existing software and other such assets as well as continuity in the use of firewalls and other network settings. It can therefore reduce the cost of system reconfiguration in the case of migration to a virtual environment.

Because a virtual environment can be constructed in the cloud while maintaining compatibility with existing systems, the building of hybrid clouds that interwork with existing systems and cloud services can further expand the range of uses of cloud services. For this use case, a multilayer virtual network is constructed for each customer, and multiple virtual servers are placed on each network and assigned arbitrary

IP (Internet protocol) addresses. We confirmed that communication with outside networks is possible in such a system (Fig. 6).

(2) Office migration

Office migration is a use case that assumes the need to temporarily move a virtual environment to a different location to maintain business continuity during a disaster, pandemic, or other anomalous situation or to reduce the cost of electrical power.

With the abovementioned overlay virtual network, a tunnel can even be constructed between remote software switches if IP connectivity is available. It is thus possible to construct a virtual network that involves a remote datacenter and create a single network environment that spans multiple locations. Doing so enables on-demand migration of virtual servers between datacenters without the need to change physical network settings. Previously, this would have required constructing individual Layer-2 VPNs and changing network settings within the datacenter.

NTT has constructed an environment in which

\*2 Demilitarized zone (DMZ): A network segment that lies between trusted networks and external networks.

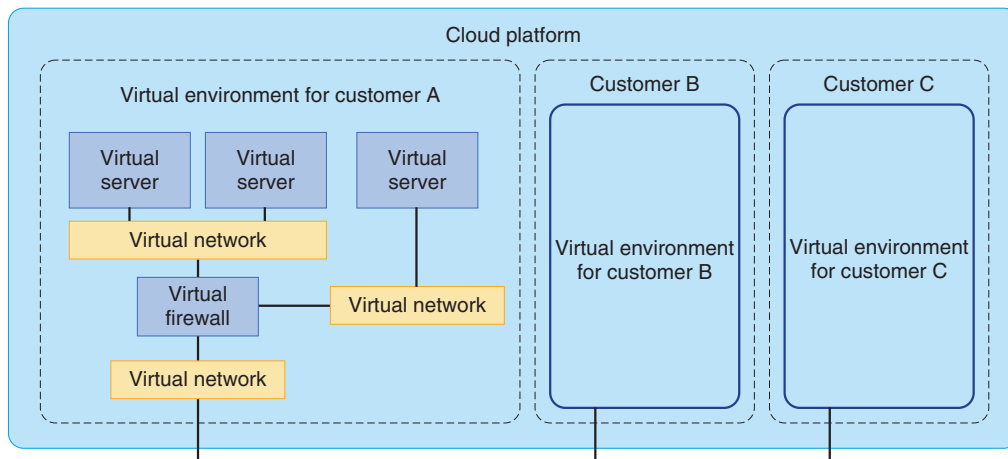
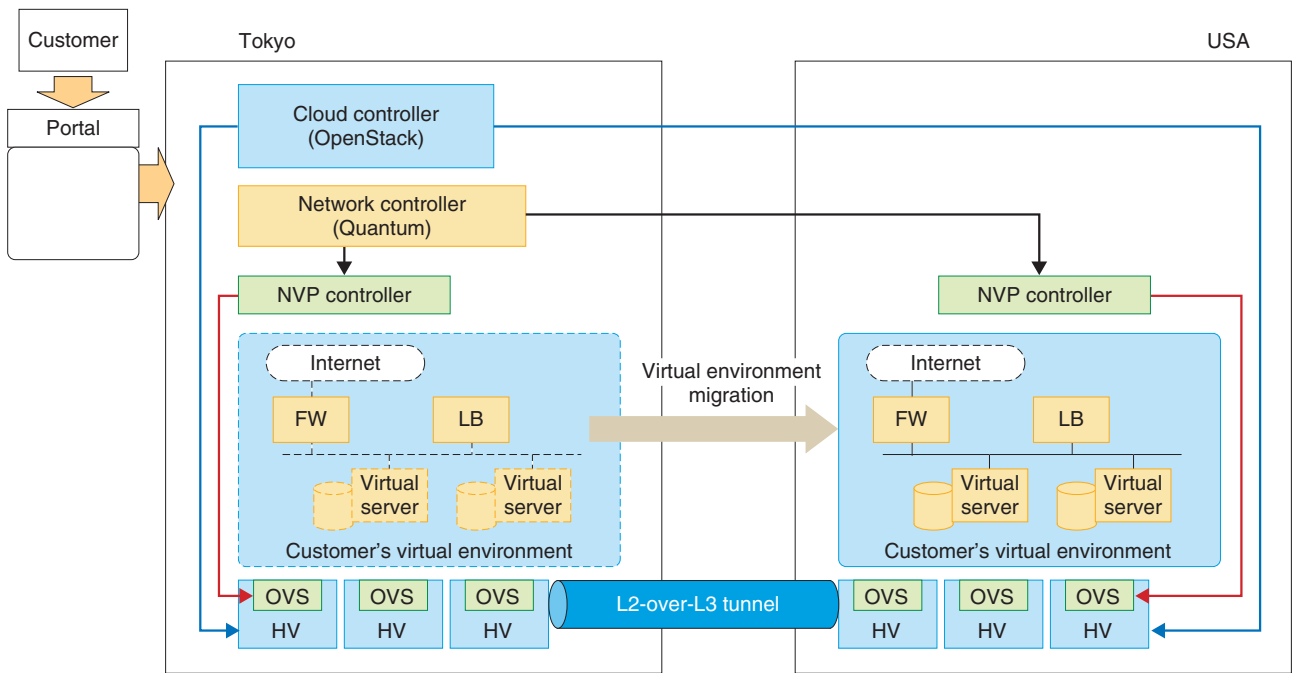


Fig. 6. Flexible configuration of virtual environment.



HV: hypervisor  
 L2: Layer 2  
 L3: Layer 3

Fig. 7. Virtual environment migration.

network controllers can be arranged to construct virtual networks between arbitrary datacenters and in which virtual servers can be freely placed and moved to enable unified control of multiple datacenters. We have confirmed that in this use case the customer can construct a virtual environment in any datacenter and

move the entire environment to a different location as required (Fig. 7).

We demonstrated these results jointly with NTT Communications at the Open Networking Summit held in April 2012. The demonstration system was constructed at three actual locations in Japan, the

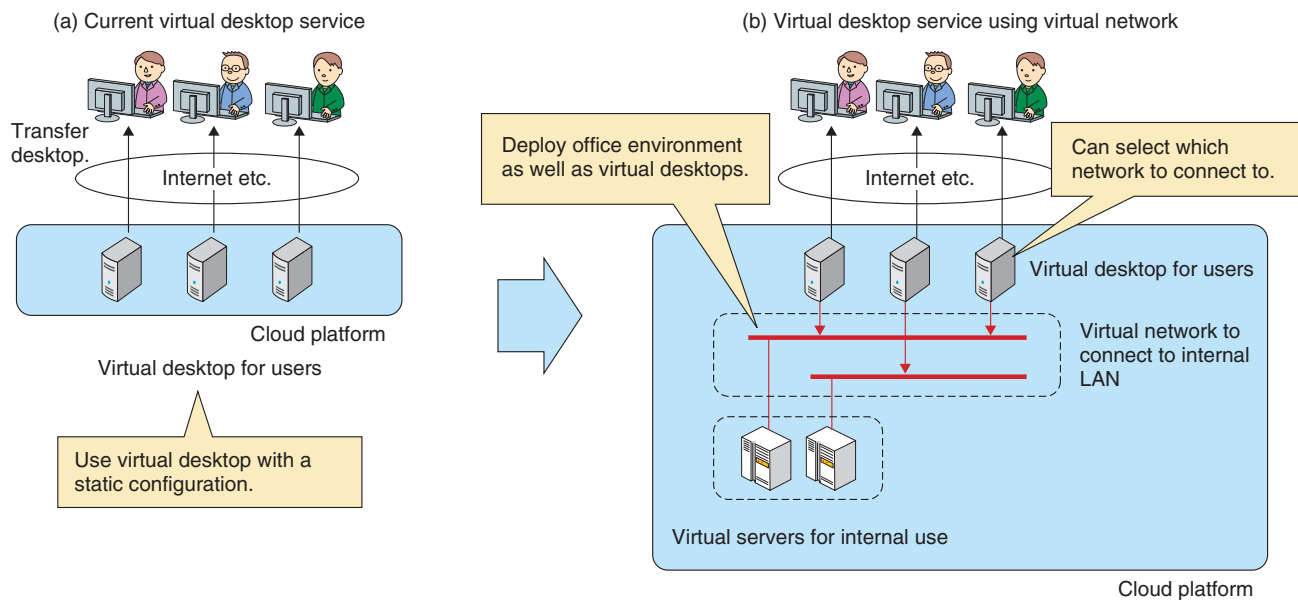


Fig. 8. Flexible virtual office environment.

USA, and Europe and migration of virtual environments among datacenters in different countries was verified.

### (3) Flexible virtual office environment

In the flexible virtual office environment use case, an IaaS platform is applied as an in-house system to provide a virtual desktop environment to company employees. A virtual desktop is a service that provides remote access to the desktop of a virtual server, thus presenting the same working environment regardless of the terminal being used.

Combining virtual network technology and a virtual desktop environment allows dynamic network construction and setup, making it possible to switch the virtual desktop to a different network without performing any extra configuration in the physical network. In this use case, the virtual-desktop users can switch networks on-demand according to their assigned tasks. We are now testing the basic operation of this use case and considering the application of OpenStack (Fig. 8).

## 9. Concluding remarks

We described the use of OpenStack, an OSS IaaS platform, in the construction of IaaS platforms. Specifically, we have verified use cases that involve interworking with a network controller and we presented examples of application to cloud services. For that task, we focused on the network, which is a core competence of the NTT Group. In the two years or so since the beginning of the OpenStack project, various proposals have been made by the active OpenStack community. We intend to continue to cooperate closely with that community to improve OpenStack functionality and raise the quality to the level required for commercial use.

## References

- [1] OpenStack. <http://www.openstack.org>
- [2] H. Kitazume, T. Koyama, T. Kishi, and T. Inoue, "Network Virtualization Technology for Cloud Services," NTT Technical Review, Vol. 9, No. 12, 2011. <https://www.ntt-review.jp/archive/ntttechnical.php?contents=ntr201112fa4.html>



**Shintaro Mizuno**

Senior Research Engineer, Cloud System SE Project, NTT Software Innovation Center.

He received the B.E. and M.E. degrees from Tokyo Institute of Technology in 1995 and 1997, respectively. Since joining NTT in 1997, he has been working on authentication systems, secure communication systems, and information security. His current research interests include cloud management system and software defined networks.



**Koji Iida**

Senior Research Engineer, Cloud System SE Project, NTT Software Innovation Center.

He received the B.E. and M.Sc. degrees from Keio University, Kanagawa, in 1993 and 1995, respectively. He joined NTT Information Platform Laboratories in 1995 and studied enterprise communication middleware and distributed object technologies. He moved to NTT Information Sharing Platform Laboratories in 2007 and researched identity management technology and cloud computing technology. As a result of organizational changes in July 2012, he is now in NTT Software Innovation Center.



**Hiroshi Sakai**

Senior Research Engineer, Supervisor, Cloud System SE Project, NTT Software Innovation Center.

Since joining NTT in April 1993, he has contributed to the development of network service systems including intelligent networks and the Next Generation Network (NGN). His current research interests include an OSS-based cloud management system and software defined networks.



**Takaaki Koyama**

Senior Research Engineer, Cloud System SE Project, NTT Software Innovation Center.

He received the B.A. and M.M.G. degrees in media and governance from Keio University, Tokyo, in 1994 and 1996, respectively. He joined NTT Software Laboratories in 1996 and studied software CALS. Since 1999, he has been studying GMN-CL, which is a kind of IP-VPN technology, and developing some network equipment. Recently, his research interests have extended to enterprise cloud network systems. He is a member of the Information Processing Society of Japan.



**Daigoro Yokozeki**

Senior Research Engineer, Cloud System SE Project, NTT Software Innovation Center.

He received the B.E. and M.Sc. degrees from Waseda University, Tokyo, in 1996 and 1998, respectively. He joined NTT Information Platform Laboratories in 1998 and studied web middleware and distributed database technologies. He moved to NTT Software Innovation Center in 2012 and is currently researching cloud computing and virtual networking technologies.