Application of Server Virtualization Technology to Communication Services

● Masaaki Muraai ● Tatsuhiro Furuya ● Takashi Imai ● Shoichi Kimura

Following the Great East Japan Earthquake of 2011, the volume of traffic on the mobile communications network spiked to about 50 to 60 times the normal busy-hour volume, thereby congesting the system and making it difficult to provide users with stable services. To provide stable services even during times of congestion caused by a disaster, Fujitsu is researching and developing the use of virtualization technology for flexibly and quickly allocating machine resources to communication services. This control technology is expected to improve the communication rate of priority services by decreasing the allocation of resources to services deemed low priority at the time of a disaster such as video and rich content delivery while increasing the allocation of resources to high-priority services such as voice calls and mail, which typically would be for obtaining disaster-related information and for checking on the well-being of people in the disaster area. This paper introduces Fujitsu's approach to establishing a server virtualization infrastructure for communication services.

1. Introduction

The Great East Japan Earthquake of 2011 was accompanied by an explosive jump in the volume of voice calls from mobile phones (to about 50 to 60 times the normal busy-hour volume). This spike in calls generated unprecedented large-scale, long-lasting congestion over a wide area. It is essential that stable communication services be provided even at the time of such a sudden jump in demand, so improvements must be made.[1)]

As part of a project overseen by the Ministry of Internal Affairs and Communications (MIC) of Japan titled "Experimental challenges for dynamic virtualized networking resource control over an evolved mobile core network—a new approach to reducing massive traffic congestion after a devastating disaster," Fujitsu is researching and developing control technology using virtualization techniques for flexibly and promptly allocating machine resources to communication services. This control technology is expected to improve the communication rate of priority services by adjusting the way in which resources are allocated. Specifically, it will decrease the allocation of resources to services deemed low priority at the time of a disaster such as video and rich content delivery while increasing the allocation of resources to high-priority services such as voice calls and mail, which typically would be for obtaining disaster-related information and for checking on the well-being of relatives, friends, and colleagues.

This paper introduces Fujitsu's approach to establishing a server virtualization infrastructure for communication services.

2. IaaS features

Infrastructure as a service (IaaS) based on virtualization technology is generally used as a means of managing and operating virtual-machine (VM) resources. It is a platform that supports infrastructure management and operation by unifying the pooling of VM resources and providing service providers with tools for simplifying infrastructure design and for increasing or decreasing VM resources in an on-demand manner. It also supplies functions for visualizing the progress of virtual-environment construction and post-construction operating conditions (**Figure 1**). The following summarizes the features of IaaS.[2)]

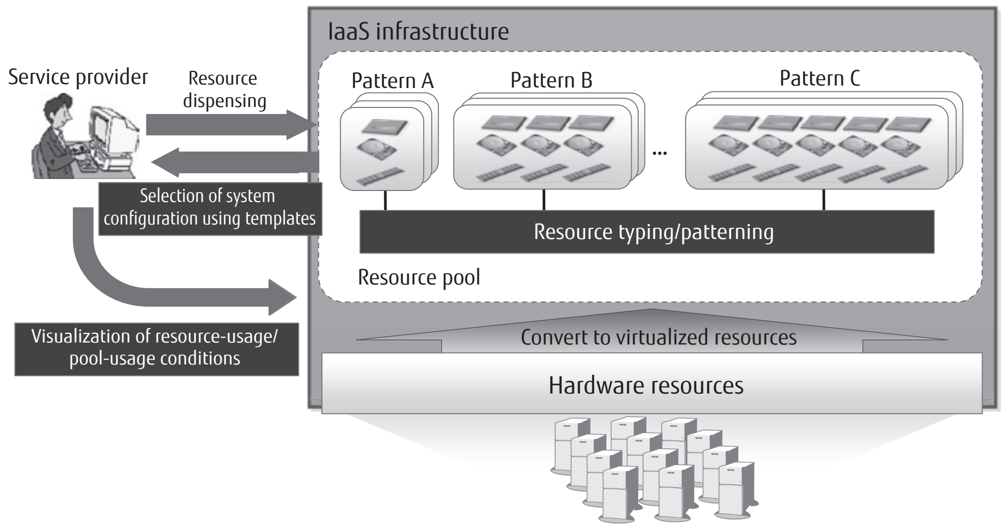

Figure 1
IaaS overview.

1) Simplification of infrastructure design

The user can select the configuration that best meets objectives by choosing an appropriate template. This has the effect of easing the load on the user by shortening the time required for designing an infrastructure and of preventing errors in defining and setting parameters.

2) On-demand dispensing of resources

This feature shortens the preparation time leading up to the actual deployment of resources by preparing and simplifying VM resources (CPUs, memory, disks) in the form of convenient resource patterns.

3) Visualization of resource-usage conditions

The user can check resource-usage conditions (e.g., number/amount of CPUs/memory/disks used) and resource-pool usage conditions in a visual manner via a dashboard screen or other means.

3. Requirements when applying IaaS to communication services

The aim of IaaS is to simplify system construction and utilize physical resources more efficiently through virtualization technology so that services can be provided at low cost. Communication services, however, make up a vital infrastructure supporting social and economic activities and must therefore provide a high level of service quality and availability beyond that of ordinary Web services. The following summarizes the requirements related to the application of IaaS to communication services.

1) VM performance guarantee

IaaS improves the usage efficiency of physical-machine resources by having multiple VM share those physical resources. A typical example is the sharing of CPU resources. When a number of VMs are sharing a CPU, however, the occurrence of a load greater than expected in one VM can deprive other VMs of a needed resource, thereby degrading VM performance.

A countermeasure to this problem is to specify the ratio of a CPU core that can be shared and to control the duration of CPU allocation to a VM in accordance with that ratio. This scheme, however, can generate overhead due to task switching between VMs, preventing the expected performance of each VM from being achieved.

At the same time, there is a need to guarantee high quality in communication services (packet transfer delay: 70 ms max.; packet loss ratio: 0.1% max.).[3] To this end, the resource dispensing process must be able to prevent overhead associated with CPU task switching caused by sharing and resource competition with other VMs to guarantee VM performance.

2) Resource dispensing that improves availability

To provide infrastructure to an indefinite number of users at low cost, IaaS dispenses resources giving priority to resource usage efficiency. Accordingly, a physical machine that allocates VMs must have many free resources available for use.

Communication services, however, are required to have a high availability factor of 99.9999% on average (corresponding to an annual down time of 30 s). To achieve this, the system must not include any single points of failure with respect to various types of faults. For example, when building a system consisting of multiple VMs, consideration must be given to the correspondence between physical machines and VMs, as in creating a redundant configuration that can accommodate faults in power supplies, physical servers, chassis, etc.

3) Support of uninterruptible maintenance and operations

A system that provides communication services must run without interruption on a 24/365 basis. This means that service provision must continue even when hardware is being replaced or maintenance work is being performed. Hardware-replacement or maintenance work to date has assumed the deployment of a redundant configuration that enables a server targeted for maintenance to be simply separated from the system. This, however, results in single-system operation during the time that maintenance is being performed, causing a temporary drop in fault resistance.

Hardware-replacement and maintenance work in IaaS is performed after moving operations to another physical server while guaranteeing the preservation of memory on VMs through a "live migration" process. This guarantee is achieved by periodically transferring difference information from the memory on the source VM to the memory on the destination VM during the migration process. The basic design of communication services, however, is to perform processing "on memory" to achieve high real-time characteristics. Furthermore, the processing performance required of one server used for communication services is about 100 000 BHCA,[note)] which calls for the processing of about 300 signals per second and the frequent updating of memory. Thus, in an ordinary live migration process, there is a high possibility that memory transfer will not complete within a fixed amount of time and that live migration will terminate abnormally.

For the above reasons, a scheme other than live migration is needed for performing hardware-replacement and maintenance work.

4. Fujitsu's approach to a communication SaaS infrastructure

A software as a service (SaaS) infrastructure for communication purposes provides functions for managing and operating the resources required for achieving the high availability, reliability, and real-time characteristics demanded of communication service based on IaaS.

A communication SaaS infrastructure can be achieved through two functional components:

1) Management of resource requirements for communication services
2) Management of resources-in-operation information for dispensing resources applicable to specified requirements.

An overview of a communication SaaS infrastructure is shown in **Figure 2**. The resource coordinator indicated by (1) manages resource requirements for communication services. When a system for providing a communication service needs to be constructed, this functional component determines which resource requirements fit the characteristics of that communication service and issues a request for resources to the resource pool manager. The resource pool manager indicated by (2) receives requests for resources from the resource coordinator, extracts the resources fitting the specified requirements from the resource pool, and dispenses those resources.

In response to the "requirements when applying IaaS to communication services" described in the previous section, this communication SaaS infrastructure has the following features for satisfying the requirements of applying virtualization technology to communication services.

1) Exclusive resource allocation

To dispense resources in a manner that guarantees the performance of communication services, the communication SaaS infrastructure is capable of exclusively allocating a physical CPU core and network interface card (NIC) to a VM. To put it another way, the resource pool manager manages the state of physical resource allocation and dispenses the appropriate resources in accordance with an "exclusive" or "share" allocation specified by the resource coordinator.

note) Initialism for "busy-hour call attempts," which is the number of times that calls are attempted during the one hour period in a day having the most network congestion.

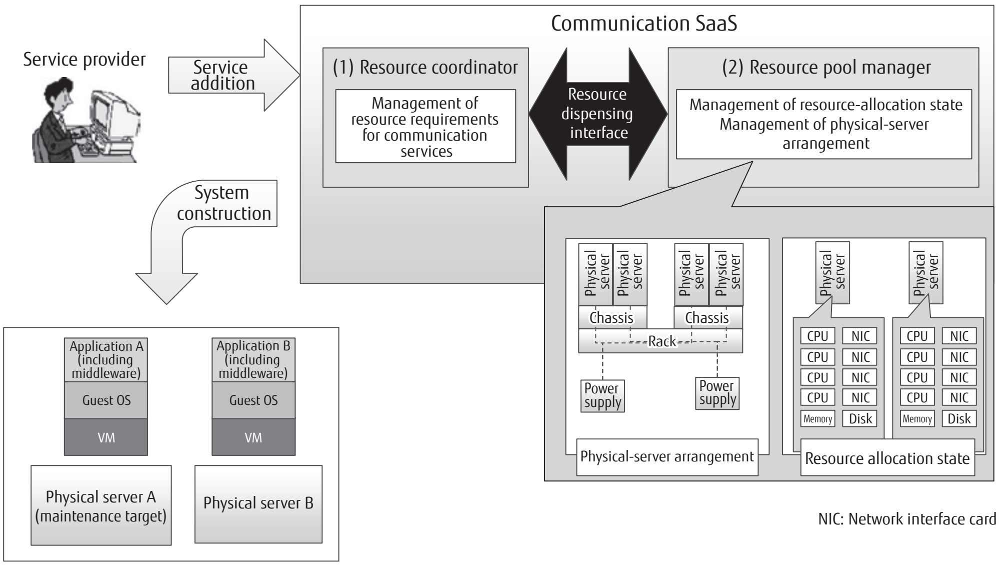

Figure 2
Communication SaaS overview.

Allocating physical resources to VMs in this way makes allocation granularity (unit of segmentation) finer and eliminates resource competition between VMs, making it possible to achieve expected processing performance. Although making allocation granularity finer increases the amount of communication between resources, this does not pose a problem because each communication service node has relatively high processing performance.

2) Arrangement of VMs to improve availability

The communication SaaS infrastructure is also capable of specifying the configuration of physical servers and creating VMs in accordance with the characteristics of the target communication service. Specifically, the resource pool manager manages information on the configuration of physical servers (racks, chassis, etc.), the redundancy of power supplies, etc. and determines the VM arrangement that can provide high availability in accordance with the physical-resource requirements of VMs specified by the resource coordinator. For example, when VMs are arranged in a redundant configuration, a system with high tolerance to power-supply and hardware failures can be constructed by selecting separate power supplies and chassis for those VMs.

3) VM mobility function for communication services

To achieve VM mobility appropriate to communication services, the communication SaaS infrastructure provides an original VM mobility function without having to use live migration prepared by a VM hypervisor. The basic configuration of this VM mobility system is shown in **Figure 3**. The communication SaaS infrastructure provides a function for routing session control signals, so when a VM has to be moved and construction of the destination VM has been completed, session signals in progress continue to be routed to the source VM while signals for initiating new sessions are transferred to the destination VM. Routing signals in this way means that the source VM can be terminated once all sessions in progress at that machine have completed. This successfully completes the VM moving process.

5. Future issues in communication SaaS

The above approach to applying server virtualization technology targets a single hub or data center and does not take into account the complete destruction of a hub or data center and the facilities within or the complete exhaustion of resources at a hub or data center. However, if we were to enable the communication processing of a hub experiencing massive call congestion to be performed by borrowing the resources of another hub with processing capacity to spare, it should be possible to eliminate such call congestion more effectively and provide stable operation. For this reason, extending the communication SaaS infrastructure

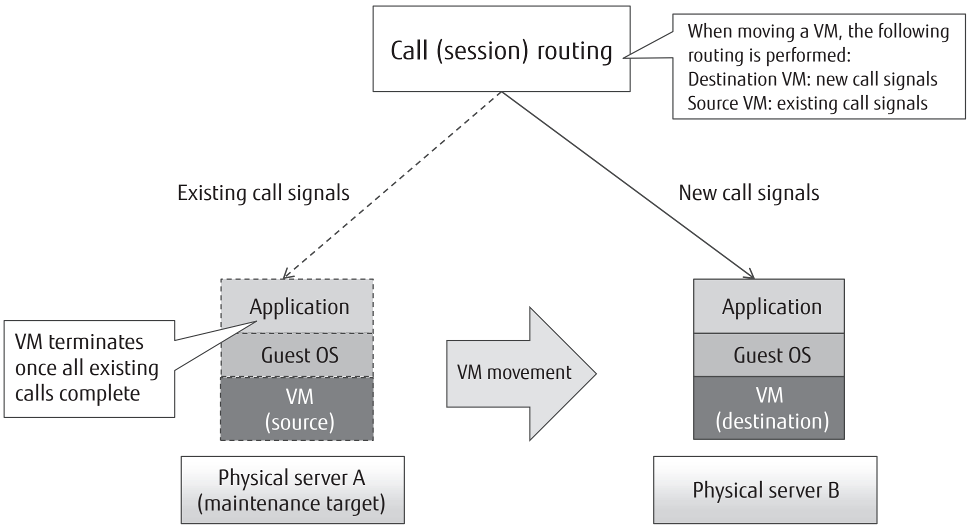

Figure 3
VM mobility system in communication SaaS.

presented here to one that spans multiple hubs is left for future study. To this end, the following issues must be addressed.

1) Selection of VM resources through inter-hub coordination

When borrowing resources for communication services across multiple hubs and building up services, it is inevitable that transmission distances will increase, thereby lengthening transmission delays. As a consequence, the performance of server resources at another hub will not necessarily be the same as those at the original hub despite being equivalent in terms of physical performance. In addition, the communication performance between resources (delay, bandwidth, packet loss rate, etc.) is always fluctuating in line with network and traffic conditions, which means that the level of resource performance originally expected can no longer be guaranteed.

2) High availability control through inter-hub coordination

Establishing redundant systems that span multiple hubs would be an effective measure against a large-scale disaster that severely damages a hub because it would accelerate resumption of operations and ensure business continuity. When switching between systems in a redundant configuration, it is important to consider exactly how accurately such switching can be performed in terms of system state and data integrity before and after the switching. Furthermore, to improve availability, the time required for synchronizing the data between the current system and the standby system must be shortened. However, the amount of data to be synchronized, the frequency of synchronization, the number of standby systems, and the location of standby systems all affect the processing overhead for performing synchronization and can have the undesired effect of actually lowering availability.

We plan to study these issues in the application of communication SaaS to multiple hubs.

6. Conclusion

Our personal experiences during the Great East Japan Earthquake made us keenly aware of the damage that can be inflicted by a major disaster. Communication services play a critical role as a lifeline during a calamity, so the stable provision of these services is essential. We consider the technology presented here to be an effective means of meeting this need. In future studies, we plan to establish a new operation model of communication services with an eye to practical application.

This technology was an outcome of research and development work performed under the MIC "Research and Development to Strengthen Disaster Resistance of the Information and Communications Network" project.

As part of these research and development efforts, a test bed emulating an actual mobile communications system is now under construction on the

Tohoku University campus. The plan is to use this test bed to evaluate the technology presented here and to uncover any problems in actual operations.

References

1) Ministry of Internal Affairs and Communications: 2011 White Paper: Information and Communications in Japan, Part 1, Information and Communications in the Aftermath of the Great East Japan Earthquake. *http://www.soumu.go.jp/johotsusintokei/whitepaper/eng/WP2011/part1.pdf*
2) M. Tajima et al.: PFU's Approach to Public Cloud Computing. *PFU TECHNICAL REVIEW*, Vol. 22, No. 2, pp. 1–5 (2011) (in Japanese).
3) T. Tsuchiya: Quality Standards and Quality Control in the NGN. *NTT Technical Review*, Vol. 7, No. 7, pp. 30–31 (2009). *https://www.ntt-review.jp/archive/ntttechnical.php?contents=ntr200907sf1.html*

Masaaki Muraai
Fujitsu Ltd.
Mr. Muraai is engaged in the development of server virtualization technology and cloud management software for communication services.

Takashi Imai
Fujitsu Ltd.
Mr. Imai is engaged in the development of server virtualization technology and software for interfacing with the virtualization infrastructure.

Tatsuhiro Furuya
Fujitsu Ltd.
Mr. Furuya is engaged in the development of server virtualization technology and cloud management software for communication services.

Shoichi Kimura
Fujitsu Ltd.
Mr. Kimura is engaged in the development of server virtualization technology and software for interfacing with the virtualization infrastructure.