Technology Reports Operational and Maintenance Efficiency Improvement OpenStack TCO Reduction

Large-scale OSS Migration Using Private Cloud and Virtualization Technologies

DOCOMO Technology, Inc. Solution Service Division Yuki Kishikawa Kentaro Ishii

Regarding OSS responsible for service and equipment monitoring, NTT DOCOMO studied facility renewals accompanying the end of the HW support used with these systems, and introduced OpenStack private cloud and virtualization technologies. This article describes issues with the conventional OSS, requirements for their renewal, the results of studying renewal methods, issues and solutions with renewal methods, function development to improve operational efficiency, effects of deployment, and future issues.

1. Introduction

To continuously provide stable mobile services, NTT DOCOMO has developed an Operations Support System (OSS)*1 that is responsible for monitoring services and equipment, and has introduced this system commercially.

As the end of support (EOL: End Of Life*2) of HardWare (HW) used with OSS approaches, we have continued to study facility renewals with a view to making improvements into the future. With this endeavor, developing systems consisting of many pieces of HW and efficiently developing SoftWare (SW) when HW and OS are upgraded have been central themes. This article discusses the construction of an OpenStack*3 private cloud*4 environment, and discusses studies and outcomes on issues and solutions with the construction and operation of OSS SW in virtual environments.

2. Conventional OSS Configurations, Issues and Renewal Requirements

1) HW Issues

Conventional OSS consists of thousands of general-purpose blade-type IA servers*5, storages and

*1 OSS: Enterprise operations support system. For communications operators, this can include some or all of fault management, configuration management, charging management, performance management, and security management for the networks and systems providing the services.

network equipment. Also, efficient disaster recovery*6 mechanisms for large-scale disasters [1] - [3] were introduced to ensure the continuity of OSS. **Figure 1** shows the conventional OSS configuration.

While it is important to ensure spare HW resources so that OSS can operate continually in emergencies such as disasters or network congestion*7, this lowers resource usage rate in times of normal operation, which in turn lowers the usage efficiency of facilities. Moreover, as an issue with conventional OSS facilities in terms of cost, these systems require facilities investment for initial deployment (CAPEX: CAPital EXpenditure*8) in addition to HW installation space rental, electricity, maintenance and other operating expenses (OPEX: OPerational EXpense*9).

2) SW Issues

In terms of SW, testing OSS OS when upgrading is problematic. For example, the CentOS*10 version currently used with OSS is not supported with the latest HW products, which means that when facilities are renewed with new HW, OS used by SW also have to be upgraded and tested to ensure proper operation. To date, we have periodically renewed facilities (including operations testing associated with OS upgrades). For this, it is also important to optimize development costs associated with OS upgrades by properly managing the SW level of degree of dependence on OS and localizing affected SW.

With the aim of solving the above issues, **Table 1** shows specific requirements for renewing facilities. We studied renewal methods that satisfy these

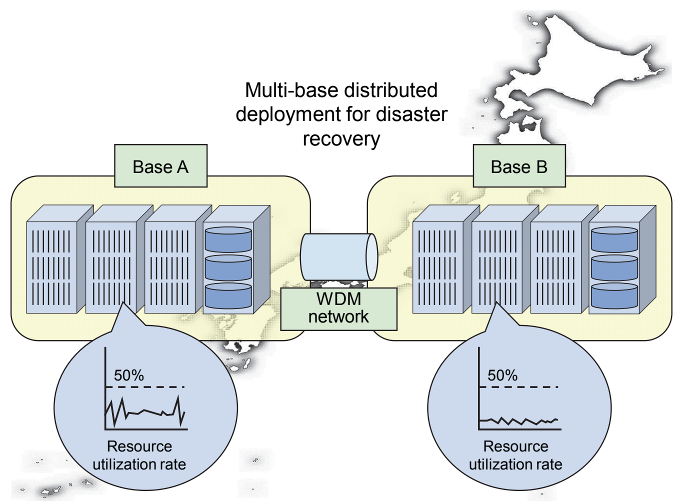

Figure 1 Conventional OSS configuration

*2 EOL: Refers to cessation of product manufacture and sales, cessation of support services for SW products, or cessation of provision of modification/upgrade programs for bug fixing and functional improvements.

*3 OpenStack: Cloud-infrastructure SW that uses server virtualization technology to run multiple virtual servers on a single physical server. It can allocate virtual servers to different cloud services used. OpenStack is open source SW.

*4 Private cloud: Refers to an in-house cloud system configured in a corporation or organization, and provided to various in-house divisions or group companies. In contrast, open cloud services that do not restrict their services to certain users are called “public cloud” services.

*5 IA server: A server equipped with an Intel microprocessor or an Intel compatible processor. Its internal structure is very similar to that of an ordinary PC, and it is less expensive than servers based on other types of microprocessor.

requirements.

3. Results of Study on Renewal Methods

To satisfy the requirements of facility renewal, it is necessary to improve the usage efficiency of facilities. Here, we made efforts to improve usage efficiency by implementing several applications with one HyperVisor (HV)*11.

1) Traditional Distributed Architectures

Conventionally, OSS applications have adopted distributed architecture which is made up of platform SW called Distributed Data Driven Architecture (D3A)*12 [4] and a group of distributed applications called multiple ELements (EL)*13 which operate on the platform. The D3A platform is responsible for transferring data between the distributed EL, and achieves simplification of EL functions. **Figure 2** shows the conventional OSS SW distributed configuration.

2) Adoption of HV-type Virtualization Technologies

Compared to conventional distributed architecture, more efficient use of facilities is enabled by using HV-type virtualization technologies and configuring and operating multiple Virtual Machines (VM)*14 on a single piece of HW.

Figure 3 shows the system configuration of the OSS SW using the HV-type virtualization technologies. With conventional D3A platform, it was not possible to implement more than two EL of the same type on one piece of HW, but because this

Table 1 Facility renewal requirements

Number	Requirements
Requirement 1	Efficiency with optimized HW resources usage
Requirement 2	Optimized SW development investments
Requirement 3	Efficient HW investments

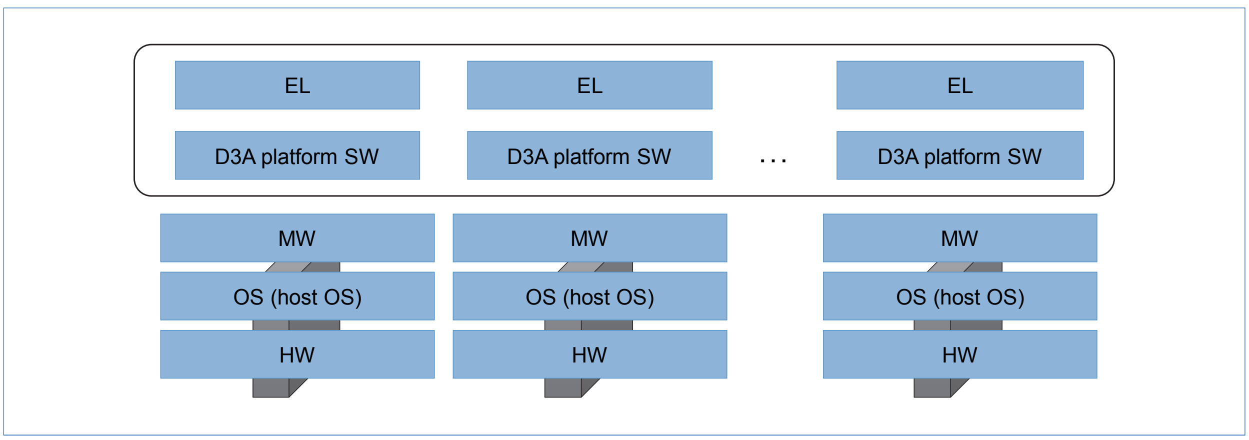

Figure 2 OSS application distributed configuration

*6 Disaster recovery: Repair and restoration of a system damaged by a natural disaster or other calamity. Also, preventive measures for minimizing damage.

*7 Congestion: Impediments to communications services due to communications requests being concentrated in a short period of time and exceeding the processing capabilities of the service control server.

*8 CAPEX: The amount of money expended on facility investments.

*9 OPEX: The amount of money expended for maintaining and operating facilities.

*10 CentOS: A free Linux distribution aiming for complete compatibility, based on the software source code included in Red Hat Enterprise Linux (see *16), but rebuilt without the Red Hat trademark and commercial packages.

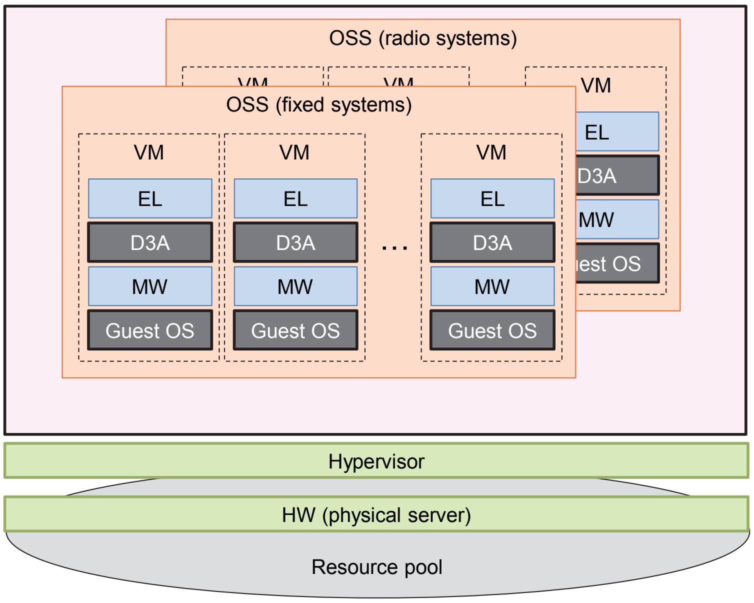

Figure 3　Configuration after implementation of HV-type virtualization

technology enables separate operating space with each VM, it's possible to implement several VMs with the same type of EL installed on the same piece of HW. Also, if there are changes to the OS or MiddleWare (MW)*15 due to HW upgrades, those changes are isolated and VMs are not affected as long as the OS is supported by HV. Basically, this means the effects of OS version upgrades can be avoided, and hence investment in SW development can be optimized for the long term. We adopted this method because the HV-type virtualization technologies can meet the full facility renewal requirements described in Table 1.

Furthermore, as HV, we adopted Linux®*16 standard free-to-use Kernel-based Virtual Machines (KVM)*17 for these facility renewals because we have a policy of building our OSS using royalty-free open source SW.

3) Adoption of the Private Cloud System

In terms of infrastructure, we have individually constructed our conventional OSS environment, but this time, we adopted a private cloud system using OpenStack to implement virtualization technologies for these renewals. In particular, to also use excess resources effectively, we can make company-wide operations more efficient by using the above system instead of building independent OSS environments or using external cloud services.

4. Challenges with Cloud Computing

1) Challenges with Using the Cloud

While conventional OSS SW was originally configured based on on-premises*18 physical server operations, we found new issues with virtualization

*11 HV: A virtual server technology that assigns and manages physical resources for SW to mount on virtual machines, and runs multiple virtual machines on physical resources.

*12 D3A: An architecture developed at NTT DOCOMO, which groups multiple IA servers to achieve high performance.

*13 EL: A basic structure of D3A, a mechanism that enables distributed operation on a number of servers by dividing SW into functional units. Because EL run on Java® VM, they are not dependent on particular HW vendors. Oracle and Java are registered trademarks of the Oracle Corporation and its subsidiaries and related companies in the United States and other countries.

*14 VM: A computer created in a virtual manner by SW.

*15 MW: Positioned between OS and SW, middleware is a collection of common basic functions and generic processing provided to all SW running on an OS.

and cloud systems with these renewals. For example, since EL with Active/Stand-by (ACT/SBY)*19 redundant configuration were separated in physical servers in the on-premises environment, it was extremely rare that both systems would go down due to HW failure, and hence the possibility that OSS services would be down was extremely low. However, in the virtual, cloud environment it's possible to implement the redundant ACT/SBY EL on the same HW, which means there is a risk of both systems failing if the HW fails. Also, with ACT/SBY switching, EL configured with a shared disk preserved data consistency with shared disk mount switching, but the current OpenStack does not always enable suitable unmount operations, which can lead to shared disk mount switching failures, thus requiring recovery time and lengthening the impact on business operations. In addition, appropriate control of VM placement is required for EL intended for the physical structure of conventional OSS, such as EL for multiple ACT configurations or EL implemented on the same HW to ensure performance. **Table 2** and **Figure 4** (a) describe specific issues.

2) Solutions

When using a cloud system, it is difficult to select HW or change cloud system settings. Therefore, solutions are required that entail changing SW specifications for virtualization. Fig. 4 (b) describes these solutions. As specific solutions to numbers 1 to 3 in Table 2, we adopted OpenStack's "availability zone*20" designation functions that enable placement of VM on specific HW and filter functions (Different Host Filter, Same Host Filter etc.) [5]. To preserve physical separation and redundancy, the Availability Zone function allows the user to specify the physical location of a VM such as a specific rack, floor or data center. Similarly, the filter functions enable VM launch rules to be set in detail. Normally these are specified manually, but with large-scale systems such as the DOCOMO OSS in which one system may contain 1,000 or more VMs, such settings would entail massive amounts of work and be prone to human error. For this reason, we implemented a virtualization controller for these technologies.

In number 4 in Table 2, we resolved the issue by changing storage configuration and revising the SW to implement data replication*21 between ACT/SBY.

Table 2 Challenges with cloud computing and solutions

Number	Issues to consider
1	EL for ACT/SBY configuration: To ensure redundancy, ACT-side EL and SBY-side EL (VM) are not implemented on the same HW
2	EL for multiple ACT configuration: To ensure redundancy, all ACT EL (VM) are not implemented on the same HW
3	Some EL correlated greatly with operations: To ensure processing performance, correlated EL (VM) are implemented on the same HW
4	EL on shared disks: Due to the OpenStack constraint that unmounts do not operate properly when a VM stops, shared disk systems are changed to individual disk systems

*16 Linux®: An open-source Unix-type OS that can be freely redistributed under GNU Public License (GPL). A registered trademark or trademark of Linus Torvalds in the United States and other countries.

*17 KVM: SW for achieving virtualization. KVM is open source SW.

*18 On-premises: Refers to an environment where HW that makes up a corporate system is possessed, operated and maintained by the company.

*19 ACT/SBY: A system configuration in which two servers perform the same function with one server in active mode (ACT) and the other in standby mode (SBY). Service interruptions are prevented by immediately continuing operations on the SBY server whenever a fault occurs on the ACT server. The SBY server is always kept in the same state as the ACT server during normal operations in preparation for switching.

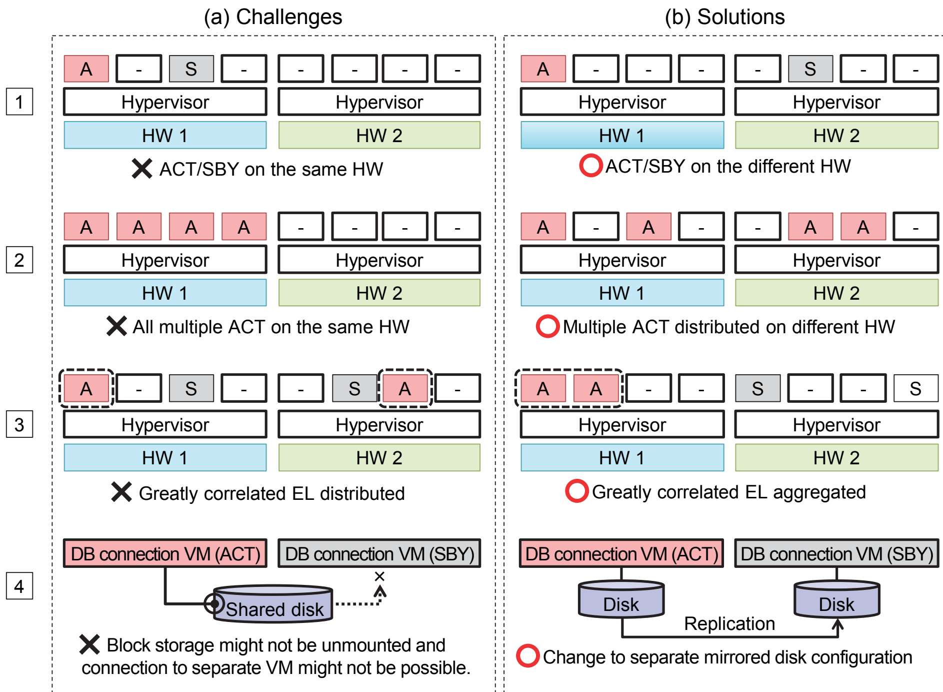

Figure 4 Challenges and solutions

This enabled migration to a virtualization and cloud environment while preserving the redundancy, performance and functions of conventional OSS SW.

5. Challenges with HV-type Virtualization Technologies and Their Solutions

Because HV-type virtualization technologies use SW to emulate HW, and because many VMs can share one HW resource, performance can degrade due to virtual layers (SW processing). For this reason, in particular, sufficient study and solutions on storage and network performance are required to solve the specific issue of performance degradation with HV-type virtualization technologies.

In these facility renewals, we confirmed that the I/O performance*22 of external storage in the OpenStack private cloud environment was lower than the conventional method. Hence, we avoided using external storage HW in the cloud environment and used higher performance internal flash storage. Although this completed solutions in terms of performance, issues remained with data persistency.

In the OpenStack virtual environment, data saved in external storage is basically persistent, and remains on the disk even if a VM stops. Hence, on restarting the VM, and connecting the external

*20 Availability zone: In OpenStack, there are areas that are geographically separate such as data centers that are referred to as “regions.” An “availability zone” refers to an independent location in a rack or power system within a region.

*21 Replication: A data-copy process between file systems performed on a one-to-one basis. Defines pairs of replication sets between file systems. Copies data sets not stored in the transfer destination and copies differences only if data sets have already been stored.

*22 I/O performance: Refers to the input/output performance of data and signal exchange between the CPU or memory and storage or disks.

storage to the VM, the stored data can be reused. However, data stored in one type of OpenStack storage, ephemeral storage*23, is not persistent, and is deleted if the VM is deleted. Hence, with OSS, because some data such as equipment monitoring and controlling logs and traffic data must be retained for a certain period, we upgraded SW data replication functions to ensure the persistency of such data. Specifically, we took measures to replicate ACT-side EL data in the SBY-side EL internal disk with ACT/SBY-related EL, as described in **Figure 5**.

6. Improving Operational and Maintenance Efficiency

When increasing equipment with operations in physical environments, it takes two to three months to procure and build HW, and install SW, depending on the scale. However, with virtualization technologies, HW required to build systems is abstracted, therefore virtual servers, networks and storage can be configured quickly with SW controls. Also, when testing requires changes to the system, the functions described in **Table 3** can be achieved by combining build and fault because virtual environments can also be stopped rapidly.

1) Developing Three Functions

The introduction of virtualization technologies enables not only improved operational efficiency due to improved HW usage efficiency and better use of excess resources, but also further improves business operational efficiency with improved maintenance efficiency and automation of operations. In view of this, we developed the following three functions with these facility renewals.

- Instantiation

  Instantiation means generating a VM. This refers to the operations from installing a VM on a physical machine, making virtual network settings, launching the VM and launching SW through other settings to achieve a usable state. We have accomplished this with a single touch.

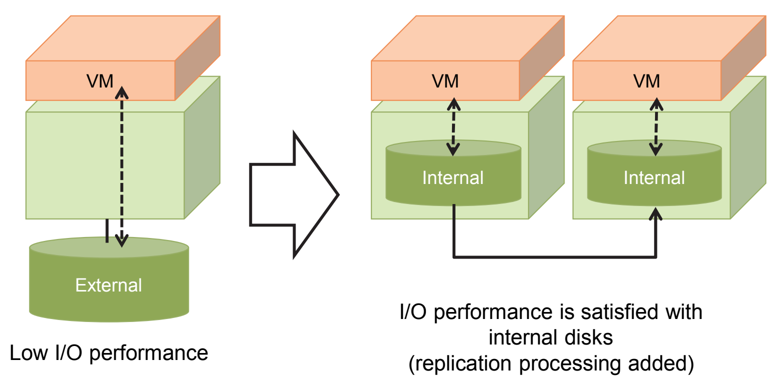

Figure 5　Persistent data handling with performance improvements

*23 Ephemeral storage: Volatile memory in OpenStack. When an instance is deleted, data in storage is also deleted.

Table 3 Developing functions for more efficiency

Number	Issues to consider
Function 1	One-touch VM creation (instantiation)
Function 2	One-touch redundancy recovery (healing)
Function 3	One-touch equipment expansion (scale out)

- Healing

  Healing means recovery of redundant VM and SW configurations. If one side of a VM with a redundant ACT/SBY configuration is lost, ACT/SBY switching takes place in the application layer and operations can continue, but the SW only runs on one side. In this situation, if a failure occurs and a VM stops leading to a system outage, recovery of redundancies becomes urgent. To counter this issue, we have achieved a one-touch redundant configuration recovery procedure.
- Scale out

  Scale out means expanding VMs. While it was possible to improve performance by increasing VMs with the same type of EL installed if the performance of operational VMs came under pressure, it was difficult to handle HW expansion for sudden performance concerns because this required HW expansion planning and works. Thus, with this development, we have achieved a one-touch facility expansion procedure.

2) Introduction of Virtualization Controllers

Figure 6 describes OSS architecture to further increase efficiency. A virtualization controller is responsible for executing and managing the functions described in Table 3, and achieves functionality interlocked with the OpenStack/D3A platform. The OpenStack and D3A functionality enables the achievement of instantiation, healing and scale out with manual work. However, with the introduction of virtualization and the private cloud, we have also implemented an automated system with execution and management functions to maximize simplicity and speed with maintenance.

This controller basically launches and stops SW and VMs. Conventional OSS consists of a combination of D3A platform SW and SW (EL), where the D3A platform launches and stops SW and manages system status. For this reason, it's not necessary to associate virtual controllers with each SW (EL), because SW and VM lifecycle management can be achieved with association to the D3A platform only. Thus, this architecture contributes to lower development costs because impacts are localized to existing SW.

In addition, it is also possible to entirely automate functions and eliminate the need for human intervention. While automation leads to a decrease in human errors, there are additional development costs associated with it. Hence, in this development, we estimated the cost of full automation and OPEX reduction effect, and chose one-touch automation operations. Generally, these developments comply with European Telecommunications Standards Institute (ETSI)*24 Network Functions Virtualisation

*24 **ETSI**: European Telecommunications Standards Institute. A European standardization body engaged in the standardization of telecommunications technologies. Headquartered in Sophia Antipolis, France.

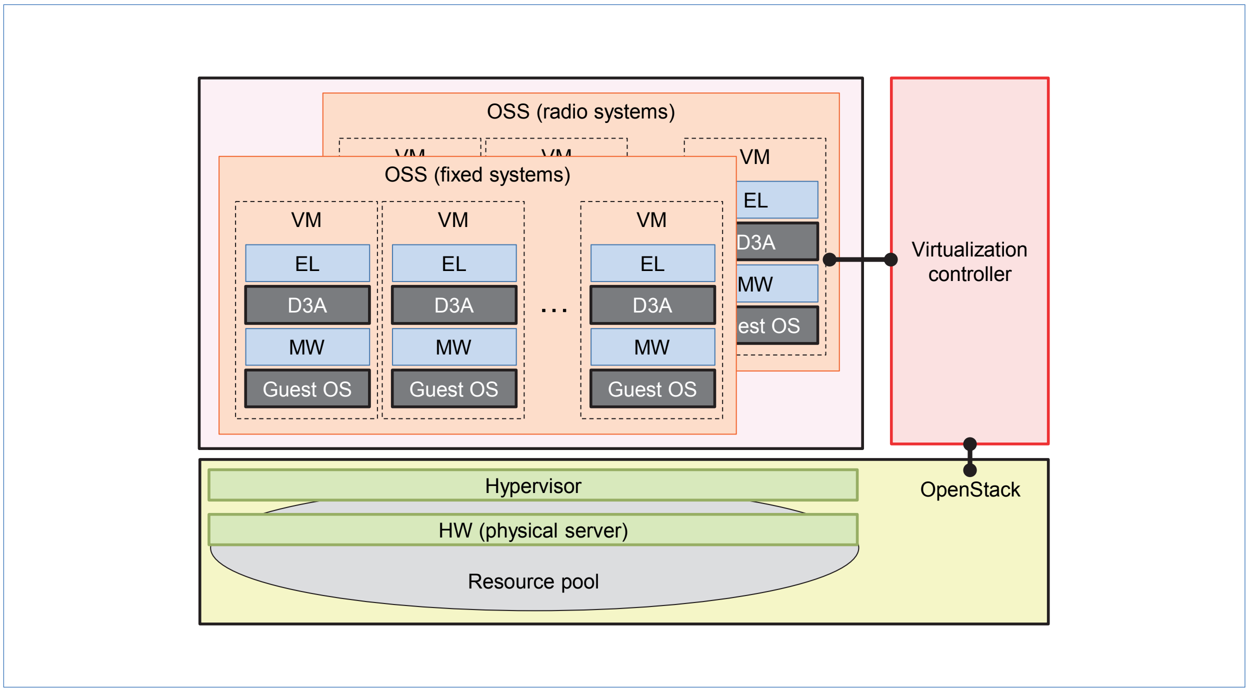

Figure 6 Architecture for achieving further efficiency

(NFV)*25 standards [6], and the virtualization controller is equivalent to Virtual Network Function Manager (VNF Manager)*26 included in Management and Orchestration (MANO)*27 in NFV standards.

7. Effects of Cloud Implementation and Future Challenges

Figure 7 describes the Total Cost of Ownership (TCO)*28 reductions calculated for the implementation of the OpenStack private cloud and virtualization technologies. This figure compares the TCO over seven years including operational costs of the current physical environment (Fig. 7 (a)) with the implementation of OpenStack private cloud and virtualization technologies (Fig. 7 (b)).

In terms of CAPEX, we achieved massive reductions in equipment investments. After renewing, optimized HW facility investments through improved HW performance and efficient uses of resources by virtualizing were largely responsible for reductions in TCO.

In terms of OPEX, both maintenance and electricity expenses were reduced by the above-mentioned optimization of HW facilities investments. Also, as combining the functions shown in Table 3 with the virtualization technologies enables planned execution of maintenance support, a future issue is the study of operational methods to maximize effectiveness with the aim of further increasing efficiency.

8. Conclusion

This article has presented studies on the im-

*25 NFV: A technology that uses virtualization technologies to implement processing for communications functionality in SW running on general-purpose HW.

*26 VNF Manager: The system that performs VNF control operations such as launching and termination as VNF lifecycle control.

*27 MANO: A mechanism that provides VNF management functions and orchestration functions for HW and SW resources in virtual environments.

*28 TCO: The total expenses incurred with initial deployment and operations management of a system.

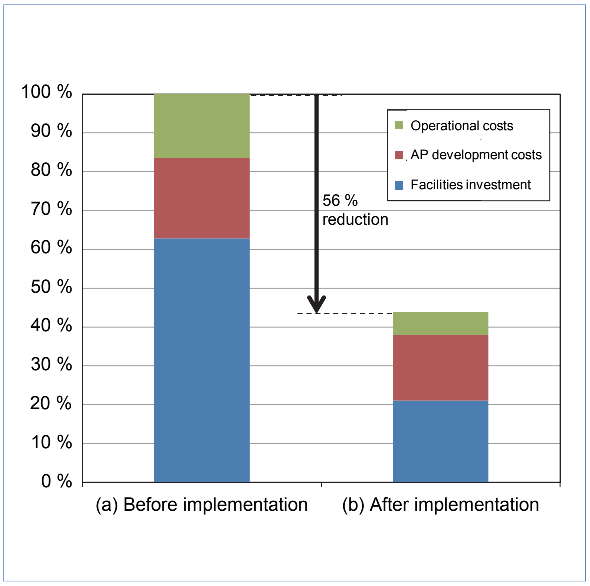

Figure 7 TCO cost reductions with cloud and virtualization technologies implementation

plementation of OpenStack private cloud and virtualization technologies in facility renewals of conventional OSS, and has shown the potential to reduce TCO and use excess resources effectively. This was NTT DOCOMO's first attempt to migrate large-scale OSS to a cloud system. As a result of various measures to achieve the desired performance and redundancy by moving non-cloud native*29 OSS SW to cloud systems, we achieved commercial implementation. While new operational processes are required with the implementation of virtualization technologies, these developments are highly advantageous because they can achieve efficient, agile and flexible operations. We intend to continue studying ways to further improve business efficiency by maximizing the effects of implementation of cloud computing and virtualization technologies.

REFERENCES

[1] K. Takahashi: "OSS configuration techniques to economically realize disaster recovery," The Institute of Electronics, Information and Communication Engineers, Technical report, NS, Network System 112 (392), pp.61-66, Jan. 2013 (In Japanese).

[2] K. Takahashi, M. Onuki: "Practical application of large-scale distributed operation support systems," Journal of The Institute of Electronics, Information and Communication Engineers, B, Communications J92-B (7), pp.970-980, Jul. 2009 (In Japanese).

[3] H. Kitahama et al.: "Improving IP-based OSS Reliability During Large-scale Disasters," NTT DOCOMO Technical Journal, Vol.18, No.2, pp.72-77, Oct. 2016.

[4] Y. Kato, S. Tamuki: "Introduction to containers learned with LXC - Technology to realize lightweight virtualization environment," Gihyo.jp ADMINISTRATOR STAGE series, Gijutsu-Hyohron, May-Jun. 2014 (In Japanese).
http://gihyo.jp/admin/serial/01/linux_containers

[5] OpenStack Administrator Guide: "Filter Scheduler."
http://docs.openstack.org/developer/nova/filter_scheduler.html

[6] ETSI: "ETSI Network Functions Virtualization completes first phase of work," Jan. 2015.
http://www.etsi.org/news-events/news/864-2015-01-pressetsi-network-functions-virtualisation-completes-first-phase-ofwork

*29 Cloud native: Refers to systems and services designed for configuration and operation on the cloud, rather than on-premises.