

## Linux and KVM Activities Providing Foundational Support for Use of OSS

*Kenji Yoshida, Kumiko Goto, Toshiaki Makita, Toshio Nakazawa, and Seiji Kihara*

### Abstract

The NTT Open Source Software Center is working on various aspects of Linux and Kernel-based Virtual Machine (KVM) to support the use of open source software (OSS) within the NTT Group. In this article, we introduce some of these activities, which include building expertise in order to use OSS with confidence, improving the quality of OSS products by cooperating with development communities, and promoting OSS virtualization technology.

*Keywords: Linux support, OSS community activity, server virtualization*

### 1. Introduction

When speaking of software, we tend to focus on applications, which are directly visible to users. However, there are other elements that act as “unsung heroes” operating behind the scenes. Operating systems (OSs) and virtualization software are examples of this type of software that operates closest to the hardware.

An OS provides the infrastructure for running various other types of software such as applications. It also provides functions such as process management, memory management, file systems, and networking. The NTT Open Source Software (OSS) Center provides technical support to NTT Group companies for Red Hat Enterprise Linux (RHEL), which is an industry standard Linux OS distribution provided by Red Hat, Inc. and used widely within our group companies. We also consult separately regarding other distributions such as CentOS when requested.

To improve the quality of this support, we investigate and test new versions of RHEL when they are released. In doing so, we actively contribute patches to the OSS development community when we find bugs in order to improve software quality.

Recently, server virtualization technology has been focused on as a way of reducing total cost of ownership (TCO) and achieving more flexible system

operations by integrating servers and systems. Kernel-based Virtual Machine (KVM) is an OSS hypervisor, the standard virtualization technology in Linux, that is included in RHEL. As such, we are actively working on it.

Initially, the operations management features of KVM were limited, so it was not used much within our group. However, Red Hat Enterprise Virtualization (RHEV) has since been introduced, and it provides more substantial operations management features. We are therefore working with RHEV to promote the use of KVM (Fig. 1).

### 2. Testing new RHEL versions and providing feedback

#### 2.1 Investigation and testing of new RHEL releases

For each major release<sup>\*1</sup> of RHEL, Red Hat offers minor releases<sup>\*1</sup> to fix bugs and security vulnerabilities approximately every six months. Red Hat conducts testing before releasing them so they can be used safely for enterprise services.

However, system administrators tend to be conservative when considering whether applying software updates is necessary because doing so can affect the

\*1 The latest version is RHEL 6.5; ‘6’ indicates the major release and ‘5’ the minor release.

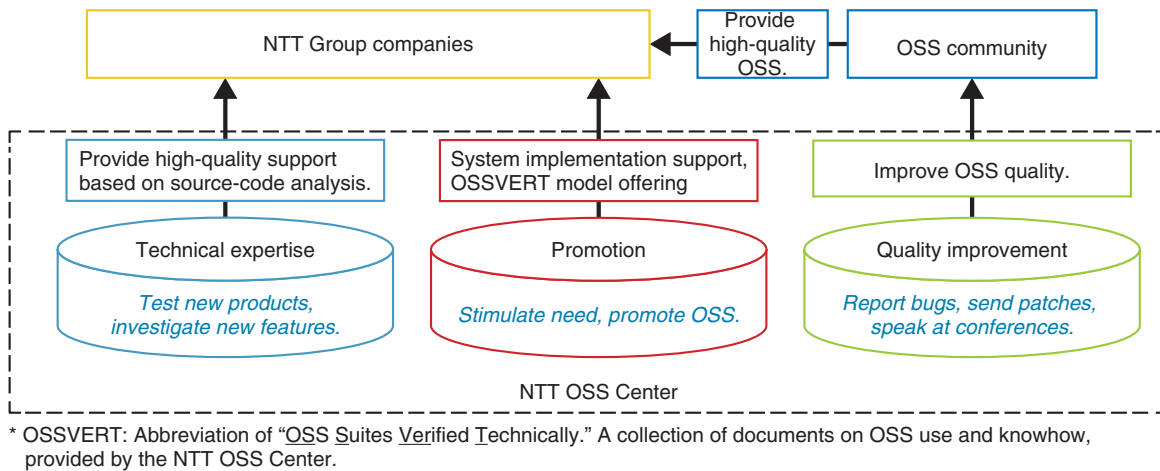


Fig. 1. NTT OSS Center initiatives for use of OSS within the NTT Group.

Table 1. Benchmark tools used in our performance testing.

Measurement items	Benchmark tools
System call, glibc performance	LMbench ( <a href="http://www.bitmover.com/lmbench/">http://www.bitmover.com/lmbench/</a> )
Network performance	netperf ( <a href="http://www.netperf.org/">http://www.netperf.org/</a> )
Task scheduler performance	ocbench ( <a href="http://linux.1wt.eu/sched">http://linux.1wt.eu/sched</a> )
Disk I/O performance	IOzone ( <a href="http://www.iozone.org/">http://www.iozone.org/</a> )

end users of their enterprise services and incur a substantial cost. This is why an even stronger guarantee of stability is required in our group.

For this reason, with the release of each minor version, certain items need to be checked on the kernel and virtualization software to prevent problems from occurring. Specifically, we check:

- (1) Whether there are significant changes such as added or removed functionality by reading release notes and related documentation.
- (2) Whether it can be installed on a PC server without difficulty or can be upgraded from the previous version.
- (3) Whether there are any significant changes in performance.

On the basis of the results of these checks, we decide whether the new version can be used safely in our group, and if so, we recommend it as a *recommended version*. As an example, for RHEL 6.4 in February 2013, we temporarily deferred making it a recommended version since it failed to pass item (3). The following is a detailed description of how we reached that decision.

## 2.2 Performance testing

For the newly released RHEL 6.4, as we always do for each minor release, we did some performance tests on a bare-metal server and virtual machines, using KVM for example, to determine whether there were any significant performance changes in the basic functions of the OS. The details of the measurements taken for these tests are given in **Table 1**.

## 2.3 Disk I/O problem with RHEL 6.4 KVM

On a KVM virtual machine, the disk I/O (input/output) performance test result showed that throughputs for a particular I/O pattern dropped 20% compared to the RHEL 6.3 environment. Since this was a significant change in performance, we temporarily deferred making RHEL 6.4 a recommended version in cases where it would be used for KVM virtualization, and we further investigated the issue.

As a result of the investigation, we reached the conclusion that what had triggered the performance regression was a bug in QEMU—the machine emulator and virtualizer—that could disable a disk I/O optimization.




Fig. 2. Disk I/O performance test results (KVM environments).

We also confirmed that by using a later version of QEMU that had received a bug fix patch, we could gain the same performance as RHEL 6.3 (Fig. 2). Finally, on the basis of this result, we decided to recommend the use of the QEMU version that included the bug fix for RHEL 6.4 KVM virtualization [1].

This is an example of how our testing can prevent problems from occurring in real systems. We plan to continue this kind of testing so that our group companies can use RHEL with confidence for enterprise services.

#### 2.4 The next major release (RHEL7)

Red Hat is now preparing to release RHEL 7, the next major version. RHEL 7 is expected to have improved scalability and many new features, so it will most likely become one of the most widely used operating systems in the future.

Accordingly, in preparation for its release, we are strengthening our know-how of RHEL 7, checking its stability, and feeding back any problems found to Red Hat and the OSS community to ensure that it is reliable.

### 3. Contributing to OSS development community and improving software quality

In addition to giving feedback to Red Hat, we also contribute to the Linux community to improve the quality of the software. Specifically, we submit patches and have discussions with developers through mailing lists and community meetings. Because RHEL is based on OSS products developed by the

Linux community, providing feedback to the community leads to improving the quality of RHEL.

We recently submitted patches to fix bugs in KVM and networking, and they were merged into the latest Linux kernel. We also gave presentations at various conferences related to Linux and KVM, for example, LinuxCon and KVM Forum, to publicly discuss problems with developers around the world. These contributions to the community led to our receiving the Japan OSS Contribution Award and the Japan OSS Incentive Award in February 2013.

We are also testing and improving new features that have not been supported on RHEL. As an example, the latest Linux kernels incorporate Open vSwitch, a software switch that supports OpenFlow<sup>\*2</sup>, which is attracting attention in the cloud computing and network virtualization fields. We have already investigated its internal structure and have tested its performance to enhance our knowledge.

We will continue contributing to the community and trying to improve the quality of OSS products.

## 4. Further promoting OSS virtualization technology

### 4.1 Issues in promoting KVM

Although we have been working on KVM, it is not widely used within our group. One of the reasons for this is that the operations management features for

<sup>\*2</sup> OpenFlow: A standard protocol that allows direct access to and manipulation of the forwarding plane of network devices. It is attracting attention as one of the core technologies for network virtualization and SDN (software-defined networking).




Fig. 3. Centralized management using a virtualization management product.

KVM virtual environments are still inadequate.

Current popular proprietary (non open-source) products that have operations management features generally provide a GUI (graphical user interface) enabling various tasks to be performed easily even if they are complex. On the contrary, there was no such product for KVM, and depending on the task, many commands had to be entered through a CUI (command-line user interface). In this way, the lack of operations management features had been an obstacle to promoting KVM.

#### 4.2 OSS virtualization management products: RHEV and oVirt

The oVirt community was created to develop an open source management product for KVM. Red Hat is testing and improving the product developed by the oVirt community to achieve the quality necessary for enterprise use and is releasing it as RHEV (Fig. 3). Many of the operations management features needed for KVM virtual environments are already implemented in RHEV, and efforts are continuing to add new functions and to improve its quality. We have confirmed that RHEV has no major bugs in its frequently used features and is ready for practical use.

To promote KVM, we will continue working on RHEV and on improving the quality of RHEV and oVirt by feeding back the know-how we gain to the oVirt community.

### 5. Promoting OSS in the future

As we have explained in this article, using OSS in our group is not just for avoiding vendor lock-in, but for gaining various benefits such as cost reductions and analyzability of software.

We will continue to carry out the functions of the NTT OSS Center by assessing the maturity of OSS technology, promoting the use of OSS within our group, and providing high-quality support. Furthermore, the experience gained through these activities will help us contribute to the OSS community extending throughout the world, which in turn will lead to the further development of OSS.

### References


- [1] QEMU-KVM bug fix update,  
<https://rhn.redhat.com/errata/RHBA-2013-0857.html>



**Kenji Yoshida**

Manager, Infrastructure Software Technology Unit, NTT Open Source Software Center.

He received the B.E. and M.E. in electronic engineering from University of Fukui in 1993 and 1995, respectively. He joined NTT in 1995 and has been engaged in research on IP service network engineering. He then moved to NTT Open Source Software Center in 2010, where he has been providing Linux technical support.



**Toshio Nakazawa**

Senior Expert, Infrastructure Software Technology Unit, NTT Open Source Software Center.

He received the B.E. in electrical and electronic engineering from Meiji University, Tokyo, in 2002. He joined NTT Software Corporation in 2002 and was involved in developing systems for mobile communication. He then moved to NTT Open Source Software Center in 2010, where he has been providing technical support for Linux. He is now interested in operations management products for KVM and is promoting them within the NTT Group.



**Kumiko Goto**

Engineer, SE Department Core Technology, Quality Management and Engineering Division, NTT Comware.


She received the B.E. in information and system engineering from Chuo University, Tokyo, in 2005. She joined NTT Comware in 2005, and then moved to NTT Open Source Software Center in 2010, where she was responsible for providing Linux technical support. She moved back to NTT Comware in 2014.



**Seiji Kihara**

Senior Manager, Infrastructure Software Technology Unit, NTT Open Source Software Center.

He received the B.S. and M.S. in information science from the Tokyo Institute of Technology in 1990 and 1992, respectively. He joined NTT in 1992. His current research interests include operating systems, computer networks, and open source software. He is a member of the Information Processing Society of Japan (IPSJ), the Institute of Electronics, Information and Communication Engineers (IEICE), and the Association for Computing Machinery (ACM).



**Toshiaki Makita**

Senior Expert, Infrastructure Software Technology Unit, NTT Open Source Software Center.

He received the B.S. and M.S. in informatics from Kyoto University in 2004 and 2006, respectively. He joined NTT WEST in 2006 and studied Ethernet VPN. Since moving to NTT Open Source Software Center in 2011, he has been interested in networking technologies and developing a Linux kernel networking subsystem.