Regular Paper

Lightweight Recovery from Kernel Failures Using Phase-based Reboot

KAZUYA YAMAKITA^{1,a)} HIROSHI YAMADA^{1,2} KENJI KONO^{1,2}

Received: July 19, 2011, Accepted: December 6, 2011, Released: March 27, 2012

Abstract: Although operating systems (OSes) are crucial to achieving high availability of computer systems, modern OSes are far from bug-free. Rebooting the OS is simple, powerful, and sometimes the only remedy for kernel failures. Once we accept reboot-based recovery as a fact of life, we should try to ensure that the downtime caused by reboots is as short as possible. This paper presents “phase-based” reboots that shorten the downtime caused by reboot-based recovery. The key idea is to divide a boot sequence into *phases*. The phase-based reboot reuses a system state in the previous boot if the next boot reproduces the same state. A prototype of the phase-based reboot was implemented on Xen 3.4.1 running para-virtualized Linux 2.6.18. Experiments with the prototype show that it successfully recovered from kernel transient failures inserted by a fault injector, and its downtime was 34.3% to 93.6% shorter than that of the normal reboot-based recovery.

Keywords: reboot-based recovery, operating system reliability, virtualization

1. Introduction

High availability is important for all ranges of computer systems from high-end enterprise systems to low-end consumer devices. High-end enterprise systems lose millions of dollars if their services are unavailable. Low-end device vendors would lose their customers if their products such as smart phones and HDD recorders were not very stable or sometimes got hung up. Upgrading iPhoneOS 3.x to iOS 4.0 on iPhone 3G causes severe performance degradation and makes iPhone 3G service nearly unavailable. Apple was criticized for delivering an inferior operating system and finally took action to investigate the series of complaints related to performance.

Operating systems are crucial for achieving high availability of computer systems. Compared with application-level failures, kernel-level failures are known to occur less frequently, but they have a considerable impact on the overall availability of software systems. Even if the applications running on the operating system are highly available, a bug inside the kernel may result in a failure of the entire software stack; no application can continue to run on the crashed kernel.

Modern operating systems are far from bug-free. Palix et al. [2] report that the rate of introduction of bugs continues to rise even in Linux 2.6. In addition, the average time between when a bug is introduced and when a fix is released is 1.8 years for Linux kernels. Our investigation of the change logs of Linux 2.6.24 and 2.6.25 also revealed that there are critical bugs inside the kernel core components. Kernel bugs are not the sole reason for kernel failures. Soft errors in high-density semiconductors are increas-

ing [3], and they cause incorrect values to be read from memory or incorrect instruction results to be produced.

For end users of computer systems, sometimes the only remedy for kernel failures is to reboot the operating system (and thus the entire software stack). For example, if a smart phone freezes due to a kernel failure, the end user reboots it in the expectation that the reboot will recover the smart phone; she does not have any skill or tools to diagnose and recover from the failure. Aside from low-end consumer devices, skillful administrators for high-end enterprise systems sometimes reboot the system to avoid or recover from failures. A Cisco Security Advisory [4] reported that their network products had a bug involving a memory leak, and the reboots were necessary to recover from it until a bug fix was released. IBM Director, a cluster management system for xSeries servers, periodically reboots (i.e., rejuvenates) the system to counteract software aging [5].

Once we accept reboot-based recovery as a fact of life, we need to try to reduce the downtime caused by reboots as much as possible. This paper proposes “phase-based” reboots that shorten the downtime caused by reboot-based recovery. In a phase-based reboot, a boot sequence is divided into three *phases*: 1) the hardware-initialization phase, 2) the kernel-boot phase, and 3) the daemon-startup phase. The key idea behind phase-based reboot is that a reboot repeats the same procedure as in the previous boot and sometimes reproduces the system state identical to the previous one; we can *reuse* a system state in the previous boot if the next boot reproduces the same state. In the phase-based reboot, a system state is saved after each boot-phase is finished. When a reboot is done for recovery, our system restores the saved state to skip the boot-phases that reproduce the same states as in the previous boot.

¹ Keio University, Yokohama, Kanagawa 223–8522, Japan

² JST CREST, Chiyoda, Tokyo 102–0075, Japan

^{a)} yama@sslabs.ics.keio.ac.jp

An earlier version of this paper appeared in IEEE/IFIP DSN [1].

To save and restore a system state, the phase-based reboot uses the *snapshot* functionality of virtual machines (VMs). At first glance, saving and restoring a system state is straightforward; the entire memory image of the VM is saved to and restored from a disk. However, this is time-consuming, especially when the memory size assigned to a VM is large. In the worst case, the phase-based reboot takes a longer time to reboot than a normal reboot. To avoid this situation, our mechanism avoids saving unnecessary memory pages that can be reconstructed after the memory image is restored.

Restoring a system state is much more complicated. The memory image saved to a disk contains a disk cache that may be updated after the snapshot is taken. In other words, the disk cache in the saved image may be out of date. If the saved image is simply restored, the out-of-date disk cache is also restored and regarded as up-to-date. To solve this problem, our mechanism refreshes in-memory file objects with the corresponding disk blocks after it restores the saved image.

A prototype of the phase-based reboot was implemented on Xen 3.4.1 running para-virtualized Linux 2.6.18. Experiments with the prototype showed that the phase-based reboot successfully recovered from kernel transient failures inserted by the kernel fault injector, and its downtime was 34.3 to 93.6% shorter than that of the normal reboot-based recovery.

The rest of this paper is organized as follows. Section 2 presents the key idea of phase-based reboot and its semantics. Section 3 overviews phase-based reboot. Sections 4 and 5 describe the design and implementation of phase-based reboot, respectively. Section 6 presents our experimental results. Section 8 discusses work related to ours. Finally, Section 9 concludes this paper.

2. Phase-based Reboot

To reduce the downtime of reboot-based recovery, the phase-based reboot skips some phases of a time-consuming boot sequence. In this section, we describe the key idea behind the phase-based reboot and its recovery semantics.

2.1 Key Idea

A boot sequence can be divided into three *phases*: 1) hardware initialization, 2) kernel boot, and 3) daemon startup. Normal reboot-based recovery executes all the boot phases in order to reconstruct a consistent system state and restart services. The normal reboot-based recovery repeats the same boot sequence as in the previous boot. Most parts of the sequence are similar to the previous boot because the system configuration is not changed in the context of reboot-based recovery. In the reboot-based recovery, the system is *not* rebooted so as to make the configuration changes effective; the boot sequence starts from the same system configuration and is thus expected to result in the same system state as the previous boot.

As illustrated in **Fig. 1**, the system starts its services after initializing the hardware, booting the kernel, and starting up every daemon. Through these operations, a consistent system state is constructed from which we can start services. When a kernel crashes, the reboot-based recovery is attempted; the system re-

Fig. 1 Key idea behind phase-based reboot. In most cases, the reboot process produces the same system state as in the previous boot.

peats the same procedures to reconstruct the same system state from which we can restart the services. In the reboot-based recovery, no system configuration is changed in most cases. Therefore, the resulting system states are expected to be the same in the previous and current boot sequences.

The key idea behind the phased-based reboot is to save and reuse consistent system states during the reboots. If the next reboot always creates the same state as in the previous one, we can simply save and restore the previous state for reboot-based recovery instead of rebooting the entire system. Unfortunately, this is an oversimplification. During service operations, an administrator may change the configuration of some daemons. In this case, we cannot reuse the system state saved in the previous boot because the different configuration may result in a different system state.

To address this reusability problem, the phase-based reboot saves system states at several points, called *restartable points*, during the boot sequence. System states saved at restartable points are called *restartable candidates*, from which the user can select the appropriate point to start the system reboot from. By default, the phase-based reboot saves the candidates every time each boot-phase is finished. Administrators can add more restartable points based on their intimate knowledge of the target system. If no configuration is changed, we can use the most recent candidate from which the system reboot starts. Note that this is the most common case in reboot-based recovery. If a configuration is changed as in the above example, the user can select an appropriate candidate from which to restart the system from. In the above example, the user restarts the system just after the kernel initialization is finished. To help the user select an appropriate restartable point, the phase-based reboot can determine which restartable point can be used for recovery.

2.2 Recovery Semantics

Figure 2 shows the failure coverage for a normal OS reboot and the phase-based reboot. The phase-based reboot handles kernel transient failures in a way similar to a normal OS reboot. Kernel transient failures include memory leaks and non-deterministic kernel panics. By rebooting an OS, we can recover from kernel transient failures even if their root causes are unknown. Rebooting an OS eliminates a corrupted memory state and returns the system state back to its initial state, which is known to be consistent, making it possible to safely restart services. The phase-based reboot inherits this advantage from the OS reboot, and it restores the system state to a clean and consistent one at a restartable point.

Unlike normal OS reboots, the phase-based reboot cannot recover from a failures caused by inconsistent hardware states. To

recover from inconsistency in hardware devices, the devices must be re-initialized; the faulting machine must be reset physically. Since the phase-based reboot skips the hardware-initialization phase, this type of reboot cannot recover from hardware inconsistency. This is not a serious shortcoming of the phase-based reboot. When a failure occurs, the user tries the phase-based reboot first. If the failure cannot be recovered, the user physically resets the entire machine. From our experience in investigating Linux change logs, most of the bugs in Linux corrupt in-memory kernel states, which can be recovered from the phase-based reboot. There are only a few bugs that make hardware devices inconsistent.

As in the normal reboot-based recovery, the phase-based reboot cannot handle all types of failures. Since the phase-based reboot is an optimization of the normal reboot-based recovery, it inherits all the shortcomings of reboot-based recovery. First, the failures that persist across reboots cannot be recovered. For example, if a hardware device is corrupted physically, reboot-based recovery is useless. If the persistent data in file systems are corrupted, we have to run `fsck` to repair the corruption. Second, the reboot-based recovery cannot handle deterministic failures that can be reproduced by executing the same path. Finally, the reboot-based recovery sometimes fails to restart user-level applications that save their states to non-volatile devices. If a kernel failure prevents an application from saving its state, the application may be confused after the OS reboot. To correctly restart a service, the application must perform the recovery operation. For example, an application using a database server must roll back the SQL transactions that were processed when the kernel crashed.

Fig. 2 Failure coverage for a normal OS reboot and the phase-based reboot.

3. Overview

The phase-based reboot leverages the snapshot function provided by system virtualization to restart the system at a restartable point. System virtualization is becoming commonplace in a computing environments. The snapshot function enables us to save/restore a virtual machine (VM) state including CPU registers, memory, and disks at an arbitrary point. The phase-based reboot uses a snapshot taken during an OS boot to restart the VM at a restartable point. We refer to the snapshot as a *restartable image*. The phase-based reboot overwrites CPU registers and memory states preserved in the restartable image to the running VM; the phase-based reboot never rolls back the disk state to save updates of disks in the service operation.

In the phase-based reboot, we treat snapshots as restartable candidates that can be used as a restartable image. The phase-based reboot appropriately selects a snapshot from the restartable candidates and restores it. To prepare restartable candidates, we take snapshots at many points during an OS boot. **Figure 3** shows a typical example of how restartable candidates are prepared. We can collect them by taking a snapshot when the kernel boot is complete, every daemon has been launched, and after a log-in prompt appears. When a phase-based reboot is conducted, we pick up a snapshot from the restartable candidates that has the same state as after the normal reboot-based recovery.

The use of system virtualization allows us to prevent kernel failures inside a VM from corrupting its restartable candidates. If their contents are modified by the propagation of kernel failures, we cannot successfully restore the restartable image. Since the virtual machine monitor (VMM) isolates VMs running on it, the kernel failures are not propagated to the other VMs or VMM, where snapshots are saved. Although commodity Oses offer a hibernation mechanism that saves its memory state to disks, the saved memory image are more easily affected by kernel failures. This is because the hibernation mechanism saves memory images to the disk that is not isolated from the OS. This is not reasonable for saving restartable candidates.

However, the phase-based reboot does not come without effort. It poses several design challenges. First, the existing snapshot restoration takes a long time if the memory size assigned to the VM is large. Next, when a snapshot is restored, the file system objects are restored as well, which leads to eliminating the disk update in the service operation. Lastly, we need a way to determine which restartable candidate is a restartable image in order to help select a proper restartable image.

Fig. 3 Comparison of normal reboot and phase-based reboot.

4. Design

For the first challenge, we avoid saving pages that are unnecessary for the VM to work correctly after the snapshot is restored. For the second challenge, we design a kernel module that updates the file system objects. To address the last challenge, we prepare a support tool that infers application states that will be built by the normal reboot. To do so, the support tool checks whether files accessed during the guest OS boot are updated in the service phase.

4.1 Snapshot Optimization

The conventional VM snapshot function saves and restores all the memory pages of a VM even if the pages are not used for the kernel and user processes. If a VM is assigned 1,024 MB of memory, the VMM saves all the memory on the disk even if the VM uses only 128 MB. As the assigned memory size is larger, the restoration of the snapshot takes longer, causing many disk I/O problems.

To shorten the time for restoring a restartable image, we *shrink* the size of VM memory checkpoints. Our technique reduces disk I/O involved in saving and restoring the memory checkpoints. We borrowed this idea from the hibernation mechanism offered by commodity OSes. Specifically, our technique avoids saving pages that are not necessary for the system to work correctly after the restore operation as shown in Fig. 4. Such pages include a free page and file cache pages. For example, if a VM is assigned 1,024 MB of memory, and 928 MB are free pages, we save only 72 MB since free pages are a soft state and can be reproduced from the disk. We believe that our technique is effective because memory usage is not heavily utilized during the boot phase where restartable candidates are taken.

In this work, we focused on a free page and a page containing soft-state kernel objects. Even if a VMM discards the contents of free pages, the guest works correctly because free pages are initialized when the kernel uses them. Soft-state kernel objects include caches for disk blocks and caches for kernel resource managers. A file cache is a typical example of soft-state kernel objects. Because a file cache contains the data on disk, the guest can reproduce it by reading the data from the disk. Likewise, caches for resource managers such as a slab cache can be reproduced from the original data structure.

To avoid saving these pages, we modify the guest kernel to explicitly inform the VMM which pages are unnecessary. When a snapshot is taken, the guest kernel examines its memory objects and sends the VMM the guest physical address of the unnecessary

pages. The VMM does not store them in a snapshot, based on the given addresses. When the snapshot is restored, the VMM compensates for the lost pages by allocating new pages to the guest. After that, the restored VM starts to run.

Shrinking the VM memory checkpoints also enables us to put the checkpoints on small and faster access devices. When a VM memory checkpoint is significantly small, we can place it on solid state drives or RAM disks whose accesses are much faster than disk drives. This accelerates the restoration of the restartable image, leading to much faster reboot-based recovery.

4.2 Update of File System Objects

We need to take into account the memory objects of file systems after restoration from a restartable image. File systems are OS core components that manage disk caches including the cache of data blocks, metadata, and file system metadata. Because file systems manage such memory objects, a restartable image naturally contains them, but the file systems fail to keep the disk updates in the service phase when a restartable image is restored. For example, the filesystems' metadata such as the super block cause this problem. The kernel only updates these metadata in the memory, and writes them back to the disks; the metadata are never read from the disks after the partition has been mounted. When a restartable image is restored, the older file system metadata are overwritten on the current metadata. This causes the file system to inconsistently manage disk blocks such as free blocks and data blocks.

The i-node cache of the file opened with the append mode causes a similar problem. When the i-node cache preserved in a restartable image does not have an append region, restoring the restartable image results in the i-node data being overwritten on the newer data on the disk. This means the appended regions of the file are eliminated.

Although remounting the disk volumes is a simple solution for this inconsistency problem, this solution is not suitable for phase-based reboots. Specifically, we take a snapshot after unmounting the disk volumes. When the snapshot is restored, we mount them. This way naturally refreshes file system objects, thus avoiding the inconsistency problem described above. However, we have to close all the files in the disk volumes to safely unmount the disk volumes. This constraint is critical for the phase-based reboot because some applications keep files open while running. For example, `syslog.d` keeps its log file open with `O_APPEND`, and `crond` keeps its pid file open. Therefore, we cannot put these applications into restartable candidates. To put such applications on a restartable image, we explore an alternative to solve the in-

Fig. 4 Snapshot optimization.

Fig. 5 Selection of the proper candidate.

consistency of file system objects.

To solve the inconsistency, our kernel module forces the kernel file system component to read such metadata again just after a restartable image has been restored. When the restoration of a snapshot is completed, our kernel module forces the file system component to read the file system metadata and i-node from the disk and it updates them. The consistency of file cache pages is kept by our snapshot shrinking because it releases the file cache pages.

4.3 Finding Restartable Images

We use as a restartable image a restartable candidate, where the application states are the same as after a normal OS reboot. Because the memory contents saved in a restartable image are overwritten to the target VM, the restartable image needs to contain applications' memory contents that will be built by the normal OS reboot. If an application memory image in a restartable image is different from that after the normal OS reboot, the wrong memory image will be built on the VM. This means we cannot produce the effect of the reboot-based recovery.

Suppose that a restartable candidate contains a running application that reads a configuration file in its boot phase. If the file is updated, the application should be launched with the new configuration after reboot-based recovery. However, restoring the restartable candidate builds an application image that is based on the older configuration. For example, as shown in Fig. 5, when the configuration file of daemon 3 is changed in the service operation, the user should select and restore the restartable image saved just after daemon 2 has been launched. Then, the user can initialize daemon 3. However, this selection depends on users' skills and the knowledge about the target system. Although users acquainted with the target system can select the proper candidate, it is difficult for unskilled users to select it correctly.

Another example is that a restartable candidate contains a running application that keeps a file open to log its state. Restoring this image may cause a log corruption if the application logs its state in the service phase. Because the file offset of the application is also restored, the application may overwrite the log contents that were logged before restoring. Although one way to solve this problem is to redesign applications to force them to reconstruct their states after a snapshot is restored, modifying all of the applications is unreasonable.

To find an appropriate restartable image, our checker infers the application states that will be built by the normal reboot. To do so, it checks whether files accessed until a restartable candidate is taken are updated in the service operation. If these files are not

updated, we evaluate whether the selected restartable candidate can be used as a restartable image. We assume that applications launch in the same way if files accessed by them are not updated. For example, some applications start to run based on their configuration files. If the configuration files are not modified, the applications start in the same way at the next OS boot. Even if an application reads files and caches their contents in memory, it builds the cache again when the file contents have not been modified. When a log file is not updated, the application does not overwrite old log contents since the file offset has not changed.

We believe that our checking tool supports to configure the VM environment suitable for the phase-based reboot. If non-skillful users want to enjoy the phase-based reboot, they can use a VM image configured by skillful users with our tool. Although the tool can be used to check restartable candidates when the phase-based reboot is performed, it causes additional downtime of the phase-based reboot. Optimizing the checking tool for using in the service operation is another challenge, which is out of scope of this paper.

We note that our checker does not cover all types of applications. For example, it does not manage applications whose behavior is defined by network conditions and time. If such applications are put into a restartable image, the phase-based reboot may build the wrong application states after network conditions and times are changed. To manage the applications, we need to extend our checker to determine whether applications' states in the restartable candidates are the same as those built by the normal OS reboot.

We prepare a mechanism for files opened with the append mode such as `O_APPEND` to aggressively omit the boot phase. When files are opened with the append mode, the kernel sets files' offset of the application to the end of the files in `write()`. This means the file offset is automatically set to the end of the file by `write()` even after the snapshot is restored. If the file contents are not updated, except for the appended region, the checker determines whether the application that opens the file with the append mode can consistently run after a restartable image is restored. In this situation, it does not issue a warning that a restartable candidate is not a restartable image.

5. Implementation

We implemented the phase-based reboot in Linux 2.6.18 running on Xen 3.4.1. Our core implementation consists of three modules; a *file access monitor*, *kernel object manager*, and *file update checker*. Both the file access monitor and kernel object manager are running inside the guest kernel in a domain U. The

Fig. 6 An Overview of phase-based reboot.

file update checker is running inside domain 0. The file access monitor logs the name and last modification time of files accessed until a restartable candidate is taken. To do so, it records information on files accessed by the processes. The kernel object manager appropriately handles disk cache being managed by the ext3 file system. In addition, it frees slab cache and tells the addresses of free pages to the underlying hypervisor in order to remove the entries of the P2M table when a restartable candidate is taken. The file update checker inspects virtual disk images mounted by the target VM and checks file updates by referring to a log produced by the file access monitor. We added about 2,000 lines of new codes to Linux 2.6.18 running on Xen 3.4.1 for three modules.

Our implementation is overviewed in Fig. 6. Figure 6(a) depicts the execution of saving restartable candidates. For ease of implementation, we run a daemon server that triggers our guest kernel-level mechanism. Since we can take a snapshot only in domain 0, the client running in domain 0 communicates with the daemon server. To take a restartable candidate, the client asks the server to run the kernel-level mechanism. The client starts taking a snapshot when it is notified of the completion of the module tasks by the daemon server. Note that there is a race condition where a process can modify files until the client starts taking a snapshot after the completion of the module tasks. To avoid this situation, we need to implement a mechanism that enables domain U to take its snapshot.

First, the file access monitor logs information on accessed files for the file update checker. Next, the kernel object manager flushes the dirty buffer and releases the disk cache and slab cache. Then, the kernel object manager tells the underlying hypervisor to remove entries of free pages in the P2M table with the balloon driver to shrink the memory checkpoint. Finally, we save the shrunken memory image as a restartable candidate.

Figure 6(b) shows the execution flow when the phase-based reboot is triggered. Xen restores the selected restartable candidate, and the kernel object manager reallocates free pages because the VM snapshot has been shrunk. Finally, the kernel object manager updates i-nodes, dentries, and super block data in the memory. If necessary, the file update checker checks whether the VM has updated the files in the file access logs and determines which restartable candidate is restartable.

5.1 File Access Monitor

The file access monitor logs information on accessed files for

the file update checker. Specifically, it memorizes the absolute path of the accessed file, its i-node number, and its last modification time. The file access monitor also memorizes whether a file has been opened with `O_APPEND`. It saves the memorized information as a file on the guest file system when the kernel object manager triggers it. The log is used by the file update checker, as will be described later.

The file access monitor monitors `sys_open`, `sys_stat`, and `sys_exec` to find out which files have been accessed. The monitoring is stopped when our system call, `pbr_ready()`, is issued in order to avoid overhead of file access monitor activities in the service phase.

5.2 Kernel Object Manager

The kernel object manager manipulates kernel objects being managed by ext3 and the slab allocator. The manipulation is carried out when restartable candidates are taken and when they are restored. When a restartable candidate is taken, the kernel object manager flushes the dirty buffer and releases the page cache corresponding to the buffer cache, i-nodes, and dentries. This is done to shrink the memory checkpoint. If some processes are using i-nodes and dentries, the kernel object manager does not release these objects. When the restartable candidate is restored, the kernel object manager updates the unreleased cache by fetching the data from the virtual disk. At the same time, it updates the super block data in the memory.

The kernel object manager also unregisters free pages from the P2M table for the Xen hypervisor to shrink the memory checkpoint of the domain. To remove the entries of free pages in the P2M table, the kernel object manager leverages a *balloon driver* [6]. When the balloon is inflated, it pins down free pages and tells the Xen hypervisor to remove their entries in the P2M table. When the balloon is deflated, it releases the pinned pages and registers their entries to the P2M table again.

The kernel object manager controls the balloon when a restartable candidate is taken and restored. After releasing the page cache and slab cache, it inflates the balloon to remove the entries of free pages in the P2M table. When the restartable candidate is restored, the kernel object manager deflates the balloon to register the free pages in the P2M table again.

5.3 File Update Checker

The file update checker helps to determine which restartable

candidate is restartable. We can run the file update checker just after a restartable candidate is taken. To detect file updates, it pairs the log produced by the file access monitor with the restartable candidate. The file update checker mounts the virtual disk used by the target domain U after the restartable candidate is taken to salvage the produced log.

To determine which restartable candidate is restartable, the file update checker detects file updates by referring to the log paired with the restartable candidate. The file update checker obtains the current state of the logged files by mounting the used virtual disks. It compares the obtained i-node number and the last modification time with the corresponding values in the log. If the current values are different from the logged ones, the file update checker judges the files to be updated, and the restartable candidate cannot be used as a restartable image.

To successfully deal with the files opened with `O_APPEND`, the file update checker calculates and logs the hash value of their contents just after a restartable candidate is taken. The file update checker gets the file contents by mounting the virtual disks. When we check whether a restartable candidate is restartable, the file update checker calculates the file contents except for the appended region again, and compares it with the logged value. If the values are the same, the file update checker does not issue a warning that the restartable candidate is not restartable.

6. Experiments

We conducted experiments to evaluate the effectiveness of the phase-based reboot. We used a machine equipped with a 3 GHz quad-core Xeon processor, 16 GB of memory, and a 73-GB SAS NHS 10,000 rpm hard disk. On this machine, we ran Xen 3.4.1 and the Linux 2.6.18 kernel in domain 0. We also ran the modified Linux 2.6.18 para-virtualized for Xen on guest domains connected to a 10-GB virtual disk. We installed Fedora Core 8 on each domain, and turned off unnecessary service daemons.

We investigated the following fundamental issues in our experiments. The first was how the phase-based reboot shortens the downtime of reboot-based recovery. The second was which restartable candidate can be a restartable image under a complicated workload. The last was whether the phase-based reboot can recover from kernel transient failures.

6.1 Downtime

We measured the downtime of the phase-based reboot to determine how the phase-based reboot shortens the downtime of reboot-based recovery. To execute the phase-based reboot, we prepared two restartable images. The first was a restartable image that was taken before the guest kernel mounted the virtual disk. We refer to the phase-based reboot that is restoring this restartable image as *pr-naive*. The other was a restartable image that was taken when a log-in prompt appeared after the kernel and all the daemons were ready. We simply refer to the phase-based reboot that is restoring this restartable image as *pr-opt*. For comparison, we also measured the downtime of a normal boot and a normal reboot on the guest domain (*guest boot* and *guest reboot*). To clarify how effective our optimization described in Section 4.1 was, we executed the phase-based reboot without our snapshot

optimization and measured its downtime (*pr without snapshot opt.*). We started measuring when each operation was triggered, and stopped when all the daemons registered in run level 3 were ready on the domain. We measured the downtime of each reboot-based recovery, varying the memory size of the guest domain.

We assumed a scenario of recovering from fail-stop failures, and compared the guest boot to *pr-opt*, *pr-naive*, and *pr without snapshot opt*. In this scenario, each reboot-based recovery contains a `fsck` execution; the Linux kernel conservatively runs `fsck` in the boot after the kernel is shut down without partitions being unmounted. Next, we assumed a scenario of recovering from gradually corrupting failures such as memory leaks, and we compared the guest reboot to *pr-opt*, *pr-naive*, and *pr without snapshot opt*. In this scenario, none of the reboot-based recoveries need a `fsck` execution because we assume the situation where the virtual disk is correctly unmounted.

Tables 1 and 2 list the average downtime of each reboot-based recovery. Table 1 indicates that the downtime of the phase-based reboot was shorter than that of the guest boot in many cases. In *pr-opt*, the downtime was 75.0% to 86.2% shorter than the guest boot, while the downtime of *pr-naive* was 34.3% to 60.6% shorter than that of the guest boot. In *pr without snapshot opt*, its downtime is shorter than the guest boot when the domain memory size was smaller than 2 GB.

Table 2 lists similar results to Table 1, where the downtime of the phase-based reboot was shorter than that of the guest reboot in many cases. In *pr-opt*, the downtime was 86.1% to 93.6% shorter than the guest reboot. The downtime of *pr-naive* was 60.1% to 77.6% shorter than that of the guest reboot.

To analyze the downtime caused by the phase-based reboot, we show the breakdown of the downtime of *pr without snapshot opt*, *pr-naive* and *pr-opt* in Fig. 7. Figure 7 (a) and (b) reveal that our snapshot optimization significantly contributed to shortening the downtime of reboot-based recovery. In *pr without snapshot opt*, the restore time was much longer than the other configurations, *pr-naive* and *pr-opt*. In particular, the restore time was about 37

Table 1 Average downtime of guest boot, *pr without snapshot opt* with `fsck`, phase-based reboot with `fsck`.

Memory size [MB]	Guest boot [sec]	<i>pr w/o opt.</i> w/ <code>fsck</code> [sec]	<i>pr-opt</i> w/ <code>fsck</code> [sec]	<i>pr-naive</i> w/ <code>fsck</code> [sec]
64	18.96	3.23	2.62	7.47
128	18.50	3.96	2.68	7.81
256	18.39	5.82	2.59	7.90
512	18.45	10.29	2.77	7.69
1,024	18.85	18.46	2.90	8.47
2,048	19.00	38.09	3.38	9.49
4,096	19.29	67.05	4.83	12.67

Table 2 Average downtime of guest reboot, *pr without snapshot opt*, phase-based reboot, *pr-naive*.

Memory size [MB]	Guest reboot [sec]	<i>pr w/o opt.</i> w/o <code>fsck</code> [sec]	<i>pr-opt</i> w/o <code>fsck</code> [sec]	<i>pr-naive</i> w/o <code>fsck</code> [sec]
64	29.01	2.42	1.87	6.51
128	28.58	3.18	1.93	6.87
256	28.27	5.01	1.84	6.98
512	28.42	9.57	2.02	6.78
1,024	28.83	17.65	2.15	7.57
2,048	28.92	37.17	2.63	8.56
4,096	29.38	66.15	4.08	11.72

Fig. 7 Breakdown of pr without snapshot opt, pr-opt, and pr-naive downtime.

and 66 seconds when the memory size of the domain was 2 and 4 GB respectively. This is because Xen's snapshot function saves and restores all the memory pages assigned to the guest domain even if the pages are not used by the kernel and user processes. The restore time in the other configuration was at most 5.8 seconds. And, the time required for syncing and updating kernel objects was much shorter than the other functions.

In addition, our snapshot optimization successfully shrank the restartable images. For example, when a VM was assigned 1,024 MB of memory, the optimized snapshot function saved only 99 MB as a restartable image, but Xen's snapshot function saved 1,050 MB. As described in Section 4.1, since we can prepare RAM disks or solid-state drives where these memory checkpoints are placed, we can shorten the downtime of the phase-based reboot.

The figure also shows that omitting the kernel and daemon boot phase is effective to shorten the downtime of reboot-based recovery (Fig. 7 (b) and (c)). In pr-naive, booting the kernel and daemons is the main part of its downtime since the impact of the restore operation and ballooning is relatively smaller. Phase-based reboot in pr-opt effectively shortens its downtime by omitting the launch phase. In fact, the downtime of pr-opt was 61.9% to 67.2% shorter than that of pr-naive, as exhibited in Table 1.

6.2 Finding Restartable Images

To confirm how the phase-based reboot performs under a complicated workload, we checked which restartable candidate can be a restartable image after running a benchmark that models a real web site. We used RUBiS [7] on the Java EE platform, which is a three-tailored auction site prototype modeled after eBay.com [8]. We prepared additional physical machines for this experiment. The specifications of these machines were the same as the machine described previously. These machines were connected via Gigabit Ethernet. We ran the RUBiS client emulator on one machine while Xen 3.4.1 was running on another machine. The Xen machine was used as a server where three guest domains were running, a web server domain (*FrontVM*), application server domain (*AppVM*), and database server domain (*DBVM*). Apache 2.2.9, Tomcat 5.5.28, and MySQL 5.0.45 were running on *FrontVM*, *AppVM*, and *DBVM*, respectively. We emulated 500 clients and checked whether or not all the restartable candidates were restartable images. Our check was carried out two ways. One is that we conducted the phase-based reboot when the emulation had finished. The other is that we conducted the phase-based reboot while the client emulator was running. We assigned 1.7 GB of memory to each guest domain. This memory size comes from the small VM configuration in Amazon Elastic

Fig. 8 Restartable points in the second experiment.

Compute Cloud [9].

We prepared restartable candidates in the following way, which is shown in Fig. 8. We took a snapshot before mounting the virtual disk (Point A), when the kernel boot was completed (Point B), when all the configured daemons were launched (Point C), when a log-in prompt was displayed (Point D), and after Apache, Tomcat, and MySQL were launched on each VM (Point E).

The results are exhibited in Tables 3 and 4. Table 3 indicates which restartable candidate is restartable or not when the emulation has finished. In this situation, the phase-based reboot does not issue warnings in *FrontVM* and *AppVM*. This means that we can use Point E to restart the VMs. On the other hand, the phase-based reboot judges that Point E is not restartable since some files have been updated. Specifically, `/var/log/mysqld.log` and `/var/lib/mysql/ib_logfile` were opened without `O_APPEND` and were updated during the RUBiS operation. `/var/log/mysqld.log` is used for MySQL to log its execution state. `/var/lib/mysql/ib_logfile` is a log file where MySQL records transactions states.

Table 4 indicates which restartable candidate is restartable or not when we conduct the phase-based reboot while the client emulator is running. *DBVM* is sent the same warnings as when we conduct the phase-based reboot after the client emulator is completed. The phase-based reboot sometimes judges that Point E is not restartable in *FrontVM* and *AppVM* since `/etc/httpd/logs/error.log` has been updated. When the workload of RUBiS is interrupted, Apache logs the error condition into `/etc/httpd/logs/error.log`. We cannot restart candidates at Point E because this file is opened by Apache without `O_APPEND`.

We found that some logs were updated frequently. For example, `auditd` records system call names issued by specified processes, opening the log file without `O_APPEND`. If this daemon is put into restartable candidates, we cannot use them for a restartable image. To consistently run such a daemon after a phase-based reboot, we carefully avoid putting it in restartable candidates. We need to configure the daemon to start after the phase-based reboot. If a user wants to skip the boot phase of

Table 3 Restartable points when executing phase-based reboot after completing the emulation.

VM	Point A	Point B	Point C	Point D	Point E
FrontVM	OK	OK	OK	OK	OK
AppVM	OK	OK	OK	OK	OK
DBVM	OK	OK	OK	OK	/var/log/mysqld.log /var/lib/mysql/ib_logfile

Table 4 Restartable points when executing phase-based reboot during the emulation.

VM	Point A	Point B	Point C	Point D	Point E
FrontVM	OK	OK	OK	OK	Depends (/etc/httpd/logs/error_log)
AppVM	OK	OK	OK	OK	Depends (/etc/httpd/logs/error_log)
DBVM	OK	OK	OK	OK	/var/log/mysqld.log /var/lib/mysql/ib_logfile

such a daemon, he or she has to redesign the daemon to be phase-based reboot-aware. In this case, we add `O_APPEND` to the open argument.

Although this experiment shows that a real Linux distribution and the RUBiS benchmark are restartable from Point D in many cases, there are the cases in which we need to use the early restartable candidate such as Point A. For example, when the VM hardware configuration such as VM memory size is changed during service operations, we need to use Point A to restart the VMs. Evaluating other test sets that restart at Point A, B or C, is important and future work.

6.3 Recovery from Kernel Failures

To confirm that the phase-based reboot can successfully recover from kernel failures, we synthetically injected faults into the running kernel. We measured the rate of successful recovery from the kernel crashes, which were caused a total of 200 times. To inject transient failures to our kernel, we used the fault injection mechanism originally developed at the University of Michigan. This mechanism has been used in other studies [10], [11], [12]. Each fault changes a single integer value on the kernel stack of a random thread, or a single instruction, or instruction operand in the kernel code. This emulates many common errors, such as stack corruption, uninitialized variables, incorrect testing conditions, incorrect function parameters, and wild writes. We assigned 1.7 GB of memory to the guest domain.

The experimental results demonstrated that the phase-based reboot successfully recovered from all the injected kernel failures. Because the phase-based reboot completely destroys the crashed memory state and constructs a fresh state from the restartable image, the kernel failures do not affect the phase-based reboot procedure. For example, when the fault injection tool changed values in the stack memory region to incorrect values and led to kernel stop error, the phase-based reboot overwrote the memory image of the crashed VM with the restartable image, and then the crashed VM continued providing services. In this experiment, the fault injection tool never injected faults that were not transient, such as a fault that writes incorrect values to the disks. Although such kernel failures happen in the real world, they are out of the scope of the phase-based reboot; The target failures of the phase-based reboot are *transient* failures that can be recovered by a normal OS reboot.

7. Discussion

7.1 Limitation

As previously described, the phase-based reboot has limitations. The effectiveness of the phase-based reboot depends on the two major factors. First, it depends on the type of kernel failures. As described in Section 2.2, the phase-based reboot handles kernel transient failures in a way similar to a normal OS reboot. So, the phase-based reboot does not manage persistent failures and deterministic failures. Moreover, the phase-based reboot cannot recover from the failures caused by inconsistent hardware states. This is because, differently from the normal OS reboot, the phase-based reboot skips hardware initialization which can recover from such failures.

Second, the effectiveness of the phase-based reboot depends on how applications whose running state is saved in restartable candidates launch. As described Section 4.3, a restartable candidate cannot be used as a restartable image when the restartable candidate includes the running state of an application and the applications' configuration files are modified during the service operation. For example, in Section 6.2, the restartable candidate saved at Point E, which includes the running state of MySQL, cannot be used as a restartable image because MySQL's log file were updated during the RUBiS operation. If a daemon is launched earlier than many other daemons and its configuration files are frequently modified during the service operation, the phase-based reboot is not effective since it can use only restartable candidates taken before the daemon is launched. This means that we cannot skip many execution phases, failing to shorten the downtime of reboot-based recovery. However, there is not such a daemon in 34 daemons of Fedora 8 and RUBiS.

Moreover, the phase-based reboot does not always automatically produce the effect of an OS reboot, as described in Section 2.1. Some applications create a file to avoid being doubly launched. For example, `vsftpd` creates its lock file in `/var/lock/subsys`. Because the file is preserved in the disk after a restartable image is restored, the system fails to start the applications when we restore a restartable image where they have not been started yet. In order to obtain the effect of an OS reboot from the phase-based reboot in this situation, we need to shut down such applications before conducting the phase-based reboot.

Furthermore, we need to explain the limitation of our current implementation of the support tool named the file update checker. Our check is conservative because the file update checker checks file update times, not sizes and contents, by referring to the log paired with the restartable candidate. However, it does not manage applications whose behavior is defined by network conditions and time. Imagine that we change the DHCP server configuration to assign a different IP address to the VM using the phase-based reboot. In this case, the file update checker does not issue a warning that the restartable candidate is not restartable because it cannot detect the change of IP address. Addressing this limitation is our future work.

Our snapshot optimization conflicts with a daemon that puts frequently used files into buffer cache to improve the performance in the service phase. For example, `readahead_early` accesses frequently used files to put them into the buffer cache of the kernel. Our optimization shortens the time for restoring a snapshot at the expense of disposing of soft-state objects that improve system performance. We can adjust how many pages of buffer cache we release, taking into account the importance of performance in the service. However, there is a trade-off between the performance and the time for restoring the restartable image; the more buffer cache we leave, the longer time it takes to restore the restartable image.

7.2 Use Case

One of the use-cases is cloud computing environments, especially in SaaS and PaaS environments. The skillful administrators create the restartable candidates and deploy them in the datacenters. In such environments, end users can use the applications and interfaces without the knowledge about the system configurations of the datacenters. By performing the phase-based reboot, we can improve the availability of the users' services.

To use our system, users need to understand whether the system state after next reboot is the same as the previously saved state. Specifically, we need to know that the boots of the kernel and daemons are the same as the previous boot. For example, suppose that the restartable candidate is saved when one daemon is launched. In this case, the user need to check whether files opened by the daemon are modified during service operations, whether the files was opened with the append mode in the previous boot, whether the contents of the files are changed during service operations if they were not opened with the append mode, and so on.

7.3 Future Work

We briefly discuss about future work of the phase-based reboot. First, we apply the phase-based reboot to other environments. For example, the phase-based reboot is effective for embedded systems since the applications used in the systems are limited. Second, we expand the coverage of the file update checker. As described in Section 4.3, the current implementation of the file update checker does not manage applications whose behavior is defined by network conditions and time. So, we should address this limitation in order to expand the coverage of the target applications. Thirdly, we optimize the method of checking file updates.

The current file update checker checks whether the files accessed in the previous boot are modified individually. It may cause I/O contention with the running VMs since it compares i-node number and last modification time between the current state and the previously saved state in the log. So, we need to explore a way to reduce I/O contention. Finally, we optimize the snapshot function furthermore. The current optimization of the existing snapshot function is to shrink the size of VM memory checkpoints. In detail, it is to avoid saving free pages and cache pages. We think that the size of VM memory checkpoints becomes smaller by avoiding saving other soft-state memory objects in the snapshot such as the kernel code region.

8. Related Work

Various approaches have been proposed to reduce the downtime stemming from a whole program restart. Microreboot [13] achieves fine-grained software reboots. To enable a microreboot, the target application is divided into small independent software components which become units for a reboot. If rebooting a small component cannot recover from a failure, a bigger component will be rebooted. The work aims at application-level failures, and thus, it cannot shorten the reboot time for recovery from kernel failures. Also, the microreboot is complementary notion to the phase-based reboot. We can say that the microreboot focuses on “components” of software systems. We benefit from this when a reboot of small components recovers from failures. On the other hand, the phase-based reboot focuses on “phases” of software systems. We benefit from this in cases where we have to reboot larger components such as OS kernels, which take a long time to restart.

Kexec [14] and Fast Reboot [15] allow us to quickly start up a kernel. When they are invoked on a running kernel, another kernel boots without any hardware reset. Since these mechanisms require kernel support, they cannot be used when the kernel is stopped due to kernel failures. The phase-based reboot can work even when the kernel has crashed.

Different approaches have been proposed to recover from kernel failures. Otherworld [10] reboots the kernel without clobbering the state of the running applications. After the kernel crashes and is rebooted, Otherworld restores the application memory spaces, open files, and other resources. However, the downtime of Otherworld is reported to be about 1 minute. To restart the service quickly, we use both Otherworld and our method as the situation demands. We should use Otherworld if the running states of the applications are critical for recovery. On the other hand, we should use the phase-based reboot if the running states of applications are not critical.

Akeso [16] is a kernel-level mechanism that is request-oriented in the sense that it handles the recovery at the request level such as system calls or interrupts. When a failure occurs in the kernel, Akeso rolls back the kernel state to the beginning of the function and makes the function return an error. However, it requires a complicated annotation in various places within the kernel code along with the context. Writing a correct annotation requires accurate knowledge of the kernel and is a laborious and error-prone task. The phase-based reboot does not require any annotation

and is basically complementary to kernel-level mechanisms; the phase-based reboot can quickly rejuvenate them to achieve more reliable services.

Previous studies have focused on a kernel component. Nooks [11], [17] pushes a device driver into a lightweight protection domain and transparently recovers device drivers when they fail. LeVasseur et al. proposed an approach to isolating device drivers using dedicated VMs [18] to limit the drivers' crash influence. Membrane [19] is a kernel-level mechanism to make file systems restartable. It periodically saves checkpoints of file system states. If the file system fails, Membrane restores the file system state from the recent checkpoint and consistently and transparently updates the stateful information to applications. These approaches focus on certain kernel components' failures, but the phase-based reboot focuses on failures in any kernel component.

Some previous studies have made better use of virtualization to improve the reliability of the system. Bresoud and Schneider proposed a hypervisor-based approach to implementing a fault-tolerant system [20]. It replicates the state of a system remotely and recovers from failures in a failover manner. Remus [21] is a failover mechanism that uses VMM. Remus replicates snapshots of an entire running OS instance between a pair of physical machines. These failover approaches basically focus on hardware failures, while the phase-based reboot focuses on software failures in the kernel.

CuriOS [22] recovers failed services transparently to clients in a microkernel OS. CuriOS stores client-specific states in client-associated but client-inaccessible memory. When OS system servers fail, the servers use the preserved client states to restart without affecting the clients. Vino [23] provides a mechanism to recover from extension failures without rebooting the OS. Vino encapsulates extensions in a transaction to spontaneously abort them and clean up their states. These can run on a microkernel or special formed kernel, while the phase-based reboot is suited to commodity OSes such as Linux.

Approaches to improving the reliability of applications and virtual machine monitors have also been proposed. Many techniques target application failures. Examples of these techniques are checkpoint-restarting methods [24], protecting the system from code injection attacks [25], [26], diagnosing failures and patching online [27], [28], and changing application execution environments [29]. Roothammer [30] achieves a fast VMM rejuvenation by preserving the running VMs in memory while rebooting the VMM. These approaches are complementary to the phase-based reboot, whose target is kernel failures.

The phase-based reboot is also complementary to other techniques in OS error detection sensors such as an SVA runtime mechanism [31] and software guards used in the XFI system [32]. These sensors can reduce our reboot recovery latency.

9. Conclusion

We proposed a "phase-based" reboot that shortens the downtime of reboot-based recovery. The key idea is to divide a boot sequence into phases. The phase-based reboot reuses a system state in the previous boot if the next boot reproduces the same state. By doing so, it skips some phases of a time-consuming

boot sequence that reproduces the same states as in the previous boot. A prototype of the phase-based reboot was implemented on Xen 3.4.1 running para-virtualized Linux 2.6.18. Experimental results showed that the prototype successfully recovered from kernel failures inserted by a kernel fault injector, and its downtime was 34.3% to 93.6% shorter than that of the normal reboot-based recovery.

Reference

- [1] Yamakita, K., Yamada, H. and Kono, K.: Phase-based Reboot: Reusing Operating System Execution Phases for Cheap Reboot-based Recovery, *Proc. 41st Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN 2011)*, pp.169–180 (2011).
- [2] Palix, N., Thomas, G., Saha, S., Calvés, C., Lawall, J. and Muller, G.: Faults in Linux: Ten Years Later, *Proc. 16th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '11)*, pp.305–318 (2011).
- [3] Baumann, R.C.: Soft errors in commercial semiconductor technology: Overview and scaling trends (2002).
- [4] Advisory, C.S.: Cisco catalyst memory leak vulnerability, ID:13618 (2001).
- [5] Castelli, V., Harper, R.E., Heidelberg, P., Hunter, S.W., Trivedi, K.S., Vaidyanathan, K. and Zeggert, W.P.: Proactive Management of Software Aging, *IBM Journal of Research and Development*, Vol.45, No.2, pp.311–332 (2001).
- [6] Waldspurger, C.A.: Memory Resource Management in VMware ESX Server, *Proc. 5th USENIX Symposium on Operating System Design and Implementation (OSDI '02)*, pp.181–194 (2002).
- [7] RUBiS, available from (<http://rubis.objectweb.org/>).
- [8] eBay.com, available from (<http://www.ebay.com/>).
- [9] Amazon.com: Amazon Elastic Compute Cloud (Amazon EC2), available from (<http://aws.amazon.com/ec2/>).
- [10] Depoutovitch, A. and Stumm, M.: Otherworld - Giving Applications a Change to Survive OS Kernel Crashes, *Proc. 5th European Conference on Computer Systems (EuroSys '10)*, pp.181–194 (2010).
- [11] Swift, M.M., Bershad, B.N. and Levy, H.M.: Improving the Reliability of Commodity Operating Systems, *Proc. 19th ACM Symposium on Operating Systems Principles (SOSP '03)*, pp.207–222 (2003).
- [12] Ng, W.T. and Chen, P.M.: The Systematic Improvement of Fault Tolerance in the Rio File Cache, *Proc. 1999 Symposium on Fault-Tolerant Computing (FTCS '99)*, pp.76–83 (1999).
- [13] Candea, G., Kawamoto, S., Fujiki, Y., Friedman, G. and Fox, A.: Microreboot - A Technique for Cheap Recovery, *Proc. 6th USENIX Symposium on Operating Systems Design and Implementation (OSDI '04)*, pp.31–44 (2004).
- [14] Nellitheertha, H.: Reboot Linux faster using kexec, available from (<http://www.ibm.com/developerworks/linux/library/l-kexec.html>).
- [15] Sun Microsystems: Using Fast Reboot on the x86 Platform (2008), available from (<http://dlc.sun.com/osol/docs/content/SYSADV1/ghsut.html/>).
- [16] Lenharth, A., Adve, V. and King, S.T.: Recovery Domains: An Organizing Principle for Recoverable Operating Systems, *Proc. 14th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '09)*, pp.49–60 (2009).
- [17] Swift, M.M., Annamalai, M., Bershad, B.N. and Levy, H.M.: Recovering Device Drivers, *Proc. 6th USENIX Symposium on Operating Systems Design and Implementation (OSDI '04)*, pp.1–16 (2004).
- [18] LeVasseur, J., Uhlig, V., Stoess, J. and Gätz, S.: Unmodified Device Driver Reuse and Improved System Dependability via Virtual Machines, *Proc. 6th USENIX Symposium on Operating Systems Design and Implementation (OSDI '04)*, pp.17–30 (2004).
- [19] Sundararaman, S., Subramanian, S., Rajimwale, A., Arpaci-Dusseau, A.C., Arpaci-Dusseau, R.H. and Swift, M.M.: Membrane: Operating System Support for Restartable File Systems, *Proc. 8th USENIX Conference on File and Storage Technologies (FAST '10)*, pp.281–294 (2010).
- [20] Bresoud, T.C. and Schneider, F.B.: Hypervisor-based Fault-tolerance, *Proc. 15th ACM Symposium on Operating Systems Principles (SOSP '95)*, pp.1–11 (1995).
- [21] Cully, B., Lefebvre, G., Meyer, D., Feeley, M., Hutchinson, N. and Warfield, A.: Remus: High Availability via Asynchronous Virtual Machine Replication, *Proc. 5th USENIX Symposium on Networked Systems Design and Implementation (NSDI '08)*, pp.161–174 (2008).
- [22] David, F.M., Chan, E.M., Carlyle, J.C. and Campbell, R.H.: CuriOS: Improving Reliability through Operating System Structure, *Proc. 8th USENIX Symposium on Operating Systems Design and Implementation (OSDI '08)*, pp.59–72 (2008).

- [23] Seltzer, M.I., Endo, Y., Small, C. and Smith, K.A.: Dealing With Disaster: Surviving Misbehaved Kernel Extensions, *Proc. 2nd USENIX Symposium on Operating Systems Design and Implementation (OSDI '96)*, pp.213–227 (1996).
- [24] Elnozahy, E.N.M., Alvisi, L., Wang, Y.-M. and Johnson, D.B.: A Survey of Rollback-Recovery Protocols in Message-Passing Systems, *ACM Computer Surveys*, Vol.34, No.3, pp.375–408 (2002).
- [25] Etoh., J.: GCC Extension for Protecting Applications from Stack-Smashing Attacks, available from <http://www.tri.ibm.com/projects/security/ssp>.
- [26] Team, P.: Address Space Layout Randomization (2003), available from <http://pax.grsecurity.net/docs/aslr.txt>.
- [27] Gao, Q., Zhang, W., Tang, Y. and Qin, F.: First-Aid: Surviving and Preventing Memory Management Bugs during Production Runs, *Proc. 4th ACM European Conference on Computer Systems (EuroSys '09)*, pp.159–172 (2009).
- [28] Sidirolou, S., Laadan, O., Perez, C.R., Viennot, N., Nieh, J. and Keromytis, A.D.: ASSURE: Automatic Software Self-healing Using REscue points, *Proc. 14th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '09)*, pp.37–48 (2009).
- [29] Qin, F., Tucek, J., Sundaresan, J. and Zhou, Y.: Rx: Treating Bugs As Allergies - A Safe Method to Survive Software Failures, *Proc. 20th ACM Symposium on Operating Systems Principles (SOSP '05)*, pp.235–248 (2005).
- [30] Kourai, K. and Chiba, S.: A Fast Rejuvenation Technique for Server Consolidation with Virtual Machines, *Proc. 37th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN '07)*, pp.245–255 (2007).
- [31] Criswell, J., Lenharth, A., Dhurjati, D. and Adve, V.: Secure Virtual Architecture: A Safe Execution Environment for Commodity Operating Systems, *Proc. 21st ACM Symposium on Operating Systems Principles (SOSP '07)*, pp.31–44 (2007).
- [32] Erlingsson, Ú., Abadi, M., Budiu, M. and Necula, G.C.: XFI: Software Guards for System Address Spaces, *Proc. 7th USENIX Symposium on Operating Systems Design and Implementation (OSDI '06)*, pp.75–88 (2006).

Kenji Kono received his B.Sc. degree in 1993, M.Sc. degree in 1995, and Ph.D. degree in 2000, all in computer science from the University of Tokyo. He is an associate professor of the Department of Information and Computer Science at Keio University. His research interests include operating systems, system software, and

Internet security. He is a member of the IEEE/CS, ACM and USENIX.

Kazuya Yamakita was born in 1987. He received his B.E. degree from Keio University in 2010. He is currently a M.E. student in Keio University. His research interests include operating systems, virtualization, and dependable systems.

Hiroshi Yamada was born in 1981. He received his B.E. and M.E. degrees from the University of Electro-communications in 2004 and 2006, respectively. He received his Ph.D. degree from Keio University in 2009. He is currently a research associate of the Faculty of Science and Technology at Keio University. His research

interests include operating systems, virtualization, and dependable systems. He is a member of ACM, USENIX and IEEE/CS.