Always Up-to-date – Scalable Offline Patching of VM Images in a Compute Cloud

Wu Zhou Peng Ning
North Carolina State University
{wzhou2, pning}@ncsu.edu

Xiaolan Zhang Glenn Ammons
IBM T.J. Watson Research Center
{cxzhang, ammons}@us.ibm.com

Ruowen Wang
North Carolina State University
rwang9@ncsu.edu

Vasanth Bala
IBM T.J. Watson Research Center
vbala@us.ibm.com

ABSTRACT

Patching is a critical security service that keeps computer systems up to date and defends against security threats. Existing patching systems all require running systems. With the increasing adoption of virtualization and cloud computing services, there is a growing number of dormant virtual machine (VM) images. Such VM images cannot benefit from existing patching systems, and thus are often left vulnerable to emerging security threats. It is possible to bring VM images online, apply patches, and capture the VMs back to dormant images. However, such approaches suffer from unpredictability, performance challenges, and high operational costs, particularly in large-scale compute clouds where there could be thousands of dormant VM images.

This paper presents a novel tool named *Nüwa* that enables efficient and scalable offline patching of dormant VM images. Nüwa analyzes patches and, when possible, converts them into patches that can be applied offline by rewriting the patching scripts. Nüwa also leverages the VM image manipulation technologies offered by the Mirage image library to provide an efficient and scalable way to patch VM images in batch. Nüwa has been evaluated on freshly built images and on real-world images from the IBM Research Compute Cloud (RC2), a compute cloud used by IBM researchers worldwide. When applying security patches to a fresh installation of Ubuntu-8.04, Nüwa successfully applies 402 of 406 patches. It speeds up the patching process by more than 4 times compared to the online approach and by another 2–10 times when integrated with Mirage. Nüwa also successfully applies the 10 latest security updates to all VM images in RC2.

1. INTRODUCTION

Patching is a basic and effective mechanism for computer systems to defend against most, although not all, security threats, such as viruses, rootkits, and worms [13, 19, 21]. Failing to promptly patch physical machines can subject the systems to huge risks, such as loss of confidential data, compromise of system integrity, and failure to provide regular system services. Unfortunately, applying security patches is a notoriously tedious task, due to the large number of patches and the high rate at which they are released — it is estimated that, in an average week, vendors and security organizations release about 150 vulnerabilities and associated patching information [15]. As a result, most software runs with outdated patches [11, 12].

The problem is exacerbated by the IT industry's recent shift to virtualization and cloud computing. Virtualization allows a complete system state to be conveniently encapsulated in a virtual machine (VM) image, which can run on any compatible hypervisor. Based on virtualization, cloud computing services (e.g., Amazon Elastic Compute Cloud (EC2) [2], NCSU Virtual Computing Lab (VCL) [20]) provide on-demand computing resources to customers' workloads, usually encapsulated in VM images. Because VM images are normal files, they can be easily copied to create new VM images. This has led to a new "VM image sprawl" problem, where a large number of VM images are created by the users and left unattended. A direct result of the VM image sprawl problem is the significantly increased management cost of maintaining these VM images, including regularly applying security patches to both active VMs and dormant VM images.

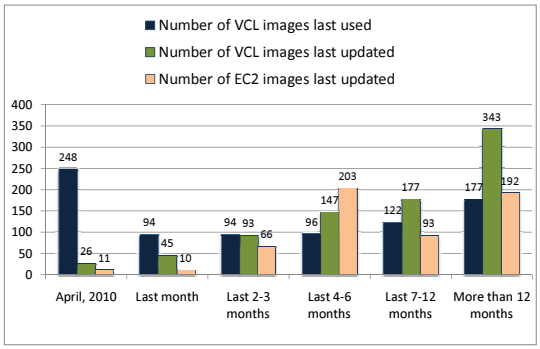

Figure 1: Counts of EC2 and VCL images, grouped by their time of last update or use (data collected on April 15, 2010)

A Glance at Two Compute Clouds: Figure 1 shows how recently VM images in two operational compute clouds, Amazon EC2 [2] and VCL [20], were updated; for VCL, the figure also shows how recently images were used. [1] There are a total of 831

[1]The VCL data was provided by the VCL management team, while the EC2 data was retrieved from the public Amazon Machine Images (AMIs) listed at Amazon's AMI page (`http://developer.amazonwebservices.com/connect/`

ACSAC '10 Dec. 6-10, 2010, Austin, Texas USA

VM images in VCL and 575 public VM images posted at EC2's AMI page. However, more than 91% of the VCL images and more than 96% of the EC2 images have not been updated for at least 1.5 months. Moreover, more than 58% of the VCL images have not been used in the last 1.5 months. Note that it is not the case that these inactive images will not be needed in the future. Indeed, based on the VCL log, VCL purged 776 VM images marked by the users as "deleted" in the past; all of the remaining 831 images were explicitly marked as needed by their owners.

Our investigation of EC2 and VCL leads to two observations:

- Most VM images in compute clouds are not properly patched. The longer a VM image remains unpatched, particularly after a major vulnerability is discovered, the more likely it is to threaten other machines in the compute cloud or in the Internet. Also, unpatched images owned by organizations or companies may not be compliant with the organizations' security policies.
- A significant portion of the VM images are mostly offline and infrequently booted. Thus, any attempt to start these VMs and install patches will be an extra cost to the image owners. The cloud service providers may certainly offer patching as a free service; however, they will have to sacrifice CPU cycles that could potentially bring in revenues.

Inadequacy of Existing Patching Utilities: Traditional patching utilities, originally designed for running systems, require the VM images to be online before the patch can be applied. There are a few recent attempts to patch offline VM images using traditional patching utilities. For example, the Microsoft Offline Virtual Machine Servicing Tool [10] brings the VM image online, applies the patches, and captures the VM back to a dormant image. Realizing that not all VM images are needed immediately by their users, a lazy patching approach was developed in [27], which injects the patch installer and patch data into the VM image in such a way that the patch process is triggered at the next booting time. This optimization can yield significant savings in the total time spent in patching in the case where only a small percentage of dormant images will ever be used. However, the tradeoff is that users will now see delays in image startup, which can be significant for images that have accumulated a long list of yet-to-be-applied patches. Our own experiences show that update time can be fairly long (in the order of 10s of minutes) for stale systems (e.g., dormant for 1 month). In modern clouds where VM instances are dynamically provisioned to meet varying demands, this delay is unacceptable. Additionally, for enterprises systems, it is often required that all IT assets (physical or virtual, dormant or online) be up to date with regard to patches for security or compliance reasons. This will apply to cloud providers as enterprises embrace the cloud computing model. Finally, in a cloud environment where customers are charged for resources used during patching, this approach imposes costs that customers might not accept.

In general, patching approaches that require VMs to be online are a poor fit for VM images in compute clouds. Note that it takes on the order of minutes just to power up and shut down a VM image. With the large number of dormant VM images that are infrequently used, these approaches add significant extra costs either for customers or for cloud service providers. In addition to these costs, bringing a VM image online necessarily runs code that has nothing to do with patching, which makes patching less predictable.

`kbcategory.jspa?categoryID=171`). There are indeed more public AMIs in EC2 (more than 7,000 in US East, US West, and EU West EC2 sites in mid April 2010) than those in this list. Amazon does not publish usage data.

Our Solution–Nüwa Offline Patching Tool: We propose an approach that is fundamentally different from the traditional online model. We argue that the only way to make the patching process scalable in a cloud environment, where the number of images can potentially reach millions [2], is to do it offline. A closer look into the patching process reveals that it can be decomposed into a sequence of actions, not all of which require a running system. In fact, most of the patching actions only depend on and have an impact on file system objects, which are already encapsulated in the VM image itself. Among the actions that do depend on or have impacts on a running system, we find that many are unnecessary when patching offline, and some can be safely replaced by other actions that do not need the running system. Based on these findings, we design and implement Nüwa [3], a scalable offline patching tool for VM images. By patching offline, Nüwa avoids the expensive VM start and stop time, and, for the majority of cases, ensures that, when a VM image is ready to be started, it has the latest patches installed.

Because Nüwa is an offline patching tool, it can leverage novel VM image manipulation technologies to further improve scalability. In particular, Nüwa is integrated with the Mirage image library [24], which stores identical files once and treats images as logical views on this collection of files. By exploiting Mirage, Nüwa can patch all images that contain a file by patching that single file and updating each image's view, thus providing efficient and scalable offline patching in batch.

Our implementation of Nüwa supports the Debian package manager [5] and the RPM package manager [8]. We evaluated Nüwa with 406 patches to a freshly installed Ubuntu-8.04. Our evaluation shows that Nüwa applies 402 of the 406 patches offline and speeds up the patching process by more than 4 times compared to the online approach. This can be further improved by another 2–10 times when the tool is integrated with Mirage, making Nüwa an order of magnitude more efficient than the online approach. We also evaluated Nüwa on real-world images from the IBM Research Compute Cloud (RC2) [25], a compute cloud used by IBM researchers worldwide. Nüwa successfully applies the 10 latest security updates to all VM images in RC2.

This paper is organized as follows. Section 2 gives background information on patching and describes our design choices and technical challenges. Section 3 presents an overview of our approach. Section 4 describes the mechanisms we use to convert an online patch into one that can be safely applied offline. Section 5 describes how we leverage efficient image manipulation mechanisms to further improve scalability. Section 6 presents our experimental evaluation results. Section 7 discusses related work. Section 8 concludes this paper with an outlook to the future.

2. PROBLEM STATEMENT

2.1 Background

Software patches, or simply patches, are often distributed in the form of software update packages (e.g., *.deb* or *.rpm* files), which are installed using a package installer, such as `dpkg` and `rpm`. In this section, we give background information on the format of software packages and the package installation process. We use the Debian package management tool `dpkg` as an example. Most software package management tools follow the same general style with only slight differences.

[2]Amazon EC2 already contains over 7,000 public VM images as of April 2010, without including private images that users choose not to share with others [18].

[3]Named after the Chinese Goddess who patches the sky.

Packages are distribution units of specific software. A package usually includes files for different purposes and associated metadata, such as the name, version, dependences, description and concrete instructions on how to install and uninstall this specific software. Different platforms may use different package formats to distribute software to their users. But the contents are mostly the same. A Debian package, for example, is a standard Unix `ar` archive, composed of two compressed tar archives, one for the filesystem tree data and the other for associated metadata for controlling purposes. Inside the metadata, a Debian package includes a list of configuration files, md5 sums for each file in the first archive, name and version information, and shell scripts that the package installer runs at specific points in the package lifecycle.

The main action in patching is to replace the old buggy filesystem data with the updated counterparts. Moreover, the package installer also needs to perform additional operations to ensure the updated software will work well in the target environment. For example, dependences and conflicts must be resolved, a new user or group might have to be added, configuration modifications by the user should be kept, other software packages dependent on this one may need to be notified, and running instances of this software may need to be restarted. Most of these actions are specified in scripts provided by the package developers. Because these scripts are intended to be invoked at certain points during the patching process, they are called *hook scripts*. The hook scripts that are invoked before (or after) file replacement operations are called *pre-installation* (or *post-installation*) *scripts*. There are also scripts intended to be invoked when relevant packages (e.g., dependent software) are installed or removed.

More details about Debian package management tools can be found in the Debian Policy Manual [6].

2.2 Design Choices and Technical Challenges

Our goal is to build a patching tool that can take *existing* patches intended for online systems and apply them *offline* to a large collection of dormant VM images in a manner that is *safe* and *scalable*. By safety we mean that applying the patch offline achieves the same effect on the persistent file systems in the images as applying it online. By scalability we mean that the tool has to scale to thousands, if not millions of VM images. In this paper we only consider dormant VM images that are completely shutdown; VM images that contain suspended VMs are out of the scope of this paper.

We made a conscious design decision to be backward compatible with an existing patch format. It is tempting to go with a "clean slate" approach, where we define a new VM-friendly patch format and associated tools that do not make the assumption of a running system at the time of patch application. While this is indeed our long-term research goal, we think its adoption will likely take a long time, given the long history of the traditional online patching model and the fact that it is an entrenched part of today's IT practices, ranging from software development and distribution to system administration. Thus, we believe that an interim solution that is backward compatible with existing patch format, and yet works in an offline manner and provides much improved scalability, would be desirable.

Several technical challenges arise in developing such a scalable offline patching tool, as discussed below:

Identifying Runtime Dependences: The current software industry is centered around running systems and so are the available patching solutions. A running system provides a convenient environment to execute the installation scripts in the patch. The installation scripts query the configuration of the running system to customize the patch appropriately for the system. Some scripts also restart the patched software at the end of the patching process to ensure its effect takes place. Some patches require running daemons. For example, some software stores configuration data in a database. A patch that changes the configuration requires the database server to be running in order to perform schema updates.

The challenge is to separate runtime dependences that can be safely emulated (such as information discovery that only depends on the file system state) or removed (such as restarting the patched software) from the ones that cannot (such as starting a database server to do schema updates). We address this challenge by a combination of manual inspection of commands commonly used in scripts (performed only once before any offline patching) and static analysis of the scripts.

Removing Runtime Dependencies: Once we identify runtime dependences that can be safely emulated or removed, the next challenge is to safely remove these dependences so that the patch can be applied to a VM image offline and in a manner that does not break backward compatibility. Our solution uses a script rewriting approach that preserves the patch format and allows a patch intended for an online system to be applied safely offline in an emulated environment.

Patching at a Massive Scale: As the adoption of virtualization and cloud computing accelerates, it is a matter of time before a cloud administrator is confronted with a collection of thousands, if not millions of VM images. Just moving from online to offline patching is not sufficient to scale to image libraries of that magnitude. We address this challenge by leveraging Mirage's capabilities in efficient storage and manipulation of VM images [24].

3. APPROACH

It seems plausible that patching VM images offline would work, given the fact that the goal of patching is mainly to replace old software components, represented as files in the file system, with new ones. Indeed, to patch an offline VM image, we only care about the changes made to the file system in the VM image; many changes intended for a running system do not contribute to the VM image directly.

Simple Emulation-based Patching: One straightforward approach is to perform the file replacement actions from another host, referred to as the *patching host*. The patching host can mount and access an offline VM image as a part of its own file system. Using the `chroot` system call to change the root file system to the mount point, the patching host can emulate an environment required by the patching process on a running VM and perform the file system actions originally developed for patching a running VM. We call this approach *simple emulation-based patching* and the environment set up by the above procedure the *emulated environment*.

Failures and Observations: Unfortunately, our investigation shows that the installation scripts used by the patching process pose a great challenge to simple emulation-based patching. For example, Figure 2 shows two segments of code from `dbus.postinst`, the post-installation script in the `dbus` package. The first segment (lines 1 to 7) detects possibly running `dbus` processes and sends a reboot notification to the system if there exists one. The second segment (lines 9 to 16) restarts the patched `dbus` daemon so that the system begins to use the updated software. Both segments depend on a running VM to work correctly. The simple emulation-based patching will fail when coming across this script.

We looked into the internals of patching scripts. After analyzing patching scripts in more than one thousand patching instances, we made some important observations. First, most commands used in the patching scripts are *safe* to execute in the emulated environment, in the sense that *they do not generate undesirable side*

```
1 if [ "$1" = "configure" ]; then
2  if [ -e /var/run/dbus/pid ] &&
3     ps -p $(cat /var/run/dbus/pid); then
4     /usr/share/update-notifier/notify-reboot-required
5     ...
6  fi
7 fi
8 ...
9 if [ -x "/etc/init.d/dbus" ]; then
10  update-rc.d dbus start 12 2 3 4 5 . stop 88 1 .
11  if [ -x "`which invoke-rc.d`" ]; then
12     invoke-rc.d dbus start
13  else
14     /etc/init.d/dbus start
15  fi
16 fi
```

Figure 2: Excerpts of the `dbus.postinst` script

effects on the persistent file system that would make the patched VM image different from one patched online except for log files and timestamps. Examples of such commands include the test commands in lines 2, 9 and 11, `cat` in line 3, `/usr/share/update-notifier/notify-reboot-required` in line 4, `update-rc.d` in line 10, and `which` in line 11. Second, some command executions have no impact on the offline patching and thus can be skipped. For example, `invoke-rc.d` in line 12 of Figure 2 is supposed to start up a running daemon, and its execution has no impact on the persistent file system. Thus, we can just skip it. We call such code *unnecessary code*. Third, there are usually more than one way to achieve the same purpose. Thus, it is possible to replace an unsafe command with a safe one to achieve the same effect. For example, many scripts use `uname -m` to get the machine architecture; unfortunately, `uname -m` returns the architecture of the patching host, which is not necessarily the architecture for which the VM image is intended. We can achieve the same purpose by looking at the file system data, for example, the architecture information in the ELF header of a binary file.

Safety Analysis and Script Rewriting: Motivated by the above observations, in this paper, we propose a systematic approach that combines safety analysis and script rewriting techniques to address the challenge posed by scripts. The safety analysis examines whether it is safe to execute a script in the emulated environment, while the rewriting techniques modify unsafe scripts to either eliminate unsafe and unnecessary code, or replace unsafe code with safe one that achieves the same purpose. Our experience in this research indicates that the majority of unsafe scripts can be rewritten into safe ones, and thus enable patches to be applied to offline VM images in the emulated environment.

However, not all scripts can be handled successfully in this way. We find some patching instances, after safety analysis and rewriting, still unsafe in the emulation-based environment. Some patches have requirements that can only be handled in a running environment. For example, the post-installation script in a patch for MySQL may need to start a transaction to update the administrative tables of the patched server. As another example, `mono`, the open source implementation of C# and the Common Language Runtime, depends on a running environment to apply the update to itself.

The Nüwa Approach: To address this problem, we adopt a hybrid approach in the development of Nüwa. When presented with a patch, Nüwa first performs safety analysis on the patching scripts included in the original patch. If all scripts are safe, Nüwa uses simple emulation-based patching directly to perform offline patching. If some scripts are unsafe, Nüwa applies various rewriting techniques, which will be discussed in detail in Section 4, to these scripts, and performs safety analysis on the rewritten scripts. If these rewriting techniques can successfully convert the unsafe scripts to safe ones, Nüwa will use simple emulation-based patching with the rewritten patch to finish offline patching. However, in the worst case, Nüwa may fail to derive safe scripts through rewriting, and will resort to online patching. In reality, we have found such cases to be rare – our results show that less than 1% of the packages tested in our experiments fall into this category (Section 6.1).

In addition to patching individual VM images, Nüwa also leverages VM image manipulation technologies to further improve scalability. In particular, Nüwa uses features of the Mirage image library [24] to enable scalable patching of a large number of VM images in batch.

To distinguish between the two variations of Nüwa, we refer to the former as *standalone Nüwa*, and the latter, which leverages Mirage, as *Mirage-based Nüwa*. In the following, we describe the novel techniques developed for offline patching in the context of both standalone and Mirage-based Nüwa.

4. SCRIPT ANALYSIS AND REWRITING

This section explains how safe patch scripts are identified and, when possible, unsafe scripts are transformed into safe scripts. The analysis is based on three concepts — impact, dependence, and command classification, which are defined in Section 4.1. Section 4.2 presents rewriting techniques that, using information from safety analyses, convert many unsafe scripts into safe scripts.

In our implementation, safety analysis and script-rewriting run immediately before the package manager (i.e., `dpkg` and `rpm`) executes a patch script. As a result, analyses and transformations have access to the script's actual environment and arguments and to the image's filesystem state.

Patch scripts are in general shell scripts. For example, patch scripts in Debian are SUSv3 Shell Command Language scripts [17] with three additional features mandated by the Debian Policy Manual [6]. Patch scripts are executed by an interpreter that repeatedly reads a command line, expands it according to a number of expansion and quoting rules into a command and arguments, executes the command on the arguments, and collects the execution's output and exit status. The language is very dynamic (for example, command-lines are constructed and parsed dynamically), which forces our analyses and transformations to be conservative. Nonetheless, simple, syntax-directed analyses and rewritings suffice to convert unsafe scripts to safe versions for 99% of the packages we considered.

4.1 Impact, Dependence, and Command Classification

The goal of command classification is to divide a script's command lines into three categories: (1) safe to execute offline, (2) unsafe to execute offline, and (3) unnecessary to execute offline. To classify command lines, we divide a running system into a "memory" part and a "filesystem" part, and determine which parts may influence or be influenced by a given command line. The intuition is that the "filesystem" part is available offline but the "memory" part requires a running instance of the image that is being patched.

Table 1: Commands w/ FS-only impacts

Command Type	Example Commands
File attribute mod.	`chown, chmod, chgrp, touch`
Explicit file content mod.	`cp, mv, mknode, mktemp`
Implicit file content mod.	`adduser, addgrp, remove-shell`

We say that a command-line execution *depends on the filesystem* if it reads data from the filesystem or if any of its arguments or inputs flow from executions that depend on the filesystem. An

execution *impacts the filesystem* if it writes data to the filesystem or if its output or exit status flow to executions that impact the filesystem. Table 1 lists some commands whose executions impact the filesystem:

We say that a command-line execution *depends on memory* if it inspects any of a number of volatile components of the system's state (perhaps by listing running processes, opening a device, connecting to a daemon or network service, or reading a file under `/proc` that exposes kernel state) or any of its arguments or inputs flow from executions that depend on memory. An execution *impacts memory* if it makes a change to a volatile component of the system's state that outlives the execution itself, or if its output or exit status flow to executions that impact the memory.

Note that all executions have transient effects on volatile state: they allocate memory, create processes, cause the operating system to buffer filesystem data, and so forth. For the purposes of classification, we do not consider these effects to be impacts on memory; we assume that other command-line executions do not depend on these sorts of effects. Table 2 lists some commands that impact or depend on memory.

Table 2: Commands w/ memory impact/dependence

Command Type	Example Commands
Daemon start/stop	`invoke-rc.d, /etc/init.d/`
Process status	`ps, pidof, pgrep, lsof, kill`
System info. inquiry	`uname, lspci, laptop-detect`
Kernel module	`lsmod, modprobe`
Others	Database update, mono gac-install

The definitions for command-line executions are extended to definitions for static command lines. A command line depends on memory (or the filesystem) if any of its executions depend on memory (or the filesystem). A command line impacts memory (or the filesystem) if any of its executions impact memory (or the filesystem).

To seed impact and dependence analysis, we manually inspected all commands used in patch scripts to determine their intrinsic memory and filesystem impacts and dependences. This might seem to be an overwhelming task but, in practice, scripts use very few distinct commands; we found only about 200 distinct commands used by more than 1,000 packages. It may be possible to derive this information by instrumenting command executions. In practice, we expect that it would be provided by package maintainers.

Table 3: Command classification

Depend on FS	Depend on Memory	Impact on Memory	Impact on FS	Safety
Yes/No	No	No	Yes/No	Safe
Yes/No	No	Yes	Yes	Unsafe
Yes/No	Yes	No	Yes	Unsafe
Yes/No	Yes	Yes	Yes	Unsafe
Yes/No	No	Yes	No	Unnecessary
Yes/No	Yes	No	No	Unnecessary
Yes/No	Yes	Yes	No	Unnecessary

Our analysis concludes that a static command-line depends on memory if one of the following holds: (1) The command is unknown; (2) the command has an intrinsic memory dependence; (3) one or more of the arguments is a variable substitution; (4) the input is piped from a command that depends on memory; or (5) the input is redirected from a device, a file under `/proc`, or from a variable substitution.

The rules for filesystem dependences and for impacts are similar. Note that the analysis errs on the side of finding spurious dependences and impacts. That is, these analyses are simple "may-depend/may-impact" analyses, which are both flow and context insensitive.

Table 3 shows how each command line's classification as safe, unsafe, or unnecessary is determined from its filesystem and memory impacts and dependences. Safe command lines do not depend on or impact memory. These are the commands that can and should be executed offline. Script rewriting preserves these commands. Unnecessary command lines have no impact on the filesystem. There is no reason to execute them offline because they do not change the image. In fact, if they depend on or impact memory, then they must be removed because they might fail without a running instance. Script rewriting removes these commands. Unsafe command lines may execute incorrectly offline because they depend on or impact memory and also impact the filesystem. In some cases, script rewriting cannot remove these command lines because their filesystem impacts are required. If any unsafe command line cannot be removed, then the patch cannot be executed offline.

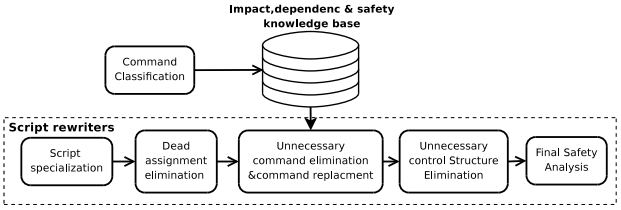

Figure 3: Flow of script analysis and rewriting

4.2 Rewriting Techniques

Figure 3 shows the rewriting techniques that Nüwa applies before executing each patch script. Rewriting a script can change the results of safety analysis, so Nüwa reruns safety analysis after applying these techniques. If safety analysis proves that all command lines in the script are safe, then the rewritten script is executed offline. Otherwise, Nüwa resorts to online patching.

Nüwa currently applies five rewriting techniques, which are described below. For clarity, the presentation does not follow the order in which the techniques are applied (that order is shown in Figure 3). The first two techniques consider command-lines, annotated by safety analysis, in isolation; the last three analyze larger scopes.

Unnecessary Command Elimination: This technique removes unnecessary commands, which, by definition, have neither direct nor indirect impact on the filesystem. Figure 4 shows an example.

```
/etc/init.d/acpid
/etc/init.d/cupsys
killall
```

Figure 4: Examples of command lines that are removed by unnecessary command elimination

Command Replacement: Some command lines that depend on memory can be replaced with command lines that depend only on the filesystem. This often happens with commands that need information about the system, in particular when the information is available both in the filesystem and, if there is a running instance, in memory.

```
uname -m
  -> dpkg --print-architecture
```

```
uname -s
  -> echo "Linux"
```

Figure 5: Memory-dependent command lines and their replacements

For example, the `uname` command prints system information; depending on its arguments, it will print the hostname, the machine hardware name, the operating system name, or other fields. `uname` gets its information from the kernel through the `uname` system call. Without a running instance, information from the kernel cannot be trusted. However, certain fields are statically known constants or available through commands that depend only on the filesystem; Figure 5 shows two examples.

Note that the command replacement technique not only removes memory-dependent commands but also ensures that the offline script uses values appropriate to the image instead of values from the host. Nüwa's implementation of command replacement consults a manually constructed table of command lines and their known replacements.

Before rewriting:

```
1  if [ -x "‘which invoke-rc.d‘" ]; then
2     invoke-rc.d dbus start
3  else
4     /etc/init.d/dbus start
5  fi
```

After rewriting:

```
All eliminated
```

Figure 6: Example of control structure analysis (from `dbus.postinst`)

Unnecessary Control-structure Elimination: This technique, a generalization of unnecessary command elimination, removes compound commands like `if` and `case` statements.

Figure 6 shows an example. Both the true branch and the false branch of the `if`-statement are unnecessary and would be eliminated by unnecessary command elimination. The conditional would not be eliminated because safety analysis conservatively assumes that all conditionals impact both memory and the filesystem through control-flow. By contrast, unnecessary control-structure elimination eliminates the entire `if`-statement because, after eliminating both branches of the `if`-statement, the conditional is unnecessary: It clearly has no filesystem impact through control-flow or any other means.

We perform unnecessary control-structure elimination in a bottom-up fashion (i.e., process inner control structures first). For each control structure being processed, we first try to eliminate all statements in each branch of the structure. If all statements in every branch can be eliminated, we consider the conditional itself: If it no longer impacts the filesystem, the entire control structure is removed.

Note that Nüwa applies unnecessary control-structure analysis to many kinds of compound commands and command lists, including the `case` construct and command lists built from the short-circuiting statements (|| and &&).

Script Specialization: This technique removes command lines and control structures that cannot execute, given the script's actual environment and arguments and the VM image's filesystem state. Recall that this context is available because safety analysis and script-rewriting run immediately before `dpkg` executes a patch script.

Figure 7 shows an example, which was extracted from the post-installation script for the `acpid` package. Except during error recovery, `dpkg` calls post-installation scripts with `configure` as the first positional parameter (`$1`). Therefore, the `case` statement can be replaced with the first branch. Next, since the rest of the script changes neither `/var/run/hald/hald.pid` nor `/etc/init.d/hal`, the conditional can be evaluated at rewriting time; in this case, the conditional is false and the false branch is empty so the entire if statement is removed.

The current implementation of script specialization is a collection of ad hoc rewriting passes, which Nüwa applies before applying any other rewriting techniques. One pass replaces positional parameters with actual parameters. Another evaluates conditionals built from filesystem tests, when the tests depend only on the initial filesystem state. A third evaluates the command line `dpkg -compare-versions`, which is used frequently and whose result can be determined from the VM image's package database.

Before rewriting:

```
1  HAL_NEEDS_RESTARTING=no
2  case "$1" in
3    configure)
4      if [ -x /etc/init.d/hal ] &&
5        [ -f /var/run/hald/hald.pid ]; then
6        HAL_NEEDS_RESTARTING=yes
7        invoke-rc.d hal stop
8      fi
9      ;;
10   reconfigure)
11     ...
12 esac
```

After rewriting:

```
HAL_NEEDS_RESTARTING=no
```

Figure 7: Example of script specialization (from `acpid.postinst`)

All passes are conservative and err on the side of missing rewriting opportunities. For example, positional-parameter replacement leaves the script unchanged if the script uses the `shift` statement, which renames the positional parameters.

Dead-assignment Elimination: This technique removes assignments to unused variables. Some dead assignments come from the original scripts; others are created by script specialization, which can convert conditionally dead assignments to dead assignments.

Figure 8 shows an example of dead assignment, extracted from `xfonts-scalable.postinst`. In this script, the command `laptop-detect` is intrinsically memory-dependent. If its result flows to a command line that impacts the filesystem, the script would be unsafe. Fortunately, the `LAPTOP` variable is unused in the rest of the script. Removing its assignment leaves the body of the inner `if` statement empty, which makes the conditional unnecessary, which in turn allows the entire inner `if` statement to be removed. The outer `if` statement is then removed in a similar fashion.

Before rewriting:

```
LAPTOP=""
if [ -n "$(which laptop-detect)" ]; then
    if laptop-detect >/dev/null; then
            LAPTOP=true
    fi
fi
```

After rewriting:

```
All eliminated
```

Figure 8: Example of dead-assignment elimination

The first assignment in Figure 7, which is conditionally dead in the original script, could be transformed into a dead assignment by script specialization.

Dead-assignment elimination depends on a syntax-directed dataflow analysis of the main body of the script. An assignment is *dead* if the assigned value cannot reach a *use* before reaching the end of the script or another assignment; the analysis conservatively judges an assignment to be dead if it it does not occur in a loop and is followed by another assignment in the same syntactic scope, with no intervening uses in any syntactic scope, or if no uses follow at all. Function bodies are not considered, except that any use of a variable within a function body is considered reachable from any assignment to that variable in the entire program.

5. SCALABLE BATCH PATCHING

A motivating assumption of this work is that, as cloud computing becomes more widely adopted, image libraries will grow to contain thousands or perhaps even millions of images, many of which must be patched as new vulnerabilities are discovered. Even with the offline patching techniques presented in Section 4, patching so many images individually would take a significant amount of time.

This section explains an approach to batch patching a large number of images offline that exploits an observation and a conjecture about patching images. The observation is that, if the same patch is applied to two similar images, then any given patch-application step is likely to have the same effect on both images. For example, the same files will be extracted from the patch both times. The conjecture is that the images that must be patched are likely to be similar to one another; this conjecture seems particularly reasonable for clouds (such as Amazon's EC2 [2]) that encourage users to derive new images from a small set of base images.

Nüwa's batch patching harnesses Mirage, a scalable VM image storage solution that exploits the similarity between images [24]. We first give a brief overview of Mirage before describing the batch patching solution.

5.1 Overview of Mirage

The Mirage image library maintains a collection of VM images and provides an image-management interface to users: users can import images into the library, list existing images in the library, check out a particular image, and check in updates of the image to create either a new image or a new version of the original image. A separate interface allows system administrators to perform system-wide operations, such as backup, virus scan, and integrity verification of all image content.

A design goal of Mirage is to support operations on images as structured data. To this end, Mirage does not store images as simple disk images. Instead, when an image is imported into the library, Mirage iterates over the image's files, storing each file's contents as a separate item in a content-addressable store (CAS); the image as a whole is represented by a manifest that refers to file-content items and serves as a recipe for rebuilding the image when it is checked out. An earlier paper [24] described this format and explained how it allows certain operations on images to be expressed as fast operations on the image's manifest. For example, creating a file, assuming that the desired contents are already in the CAS, reduces to adding a few hundred bytes to the manifest.

Mirage's new *vmount* feature, which was not described in the earlier paper, allows users to mount library images without rebuilding them. Vmount is implemented as a FUSE [26] daemon and fetches data from the CAS as it is demanded; by contrast, checking out an image requires fetching every file's contents from the CAS. Vmount also implements a new extended filesystem attribute that allows direct manipulation of the manifest. For each regular file, the value of this attribute is a short, unique identifier of the file's contents. Setting the attribute atomically replaces the file's contents with new contents.

After modifying an image through Vmount, the user can check in the changes as a new image or as a new version of the original image. The original image is not disturbed, and the time to check in is proportional to the amount of new data instead of to the size of the image.

Vmount has three benefits for batch patching. First, there is no need to rebuild each image. Arguably, this is merely a workaround for a problem created by the decision to store images as manifests. Second, if two images share data in the CAS and are patched sequentially through Vmount, then reading the shared data the second time is likely to be fast, because the data will be in the host's buffer cache. By contrast, if two disk images are patched sequentially, then the fact that they share data is effectively hidden from the host's operating system. The largest benefit is that Vmount allows batch patching to operate on manifests without major modifications of system tools like `dpkg`. Time-critical patching steps can be changed to use the new filesystem attribute, without creating a dependence on the manifest format, while less profitable steps continue to use the normal filesystem interface.

5.2 Batch Patching via Mirage

A straightforward way to patch a batch of images is to iterate the patching process for individual images. For images in Mirage, each iteration mounts an image with Vmount, applies the patch [4], and checks in the modified image.

Our approach optimizes this straightforward approach by moving invariant operations out of the loop that visits each image. Currently, Nüwa optimizes one source of invariant operations: unpacking the patch, which copies the patch's files to the image and, ultimately, adds their contents to the Mirage CAS. These copies and CAS additions are good operations to move out of the loop because they consume most of the time of applying most patches; in future work, we plan to hoist more invariants out of the loop.

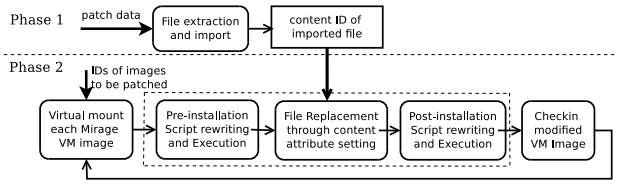

Figure 9: Batch patching VM images via Mirage

Figure 9 shows the two phases of batch patching via Mirage. Phase 1 performs the loop-invariant operation: Nüwa extracts the patch's files and imports them into Mirage. The result is a list of content identifiers, one for each file. In phase 2, Nüwa iterates over the images. For each image, Nüwa mounts the image with Vmount, rewrites and executes the pre-installation scripts, emulates the "unpack" step of the package manager (e.g., `dpkg`), using the Mirage filesystem attribute to set the contents of the patch's files, rewrites and executes the post-installation scripts, and checks in the modified VM image. If script rewriting ever fails to produce a safe script, then Nüwa resorts to online patching.

6. EXPERIMENTAL EVALUATION

We have implemented both standalone Nüwa and Mirage-based Nüwa by extending the Debian Almquist Shell (i.e., `dash`) [4]. (Our script rewriting was performed based on the syntax tree generated by `dash`.) Our implementations assume a Linux host sys-

[4]If the patch must be applied online, then the image must be rebuilt.

tem. We have tested the standalone Nüwa on patching hosts running CentOS 5.2, Ubuntu 9.0.4 and OpenSuSE 11.1. Our implementations currently support VM images of any Linux distributions based on Debian package management tools (e.g., Debian, Ubuntu, Knoppix) or RPM package manager (e.g., RHEL, CentOS). However, Mirage-based Nüwa currently only works on the Debian package manager, as the optimizations have not been completely ported to RPM yet.

We performed three sets of experiments to evaluate Nüwa, including (1) patching individual VM images offline, (2) Mirage-based offline patching in batch, and (3) patching VM images in a real-world compute cloud RC2. The first two sets of experiments were performed on a DELL OptiPlex 960 PC, with a 3GHz Intel Core 2 Duo CPU and 4GB DDR2 memory. The third set of experiments were performed in RC2. Unless otherwise noted, we used the x86-64 version of OpenSuSE 11.1 version as the patching host OS in all experiments. For compatibility reasons, we updated its kernel to version `2.6.31.11-0.0.0.2.9c60380-default`.

6.1 Patching Individual VM Images

The objective of this set of experiments is two-fold: First, we would like to evaluate the correctness of the offline patching approach used in Nüwa (i.e., whether the offline patching approach has the same effect on the VM images as online patching). Second, we would like to see the efficiency of our offline patching approach in Nüwa compared with the online patching approach.

In this set of experiments, we used the Linux Kernel-based Virtual Machine (KVM) [7] to start instances of VM images for online patching. For offline patching, we used the VMware disk library to mount the VM images in the host environment. Our tool can be logically decomposed into two parts: the script rewriter and the patch applier. We copied both components into the mounted VM image, with the patch applier replacing the original package installer inside the target VM image.

To perform the evaluation, we first created an empty disk image in the flat VMDK disk format with the `kvm-img` image creation tool, then brought this disk image online through KVM, and installed a default configured 64-bit Ubuntu-8.04 inside. This was used as the base VM image for both offline and online patching.

We gathered all 406 patches available for the base VM image (64-bit Ubuntu-8.04) on October 26, 2009. The correctness of offline patching is verified by a file-by-file comparison of the results of online and offline: If two VM images, which are obtained through patching the base VM image online and offline, respectively, differ only in log files and timestamps, we consider the offline patching to be correct. To further evaluate the effectiveness of the rewriting techniques, we used the simple emulation-based patching mentioned in Section 3 as a reference.

Table 4: Comparison of offline patching methods

	# successes	# failures	success ratio
Simple emulation	369	37	90.9%
Nüwa	402	4	99.0%

Table 4 shows the experimental results for evaluating the correctness of our techniques. Nüwa can successfully apply 402 out of the 406 patches offline, achieving a 99.0% success ratio. The results also show that the rewriting techniques contributed significantly to the success; they helped improve the success ratio by about 10%. Note that the failure cases are failures of offline patching, not of Nüwa; Nüwa automatically detects all of these failures and can cope with them through automatic online patching, as discussed in Section 3.

The four failure cases are the `mono-gac` package[5] and three other packages that depend on `mono-gac`. Through further analysis, we found that `mono-gac` failed because the installer needed to access some kernel information (e.g., `/proc/self/map` and `/proc/cpuinfo`) in order to work correctly. This information cannot be retrieved in the emulated environment.

To compare the efficiency of Nüwa's offline patching techniques with that of online patching, we performed another set of experiments. We assumed the most efficient form of online patch, automated online patching. Specifically, we insert the patch data into the VM image through the emulated environment and then schedule a patching process at boot time by modifying the booting script in the VM image. We then boot the VM, perform online patching, and shut down the VM automatically once the patching is complete.

We collected two sets of data from these experiments. The first is the time (in seconds) required to apply each applicable patch to the base VM image through the offline patching approach in Nüwa, and the second is the time needed to apply the same set of patches through automated online patching.

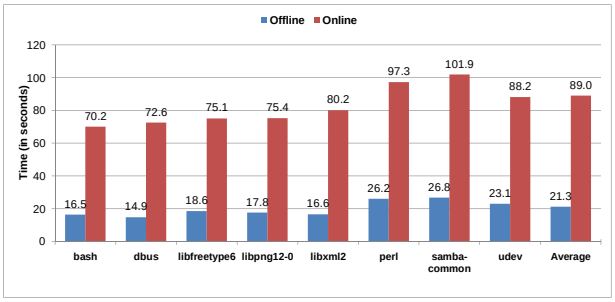

Figure 10: Time used by offline and online patching ("Average" is computed over 402 applicable packages)

Figure 10 shows the time (in seconds) required to apply some applicable patches to the base VM image through the Nüwa offline patching and the automated online patching, respectively. Due to the limited space, we only show the timing results for eight selected patches and the average for all 402 applicable patches. On average, the Nüwa offline approach takes only 23.9% of the time required by automated online patching (a factor of 4 speedup). This improvement, combined with the fact that Nüwa needs much less human intervention and physical resources, shows that it brings significant benefits to patching VM images.

This set of experiments demonstrates that Nüwa's offline patching techniques, particularly the rewriting techniques, are effective and that offline patching using Nüwa can significantly reduce the overhead involved in patching.

6.2 Batch Patching via Mirage

The primary objective of this set of experiments is to measure the scalability offered by Mirage-based Nüwa by comparing the performance of Mirage-based batch patching with that of one-by-one patching.

We generated 100 VM images using 32-bit Ubuntu 8.04 as the base operating system for this set of experiments. The Ubuntu installer can install a support for a number of basic, predefined tasks; some of these tasks are for running specific servers, while others are for desktop use. We generated test VM images from 100 randomly selected subsets of 12 of these tasks (listed in Table 5).

[5] `mono-gac` is a utility to maintain the global assembly cache of mono, an open source implementation of C# and the CLR.

Table 5: Basic Ubuntu tasks

#	Task Name	#	Task Name
1	lamp-server	2	mail-server
3	dns-server	4	openssh-server
5	print-server	6	samba-server
7	postgresql-server	8	ubuntustudio-audio
9	ubuntustudio-audio-plugins	10	ubuntustudio-graphics
11	ubuntustudio-video	12	ubuntu-desktop

We retrieved 154 security updates (i.e., security patches) for 32-bit Ubuntu 8.04 from Ubuntu Security Notices [28]. We also retrieved the ranking of Ubuntu packages given by Ubuntu's popularity contest [9], and sorted the 154 security patches accordingly. For our performance evalution, we selected the security updates corresponding to the eight most popular packages (as of January 18, 2010), including `dash`, `libdbus`, `libglib-2.0`, `udev`, `libfreetype`, `libpng`, `libxml2`, and `dbus`.

For each of the eight patches, we measured the time to apply the patch to the test VM images one-by-one and the time to apply the patch to the test VM images as batches of increasing sizes. Figure 11 shows that for all eight security patches, Mirage-Nüwa achieves considerable speedup over individual patching. Moreover, the speedup also increases as the number of images patched in a batch increases, and plateaus between 80 and 100 images.

For seven of the eight security patches (`udev` is the exception), the average speedup over one-by-one patching increases from 5.1 times to 8.5 times as the number of images in a batch increases from 10 to 100. Note that this speedup is on top of the factor of 4 speedup achieved over traditional online patching, thus bringing the total speedup over traditional online patching to about 30 when patching 100 images in a batch.

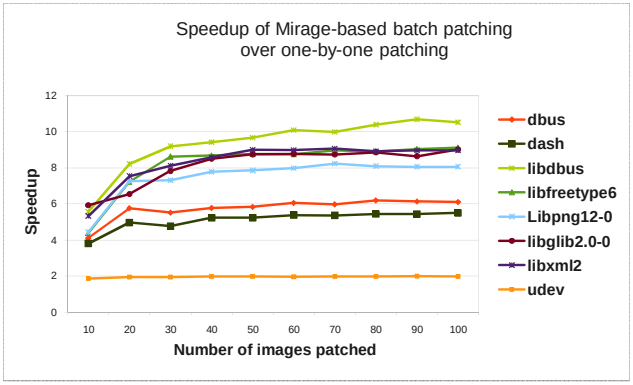

Figure 11: Scalability of Mirage-based Nüwa

However, the speedup for `udev` is much smaller, compared with the other seven patches. In fact, the speedup for `udev` is only around 2. Further investigation showed that the `udev` patch spends more time in pre-installation and post-installation scripts than the others; thus, the file replacement operations constitute a smaller portion of the entire patching process.

This set of experiments demonstrates that Mirage-based Nüwa is scalable and can improve the performance of offline patching significantly. Overall, Nüwa offline patching is an order of magnitude more efficient than online patching.

6.3 Patching a Real Cloud

In this experiment we assess the performance of Nüwa in a real production cloud. We patch the entire repository with the latest security updates published in the OS distributor's website. We set out to answer two questions: 1) how many of the images can be successfully patched offline using Nüwa, and 2) whether it is feasible to patch the *entire* image repository on a *daily* basis.

Our experiments are based on RC2 [25], a compute cloud very similar to Amazon's EC2, that is used by IBM researchers worldwide. Although small compared to EC2, RC2 is a production cloud that is used daily by IBM researchers. We created a replica of the RC2 image repository in our own testbed, so as to have a controlled experimental environment. The replica contains a total of 278 images, to which we apply the security patches from Red Hat's security advisories website [23]. All 278 images are running Red Hat 5.3 with the exception of one that is running CentOS 5.3. We used the RPM-based implementation of Nüwa since all Red Hat distributions use RPM for package installation. For this experiment we did not use the Mirage batch optimization because this feature has not yet been implemented in the RPM-based Nüwa.

We set up a dedicated host to run the patch process. The host is a blade with 4 Xeon 3.16GHz processors and 8GB RAM, running OpenSuse 11.1. The image repository is on a different, similar blade and the host accesses the repository via an NFS mount through a SAN network.

Table 6: Latest applicable security updates from RedHat rated "important" and higher

#	Update	Severity	Advisory
1	krb5	critical	RHSA-2010:0029
2	nspr/nss	critical	RHSA-2009:1186
3	openssl	important	RHSA-2010:0162
4	sudo	important	RHSA-2010:0122
5	acpid	important	RHSA-2009:1642
6	elinks	important	RHSA-2009:1471
7	dnsmasq	important	RHSA-2009:1238
8	bind	important	RHSA-2009:1179
9	cups	important	RHSA-2009:1082
10	freetype	important	RHSA-2009:1061

Patching the entire repository of 278 available images with the latest critical security update (krb5 [22]) takes about 45 seconds per image, totaling about 3.5 hours. All images were patched successfully, completely offline. Note that the patching time includes all time to set up the image for patching, download the update from a remote Red Hat Network server, and install the downloaded packages. We believe this number can be reduced 10-fold with an optimized storage configuration (e.g., a repository on a local disk or on direct-attached SAN storage), a local package server, and the Mirage batch-patching optimization, thus potentially allowing an average compute node (which can itself be a VM in the compute cloud) to apply a single security patch to about 19,200 (24*3600S/(45S/10)) images on a daily basis.

To test the robustness of Nüwa, we took the latest applicable security updates (shown in Table 6) from Red Hat's security advisories [23] that are rated "important" or "critical" and applied them across the entire repository. There are 10 updates which consist of 24 individual packages. All updates were successful on all 278 images, suggesting that Nüwa is robust enough to be offered as a real service in a production cloud.

7. RELATED WORK

Several available commercial tools [10,27,29] can apply patches to dormant VM images. But that does not mean the patches are applied in an *offline* manner. As a matter of fact, all of them require the image to be running when the patches are actually installed. Microsoft's Offline VM Servicing Tool [10] first "wakes" up the virtual machine (deploys it to a host and starts it), then triggers the appropriate software update cycle to apply the patches, and finally shuts down the updated virtual machine and returns it to the image

library. In the cases of VMware Update Manager [29] and Shavlik NetChk Protect [27], patches are first inserted into image at some specified locations, then applied when the image is powered up. We resort to this approach when Nüwa identifies patches that contain unsafe commands.

In some cases, it is preferable to apply patches online. In general, systems that tend to stay online for a long period of time, such as highly available servers, fall into this category. In those cases, "dynamic update" techniques [1, 3, 14, 16] are used to apply patches to the target software without shutting them down. In contrast, Nüwa targets VM images that have already been shut down and may stay in dormant state for an extended period of time. Thus, these approaches are complimentary to Nüwa.

8. CONCLUSION

In this paper, we developed a novel tool named Nüwa to enable efficient patching of offline VM images. Nüwa uses safety analysis and script rewriting techniques to convert patches, or more specifically the installation scripts contained in patches, which were originally developed for online updating, into a form that can be applied to VM images offline. Nüwa also leverages the VM image manipulation technologies offered by the Mirage image library [24] to provide an efficient and scalable way to patch VM images in batch. We implemented a standalone Nüwa and a Mirage-based Nüwa; standalone Nüwa supports two popular package managers, the Debian package manager [5] and the RPM package manager [8], while Mirage-based Nüwa supports only the former. In addition to evaluating Nüwa with security patches and VM images configured with popular packages according to Ubuntu popularity contest, we also applied Nüwa to a real cloud RC2. Our experimental results demonstrate that 1) Nüwa's safety analysis and script rewriting techniques are effective – Nüwa is able to convert more than 99% of the patches to safe versions that can then be applied offline to VM images; 2) the combination of offline patching with additional optimization made possible through Mirage allows Nüwa to be an order of magnitude more efficient than online patching; and 3) Nüwa successfully patched 278 images in a real compute cloud.

A limitation of Nüwa is that it currently does not support offline patching of a suspended VM image, which includes a snapshot of the system memory state in addition to the file system. In our future research, we will investigate techniques to patch suspended VM images and how to perform offline patching on Windows platforms.

Acknowledgement

This work is supported by the U.S. National Science Foundation (NSF) under grant 0910767, and by an IBM Open Collaboration Faculty Award. The contents of this paper do not necessarily reflect the position or the policies of the U.S. Government or IBM.

9. REFERENCES

[1] Gautam Altekar, Ilya Bagrak, Paul Burstein, and Andrew Schultz. Opus: online patches and updates for security. In *SSYM'05: Proceedings of the 14th conference on USENIX Security Symposium*, pages 19–19, Berkeley, CA, USA, 2005. USENIX Association.

[2] Amazon. Amazon elastic compute cloud (EC2). `http://aws.amazon.com/ec2/`.

[3] Jeff Arnold and M. Frans Kaashoek. Ksplice: automatic rebootless kernel updates. In *EuroSys '09: Proceedings of the 4th ACM European conference on Computer systems*, pages 187–198, New York, NY, USA, 2009. ACM.

[4] Debian community. Debian Almquist shell. `http://en.wikipedia.org/wiki/Debian_Almquist_shell`.

[5] Debian community. Debian package manager. `http://www.debian.org/dpkg`.

[6] Debian Community. Debian policy manual. `http://www.debian.org/doc/debian-policy/`, 2009.

[7] KVM community. Linux kernel-based virtual machine. `http://www.linux-kvm.org/`.

[8] RPM community. RPM package manager. `http://www.rpm.org/`.

[9] Ubuntu Community. Ubuntu popularity contest. `http://popcon.ubuntu.com/`.

[10] Microsoft Corporation. Offline virtual machine servicing tool 2.1. `http://technet.microsoft.com/en-us/library/cc501231.aspx`.

[11] Forbes. Cybersecurity's patching problem. `http://www.forbes.com/2009/09/14/sans-institute-software-technology-security-cybersecurity.html`. Visited on 2009-11-06.

[12] Stefan Frei, Thomas Duebendorfer, Gunter Ollmann, and Martin May. Understanding the Web browser threat. Technical Report 288, TIK, ETH Zurich, June 2008. Presented at DefCon 16, Aug 2008, Las Vegas, USA. `http://www.techzoom.net/insecurity-iceberg`.

[13] Thomas Gerace and Huseyin Cavusoglu. The critical elements of the patch management process. *Commun. ACM*, 52(8):117–121, 2009.

[14] Deepak Gupta and Pankaj Jalote. On line software version change using state transfer between processes. *Softw. Pract. Exper.*, 23(9):949–964, 1993.

[15] Huseyin Cavusoglu Hasan, Hasan Cavusoglu, and Jun Zhang. Economics of security patch management. In *The Fifth Workshop on the Economics of Information Security (WEIS 2006)*, June 2006.

[16] Michael Hicks and Scott M. Nettles. Dynamic software updating. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 27(6):1049–1096, November 2005.

[17] The IEEE and The Open Group. The single UNIX specification, version 3. `http://www.unix.org/version3/online.html`, 2004.

[18] Cloud Market. The cloud market: EC2 statistics. `http://thecloudmarket.com/stats`.

[19] Microsoft. The microsoft security update release cycle. `http://www.microsoft.com/security/msrc/whatwedo/updatecycle.aspx`.

[20] NC State University. NC State University virtual computing lab (VCL). `http://vcl.ncsu.edu/`.

[21] United States General Accounting Office. Effective patch management is critical to mitigating software vulnerabilities. gao-03-1138t, September 2003.

[22] RedHat. Critical: krb5 security update.

[23] RedHat. RedHat Security Advisories. `http://rhn.redhat.com/errata/rhel-server-errata-security.html`.

[24] D. Reimer, A. Thomas, G. Ammons, T. Mummert, B. Alpern, and V. Bala. Opening black boxes: using semantic information to combat virtual machine image sprawl. In *VEE '08: Proceedings of the fourth ACM SIGPLAN/SIGOPS international conference on Virtual execution environments*, pages 111–120, 2008.

[25] Kyung Dong Ryu, Xiaolan Zhang, Glenn Ammons, Vasanth Bala, Stefan Berger, Dilma M. Da Silva, Jim Doran, Frank Franco, Alexei Karve, Herb Lee, James A. Lindeman, Ajay Mohindra, Bob Oesterlin, Giovanni Pacifici, Dimitrios Pendarakis, Darrell Reimer, and Mariusz Sabath. RC2 – A Living Lab for Cloud Computing. In *Lisa '10: Proceedings of the 24th Large Installation System Administration*, 2010. Earlier version available as an IBM technical report at `http://domino.watson.ibm.com/library/CyberDig.nsf/Home`.

[26] Miklos Szeredi. Fuse: Filesystem in userspace. `http://fuse.sourceforge.net/`, 2010.

[27] Shavlik Technologies. Offline virtual machine image quick start guide. `http://www.shavlik.com/documents/qsg-prt-6-1-offline_vm.pdf`.

[28] Ubuntu. Ubuntu security notices. `http://www.ubuntu.com/usn/`.

[29] VMware. VMware vcenter update manager. `http://www.vmware.com/products/update-manager/`.