Splitting Interfaces: Making Trust Between Applications and Operating Systems Configurable

Richard Ta-Min   Lionel Litty*   David Lie
Department of Electrical and Computer Engineering
University of Toronto
{*tamin,llitty,lie*}@*eecg.toronto.edu*

Abstract

In current commodity systems, applications have no way of limiting their trust in the underlying operating system (OS), leaving them at the complete mercy of an attacker who gains control over the OS. In this work, we describe the design and implementation of Proxos, a system that allows applications to configure their trust in the OS by partitioning the system call interface into trusted and untrusted components. System call routing rules that indicate which system calls are to be handled by the untrusted commodity OS, and which are to be handled by a trusted *private OS*, are specified by the application developer. We find that rather than defining a new system call interface, routing system calls of an existing interface allows applications currently targeted towards commodity operating systems to isolate their most sensitive components from the commodity OS with only minor source code modifications.

We have built a prototype of our system on top of the Xen Virtual Machine Monitor with Linux as the commodity OS. In practice, we find that the system call routing rules are short and simple – on the order of 10's of lines of code. In addition, applications in Proxos incur only modest performance overhead, with most of the cost resulting from inter-VM context switches.

1 Introduction

While significant effort has been invested into making our computing infrastructure more secure, the number of security incidents continues to increase at an alarming pace. The CERT Coordination Center reports that the number of security incidents increased approximately six-fold in the three years between 2000 and 2003, after which they indicate that incidents had become so commonplace that they were not even worth reporting [4]. Despite these statistics, businesses and individuals continue to put increasing trust in computers to store and secure sensitive information, such as financial data, health records, and recently, votes for government elections [15].

Though a great deal of work goes into making operating system kernels more secure, in the vast majority of cases the vulnerabilities being exploited are not in the kernel, but in privileged applications running as user processes. The problem lies not in the reliability of kernel code, but in the overly permissive interface that commodity operating systems (OSs) export, which a privileged application can abuse to make the operating system kernel read or modify the state of any other application. On the other hand, many applications require such privileges to run on a commodity operating system, providing the attacker with many opportunities to take control of the operating system interface. As a result, it seems appropriate that applications that perform security-sensitive operations should have little or no trust in the kernel that lies on the other side of a commodity OS interface.

There have been several attempts to address this situation. One solution is to use a microkernel [1], which minimizes the amount of code running in supervisor mode. However, changing the underlying architecture of the OS kernel without changing the interface that applications use will not give applications any more protection than they currently have. On the other hand, narrowing the application-OS interface requires a large amount of effort to port or rewrite applications currently targeted towards a broad commodity OS interface [23]. There have also been attempts to restrict the interface in existing commodity OSs such as Linux with fine-grained access controls [16]. While effective in principle, the ability to have such controls means that the policy description must be equally fine-grained, making it very complex and time consuming to configure such systems correctly [14]. A third solution is to run the security-sensitive application in its own *private OS* on a virtual machine (VM) executing on top of a virtual machine monitor (VMM), and thus

*Department of Computer Science, University of Toronto

completely remove all other applications from the trusted computing base (TCB) of the system [10]. This private OS would only support the one application and be specially tailored to its needs. The problem that arises is that applications typically share data and interact with other applications through operating system facilities such as files and pipes. Therefore, short of changing the way applications communicate, we are forced to move the other applications into the private OS as well. As a result, the security-sensitive application is made to tolerate other applications in its TCB that it needs to interact with, but does not necessarily trust.

In this work, we attempt to address these issues by building a system that allows an application developer to choose what operating system facilities should be provided by an untrusted commodity OS, and what facilities need to be provided by a trusted private OS. In this way, applications may continue to use functionality in the commodity OS to communicate with other programs, and avoid having to duplicate functionality in the private OS that does not have to be trusted. This ability is provided by running both commodity and private OSs on a VMM, and using a thin operating system proxy, called *Proxos*, which we have designed. Proxos is a small library that mimics an operating system by handling system calls made by the application.

Proxos takes a novel approach to allowing applications to specify their trust in an operating system. Rather than requiring that the application developer partition the application code into components based on whether they trust the commodity OS or not [23], Proxos only requires the developer to partition the system call interface into system calls that must be trusted and those that need not. Using high-level system call routing rules specified by the application developer, Proxos transparently routes each system call made by the application to the commodity OS if the request does not need to be trusted, or to the private OS if it does. Specifying trust by partitioning the system call interface has the benefit that applications currently implemented for commodity OSs can be easily ported to Proxos with very little effort (typically by only modifying on the order of several hundred lines of code). Consequently, the application developer is able to remove the entire commodity OS from the TCB of their application while maintaining reasonable performance.

In this paper, we make three main contributions. First, we have designed a language that allows developers to configure trust relationships using short and simple system call routing rules. In practice, we find that routing rules can usually be specified in 50 lines or less. Second, we have designed and implemented a prototype of Proxos on top of the Xen VMM [2], with Linux as the commodity OS. We describe the modifications we made to Xen and Linux and evaluate the amount of code that these modifications add to each component. Finally, we demonstrate the utility of our system by porting three existing applications: a web browser that protects user privacy, an SSH authentication server, and an SSL certificate service used by the Apache web server. We describe the security of the new applications, the issues we encountered in porting them to Proxos, as well as the performance impact of moving to Proxos.

We will start by giving a high-level description of the system architecture needed to run Proxos applications, as well as a description of the Proxos routing language in Section 2. Section 3 follows with an explanation of our prototype, and gives details on modifications we made to Xen and Linux, details on our Proxos implementation, and details on some example private OS functions we have written. To show what applications one might run on Proxos, we describe three representative applications we have ported to Proxos in Section 4, and evaluate the performance impact of Proxos against a vanilla Xen/Linux system in Section 5. Finally, we finish with related work in Section 6, and give our conclusions in Section 7.

2 Overview

In this section, we describe the overall architecture of the system, as well as a description of the security guarantees our system provides. Then, we give a description of the Proxos system call routing language.

2.1 System Architecture

The architecture of our system is illustrated in Figure 1. The system consists of several VMs running on top of a VMM that enforces memory isolation between the VMs and allocates CPU execution time to the VMs. VMs can make *hypercalls* to the underlying VMM to access resources such as disks and other devices, or to signal or create other VMs. A *commodity OS VM* runs a commodity OS that provides the facilities usually found in a standard operating system, such as file system implementations, a network stack and a user interface. An *administrative VM* (not shown in the diagram) contains management tools used to create and manage other VMs. Applications that want to be isolated from the commodity OS are run inside their own *private VM* along with a Proxos instance. We call such applications *private applications*. A set of methods inside the private VM implement a *private OS*, whose purpose is to handle system calls the private application does not trust the commodity OS with.

Proxos handles all system calls made by the application. Depending on the routing rules configured by the

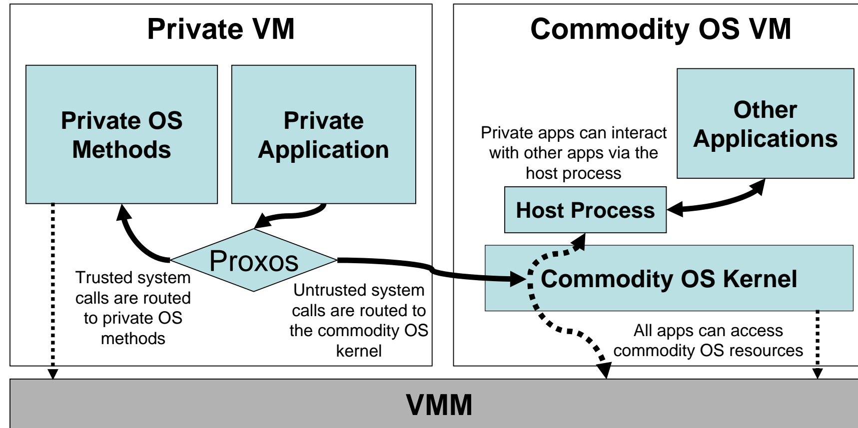

Figure 1: The Proxos System Architecture. Proxos handles all system calls the private application makes by routing them to either the commodity OS or the private OS.

application developer, Proxos will route non-security-sensitive system calls to the commodity OS via inter-VM remote procedure calls (RPCs), and security-sensitive system calls to methods in the private OS. Both Proxos and the private OS are implemented as libraries that are statically linked with the application. As a result, all system calls are converted into subroutine calls to Proxos. The application, along with Proxos and the private OS run on the bare VMM. Since only one application runs in each private VM, all code in a private VM runs in the same protection domain.

To run a commodity application as a private application, the developer first identifies which operating system objects the application uses and that need to be protected from a compromised commodity OS. With this knowledge, the developer identifies the system calls that access these objects and specifies that they are to be forwarded to the private OS using the routing language described in Section 2.3. The private OS methods can be implemented especially for the application by the developer, or even obtained from a library of generic private OS methods provided by a third-party. Section 3.3 describes some private OS methods that we have implemented.

The developer may then have to perform application source code modifications in order to compile it statically, and have it use the facilities that Proxos provides. However, since Proxos exports the same system call interface as the commodity OS, these changes are generally minor. For instance, we were able to port the Glibc library (version 2.3.3) to Proxos with only 218 lines of source code modifications. Next, the private application, the private OS methods, the routing rules and Proxos are all compiled into a single binary, which can be loaded into an empty VM. The developer gives this binary image to the VMM administrator, who registers the new private application with the VMM using the administrative VM. Because the private application binaries are stored directly on the VMM, they are safe from tampering by an adversary who has subverted the commodity OS.

To run a private application, a user on the commodity OS invokes a *host process*, which requests the VMM to instantiate a new VM containing the private application. From this point on, the host process becomes the embodiment of the private application on the commodity OS. The commodity OS attributes any forwarded system call it receives from the private application to the host process that instantiated it. The commodity OS uses the user ID of this host process to make decisions about what operating system objects (such as files or sockets) the application is allowed to access, and also attributes resources used by the forwarded system calls to the host process. In this way, the commodity OS ensures fairness and security between requests made by private applications and requests made by applications running natively on the commodity OS.

Through its host process, a private application can interact with other applications running in the commodity OS by using facilities provided by the commodity OS. For example, by configuring Proxos to forward `mknod` and `open` system calls to the commodity OS, a private application can create a named pipe between it and a commodity OS application. Then, by routing `read` and `write` system calls to the commodity OS, it can communicate with the commodity OS application by making those system calls on the named pipe. For a communication channel to be created, cooperation is required from both applications, who must agree to communicate, and from the commodity OS, who must agree to fulfill the system call requests made by both applications.

2.2 Security Guarantees

While the commodity OS may at some point become under the complete control of an attacker, we assume that the underlying VMM cannot be subverted and that it continues to enforce isolation between VMs. We also rely on the application developer to properly specify what sensitive components of the interface between the application and the operating system must be protected from the commodity OS. Based on these assumptions, our system maintains the confidentiality and integrity of sensitive private application data even in the face of a compromised commodity OS. The isolation property of the VMM prevents the compromised OS from directly interfering with the private application. The compromised commodity OS can only tamper with system calls that are routed to it by Proxos. However, since these system calls were identified as non-security-critical by the developer, the compromised OS should not be able to affect the private application in any security-critical way. We point out that if the routing rules are specified incorrectly, or if a bug in the application causes it to send sensitive data to an interface that the developer believes should have only held non-sensitive data, then sensitive data could be leaked to the commodity OS. In addition, while the confidentiality and integrity of sensitive private application data are maintained, a compromised OS can impact the availability of a private application by not performing the system calls that are forwarded to it.

So far, we have considered protecting the private application from a potentially malicious OS. However, one could envision the case of a buggy private application that could negatively affect the commodity OS through the system calls it forwards to the OS. However, our design restricts the capabilities of the private application within the commodity OS to that of its host process. Since the private application only has the rights of the user who invoked it, our system does not weaken any existing mechanism that guarantees fairness among users and processes running on the commodity OS.

2.3 The Proxos Routing Language

Proxos may route each invocation of a particular system call differently depending on rules specified by the application developer. For example, Proxos may route `read` system calls differently depending on what file is being read. We wish to provide a simple and intuitive way for an application developer to partition the system call interface. In principle, one could specify a routing rule for each of the over 200 system calls that a commodity OS like Linux provides, but this would be complex and time consuming. Further, we do not believe it necessary in most cases to have such fine-grained control over system call routing. We organize system calls by the resources they access and create a Proxos routing language with which the developer can specify routes for those resources. In this language, the operating system provides six resource classes to an application: persistent storage (disk), user interface, network, randomness, system time, and memory. Peripheral devices such as printers, USB devices, etc, are abstracted by the OS into file objects and are thus part of the persistent storage category.

While it is possible to provide routing rules for all six resources, we have found that this is unnecessary. An application may choose to forward system requests to the commodity OS for two reasons: either it wants to use the resource as a communication channel with another application, or it does not need the resource to be trusted and thus wishes to include the resource outside of its TCB. As a result, persistent storage, user interface and the network are routed by Proxos because these are resources that applications either use to communicate, or may not need to trust. System time and randomness are never routed because they cannot be used as communication channels, and are provided by the underlying VMM without increasing the application's TCB. Finally, memory related system calls (such as `brk` and `mprotect`) are used to indirectly manipulate page table entries. However, a private application would never trust a commodity OS with control of its page tables since this would imply granting the commodity OS access to the private application's memory. Therefore, it does not make sense to route memory-related system calls. All non-routable system calls are directed to functions provided by Proxos.

Based on this model of operating system resources, we have designed a simple language that allows the application developer to specify which system calls will be routed to the commodity OS, and which to the private OS. Figure 2 shows a stripped-down example of a routing specification in our language. Lines prefixed with a "#" are comments. The Rules section consists of three declarations, one for each of the routable resource classes. The specifications for the disk and network resource classes are a list of tuples, where each tuple describes the particular resource, and a table of function pointers used to access the resource. In this case, the specification for the user interface (`UI`) has "`*`" as a resource description because the application wants to route all three standard I/O streams (i.e. `stdin`, `stdout`, and `stderr`) to the private OS. The example also specifies that access to any file with name `/etc/secrets` should be handled by methods in the private OS. The same is true for system calls to any UNIX domain socket bound to `/tmp/socket` and to any TCP socket with a peer IP address of 192.100.0.4 on port 1337. By default, Proxos will route all system calls to resources that do not match

```
# Rules Section
# route accesses to /etc/secrets to private OS
DISK:("/etc/secrets", priv_fs)
# route accesses to UNIX domain socket bound
# to /tmp/socket and TCP socket bound to peer
# 192.100.0.4 port 1337 to private OS
NETWORK:("unix:/tmp/socket", priv_unix),
        ("tcp:192.100.0.4:1337", priv_tcp)
# route all accesses to stdin, stdout
# and stderr to private OS
UI: (*,priv_ui)

# Methods Section
# individual methods in the private OS
# that are bound to system calls
priv_fs = {
  .open = priv_open,
  .close = priv_close,
  .read = priv_read,
  .write = priv_write,
  .lseek = priv_lseek
}
```

Figure 2: Routing Example. This example shows a simple set of routing rules that protects operations on a particular file, two network sockets, and the standard I/O streams.

any rule to the commodity OS.

The Methods section defines which methods in the private OS will handle system calls from the application. When the application attempts to open the file `/etc/secrets`, Proxos will call the `priv_open` method in the private OS to handle the request and return a file descriptor. All subsequent system call operations (such as `close`, `read`, `write` and `lseek`) on the file descriptor associated with that file will also be forwarded to the associated private OS method in the table. On the other hand, any system call on the file that is not in `priv_fs` will be forwarded to the commodity OS. Method tables for `priv_ui`, `priv_unix` and `priv_tcp` are not shown in the figure, but must also be specified by the application developer.

Rather than specifying trust policies by partitioning code, or by restricting abilities, specifying policies by partitioning interfaces to resources results in a more compact and intuitive policy description. Further, our specification language allows the application developer to use the same names for resources as those in the source code, making the routing rules easier to write and understand.

3 Prototype Implementation

There were several requirements that dictated which underlying system we chose to implement our Proxos prototype on. First, we needed a way of "hoisting" a commodity OS to a lower privilege level and inserting our own privileged code beneath it. Second, the system had to provide isolation between the private applications and the commodity OS, but at the same time allow some controlled communication between them. In light of these requirements, we eventually settled on using the Xen VMM [2] and Linux as our commodity OS for our experimental substrate. However, we believe that the features required by our system could be provided by any VMM or microkernel.

In this section, we describe the three main components we implemented in building our prototype. First, we describe our modifications to Xen and Linux to provide support for starting private applications, and to forward system calls between VMs. Second, we describe our Proxos operating system proxy prototype, which routes system calls to either the commodity Linux kernel or to private OS methods. Finally, we describe some private OS methods that we have implemented.

3.1 Modifications to the VMM and the Commodity OS

Modifications made to Xen and the Linux kernel can be categorized into three components: the start-up and shutdown of private applications, a facility for forwarding system calls between VMs, and a trusted path facility.

Since the Linux kernel and private applications do not trust each other, the private application start-up process must make several guarantees. First, the private application must not be able to gain any privileges beyond those of its host process. This implies that the Linux kernel must always be able to attribute system calls routed to it to the host process that initiated the private application forwarding the system call. Second, a compromised commodity OS should not be able to initialize a private application in an unsafe state. Finally, the private application should not be able to access any Linux kernel memory that the kernel has not authorized it to.

The VMM administrator registers private applications with the VMM via a configuration file. This file assigns a name to each private application and sets start-up parameters for each private VM. Later, when the host process starts a private application, it will use this name to indicate to the VMM which private application to start. Figure 3 describes the private application start-up process used in our prototype in detail.

In Step 1, private applications are started using the `pr_execve` system call that we added to the Linux kernel. `pr_execve` is the private application analog to the `execve` system call and, like `execve`, takes the name of the private application to be started as its argument. `pr_execve` causes the current process to become the host process for the private application.

In Step 2, the Linux kernel allocates a shared buffer that is used later for system call arguments forwarded to

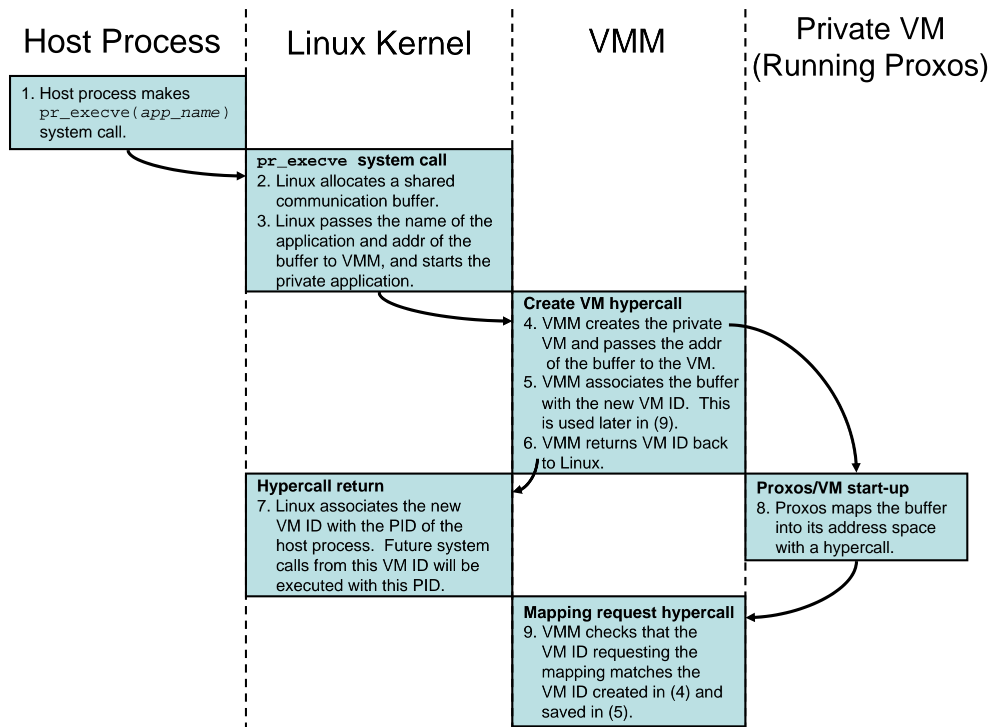

Figure 3: Private Application Start-up Sequence. The steps are arranged into columns with the titles at the top indicating what system component each step takes place in.

it from the private application. The kernel passes the address of this buffer to the VMM in Step 3, and at the same time signals the administrative VM to start a new VM for the private application with a hypercall we introduced. The administrative VM will only start the private VM with parameters set by the system administrator, ensuring that even a compromised Linux OS can only start private applications from a known, safe state. Note that a compromised Linux OS may start a private application different from the one the host process requested and attempt to get the user to use the wrong private application. To detect this, Proxos relies on application level safeguards such as the trusted path used in our web browser or cryptographic keys used in our SSH private server. We will discuss both of these applications in Section 4.

In Steps 4 to 7, the VMM creates a new private VM for the application, informs the Proxos in the private VM of the location of the shared communication buffer, and informs the Linux OS of the identity of the new VM (by giving it a VM ID). Then, in Step 8, Proxos tries to map the shared buffer into its address space via another hypercall. Originally a privileged hypercall, we modified this hypercall so that private VMs may use it. However, we also added an extra check to ensure that the VM making the mapping request in Step 9 is the same as the one to which the VMM originally passed the shared buffer address in Step 4.

Private application shutdown is much simpler as there are no security guarantees to be made. If the private application initiates the shutdown, then it informs the VMM via a standard hypercall. We extended this hypercall to notify the Linux kernel, so that the kernel may terminate the host process accordingly. Even if the private application has not terminated, the kernel may still forcibly destroy the host process (by killing the process). However, the kernel does not have the privileges to force the private application to shutdown, so by killing the host process, the kernel can only revoke the private application's ability to access commodity OS resources.

Another set of modifications allow private applications to forward system calls to the Linux OS. The goal is to reduce the latency of forwarded system calls by reducing the number of domain crossings. Xen already provides a facility that allows VMs to send events to each other. By combining this with the shared buffer between the Linux OS and the private application, we were able to add a simple RPC mechanism to Xen. We then made modifications to the Linux kernel to allow it to efficiently execute forwarded system calls. When the Linux kernel receives the system call arguments, it determines the appropriate host process to wake up by examining the source of the RPC and comparing that to information it recorded in

Step 7 of the start-up sequence. As the host process is about to be scheduled, a trip into user-space can be saved by placing the system call arguments in the appropriate registers and transferring control directly to the system call handler in the kernel. When the system call handler completes, the kernel sends the return value back to the private application via an RPC response message and returns the host process to the queue it was in before the forwarded system call arrived. As a result, our prototype handles forwarded system calls without any domain crossings in the Linux OS.

Finally, we also needed the VMM to provide a trusted path facility so that private applications can communicate directly with the user without having to trust the commodity Linux OS. This would prevent a compromised Linux OS from masquerading as a private application, as well as prevent a compromised Linux OS from eavesdropping on communication between a user and a private application. To support this, the VMM provides user interface facilities such as a console driver and graphical window system. If the private application wants to use these facilities, it routes system calls on standard I/O streams (i.e. `stdin`, `stdout` and `stderr`) to private OS methods, which will forward the requests to the VMM console driver. Similarly, it routes X window operations to private OS methods that will translate them into the appropriate operations on the VMM window system. The implementation of minimal trusted window systems on secure kernelized systems has been studied in the literature [9, 22]. Rather than reimplement these in our prototype, we simply provided an emulation of their functionality, but do not make any effort to reduce the amount of code that is added to the VMM. We did this by running an X server on Xen's administrative VM and using nested X servers to give each VM its own separate X interface.

We found that modifying Xen and Linux to allow private application start-up and shutdown, as well as forwarded system calls, had very little impact on the size of the Xen TCB. Many of the facilities needed were already present in the Xen VMM and we only had to make these accessible to unprivileged VMs and add checks to make sure they could not be abused. The only component that increases the code base of the VMM significantly is the graphical user interface. A significant portion of this component can be implemented outside of the trusted computing base of the VMM [9, 22], but exploring the design of trusted window systems was not a goal of our prototype.

3.2 The Proxos Prototype

Our prototype is derived from the *Minimal OS* example that comes with the Xen 2.0 source code. Proxos runs in a single address space and supports only one private application. Our current implementation is also single-threaded, although we plan to support threads in the future. Apart from providing basic memory and page table management, Proxos also contains: a block driver that supports raw accesses to a private block device exported by the VMM; and a console driver that provides direct access to the Xen console. Our prototype does not provide a TCP/IP stack or a network driver. We found these unnecessary as many security-sensitive applications already assume the network is not trustworthy and employ cryptographic safeguards such as SSL to protect network communications. This allows us to safely reuse the network services of the commodity OS.

Proxos uses operating system abstractions to determine where to route system calls at run time. In the case of Linux, the abstraction used by applications to access resources is a file descriptor. Initially a file descriptor is bound to a resource via a system call such as `open` or `socket`. Subsequent operations on that resource are then performed by naming the descriptor in the system call.

The design of Proxos is very simple, and is similar to the way virtual file system methods are implemented in Linux. Routing rules for the application are converted into lookup tables, which are then compiled into the Proxos library and linked with the private application. When descriptors are created, Proxos compares the name of the resource they are being bound to with the routing rules specified for the application. For example, if a file descriptor is being created via an `open` system call to a file, Proxos compares the name of the file being opened with the list of tuples provided in the `DISK` resource class. If a match is found, Proxos uses methods from the method table specified in the matching routing rule to handle subsequent system calls on the descriptor. Proxos provides a set of default methods which route untrusted system calls to the Linux OS. If a routing rule specifies a private OS method to be called, Proxos transfers control to the appropriate location in the private OS.

The private application uses file descriptors to name objects in both the private OS and the commodity OS. File descriptors in the commodity OS are allocated from a name space independent of the one the private application is using. Upon opening a new file in the commodity OS, Proxos may find that the commodity OS has assigned a file descriptor number that the private application is already using to name another object in the private OS. As a result, Proxos translates between the file descriptors used by the private application, and those used in the commodity and private OSs.

Most routable system calls can be routed transparently to the Linux OS. However, the `fork`, `execve` and `select` system calls have slightly different semantics.

When forwarded, the fork system call will cause the host process in the Linux OS to fork. The forwarded fork creates concurrency on the Linux OS side, but the application in the private VM will still contain only a single thread of execution, so parent and child code must be executed sequentially. After the fork, the private application specifies whether system calls it forwards to the Linux OS should be executed by the parent process or the child process. This is done by setting the *target PID flag* in Proxos to indicate the process ID (PID) of the process that should be the recipient of system calls forwarded to the Linux OS. The value of this flag accompanies every system call Proxos forwards to the Linux OS. The Linux OS checks that the PID specified by the flag belongs to either the host process, or a child of the host process. These semantics imply that forwarding fork system calls requires the developer to make any concurrent code sequential in the private application. To support standard fork semantics, the underlying VMM needs to be capable of duplicating the address space of the private application (preferably using copy-on-write for efficiency). While we did not support this in our prototype, we note that others have proposed adding such functionality to Xen [24].

The semantics of forwarded execve system calls are also slightly different. If the execve system call is made without a fork, the host process will terminate and a new program will take its place. If the new process is not willing to host system calls forwarded to it, the private application will be unable to forward system calls to the Linux OS. More commonly, a recently forked process will execute execve. In this case, the private application will lose the ability to forward system calls to the child, but retain the parent as the host process. More details on how fork and execve are used in private applications will be given in the description of our port of the SSH server in Section 4.2.

Finally, select has a slightly different behavior under Proxos than its Linux counterpart. select allows applications to listen on several file descriptors simultaneously and notifies them when there is activity on any of the descriptors. In Proxos, an application may execute a single select on file descriptors from both the commodity OS and the private OS. However, Proxos forwards system calls by making *synchronous* inter-VM RPCs. This limitation of our current prototype prevents Proxos from routing select system calls to both OSs simultaneously, so it serializes them and imposes a timeout on each select call. Proxos will alternate between which OS to execute select on first to ensure no file descriptor is starved. The poll system call has the same behavior as select in our system. The consequence of this is that events on file descriptors that happen close together may not be delivered to the private application in the same order that they occurred because Proxos may be polling the other OS instance when the first event occurs. However, we have not seen this to be an issue and, to the best of our knowledge, Linux makes no such ordering guarantees either.

3.3 Private OS Methods

In our prototype, we have implemented two example private OS components: one that implements a private file system, and one that implements a trusted path by forwarding standard I/O streams and X window messages to the VMM.

A private file system allows the private application access to persistent storage that is protected from tampering by the Linux OS. We wanted to implement this by adding as little code to the private VM as possible, as any code we add increases the TCB of the application. Rather than implement an entire file system, our private file system outsources most of its functionality to the commodity Linux OS through forwarded system calls, but maintains the secrecy of any information stored by encrypting all data before writing it to the Linux file system [12]. To protect the data from tampering and replay, hashes of all files stored on Linux by the private file system are kept on a private block device available directly from the underlying VMM. Doing this significantly simplifies the file system implementation, as all that is needed are the cryptographic functions, some code to manage file system buffers, and block device drivers to store the file system hashes. The drawback is that a compromised Linux OS could potentially deny the private application access to files that the private file system has saved. However, our applications typically depend on the Linux OS for other services as well, so no forward progress guarantees are broken by this.

In our prototype, the private OS implements a trusted path by routing operations on standard I/O streams and the X server's socket to the VMM. The private OS methods translate system calls on standard I/O streams to operations on Xen's console driver and route system calls on the X server to the administrative VM. A host process on the administrative VM then executes the routed system calls on a socket connected to a nested X server instance that is separate from the one that the commodity Linux OS is using.

3.4 Discussion

With the exception of modifications to the Linux kernel, all components implemented in our prototype will be part of the application TCB. As a result, we placed a lot of emphasis on keeping the impact on code size and complexity small, especially with respect to the VMM. One

Component	Lines of Code
VMM modifications	656
Linux modifications	4380
Proxos	7348
Private File System	1817
Trusted Path	1313

Table 1: Number of lines of code in each component in our Proxos prototype. The VMM modifications do not include the X server running in the administrative VM.

caveat is that Proxos does not need to support every system call that Linux exports. For example, Proxos does not support system administration calls, such as those to control swap devices, or load and unload kernel modules, as private applications will not need to make such requests. Out of the 289 system calls of the Linux 2.6.10 kernel, our Proxos prototype only needs to support (either internally or by forwarding) 56 of them to run most applications. However, we fully expect this proportion to increase as Proxos matures. The size of the components in our prototype are given in Table 1.

4 Applications

In this section, we describe three applications that we have ported to Proxos. We selected applications that will benefit from partially trusting a commodity OS, and illustrate interesting issues that arose when porting them. Our first application is a secure web browser that protects user information. Our second application is an SSH server that protects system-critical information such as passwords and host keys even if the commodity OS is compromised, but still allows users who login to gain a full shell on the commodity Linux OS. Our final application is an SSL certificate service that we use with an Apache web server to implement SSL transactions. In this case, the private keys corresponding to the certificate are protected.

4.1 Secure Web Browser

A serious threat to the security and privacy of users is spyware, which is malicious software that is surreptitiously installed on machines and monitors the web surfing habits of users. While the goal of most spyware is to collect usage data for marketing, spyware has been shown to decrease the security of user system by recording and transmitting confidential information that it has collected [18, 20].

Spyware collects information by either monitoring the user's keystrokes, or by scraping files where web browsers have recorded user information. We ported Dillo [5], a simple graphical web browser, to Proxos and configured the routing rules so that all user I/O is directed to the trusted VMM user interface, thus creating a trusted path, and all sensitive data that Dillo reads from disk or writes to disk is directed to our private file system. No other rules are specified, and thus other network operations such as HTTP requests are routed to the Linux OS by default (for extra security, the user should use HTTPS to encrypt traffic between the browser and the web server to prevent any spyware on the Linux OS from observing or tampering with it). Similarly, any documents or executables that the user downloads from the Internet are saved to the Linux file system. In addition, any external helper applications that Dillo invokes will be transparently created and executed on the Linux OS.

For the most part, no source code modifications were required to port Dillo. The only necessary modifications were due to Dillo's use of graphical themes, which are implemented as code that is dynamically loaded at run time based on the theme the user selects. In our prototype, it is not safe to load code from the Linux OS, since an adversary may have tampered with it. To support the default theme, we removed the code that loads themes at run time and statically linked the default theme into the Dillo private application. In theory, code could be safely loaded from the Linux OS if encrypted and accompanied by a valid signature that the private application could verify, but our current prototype does not support this.

4.2 SSH Authentication Server

Often when attackers compromise a system, the system administrator is not only forced to rebuild the entire system from scratch to ensure that any malicious software has been removed, but also to perform the arduous task of tracking down every user and ensuring that they change their passwords in case the attacker has been able to learn some of the old passwords. Similarly, she must change any cryptographic host keys, which the machine uses to authenticate itself, and distribute new keys to all parties that the machine interacts with. Being able to ensure the secrecy of user passwords and the host keys of a system after a security compromise would save the administrator significant time and effort.

To demonstrate the utility of Proxos in protecting the secrecy of sensitive data, we ported the OpenSSH authentication server (version 3.9p1) to Proxos. The OpenSSH server accesses several sensitive resources including configuration files, the password file and the host key file. We wrote routing rules to store the password file, host key, and global configuration files on the private file system. The SSH server also performs network operations, but no rules are specified for `NETWORK` resources since OpenSSH is designed to function with an untrusted

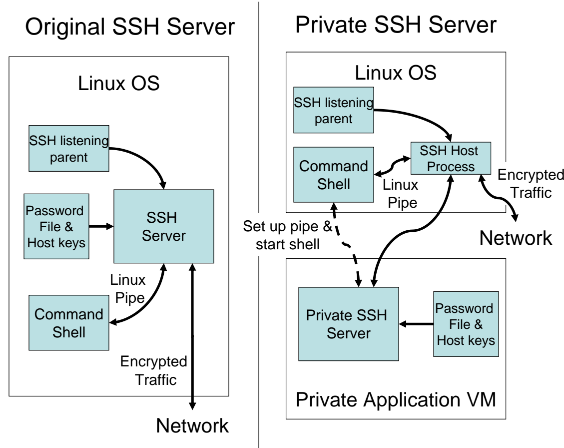

Figure 4: Comparison of the original SSH server and our private SSH server.

network. Other than the routing rules, only two modifications involving the fork system call were required to implement a private SSH server application. The architecture of our new private SSH server is shown alongside the architecture of the original SSH server in Figure 4. One modification arose because the SSH server requires some concurrency to allow multiple users to authenticate simultaneously. The native version of SSH handles this by having a parent process listen on the SSH port, and then spawning a child for every connection the parent receives. Our private SSH server still has the listening parent as a native Linux application, but implements the children as private applications. When the listening parent detects a new connection, it forks a child (on the Linux OS), which then uses pr_execve to instantiate a private SSH server VM, and in doing so becomes the host process for the new VM.

The private SSH server starts-up and reads the sensitive data from the private file system, and then proceeds with user authentication. If a user logs in using private key authentication, the private SSH server will need to access the public keys the user has placed in a file in their home directory on the Linux OS. Proxos provides access to the user's keys without any extra configuration – since the user's key files do not match any routing rules, requests to them will be forwarded to the Linux OS by default. If the user authenticates successfully, the native SSH server forks a child that will execute a command shell. Before the child starts the command shell, the native SSH server creates a pipe between itself and the command shell redirecting all input and output from the shell to itself, so that it can encrypt any shell output before sending it to the network, and decrypt any shell input coming from the network. In our version, the private SSH server changes the Proxos target PID flag to point to the new child after the fork, and then executes the child code, forwarding the system calls required to set up the pipe and start the command shell. After this, it changes the target PID flag back to the parent and executes the SSH server code. The shell will pipe all input and output through the host process to the private SSH server, which encrypts and decrypts data as appropriate between the shell and the network.

4.3 SSL Certificate Service and Apache

Next, we explored the performance impact of Proxos on Apache with SSL. As in the SSH server, the Apache server relies on concurrency so we only ported the crypto library portion of the OpenSSL library to Proxos, and left the Apache web server on the Linux OS. The crypto library uses confidential private keys stored in the SSL certificate, which would be protected if the web server was compromised. Our port uses Apache version 2.0.52 and version 0.9.7g of the OpenSSL library.

To setup SSL sessions, Apache makes calls to the OpenSSL library, which uses the OpenSSL "engine" interface to invoke the crypto library. We modified the engine interface to spawn a private application that will use the private key of the server's SSL certificate to sign challenges during an SSL handshake. Unfortunately, this operation is called on every HTTP request that uses SSL (i.e. an HTTPS request), and would give very poor performance because each request results in the instantiation and shutdown of a private VM. To remove the frequent instantiation and shutdown of the private VM, we modified Apache to spawn a process when it starts-up, which will act as the host process for a single private SSL certificate application. Apache was also modified so that a portion of the shared buffer between the host process and the private application is mapped into the address space of each Apache thread. To process an HTTPS transaction, a thread enqueues the signing request on the shared buffer, sends a signal to Proxos for processing and sleeps until the request has been processed. Since multiple Apache threads will be accessing the shared buffer, we also added the appropriate synchronization between the threads to prevent races.

4.4 Discussion

The size of our routing rule descriptions, along with the lines of code that were modified for each of the applications, as well as Glibc (version 2.3.3), is given in Table 2. In porting these applications we found that what often takes some time are modifications to application source code that are required to support operations like fork and execve in the private SSH server, or to statically

Application	Rules	LOC Modified
Dillo	53	22
SSH Server	35	108
Apache & OpenSSL	28	667
Glibc		218

Table 2: Size of routing rules and number of LOC modified for each application.

link in dynamic code in Dillo. Apache required more effort since several threads could make challenge-signing requests simultaneously, and this required careful arbitration and synchronization to preserve performance. We find these results encouraging – Proxos enables the application developer to remove the entire commodity OS kernel and privileged applications from the TCB of the private application by modifying on the order of several hundred lines of code in the application, and writing around 50 lines of routing rules.

5 Performance Evaluation

The performance of VMMs versus native operating systems has been well studied in the literature [2, 3]. To ascertain the overhead introduced by Proxos, we compare the performance of our system against a system running an unmodified Linux kernel executing on an unmodified Xen VMM. We first use microbenchmarks to better understand the components that contribute to the cost of making forwarded system calls from a private application. Then we evaluate the performance of the SSH and Apache/SSL Certificate applications described in Section 4 on our system. All tests were performed on a machine with a 3GHz Intel Pentium 4 processor, 1GB of RAM, a 7200 RPM Serial-ATA disk with 8.9 ms seek time, and a 100Mb Ethernet NIC. Our prototype is built on Xen 2.0, with Fedora Core 3 Linux running a 2.6.10 kernel as the commodity OS, and its performance is compared against vanilla versions of the same software. For our runs, 768MB of RAM were allocated to the commodity OS, and the rest was used for Xen, the administrative VM, and private applications.

5.1 Microbenchmarks

To analyze the overhead of a system call forwarded from the private application to the Linux OS, we must first understand the individual components that make up a forwarded system call. These costs are illustrated in Figure 5. In Step 1, Proxos sends an event to the Linux kernel, notifying it of the forwarded system call, and then yields the processor, causing a VMM context switch into the Linux OS VM. In the Linux kernel, a virtual interrupt

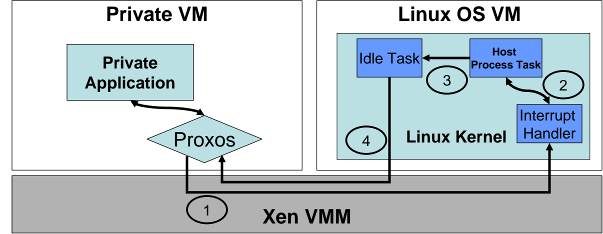

Figure 5: Breakdown of costs incurred in a forwarded system call.

Benchmark	Linux	Proxos	Overhead
NULL system call	0.37	12.88	12.51
fstat	0.57	14.28	13.71
stat	8.76	25.98	17.22
open & close	14.57	47.18	32.61
read	0.45	13.51	13.06
write	0.42	13.24	12.82

Table 3: Forwarded system call latencies on LMbench microbenchmarks. All measurements are given in μs.

handler receives the event and enqueues the system call request on the process descriptor of the host process. In Step 2, we wait until the host process is scheduled. On a lightly loaded system, this incurs only the cost of another context switch within the Linux kernel, but may take more time if the Linux kernel is heavily loaded. After the Linux kernel executes the system call, it will not yield the processor back to the VMM until either the VMM scheduler decides to preempt the Linux OS VM, or the kernel runs out of runnable processes and schedules the idle task in Step 3. Finally, in Step 4, another VM context switch occurs and Proxos can receive the result of the system call. While this accounts for four context switches, there is actually a fifth context switch because Xen will schedule the administrative VM in either Step 1 or Step 4.

We ran the system call latency benchmarks in the LMbench 2.5 microbenchmark suite [17] in a private VM configured to forward all system calls to an idle Linux OS VM, and summarize our results in Table 3. We also used the context switch microbenchmark in LMbench and measured the minimum cost of a context switch to be 2.88μs on our machine. As a result, the expected five context switches would take approximately 14μs, which tracks well with the measured results. This cost is added to every system call except for `stat` and `open`, whose larger overhead can be explained by the fact that each context switch changes virtual to physical page mappings, and causes a TLB flush. Since both `stat` and `open` take a filename as an argument, the Linux kernel

must make several queries to the buffer cache to find the correct inode (LMbench ensures that the inodes required to access the files are cached in memory), which will result in TLB misses. These misses do not occur when the benchmarks are run directly on Linux because the kernel never switches to another process, so no context switches occur.

5.2 Application Benchmarks

We now evaluate the overhead imposed on our private SSH server and SSL certificate service. Like our microbenchmarks, applications incur overhead when system calls are routed to the commodity OS. To evaluate the average overhead a forwarded system call experiences, we used an SSH client to login to our private SSH server over the loopback device and measured the time taken to copy files ranging from 32MB to 256MB over the SSH connection. Each file transfer was performed five times on both the private SSH server and a native SSH server running on Xen. The standard deviation was less than 1.5% across our measurements. Figure 6 plots the average difference in time taken by the private SSH server over the native SSH server to transfer a file, against the number of forwarded system calls the private SSH server made. We perform linear regression on the average values and found a correlation of 0.92, indicating that the overhead is well correlated with the number of system calls. We then estimate the start-up component to be 0.72s and the per-system call cost to be 15.7μs. For large files, where the cost of start-up has been amortized, the private SSH server only takes 6.0% longer to transfer the same file as the native SSH server. Note that since this overhead is comparable to the variance in our measurements, the estimated system call overhead should not be taken too literally, and is merely a rough approximation.

We suspected that a large part of the start-up cost for the private SSH server is due to VM creation. We confirmed this by measuring the time to start an empty private VM, which is approximately 0.35s. Starting a Xen VM requires the use of several user-space scripts in the administrative VM, making it very expensive, and we have not made any effort to optimize this operation. The remaining 0.37s is the time the private SSH server uses for initialization, which includes the time it takes to read in sensitive data from the private file system. This operation requires several cryptographic operations to decrypt the data and verify the authenticity of files stored on the commodity OS file system.

To evaluate the performance impact of our private SSL certificate application, we used Mindcraft's Webstone benchmark [25] extended with SSL. We configured the benchmark with 150 SSL clients, which was enough to fully load an Apache server on Xen. The same number of clients was used to measure the amount of bandwidth our Proxos-enabled web server could support. We expected the Proxos-enabled web server to introduce low overhead because HTTPS transactions mainly perform computation and make very few system calls. Our experiments show that there is actually a slight increase in throughput – 5.04Mb/s for the Proxos-enabled web server as compared to the native web server's throughput of 4.75Mb/s. From this, we surmise that the overhead Proxos introduces is very low and that the changes we made porting the system likely perturbed the system in such a way as to produce a slight performance gain.

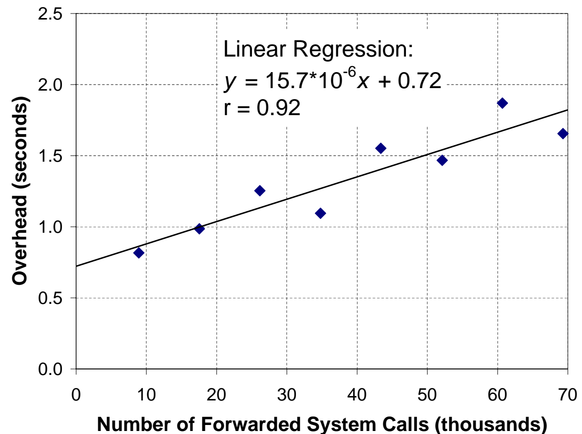

Figure 6: SSH benchmark. We plot the overhead of the private SSH server versus the number of system calls forwarded to Linux.

6 Related Work

While isolating processes from operating system components bears some similarities to multi-server microkernels [1, 11], Proxos is more similar in structure to Exokernels [8] and the Denali kernel [26], with their statically linked, single user LibOSs. Where Proxos differs from these systems is in its objective – while the goal of the former systems is to give applications some ability to customize OS management of their resources, Proxos gives applications the ability to customize the trust relationship between applications and the OS kernel.

Various systems aim to limit the damage a compromised application can cause. SELinux [16] allows the administrator to set a fine-grained mandatory access control policy for the system, thus limiting the privileges an attacker would gain by hijacking an application. However, fine-grained control has its costs. SELinux policies are large and complex – the size of the default pol-

icy set for the Fedora Core 3 Linux distribution has over 290,000 rules and consumes more than 7MB of kernel memory. In contrast, our interface routing configurations are typically around 50 lines long or less. In addition, because SELinux works by restricting the abilities of applications, its policy rules must define all the permitted behaviors of every application on the system. Since Proxos operates by isolating security-sensitive applications from the rest of the OS, Proxos policy rules only need to be defined for the applications being protected. Asbestos [6], Eros [21] and Singularity [13] also limit information flow and privileges, but through mechanisms significantly different from Proxos. Asbestos uses process labels that are updated dynamically, combining aspects of capabilities and information flow policies. Eros is a pure capability-based microkernel and Singularity only permits communication between processes through strongly typed and formally verified channels. All of these paradigms require applications to be ported to fundamentally different application interfaces. By keeping the same application interface as a commodity OS, Proxos does not require any extensive porting for existing applications. Further, applications that do not require Proxos can remain in the commodity OS and suffer no overhead.

Terra [10], Nizza [23] and Microsoft's NGSCB [7] are projects that propose new operating system models to increase the security of applications. Terra provides coarse-grained isolation by enclosing security-sensitive applications along with their own operating system in a "closed-box" VM. Applications may only communicate with applications on other VMs via the network interface. Similarly, NGSCB runs specialized "agents" in a high-assurance OS called the Nexus, which is isolated from a standard Windows OS by a VMM. In both NGSCB and Terra, each OS must contain all the functionality required by the application, even if the functionality does not have to be trusted by the application, while Proxos can reuse untrusted functionality in the commodity OS. On the other hand, Nizza, along with projects μ-Sina [12] and Perseus [19], take a fine-grained approach to minimize the amount of code in an application's TCB. They propose heavily modifying the application source code to extract and port the security-sensitive components of an application to a microkernel. The resulting applications often have reduced functionality. In contrast, by partitioning trust along a commodity OS interface, Proxos allows applications to retain the ability to communicate with untrusted applications through standard OS facilities, which is lost in Terra. At the same time, it avoids the effort to modify application source code or the reduced application functionality that Nizza entails.

Finally, nothing in this work is Xen-specific. While many features of Xen made the prototype easier to build, we believe that the Proxos infrastructure is applicable to other hypervisor based VMM systems, such as VMware ESX Server and Microsoft's forthcoming Viridian.

7 Conclusions

Current commodity OSs export an interface that is too permissive to privileged applications, allowing compromised applications to gain control of the operating system kernel and attack other applications. Proxos allows applications to partition the interface between them and the commodity OS kernel into trusted and untrusted components by specifying system call routing rules. The end result is that Proxos allows application developers to protect applications from a compromised kernel without having to make major source code modifications.

By building a Proxos prototype and porting several representative applications, we have found that specifying trust at the system call interface is a powerful and simple way of isolating applications from the operating system. Proxos routing rule specifications are short and simple, and can be expressed in 10's of lines of code. Minor source code modifications are also required to support applications, mainly due to the semantics of the `fork` system call, and to remove any instances of dynamically loaded code that cannot be eliminated by static linking. In cases such as our web server, where expensive VM start-up and shutdown may become very frequent, further modifications are necessary to preserve performance. We expect that in most cases, a single graduate student who is familiar with an application can port it to Proxos in a day or two. With a modest cost in engineering time and a reasonable impact on application performance, system call routing enables the developer to protect the secrecy and integrity of applications from a compromised operating system.

Acknowledgements

We are grateful to Nagendra Modadugu, who generously provided us code for his SSL-enabled version of Webstone. Chandu Thekkath provided insightful discussions that helped initially in this work. Various iterations of this paper were improved immensely by comments from Tom Hart, Ian Sin, Jesse Pool, Michael Stumm, Ashvin Goel, Reza Azimi, Troy Ronda, and others in the SSRG group at the University of Toronto. We would also like to thank Derek McAuley, our shepherd, for his helpful suggestions. This work was supported in part by an NSERC Discovery Grant and a MITACS seed grant.

References

[1] M. J. Accetta, R. V. Baron, W. Bolosky, D. B. Golub, R. F. Rashid, A. Tevanian, and M. W. Young. Mach: A new kernel foundation for UNIX development. In *Proceedings of Summer Usenix*, July 1986. 1, 6

[2] P. Barham, B. Dragovic, K. Fraser, S. Hand, T. Harris, A. Ho, R. Neugebauer, I. Pratt, and A. Warfield. Xen and the art of virtualization. In *Proceedings of the 19th ACM Symposium on Operating Systems Principles (SOSP 2003)*, pages 164–177, Oct. 2003. 1, 3, 5

[3] E. Bugnion, S. Devine, and M. Rosenblum. Disco: running commodity operating systems on scalable multiprocessors. In *Proceedings of the 16th ACM Symposium on Operating Systems Principles (SOSP 1997)*, pages 143–156, Oct. 1997. 5

[4] CERT Coordination Center, 2006. `http://www.cert.org`. 1

[5] Dillo web browser, 2006. `http://www.dillo.org`. 4.1

[6] P. Efstathopoulos, M. Krohn, S. Van De Bogart, C. Frey, D. Ziegler, E. Kohler, D. Mazières, F. Kaashoek, and R. Morris. Labels and event processes in the Asbestos operating system. In *Proceedings of the 20th ACM Symposium on Operating Systems Principles (SOSP 2005)*, pages 17–30, Oct. 2005. 6

[7] P. England, B. Lampson, J. Manferdelli, M. Peinado, and B. Willman. A trusted open platform. *Computer*, pages 55–62, July 2003. 6

[8] D. R. Engler, M. F. Kaashoek, and J. O'Toole. Exokernel: An operating system architecture for application-level resource management. In *Proceedings of the 15th ACM Symposium on Operating Systems Principles (SOSP 1995)*, pages 251–266, Dec. 1995. 6

[9] N. Feske and C. Helmuth. A Nitpicker's guide to a minimal-complexity secure GUI. In *Proceedings of the 21st Annual Computer Security Applications Conference (ACSAC 2005)*, pages 85–94, Dec. 2005. 3.1

[10] T. Garfinkel, B. Pfaff, J. Chow, M. Rosenblum, and D. Boneh. Terra: A virtual machine-based platform for trusted computing. In *Proceedings of the 19th ACM Symposium on Operating Systems Principles (SOSP 2003)*, pages 193–206, Oct. 2003. 1, 6

[11] H. Härtig, M. Hohmuth, J. Liedtke, S. Schönberg, and J. Wolter. The performance of μ-kernel-based systems. In *Proceedings of the 16th ACM Symposium on Operating Systems Principles (SOSP 1997)*, pages 66–77, Oct. 1997. 6

[12] M. Hohmuth, M. Peter, H. Härtig, and J. S. Shapiro. Reducing TCB size by using untrusted components - small kernels versus virtual-machine monitors. In *Proceedings of the 11th ACM SIGOPS European Workshop*, Sept. 2004. 3.3, 6

[13] G. C. Hunt, J. R. Larus, M. Abadi, M. Aiken, P. Barham, M. Fahndrich, C. Hawblitzel, O. Hodson, S. Levi, N. Murphy, B. Steensgaard, D. Tarditi, T. Wobber, and B. Zill. An overview of the Singularity project. Technical Report MSR-TR-2005-135, Microsoft Research, Microsoft Corporation, Redmond, WA, Oct. 2005. 6

[14] T. Jaeger, R. Sailer, and X. Zhang. Analyzing integrity protection in the SELinux example policy. In *Proceedings of the 12th USENIX Security Symposium*, pages 59–74, Aug. 2003. 1

[15] T. Kohno, A. Stubblefield, A. D. Rubin, and D. S. Wallach. Analysis of an electronic voting system. In *Proceedings of the 2004 IEEE Symposium on Security and Privacy*, pages 27–42, May 2004. 1

[16] P. Loscocco and S. Smalley. Integrating flexible support for security policies into the Linux operating system. In *FREENIX Track of the 2001 USENIX Annual Technical Conference (FREENIX'01)*, pages 29–42, June 2001. 1, 6

[17] L. W. McVoy and C. Staelin. LMbench: Portable tools for performance analysis. In *Proceedings of the 1996 Annual Usenix Technical Conference*, pages 279–294, Jan. 1996. 5.1

[18] A. Moshchuk, T. Bragin, S. D. Gribble, and H. Levy. A crawler-based study of spyware in the web. In *Proceedings of the 13th Annual Symposium on Network and Distributed System Security (NDSS 2006)*, Feb. 2006. 4.1

[19] B. Pfitzmann, J. Riordan, C. Stueble, M. Waidner, and A. Weber. The PERSEUS system architecture. Technical Report RZ 3335, IBM Research Division, Sept. 2001. 6

[20] S. Saroiu, S. D. Gribble, and H. M. Levy. Measurement and analysis of spyware in a university environment. In *1st Symposium on Networked Systems Design and Implementation (NSDI 2004)*, pages 141–153, Mar. 2004. 4.1

[21] J. S. Shapiro, J. M. Smith, and D. J. Farber. EROS: a fast capability system. In *Proceedings of the 17th ACM Symposium on Operating Systems Principles (SOSP 1999)*, pages 170–185, Dec. 1999. 6

[22] J. S. Shapiro, J. Vanderburgh, E. Northup, and D. Chizmadia. Design of the EROS trusted window system. In *Proceedings of the 13th USENIX Security Symposium*, pages 165–178, Aug. 2004. 3.1

[23] L. Singaravelu, C. Pu, H. Härtig, and C. Helmuth. Reducing TCB complexity for security-sensitive applications: Three case studies. In *Proceedings of EuroSys 2006*, Apr. 2006. 1, 6

[24] M. Vrable, J. Ma, J. Chen, D. Moore, E. Vandekieft, A. C. Snoeren, G. M. Voelker, and S. Savage. Scalability, fidelity, and containment in the Potemkin virtual honeyfarm. In *Proceedings of the 20th ACM Symposium on Operating Systems Principles (SOSP 2005)*, pages 148–162, Oct. 2005. 3.2

[25] Webstone: The Benchmark for Web Servers, 2006. `http://www.mindcraft.com/benchmarks/webstone/`. 5.2

[26] A. Whitaker, M. Shaw, and S. D. Gribble. Scale and performance in the Denali isolation kernel. In *Proceedings of the 5th Symposium on Operating Systems Design and Implementation (OSDI 2002)*, pages 195–209, Dec. 2002. 6