

Orchestra: Intrusion Detection Using Parallel Execution and Monitoring of Program Variants in User-Space

Babak Salamat Todd Jackson Andreas Gal Michael Franz

Department of Computer Science
School of Information and Computer Sciences
University of California, Irvine
{bsalamat, tmjackso, gal, franz}@uci.edu

Abstract

In a Multi-Variant Execution Environment (MVEE), several slightly different versions of the same program are executed in lockstep. While this is done, a monitor compares the behavior of the versions at certain synchronization points with the aim of detecting discrepancies which may indicate attacks.

As we show, the monitor can be implemented entirely in user space, eliminating the need for kernel modifications. As a result, the monitor is not a part of the trusted code base.

We have built a fully functioning MVEE, named Orchestra, and evaluated its effectiveness. We obtained benchmark results on a quad-core system, using two variants which grow the stack in opposite directions. The results show that the overall penalty of simultaneous execution and monitoring of two variants on a multi-core system averages about 15% relative to unprotected conventional execution.

Categories and Subject Descriptors D.4.6 [*Operating Systems*]: Security and Protection — *Security kernels*; K.6.5 [*Management of Computing and Information Systems*]: Security and Protection — *Invasive software*.

General Terms Design, Experimentation, Performance, Reliability, Security

Keywords process monitoring, intrusion detection, multi-variant execution, multi-core processors, vulnerability

1. Introduction

Software vulnerabilities have been a major threat for decades. The use of safe programming languages, such as Java and C#, in recent years has alleviated the problem, but there are

still many software packages written in C and C++. High performance and low-level programming provisions have made C/C++ indispensable for many applications, but writing safe and secure programs using these languages is often difficult. As a result, software vulnerabilities continue to exist in software and finding mechanisms to spot and remove them automatically continues to be a major challenge.

Many techniques have been developed to eliminate vulnerabilities, but none of them provide an ultimate solution. Modern static-analysis tools are capable of finding many programming errors, but lack of run-time information limits their capabilities, preventing them from finding all errors. They also produce a relatively large number of false positives, making them expensive to deploy in practice. Dynamic and run-time tools are often not effective either, because they do not have a reference for comparison in order to detect misbehavior. Moreover, the performance overhead of sophisticated detection algorithms utilized by such run-time tools is often prohibitively high in production systems [Hastings 1992, Nethercote 2003].

Multi-variant code execution [Cox 2006, Berger 2006, Salamat 2008a;b] is a run-time monitoring technique that prevents malicious code execution and addresses the problems mentioned above. Vulnerabilities that allow the injection of malicious code are among the most dangerous form of security flaws since they allow attackers to gain complete control over the targeted system. Multi-variant execution protects against malicious code execution attacks by running two or more slightly different variants of the same program in lockstep. At certain synchronization points, their behavior is compared against each other. Divergence among the behavior of the variants is an indication of an anomaly in the system and raises an alarm.

Extra computational overhead imposed by multi-variant execution is in the range afforded by most security sensitive applications where performance is not the first priority, such as government and banking software. Besides, the large amount of parallelism which inherently exists in multi-variant execution helps it take advantage of multi-core pro-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EuroSys'09, April 1–3, 2009, Nuremberg, Germany.
Copyright © 2009 ACM 978-1-60558-482-9/09/04...\$5.00

processors. The number of cores in multi-core processors is increasing rapidly. For instance, Intel has promised 80 cores by 2011 [Intel 2006]. Many of these cores are often idle due to the lack of extractable parallelism in most applications or due to the bottlenecks imposed by memory or I/O devices [Hsu 2003]. A multi-variant execution environment (MVEE) can engage the idle cores in these systems to improve security with little performance overhead.

Unlike many previously proposed techniques to prevent malicious code execution [Kc 2003, Barrantes 2003, Cowan 1998] that use random and/or secret keys in order to prevent attacks, multi-variant execution is a secret-less system. It is designed on the assumption that program variants have identical behavior under normal execution conditions (“in-specification” behavior), but their behavior differs under abnormal conditions (“out-of-specification” behavior). Therefore, the choice in what to vary, e.g., heap layout or instruction set, has a vital role in protecting the system against different classes of attacks. It is important that every variant be fed identical copies of each input from the system simultaneously. This design makes it impossible for an attacker to send individual malicious inputs to different variants and compromise them one at a time. If the variants are chosen properly, a malicious input to one variant causes collateral damage in some of the other instances, causing them to deviate from each other. The deviation is then detected by a monitoring agent which enforces a security policy and raises an alarm.

In contrast to previous work, our MVEE is an unprivileged user-space application which does not need kernel privileges to monitor the variants and, therefore, does not increase the trusted computing base (TCB) for processes not running on top of it. Increasing the size of the TCB is detrimental to the overall security of a system. This is because larger code bases are more prone to errors and are harder to validate. This has raised concerns in recent years and many researchers have started investigating methods to reduce the TCB size [McCune 2008, Kauer 2007, Murray 2008].

Our proposed architecture allows running conventional applications without engaging the MVEE (see Figure 1). Thus, normal applications may run conventionally on the system and in parallel with security sensitive applications which are executed on top of the MVEE.

In particular, this paper contributes the following:

- A novel technique to build a user-space MVEE that does not need any OS kernel modification. Our MVEE supervises the execution of parallel instances of the subject application using the debugging facilities of a standard Linux kernel.
- A new technique for variant generation based on reversing the direction of stack growth. Utilizing this technique with multi-variant execution defends against known stack-based buffer overflow attacks in real time.




Figure 1. Our proposed architecture does not grow trusted code of the operating system and allows execution of conventional applications without engaging the MVEE.

- A solution to the problem of preventing false positives caused as a result of inconsistent scheduling of threads and processes in multi-threaded and multi-process applications.
- A solution to the problem of preventing false positives caused by asynchronous signal delivery.
- Solutions to support a wider range of system calls in multi-variant execution, including the `exec` family.

Section 2 describes the monitoring mechanism used in our MVEE. Section 3 discusses the sources of inconsistencies among the variants which could cause false positives in MVEEs and presents methods to handle them. Section 4 presents reverse stack execution as a variant generation technique that allows stopping stack-based buffer overflow attacks when used in an MVEE. Section 5 evaluates security and performance of our implementation. Section 6 presents related work and Section 7 concludes the paper.

2. The Monitor

The monitor is the main component of our multi-variant execution environment. It is the process invoked by the user and receives the paths of the executables that must be run as variants. The monitor creates one child process per variant and starts execution. It allows the variants to run without interruption as long as they are modifying their own process space. Whenever a variant issues a system call, the request is intercepted by the monitor and the variant is suspended. The monitor then attempts to synchronize the system call with the other variants. All variants need to make the exact same system call with equivalent arguments (explained below) within a small time window. The invocation of a system call is called a *synchronization point*.

Our monitor has a set of rules for determining if the variants are synchronized with each other. If p_1 to p_n are the variants of the same program p , they are considered to be in conforming states if at every synchronization point the following conditions hold:

1. $\forall s_i, s_j \in S : s_i = s_j$
 where $S = \{s_1, s_2, \dots, s_n\}$ is the set of all invoked system calls at the synchronization point and s_i is the system call invoked by variant p_i .
2. $\forall a_{ij}, a_{ik} \in A : a_{ij} \equiv a_{ik}$
 where $A = \{a_{11}, a_{12}, \dots, a_{mn}\}$ is the set of all the system call arguments encountered at the synchronization point, a_{ij} is the i^{th} argument of the system call invoked by p_j and m is the number of arguments used by the encountered system call. A is empty for system calls that do not take arguments. Note that argument equivalence does not necessarily mean that argument values themselves are identical. For example, when an argument is a pointer to a buffer, the contents of the buffers are compared and the monitor expects them to be the same, whereas the pointers (actual arguments) themselves can be different. Formally, the argument equivalence operator is defined as:

$$a \equiv b \Leftrightarrow \begin{cases} \text{if type} \neq \text{buffer} : a = b \\ \text{else} : \text{content}(a) = \text{content}(b) \end{cases}$$

with *type* being the argument type expected for this argument of the system call. The content of a buffer is the set of all bytes contained in it:

$$\text{content}(a) := \{a[0] \dots a[\text{size}(a) - 1]\}$$

with the *size* function returning the first occurrence of a zero byte in the buffer in case of a zero-terminated buffer, or the value of a system call argument used to indicate the size of the buffer in case of buffers with explicit size specification.

3. $\forall t_i \in T : t_i - t_s \leq \omega$
 where $T = \{t_1, t_2, \dots, t_n\}$ is the set of times when the monitor intercepts system calls, t_i is the time that system call s_i is intercepted by the monitor, and t_s is the time that the synchronization point is triggered. This is the time that the first system call invocation is encountered at this synchronization point. ω is the maximum amount of wall-clock time that the monitor will wait for a variant. ω is specified in the policy and is application and hardware dependent. For example, on an n -processor system ω may be small because the expectation is that the variants are executed in parallel and should reach the synchronization point almost simultaneously. Once ω has elapsed, those variants that have not invoked any system call are considered non-complying.

If any of these conditions is not met, an alarm is raised and the monitor takes an appropriate action based on a configurable policy. In our current prototype, we terminate and restart all the variants, but other policies such as voting among the variants and terminating the non-conforming ones are possible.

Care should be taken when using majority voting, as behavior of the majority does not necessarily indicate correct

behavior. If the majority are susceptible to a particular type of attack, the system could incorrectly terminate the legitimate minority and continue with the compromised variants. Therefore, the choice of variation mechanisms and the number of the variants have a vital role in the correctness of the system when majority voting is used to tolerate attacks. While the focus of this paper is on intrusion detection and prevention, we plan to publish mechanisms designed to tolerate attacks and repair compromised variants in our future work.

2.1 Monitor Security

The monitor isolates the variants from the OS kernel and monitors all communications between them and the kernel. As mentioned before, the monitor is implemented as an unprivileged process that uses the process debugging facilities of the host operating system (Linux) to intercept system calls. This mechanism simplifies maintenance as patches to the OS kernel need not be re-applied to an updated version of the kernel. Moreover, errors in the monitor itself are less severe since the monitor is a regular unprivileged process, as opposed to a kernel patch or module running in privileged mode. If the monitor was compromised, an attacker would be limited to user-level privileges and would need a privilege escalation to gain system-level access.

The monitor is a separate process with its own address space and no other process in the system, including the variants, can directly manipulate its memory space. Therefore, it is difficult to compromise the monitor by taking control of a program variant.

Conventional system call monitors are susceptible to mimicry attacks, e.g., [Paramalli 2008]. These monitors expect certain sequences of system call invocations; if the monitored program does not follow any of the known sequences, they raise an alarm and stop execution. The conventional monitors cannot check and verify all the arguments passed to the system calls, especially contents of buffers written to output devices. This is because input data and OS behavior varies between sequences of system calls, changing the arguments and making them unpredictable. Mimicry attacks can remain undetected by keeping system calls the same as those that would have been invoked by the legitimate program, while only changing some of the system call arguments. For example, assume a legitimate Apache server opens an HTML file and sends its contents over the network. A mimicry attack could keep the `open` system call intact and pass the path of a file that contains sensitive information instead of the HTML file to the system call. In this scenario the Apache server would send sensitive information over the network and a naive system call monitor would not be able to detect the attack. Mimicry attacks are not effective against our monitor because the MVEE checks both system calls and their arguments.

2.2 System Call Monitoring

A multi-variant environment and all the variants executed in this system must have the same behavior as that of running any one of the variants conventionally on the host operating system. The monitor is responsible for providing this characteristic by performing the I/O operations itself and sending the results to the variants. When the variants try to read input data, the monitor suspends them, intercepts the input, and then sends identical copies of the data to all the variants. This is not only required to mimic the behavior of single application, but it is also essential to prevent attackers from compromising one variant at a time. Similarly, all output operations are solely performed by the monitor after making sure that all the variants agree on the output data.

Depending on the system call and its arguments, the monitor determines whether the variants should run the system call or it should be executed inside the monitor. System calls that generate immutable results (e.g., `uname`) are allowed to be executed by all the variants. If the system call result is not expected to be the same among all variants and the system call does not change the system state (e.g., `gettimeofday`, `getpid`), the call is executed by the first variant. If it changes the system state (e.g., `write`, `mkdir`), it is executed by the monitor. In either case, the results are copied to all other variants.

The monitor is notified twice per system call, once at the beginning of the call and once when the kernel has finished executing the system call handler and has prepared return values. After ensuring that the variants have invoked the same system call with equivalent parameters, the system call is executed. The Linux `ptrace` implementation requires us to perform a system call once a system call has been initiated by a program variant. However, if the system call is executed only by the monitor, the variants must skip the call. In this case, the monitor replaces the system call by a low-overhead call that does not modify the programs' states (e.g., `getpid`). The debugging interface of the OS allows the monitor to do this by changing the registers of the variant at the beginning of the system call invocation. Changing the registers causes a different system call to be executed than the one initially requested.

File, socket, and standard I/O operations are performed by the monitor and the variants only receive the results. When a file is opened for writing, for example, the monitor is the only process that opens the file and sets the registers of the variants so that it appears to them that they succeeded in opening the file. All subsequent operations on such a file are performed by the monitor and the variants are just recipients of the results. This method fails if the variants try to map a file to their memory spaces using `mmap`. The file descriptor they received from the monitor was not actually opened in their contexts and, hence, `mmap` would return an error. This causes a major restriction because shared libraries are mapped using this approach. We solve the problem by

allowing the variants to open files locally if requested to be opened read-only. This solution solves the problem of mapping shared libraries, but if a program tries to map a file opened for writing, it will fail. This is still an open problem, although our experiences indicate that `mmap` is rarely used in this manner.

When the `mmap` system call is used to map a file into the address space of a process, reads and writes to the mapped memory space are equivalent to reads and writes to the file, and can be performed without calling any system call. This allows an attacker to take control over one variant and compromise the other variants using shared memory. To prevent this vulnerability, we deny any `mmap` request that can create potential communication routes between the variants and only allow `MAP_ANONYMOUS` and `MAP_PRIVATE`. `MAP_SHARED` is allowed only with read-only permission. In practice, this does not seem to be a significant limitation for most applications.

Our platform also puts certain restrictions on the `exec` family of system calls. These system calls are allowed only if the files that are required to be executed are in a white-list passed to the monitor. The full path of all executables that each variant is allowed to execute is provided to the monitor and the monitor ensures that the variants do not execute any program other than those provided. It is obvious that the variants and all the executables that they can execute must be properly diversified.

2.3 Monitor-Variant Communication

The monitor spawns the variants as its own children and traces them. Since the monitor is executed in user mode, it is not allowed to directly read from or write to the variants' memory spaces. In order to compare the contents of indirect arguments passed to the system calls, the monitor needs to read from the memory of the variants. Also, it needs to write to their address spaces, if a system call executed by the monitor on behalf of the variants returns results in memory.

One method to read from the memory of the processes is to call `ptrace` with `PTRACE_PEEKDATA` when the variants are suspended. `PTRACE_POKEDATA` can similarly be used to write to the variants. Because `ptrace` only returns four bytes at a time, `ptrace` has to be called many times to read a large block of memory from the variants' address spaces. Every call to `ptrace` requires a context switch from the monitor to the OS kernel and back, which makes this technique inefficient for reading large buffers. To improve performance, we create a shared memory block per variant which is shared by the monitor and one variant.

Shared memory is chosen over named pipes (FIFOs) for performance reasons. Anonymous pipes cannot be used because they can only be created between a child process and its parent, while not all the variants in our system are children of the monitor. Because the variants may create new child processes, the monitor is not the parent of these children. Therefore, they cannot be connected to the monitor

through anonymous pipes. Named pipes work well in connecting these processes to the monitor, but they are not as efficient as shared memory. The downside of using both shared memory and FIFOs is the security risk, since any process can connect to them and try to access their contents. However, each shared memory block has a key and processes are allowed to attach a block only if they have the correct key. When we create shared memory blocks, their permissions are set so that only the user who has executed the monitor can read from or write to them. Therefore, the risk is limited to the case of a malicious program that is executed in the context of the same user or a super user. Both cases would be possible only when the system is already compromised. Also note that a compromised variant cannot access another variant’s shared memory even if it somehow found the other variant’s shared memory key, because attaching a shared memory block needs a system call invocation which is caught by the monitor.

Attaching the shared memory blocks to variants, as well as reading from and writing to them is not built into the applications executed in the MVEE. It is the monitor that has to force the variants to perform these operations. The creation of the shared memory blocks is postponed until they are needed. They are created by the monitor, but attaching to them has to be performed by the variants. Our method of forcing the variants to perform the required operations is based on the fact that the monitor only needs to read from or write to the address spaces of the variants when they are suspended at a system call. At such a point, the monitor makes a backup of the registers of the variants and replaces the original system call with an appropriate one (e.g., `ipc` or `shmget`). This makes the variants run the new system call instead of the original one and enables them to attach the appropriate shared memory block. After performing the operation, if the original system call needs to be executed by all variants (e.g., `mmap`), the variants’ registers are restored by the monitor and the original system call is executed.

Reading to or writing from shared memory does not need a system call. In order to perform these operations, the monitor makes each variant to allocate a block of memory using the same system call replacement method explained above. The monitor uses this memory block to inject a small piece of code that copies the contents of a buffer to another one (similar to `memcpy`). The injected code receives the addresses of source and destination buffers and their lengths in registers. Reading from or writing to the shared memory blocks is done by this code. When the monitor needs to access a variant’s memory space, it backs up the variant’s registers and sets the instruction pointer of the variant to the injected code and makes the variant write to its shared memory. A system call invocation instruction (i.e., `int 0x80`) at the end of the injected code notifies the monitor as soon as the variant finishes copying.




Figure 2. Comparison of the performance of transmitting data using shared memory, FIFOs or `ptrace`. The vertical axis shows the elapsed time, in microseconds, to transfer a buffer and the horizontal axis shows the size of the buffer. Both axes are logarithmic scale.

In order to protect this piece of code from being overwritten, the monitor forces the variant to mark it write-protected immediately after the monitor injects the code. A malicious variant cannot mark it writable without being detected by the monitor, because it has to invoke a system call to do so.

Our experiments show that the time spent to transfer a buffer using `ptrace` increases linearly with the buffer size, but it is almost constant using FIFOs when the buffer is smaller than 4KB (see Figure 2). The size of FIFOs is not configurable without recompiling the OS kernel and is set to 4KB in the Linux distribution we use for our experiments. As a consequence, large buffer sizes need multiple FIFO iterations, requiring multiple context switches. These context switches significantly increase the overhead of FIFOs when transmitting large buffers. Shared memory has the least overhead when the buffer size is larger than 40 bytes and for buffers fewer than 40 bytes in length, `ptrace` is the most efficient mechanism. Therefore, the monitor uses `ptrace` to transfer buffers smaller than 40 bytes and uses shared memory for transferring larger ones. For a 128KB buffer, shared memory is more than 900 times faster than `ptrace` and 20 times faster than FIFOs. Hence, using shared memory greatly improves the monitoring performance for applications that frequently pass large buffers to the system calls.

3. Inconsistencies Among the Variants

There are several sources of inconsistencies among the variants that can cause false positives in multi-variant execution. Scheduling of child processes and threads, signal delivery, file descriptors, process IDs, time and random numbers must be handled properly in a multi-variant environment to prevent false positives.

3.1 Scheduling

Scheduling of child processes or threads created by the variants can cause the monitor to observe different sequences of system calls and raise a false alarm. To prevent this situation, corresponding variants must be synchronized to each other. Suppose p_1 and p_2 are the main variants and p_{1-1} is p_1 's child and p_{2-1} is p_2 's child. p_1 and p_2 must be synchronized to each other and p_{1-1} and p_{2-1} must also be synchronized to each other. We may choose to use a single monitor to supervise the variants and their children or we can use several monitors to do so. Using a single monitor can cause unnecessary delays in responding to their requests. Suppose p_1 and p_2 invoke a system call whose arguments take a large amount of time to compare. Just after the system call invocation and while the monitor is busy comparing the arguments, p_{1-1} and p_{2-1} invoke a system call that could be quickly checked by the monitor, but since the monitor is busy, the requests of the children cannot be processed immediately and they have to wait for the monitor to finish its first task.

To tackle this problem, a new monitoring thread is spawned whenever the variants create new child processes. Monitoring of the newly created children is handed over to the new monitor. Figure 3 shows the hierarchy of the variants and their children and also the monitoring processes that supervise them. p_1 and p_2 are the main variants that are monitored by Monitor 1. p_{1-1} and p_{2-1} are the first children of the main variants that are monitored by Monitor 1-1 which is a child of Monitor 1 and so on.

As mentioned before, we use `ptrace` to synchronize the variants. Unfortunately, `ptrace` is not designed to be used in a multi-threaded debugger. As a result, handing the control of the new children over to a new monitor is not straightforward. The new monitor is not allowed to trace the child variants unless the parent monitor detaches from the variants first and lets the new monitor attach to them. When the parent monitor detaches from the variants, the kernel sends a signal to the variants and allows them to continue execution normally, without notifying the monitor at system call invocations. This would cause some system calls to escape the monitoring before the new monitor is able to attach to the variants.

We solved the problem by letting the parent monitor start monitoring the new child variants until they invoke the first system call. For example, in Figure 3 the Monitor 1 starts monitoring p_{1-1} and p_{2-1} until they call the first system call. Monitor 1 saves the system call and its arguments and replaces it with a `pause` system call. Then, Monitor 1 detaches from p_{1-1} and p_{2-1} . The variants receive a `continue` signal, but immediately run `pause` and get suspended. Monitor 1 spawns a new monitoring thread, which would be Monitor 1-1, and passes the process IDs of p_{1-1} and p_{2-1} to it. Monitor 1-1 attaches to the children, restores the original

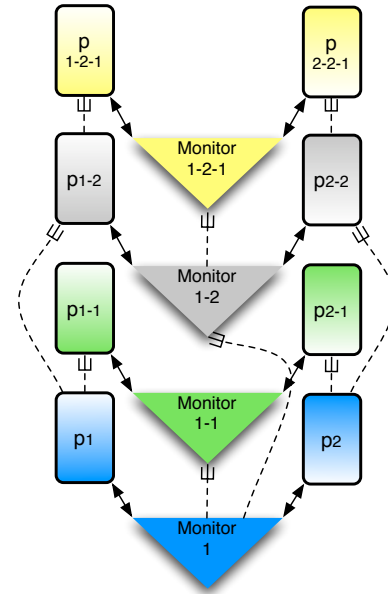


Figure 3. A new monitoring thread is spawned whenever the variants create new child processes. Monitoring of the newly created children is handed over to the new monitoring thread. The dashed lines in the figure connect parent processes to their children.

system call replaced by `pause`, and starts monitoring p_{1-1} and p_{2-1} without missing any system call.

3.2 Signal Delivery

Asynchronous signal delivery can also cause divergence among variants. For example, assume variant p_1 receives a signal and starts executing its signal handler. p_1 's signal handler then invokes system call s_1 , causing the monitor to wait for the same system call from p_2 . Meanwhile, variant p_2 has not received the signal and is still running its main program code. When p_2 calls system call s_2 , the monitor will detect the difference between s_1 and s_2 and raise an alarm.

Whenever a signal is delivered to a variant, the OS pauses the variant and notifies the monitor. The monitor has the choice to deliver the signal to the variant or ignore it. The monitor immediately delivers signals that terminate program execution, such as `SIGTERM` and `SIGSEGV`. Other signals are delivered to all variants synchronously, meaning that signals are delivered to all variants either before or after a synchronization point. If at least half of the variants receive a signal before making a system call, and the rest invoke the system call, the monitor makes the latter variants skip the system call by replacing it with a non-state-changing call and forces them to wait for the signal. The monitor then delivers the signal to all variants and restores the system call in those variants that have been made to skip it. The variants that are forced to wait for a signal and do not receive

it within a configurable amount of time are considered as non-complying.

If fewer than half of the variants receive a signal and the rest invoke a system call, the signal is ignored and the variants which are stopped by the signal are resumed. The monitor keeps a list of pending signals for each variant. The ignored signals are added to these lists. As more variants receive the signal, the monitor checks the lists and when half of the variants have received the signal, the signal is delivered using the method mentioned above. The only difference is that the signal has to be sent again to the variants that ignored it. The monitor sends the signal to these variants and removes it from the variants' pending signal lists.

We use majority voting to decide when to deliver signals despite the fact that majority voting without further consideration about the type and number of variants could introduce potential risks in the system (see Section 2). Since signal delivery time by itself cannot damage the system, using majority voting in signal delivery is not problematic. In fact the protection mechanisms of all modern operating systems prevents processes to have any outside effect unless they invoke a system call. Therefore, if any variant is compromised, it has to eventually invoke a system call to cause damage. Since the system call monitoring does not use majority voting, the attack will be eventually caught before it can cause any damage to the system.

3.3 File Descriptors and Process IDs

As mentioned in the previous section, the monitor allows the variants to open files with read-only permission. Also, anonymous pipes that connect the variants to their children are created by the variants. The file descriptors assigned to these files or pipes are not necessarily the same in different variants and can cause discrepancies among them. Therefore, the monitor replaces the assigned file descriptors by a replicated file descriptor and sends this replicated file descriptor to all the variants running the system call. The monitor keeps a record of the replicated file descriptor and the real file descriptors assigned to the variants by the OS. When a subsequent system call that operates on one of these files is encountered, the monitor restores the original file descriptors before letting the system call execute. As a result, the OS receives the correct file descriptor and operates on the intended file.

A similar approach is taken for process and group IDs. The monitor tracks the process identifiers (PIDs) of the variants. All PIDs of variants monitored by a monitoring thread are mapped to a unique value. Whenever a system call that reads the PID of a variant (`getpid`) is called, its result is replaced by the unique value and, consequently, all the variants receive the same PID. System calls that use these PIDs, such as `kill`, are also intercepted before their execution and the real PIDs of the variants are restored by the monitor. Therefore, the OS receives the correct values when running

the system call. The same approach is taken for the group, parent, and thread group IDs.

3.4 Time and Random Numbers

Time can be another source of inconsistency in multi-variant execution. The solution for this problem is simple. Whenever a time-reading system call is encountered, the monitor invokes the same system call only once and sends the result that it has obtained to all the variants.

Random numbers that are generated by the variants would be different if the variants used different random seeds. Removing the sources of inconsistencies makes all the variants use the same seed and generate the same sequence of random numbers. Reading from `/dev/urandom` is also monitored. The variants are not allowed to read this pseudo file directly. The monitor reads the file and sends the result to all the variants.

3.5 False Positives

We have addressed removing most sources of inconsistency among the variants, but there are still a few cases that can cause false positives. Although the variants are synchronized at system calls, the actual system calls are not usually executed at the exact same time. As mentioned above, files that are requested to be opened as read-only are opened by the variants. If any of these files is changed by a third application after one variant has read it and before it is read by the other variants, there is a race condition and the variants will receive different data which will cause divergence among them.

Another false positive can be triggered if variants try to read the processor time stamp counters directly, e.g., using the `RDTSC` instruction available with x86 processors. Reading the time stamp counters is performed without any system call invocation, so the monitor is not notified and cannot replace the results that the variants receive. Using system calls (e.g., `gettimeofday`) to read the system time solves this problem, although it has higher performance overhead.

Applications that output their memory addresses, such as printing the address of objects on the stack or heap, may trigger a false positive.

4. Reverse Stack Execution

Reverse stack execution is a compiler-driven technique to generate variants. We use this method to evaluate the effectiveness of our MVEE. We generate two variants that write the stack in opposite directions; one variant writes the stack conventionally (downward on x86), and the other one writes it in the reverse direction. The variant that writes the stack upward is resilient against activation record overwrites. As Figure 4 shows, when a stack-based buffer overflow vulnerability is exploited, the injected data overwrites the return address of the function in the conventional variant, but the return address remains intact in the reverse stack variant. This

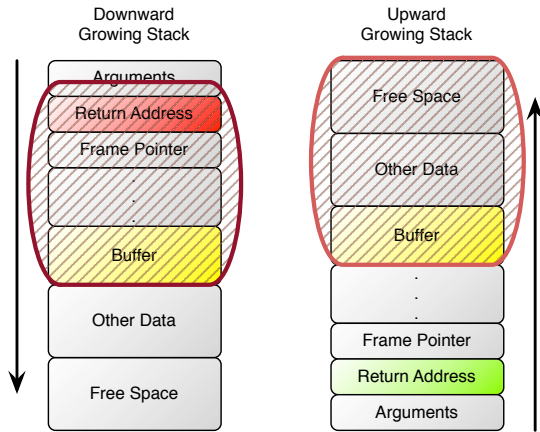


Figure 4. The two boxes with thick borders show the areas that buffer overflows overwrite. The return address and frame pointer of the current function cannot be overwritten by exploiting buffer overflow vulnerabilities when the stack grows upward.

causes the variants to run two different sets of instructions which will cause divergence and is detected by the monitor.

Multi-variant execution of these two variants allows us to prevent known stack-based buffer overflow attacks, including activation record overwrites, return-to-*libc* [Nergal 2001, Pincus 2004] and function pointer overwrites.

We modified GCC version 4.2.1 [GNU] to generate variants that write to the stack in the reverse direction. We also modified *diet libc* [Diet libc] to be able to generate reverse-stack executables. Modifying the stack growth direction is not trivial and involves many challenges. Those interested in compiler techniques for changing the stack growth direction are referred to [Salamat 2008a].

4.1 Effectiveness of Reverse Stack Execution

At first glance, it might seem that a reverse stack executable is inherently immune to stack smashing attacks and there is no need to run a reverse stack executable in an MVEE. Although a reverse stack executable is resilient against many of the known stack-based buffer overflow vulnerabilities, it cannot protect against all possible cases. As an example, consider the following C function:

```
void foo() {
    char buf[100];
    strcpy(buf, user_input_longer_than_buf);
}
```

A user input larger than `buf` can overwrite the return address of `strcpy` and hijack the reverse stack version of the application, since this address is located above the `buf` on the stack. This is shown in the right side of Figure 5.

Now compare how effective a reverse stack executable is when it runs alongside a conventional executable in the MVEE. As Figure 5 shows, exploiting the buffer overflow

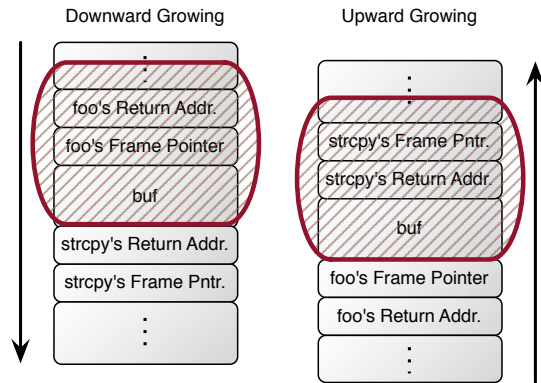


Figure 5. Buffer overflow vulnerabilities can be exploited to overwrite return addresses even in an upward growing stack.

vulnerability in the above code enables an attacker to simultaneously overwrite the return addresses of `strcpy` and `foo` in the reverse and normal executables, respectively. Since no system call is invoked between the point that `strcpy` returns and the point that `foo` returns, the MVEE does not detect any anomaly and lets the variants continue. Therefore, both variants could be diverted to an address where the attack code would be stored.

Since all inputs are identically given to all the variants, the buffer containing the attack code would have the same content in both variants. This means that the addresses used by the instructions in the attack code would be the same in the two variants. For example, suppose that the attack code includes a call to `exec` and passes the address of a small buffer that contains `"/bin/sh"` to `exec`. Also, suppose that `"/bin/sh"` is on the white-list and allowed by the MVEE. Almost all modern OS kernels randomize the beginning of the heap and as a result, the addresses of the corresponding buffers on the heaps of the two variants are not the same. Also, addresses of stack objects are also totally different. Therefore, the address of this small buffer passed to `exec` is different in each variant, but the attack code would have the same address and would fail.

In order to prevent the failure, the attacker would have to divert each variant to a different address that contains attack code valid for that particular variant. This is a high barrier to overcome and the attacker would have to know the exact location of the return addresses on the stack and also the buffers that contain attack code for each variant.

In very high security applications, one might want to add other variation mechanisms or other variants to increase the level of provided security. Instruction set randomization (ISR) [Kc 2003], heap layout randomization [Berger 2006, Bhatkar 2005] and system-call number randomization [Chew 2002] are among possible variation methods that can be used to add extra security. When ISR is used, the injected malicious code would be valid only on one of the

variants and would cause collateral damage on the others. ISR requires a software layer that decodes instructions back to those understandable by the processor at run-time. ISR cannot prevent all kinds of buffer overflow attacks, for example, return-to-*lib(c)* would still be possible. Heap layout randomization, when used in a multi-variant environment, can protect the system against heap-based buffer overruns. System call number randomization causes different system calls to be invoked when the same injected code is executed in different variants. This deviation is easily detectable by the monitor. System call number randomization is easy to deploy, but like ISR, return-to-*lib(c)* would be possible.

5. Evaluation

To demonstrate the effectiveness of the multi-variant execution environment, we create a customized test suite which includes common benchmarks and frequently used applications. This suite allows us to evaluate the security claims and assess the computational tradeoff in CPU- and I/O-bound operations. While our MVEE is capable of running many different variants, we evaluate it with two variants: standard and reverse stack.

5.1 Security

An MVEE is well-suited for network-facing services, and we use documented past exploits of Apache 1.3.29 and Snort 2.4.2 as test vectors. The vulnerabilities and their corresponding exploits are documented with specific environments. Details of these environments include versions of the compiler, operating system, as well as supporting libraries. Changes in one or many of these components of the environment can prevent an exploit from working. As a result, we reconstruct three representative exploits for Apache and Snort in our testing environment, a process that replicates the steps that an attacker would take. Other than these vulnerabilities that exist in real-life applications, we also write small programs with intentional buffer overflow vulnerabilities to test our MVEE.

All vulnerabilities used for testing are stack-based buffer overflow exploits and can be exploited using the techniques described in Aleph One's stack smashing tutorial [Aleph One 1996]. They are chosen because they are representative of a large number of stack-based buffer overflow errors that are present in software, and because these exploits have been available publicly and likely to have been used to obtain illicit access to Apache servers or systems charged with protecting networks. These exploits simulate real-world conditions, as it is likely that other server programs still contain similar implementation errors [Taschner 2007]. Finally, these vulnerabilities are chosen because they are part of the main source package and not dependent on third party libraries or plugins.

5.1.1 Apache mod_rewrite Vulnerability

The Apache mod_rewrite vulnerability was first reported by Jacobo Avariento. It affects all versions prior to Apache 1.3.29 [Dowd 2006]. The vulnerability involves an array of five `char*` variables called `token` in a parsing function called `escape_absolute_uri()`, which can be overflowed given the correct input. In this case, the input required more than five question marks in order to effect the overflow. Avariento's proof-of-concept exploit code [Avariento 2006] is a customized version of Taeho Oh's bindshell shellcode [Oh 2000] and was further modified in order to make Apache exploitable when compiled with GCC 4.2.1. The extra modifications are needed because this version of GCC arranges data on the stack differently than the versions available when Avariento discovered the vulnerability.

5.1.2 Apache mod_include Vulnerability

An anonymous author with the pseudonym "Crazy Einstein" discovered a vulnerability in Apache's mod_include module in 2004 [Einstein 2004]. The vulnerability describes an overflow in a static 8 kB array located on the stack created by the function `handle_echo()`. The array is passed as an argument to `get_tag()`, and when `get_tag()` is given an input longer than 8 kB, `get_tag()` overwrites the return address of `handle_echo()`. The exploit is successful when `handle_echo()` returns and jumps to the shellcode address. In order to make Crazy Einstein's exploit program [Einstein 2006] work in our testing environment, the program was modified to provide extra padding and proper return addresses for shellcode.

5.1.3 Snort BackOrifice Preprocessor Vulnerability

A stack-based buffer overflow vulnerability in the Snort intrusion detection system was discovered by Neel Mehta of ISS X-Force in 2005 [Mehta 2005]. Because of the trusted nature of Snort and the permissions required in order to make it effective, this vulnerability was considered extremely serious since it can give elevated or system-level privileges on a target system and the victim computer does not need to be targeted directly [Manion 2005]. The vulnerability involves a 1 kB array of `char` variables in `BoGetDirection()`, which is used to decode and decrypt BackOrifice packets. A carefully crafted packet, as described by an anonymous author named "rd", can be used to overwrite the return address of `BoGetDirection()`'s caller, `BoFind()` [rd 2005]. In order to make rd's exploit program work, it was modified with proper padding lengths and addresses corresponding to the GCC 4.2.1-based environment.

5.1.4 Effectiveness of the MVEE

For all vulnerabilities, when the variants with a downward growing stack are given the exploit code the exploits succeed and an attacker is able to obtain illicit access to the target computer. When an upward growing stack variant is presented with the same exploit code, the variant continues

to run since the buffer overflow writes into unused memory. When variants of each direction are run in parallel and under supervision of our monitor, the attempted code injection is detected and execution is terminated because shellcode executed by the downward growing stack variant contains system calls. All the buffer overflow attacks on our test programs are also detected by the MVEE, because the attack vectors either cause divergence between the variants or cause one or both variants to be terminated by the OS.

5.2 Performance

The second component of our test suite includes tests designed to assess performance of the MVEE. In order to run these tests, we compile and build executables of *find* 4.1, *tar* 1.12, a MD5 sum generation program (*md5deep* 2.0.1-001), *apache* 1.3.29 and *SPEC CPU2000* [Standard Performance Evaluation Corporation (SPEC)] with both downward and upward growing stacks, and then measure the performance penalty of these applications while running on the MVEE. Although the MVEE concept is targeted towards running security sensitive or network-facing applications, the chosen set of benchmark programs are representatives of I/O- and CPU-bound applications that might be executed in such an environment.

All performance evaluations are performed on an Intel Core 2 Quad Q9300 2.50 GHz system running Ubuntu Linux 8.10 and Linux kernel 2.6.27-9. All benchmark applications are run under two conditions: (a) with the highest scheduling priority (*nice -20*) on an otherwise unloaded machine (Figure 6) and (b) with normal scheduling priority when three other CPU intensive applications run in parallel with the benchmark programs (Figure 7).

Disk-based tests are run several times to remove disk caching effects from skewing the results, and then run again several times to collect data. Once the data is collected, the highest and lowest times are discarded and the average of the remaining times is computed.

Find: *find* is used as an I/O-bound test. In this test, we search the whole disk partition of our test platform for all C source code files (files ending in “.c”). To eliminate effects caused by *find* printing to the screen, the standard output is redirected to `/dev/null`.

Tar: *tar* is selected as a test to show the effects of the MVEE on I/O-bound applications. In this test, we check out the source code of the Eclipse development platform and create a tar archive of the data. The source code is composed of many subdirectories, each of which contains many small text and JAR files. Because of this property, the *tar* test is not reduced to a sequential read operation, which would have occurred if we had used a DVD ISO image. The size of the data set is 3 GB.

md5deep: *md5deep* is a program that generates MD5 sums for files and directories of files. It provides a good mix of I/O- and CPU-bound operations, as the program computes

the MD5 sum while reading each file. *md5deep* has been run over two CD ISO images, totaling 1.5 GB of data.

Apache: The version of Apache that is used for security testing is the same as that of the one used as a performance test. In order to see what effect the monitor has on Apache, we use the provided version of ApacheBench [Apache Software Foundation] to request a 27 KB HTML document. ApacheBench requests the file 10,000 times from a separate computer connected to the target server via an unloaded gigabit ethernet connection.

SPEC CPU2000: SPEC CPU2000 is an industry standard benchmark for testing the computational ability of a system. It is composed of various tools that have heavy CPU-bound characteristics. All of the SPEC tests are used when evaluating the performance of the MVEE, except the FORTRAN and C++ tests, because we currently only have a C library that operates in the reverse-stack mode.

5.3 Analysis

Figure 6 presents the results of the performance evaluation of the MVEE on an otherwise unloaded system and Figure 7 shows these results on a loaded system that runs three other CPU intensive applications in parallel to the benchmarks. These CPU intensive programs fully utilize three cores of the processor and are run to simulate an environment where competition to acquire the CPU is high.

The results in Figure 6 show that the monitor imposes an average performance penalty of 16% and 14% for running both upward and downward growing variants (DU execution) and two downward growing variants (DD execution), respectively. Note that the baseline of the comparison (100% performance) is a normal executable that writes the stack downward. Therefore, in cases where the benchmark is not multi-threaded or multi-process, only one of the processor cores is used when running the baseline and the other cores are idle. The figure also shows that upward-growing stack variants have an average performance penalty of 2%.

The average static size of upward growing stack benchmarks is 10% larger than that of their downward growing stack counterparts. More importantly, the number of dynamic instructions executed by the upward growing stack executables is on average 7% more than those executed by the executables that grow the stack downward (not shown in the figure). The primary reason why the runtime overhead of reverse stack execution (U execution in the figures) is small is that the difference between downward and upward growing stack variants is the addition of some ADD and SUB instructions to manage the stack pointer. Superscalar processors parallelize these instructions with other instructions in the program and execute them with almost no overhead.

In some cases, such as *mcf*, *equake*, *art*, and *tar*, we experience a small speedup when the test is run with a reverse stack. This is likely due to the fact that growing the stack upward better matches the default cache pre-fetching

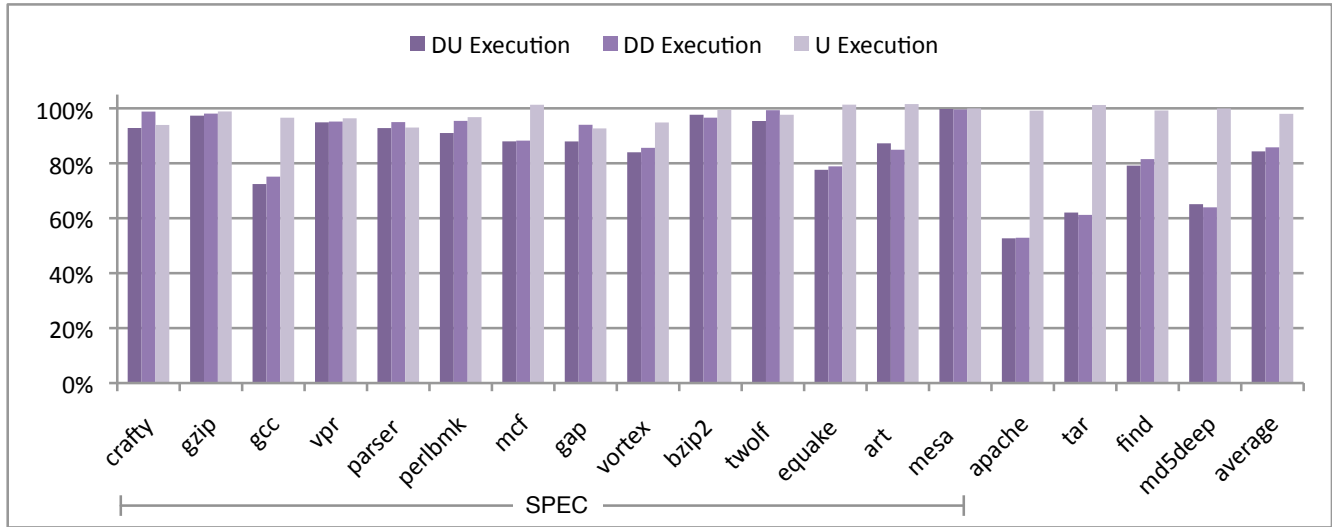


Figure 6. Comparison of the performance of program variants and the MVEE relative to conventional programs when run on an otherwise unloaded system. The performance of two identical copies of a program with a downward growing stack (DD execution), a mix of downward and upward growing stacks (DU execution), and a version of a program compiled for reverse stack execution (U execution) relative to the conventional version are shown. DU and DD execution tests are run in the MVEE, while the U execution tests are run normally.




Figure 7. Comparison of the performance of program variants and the MVEE relative to conventional programs when run on a loaded system which runs three other CPU intensive applications in parallel to the benchmark programs.

heuristics, which results in a slightly better cache hit rate and improves overall performance.

When the tests are run in the MVEE (DU execution and DD execution in Figure 6), the results show that the mostly CPU-bound SPEC tests experience little performance penalty. The two main exceptions to this are *gcc* and *quake*. *gcc* invokes more than 7000 system calls per second, which is very high compared to other benchmarks. Monitoring these system calls causes the performance degra-

dation. The performance degradation of *quake* is caused by memory bandwidth. *quake* is a memory intensive benchmark and memory bandwidth becomes the bottleneck when running two instances of *quake* in parallel.

Performance overhead of the MVEE increases significantly when the SPEC benchmarks are executed on a loaded system (see Figure 7). This is expected, because these are mostly CPU-bound benchmarks. Running two parallel instances of a CPU-bound benchmark on a system that has

only one available processor core causes a performance penalty approaching 50%.

The I/O-bound tests, especially *apache* and *tar*, experience a larger performance penalty. In the case of *apache*, the monitor does all of the socket operations and has to examine all the data sent or received via the network. This means that data that is to be sent has to be transferred from the variants to the monitor, checked for equality, and then sent over the network by the monitor. Also, all requests from the network are received by the monitor and then copied to all the variants. The performance degradation for *apache* is less than 50%. *apache* is a multi-process benchmark and is expected to engage all available processing cores when executed conventionally. As a result, we expected to see a much larger performance drop when running *apache* in the MVEE on a loaded system. Instead, we found that the network interface is the bottleneck which prevents the server from fully utilizing the available processing units [Pinheiro 2001]. Consequently, we do not see a performance drop for *apache* on a loaded system comparing to the results obtained on an unloaded system.

The performance penalty encountered in the *tar* benchmark is partially due to the monitor examining the relative path names of over 300,000 files. Moreover, the output of *tar* is a large file which is written by the monitor. All data that the variants write to the file must be transferred to the monitor, compared, and written to the file by the monitor.

Since the CPU is not the bottleneck in the I/O-bound applications, including *apache*, *tar* and *find*, we do not expect significant changes in the results when these tests are run on a loaded system. However, *tar* results show an improvement when run on a loaded system. Surprisingly, this is not only a relative performance improvement, but also an actual performance improvement; *tar* runs faster on a loaded system even when it is executed conventionally and in the absence of the MVEE. While explaining this phenomenon is not easy, we guess that assigning the same processor core for running the benchmark could be the reason. On a loaded system where processor cores are running other processes and there is only one available core, it is likely that the same core be used to run the benchmark after context switches. This could result in a higher L1 cache hit rate and better overall performance.

6. Related Work

Software security is extremely important, and hence there is a much larger body of related work than space constraints permit us to cite. We apologize for the necessity to select a subset and present the following pioneering earlier work that our research builds upon:

The idea of using diversity to improve robustness has a long history in the fault tolerance community [Avizienis 1977]. The basic idea has been to generate multiple independent solutions to a problem (e.g., multiple versions of a program, developed by independent teams in independent

locations using even different programming languages), with the hope that they will fail independently.

Along with a rising awareness of the threat posed by an increasingly severe computer monoculture, replication and diversity have also been proposed as a means for improving security. [Joseph 1988] proposed the use of n-version programming in conjunction with control flow hashes to detect and contain computer viruses. [McDermott 1997] proposed the use of *logical* replication as a defense tool in an n-version database setting. Rather than merely replicating data across databases, they re-executed commands on each of the replicated databases. This made it much more difficult for an attacker to corrupt the database in a consistent manner by way of a Trojan horse program. [Cohen 1993] proposed the use of obfuscation to protect operating systems from attacks by hackers or viruses, an idea that has reappeared in many variants. [Pu 1996] described a toolkit to automatically generate several different variants of a program, in a quest to support operating system implementation diversity. [Forrest 1997] proposed compiler-guided variance-enhancing techniques such as interspersing non-functional code into application programs, reordering the basic blocks of a program, reordering individual instructions via instruction scheduling, and changing the memory layout. [Chew 2002] proposed automated diversity of the interface between application programs and the operating system by using system call randomization in conjunction with link-time binary rewriting of the code that called these functions. They also proposed randomizing the placement of an application's stack.

Recently, researchers have started to look at providing diversity using *simultaneous* n-variant execution on the same platform, rather than merely creating diversity across a network of computers; our method falls into this category. [Cox 2006] proposed running several artificially diversified variants of a program on the same computer. Unlike our method, their approach requires modifications to the Linux kernel, which increases the maintenance effort and related security risks. They addressed a limited set of the sources of inconsistencies among the variants and their platform did not support certain classes of system calls, including `exec` family.

Also closely related, [Berger 2006] proposed redundant execution with multiple variants that provided probabilistic memory safety by way of a randomized layout of objects within the heap. Their proposed replicated execution mechanism was limited to monitoring the standard I/O. The focus of the work was on reliability (in particular resilience against memory errors) rather than on attack prevention.

A large body of existing research has studied the prevention of buffer overflow attacks at run-time through software only [Kuperman 2005, Wilander 2003]. Several existing solutions are based on obfuscating return addresses and other code and data pointers that might be compromised by an attacker [Bhatkar 2003]. The simplest of these uses an XOR mask to both "encrypt" and "decrypt" such values with low

overhead. [Cowan 1998] takes an alternative approach and places an extra value called a *canary* in front of the return address on the stack. The assumption is that any stack smashing attack that would overwrite the return address would also modify the canary value, and hence checking the canary prior to returning would detect such an attack. StackGuard does not protect against overflows in automatically allocated structures which overwrite function pointers.

[PaX] and [Solar Designer] implement non-executable stacks. This technique does not allow control transfer to the stack by marking the stack memory space as non-executable. Therefore, it prevents attackers from executing code injected to the stack. While many new microprocessors have implemented the necessary hardware support for a non-executable stack, it does not provide protection against return-to-*lib(c)* attacks [Nergal 2001]. This technique also causes compatibility issues. For instance, just-in-time compilers which generate and execute dynamic code may not work properly with non-executable stacks.

7. Conclusions and Outlook

We have presented a new defense against stack-based attacks and a new technique to build multi-variant execution environments that run as unprivileged user-space processes, limiting the repercussions of potential programming errors in building the MVEE. We have also addressed many challenges in developing such environments, including how to deal with sources of inconsistencies among the variants, and have implemented mechanisms to improve performance of the MVEE.

Our results show that deploying the MVEE on parallel hardware provides extra security with modest performance degradation. Our method uses user-space techniques to create the perception of a virtual OS kernel without requiring changes to the OS kernel proper. We have shown that the performance overhead for this approach is acceptable for many applications, in particular considering the beneficial effect of not having to modify kernel code.

Many everyday applications are mostly sequential in nature. At the same time, automatic parallelization techniques are not yet effective enough on such workloads. Even in parallel applications, such as web servers, limited I/O bandwidth prevents us from putting all available processing resources into service. As a result, parallel processors in today's computers are often partially idle. By running programs in MVEEs on such multi-core processors, we put the parallel hardware in good use and make the programs much more resilient against code injection attacks.

As far as future work is concerned, we are interested in ways to *repair* corrupted instances instead of having to terminate them. Such a system would automatically quarantine, re-initialize, and *resume* processes that have become corrupted.

Acknowledgments

We would like to thank Gernot Heiser for his valuable advice in preparing the final version of the paper. We also thank Christian Wimmer for his thoughtful feedback on this paper.

This research effort is partially funded by the Air Force Research Laboratory (AFRL) under agreement number FA8750-05-2-0216. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of AFRL or any other agency of the United States Government.

References

- [Aleph One 1996] Aleph One. Smashing the stack for fun and profit. *Phrack*, 7(2), 1996.
- [Apache Software Foundation] Apache Software Foundation. ab - Apache HTTP Server Benchmarking Tool.
- [Avariento 2006] J. Avariento. Exploit for Apache mod_rewrite off-by-one, 2006. URL http://ciberjacobito.com/sec/mod_rewrite.html.
- [Avizienis 1977] A. Avizienis and L. Chen. On the implementation of n-version programming for software fault tolerance during execution. In *IEEE International Computer Software and Applications Conference (COMPSAC)*, volume 77, pages 149–155, 1977.
- [Barrantes 2003] E.G. Barrantes, D.H. Ackley, T.S. Palmer, D. Stefanovic, and D.D. Zovi. Randomized instruction set emulation to disrupt binary code injection attacks. In *Proceedings of the 10th ACM Conference on Computer and Communications Security (CCS)*, pages 281–289, 2003.
- [Berger 2006] E.D. Berger and B.G. Zorn. Diehard: Probabilistic memory safety for unsafe languages. In *Proceedings of the 2006 ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, pages 158–168, 2006.
- [Bhatkar 2003] S. Bhatkar, D. C. DuVarney, and R. Sekar. Address obfuscation: An efficient approach to combat a broad range of memory error exploits. In *Proceedings of the 12th USENIX Security Symposium*, pages 105–120, 2003.
- [Bhatkar 2005] Sandeep Bhatkar, R. Sekar, and Daniel C. DuVarney. Efficient techniques for comprehensive protection from memory error exploits. In *Proceedings of the 14th USENIX Security Symposium*, pages 271–286, 2005.
- [Chew 2002] M. Chew and D. Song. Mitigating buffer overflows by operating system randomization. Technical report, Department of Computer Science, Carnegie Mellon University, 2002.
- [Cohen 1993] F. Cohen. Operating system protection through program evolution. *Computers and Security*, 12(6):565–584, October 1993.
- [Cowan 1998] C. Cowan, C. Pu, D. Maier, J. Walpole, P. Bakke, S. Beattie, A. Grier, P. Wagle, Q. Zhang, and H. Hinton. Stack-guard: Automatic adaptive detection and prevention of buffer-overflow attacks. In *Proceedings of the 7th USENIX Security Symposium*, pages 63–78, 1998.
- [Cox 2006] B. Cox, D. Evans, A. Filipi, J. Rowanhill, W. Hu, J. Davidson, J. Knight, A. Nguyen-Tuong, and J. Hiser. N-

- variant systems: A secretless framework for security through diversity. In *Proceedings of the 15th USENIX Security Symposium*, 2006.
- [Diet libc] Diet libc. URL <http://www.fefe.de/dietlibc/>.
- [Dowd 2006] M. Dowd. Apache Mod.Rewrite Off-By-One Buffer Overflow Vulnerability, 2006. URL <http://www.securityfocus.com/archive/1/441487/30/0/threaded>.
- [Einstein 2004] Crazy Einstein. Apache mod_include Local Buffer Overflow Vulnerability, 2004. URL <http://www.securityfocus.com/bid/11471>.
- [Einstein 2006] Crazy Einstein. Apache \leq 1.3.31 mod_include Local Buffer Overflow Exploit, 2006. URL <http://milw0rm.com/exploits/587>.
- [Forrest 1997] S. Forrest, A. Somayaji, and D. Ackley. Building diverse computer systems. In *6th Workshop on Hot Topics in Operating Systems (HotOS)*, 1997.
- [GNU] GNU. GNU Compiler Collection (GCC). URL <http://gcc.gnu.org>.
- [Hastings 1992] R. Hastings and B. Joyce. Purify: Fast detection of memory leaks and access errors. In *Proceedings of the Winter USENIX Conference*, volume 136, 1992.
- [Hsu 2003] W. Hsu and A.J. Smith. Characteristics of I/O traffic in personal computer and server workloads. *IBM Systems Journal*, 2003.
- [Intel 2006] Intel. Paul Otellini Keynote. *Intel Developer Forum*, September 2006.
- [Joseph 1988] M.K. Joseph and Avizienis. A. A fault tolerance approach to computer viruses. In *1988 IEEE Symposium on Security and Privacy*, pages 52–58, 1988.
- [Kauer 2007] B. Kauer. Oslo: Improving the security of trusted computing. In *Proceedings of the 16th USENIX Security Symposium*, pages 229–237, 2007.
- [Kc 2003] G.S. Kc, A.D. Keromytis, and V. Prevelakis. Countering code-injection attacks with instruction-set randomization. In *Proceedings of the 10th ACM Conference on Computer and Communications Security (CCS)*, pages 272–280, 2003.
- [Kuperman 2005] B. A. Kuperman, C. E. Brodley, H. Ozdoganoglu, T. N. Vijaykumar, and A. Jalote. Detection and prevention of stack buffer overflow attacks. *Communications of the ACM*, 48(11):50–56, 2005.
- [Manion 2005] A. Manion and J. Gennari. US-CERT Vulnerability Note VU#175500, October 2005. URL <http://www.kb.cert.org/vuls/id/175500>.
- [McCune 2008] J.M. McCune, B.J. Parno, A. Perrig, M.K. Reiter, and H. Isozaki. Flicker: An execution infrastructure for TCB minimization. In *Proceedings of the 3rd European Conference on Computer Systems (EuroSys)*, pages 315–328, 2008.
- [McDermott 1997] J. McDermott, R. Gelinias, and S. Ornstein. Doc, wyatt, and virgil: Prototyping storage jamming defenses. In *13th Annual Computer Security Applications Conference (ACSAC)*, pages 265–273, 1997.
- [Mehta 2005] N. Mehta. Snort Back Orifice Parsing Remote Code Execution, 2005.
- [Murray 2008] D.G. Murray, G. Milos, and S. Hand. Improving Xen security through disaggregation. In *Proceedings of the fourth ACM SIGPLAN/SIGOPS Conference on Virtual Execution Environments*, pages 151–160, 2008.
- [Nergal 2001] Nergal. The advanced return-into-lib(c) exploits: PaX case study. *Phrack*, 2001.
- [Nethercote 2003] N. Nethercote and J. Seward. Valgrind: A Program Supervision Framework. *Electronic Notes in Theoretical Computer Science*, 2003.
- [Oh 2000] T. Oh. Advanced Buffer Overflow Exploit, 2000. URL <http://www.windowsecurity.com/uplarticle/1/advanced.txt>.
- [Parampalli 2008] C. Parampalli, R. Sekar, and R. Johnson. A practical mimicry attack against powerful system-call monitors. In *ACM Symposium on Information, Computer & Communication Security (ASIACCS)*, pages 156–167, 2008.
- [PaX] PaX. URL <http://pax.grsecurity.net>.
- [Pincus 2004] J. Pincus and B. Baker. Beyond stack smashing: Recent advances in exploiting buffer overruns. *IEEE Security and Privacy*, pages 20–27, 2004.
- [Pinheiro 2001] E. Pinheiro, R. Bianchini, E.V. Carrera, and T. Heath. Load balancing and unbalancing for power and performance in cluster-based systems. In *Workshop on Compilers and Operating Systems for Low Power*, pages 182–195, 2001.
- [Pu 1996] C. Pu, A. Black, C. Cowan, and J. Walpole. A specialization toolkit to increase the diversity of operating systems. In *ICMAS Workshop on Immunity-Based Systems*, 1996.
- [rd 2005] rd. THCSnortbo 0.3 - Snort BackOrifice PING exploit, October 2005. URL <http://milw0rm.com/exploits/1272>.
- [Salamat 2008a] B. Salamat, A. Gal, and M. Franz. Reverse stack execution in a multi-variant execution environment. In *Workshop on Compiler and Architectural Techniques for Application Reliability and Security (CATARS)*, 2008.
- [Salamat 2008b] B. Salamat, A. Gal, T. Jackson, K. Manivannan, G. Wagner, and M. Franz. Multi-variant program execution: Using multi-core systems to defuse buffer-overflow vulnerabilities. In *Proceedings of the International Conference on Complex, Intelligent and Software Intensive Systems (CISIS'08)*, pages 843–848, March 2008.
- [Solar Designer] Solar Designer. Non-executable user stack. URL <http://www.openwall.com>.
- [Standard Performance Evaluation Corporation (SPEC)] Standard Performance Evaluation Corporation (SPEC). URL <http://www.spec.org>.
- [Taschner 2007] C. Taschner and A. Manion. US-CERT Vulnerability Note VU#196240, February 2007. URL <http://www.kb.cert.org/vuls/id/196240>.
- [Wilander 2003] J. Wilander and M. Kamkar. A comparison of publicly available tools for dynamic buffer overflow prevention. In *Proceedings of the 10th Annual Symposium On Network And Distributed System Security*, 2003.