Libsafe: Protecting Critical Elements of Stacks

Arash Baratloo, Timothy Tsai, and Navjot Singh
Bell Labs, Lucent Technologies
600 Mountain Ave
Murray Hill, NJ 07974 USA
{arash,ttsai,singh}@research.bell-labs.com
http://www.bell-labs.com/org/11356/libsafe.html

December 25, 1999

Abstract

The exploitation of buffer overflow vulnerabilities in process stacks constitutes a significant portion of security attacks. We present a new method to detect and handle such attacks. In contrast to previous methods, this new method works with any existing pre-compiled executable and can be used transparently, even on a system-wide basis. The method intercepts all calls to library functions that are known to be vulnerable. A substitute version of the corresponding function implements the original functionality, but in a manner that ensures that any buffer overflows are contained within the current stack frame. This method has been implemented on Linux as a dynamically loadable library called *libsafe*. Libsafe has been shown to detect several known attacks and can potentially prevent yet unknown attacks. Experiments indicate that the performance overhead of libsafe is negligible.

1 Introduction

As the Internet has grown, the opportunities for attempts to access remote systems improperly have increased. Several security attacks, such as the 1988 Internet Worm [7, 18, 19], have even become entrenched in Internet history. Some attacks, such as the Internet Worm, merely annoy or occupy system resources. However, other attacks are more insidious because they seize root privileges and modify, corrupt, or steal data.

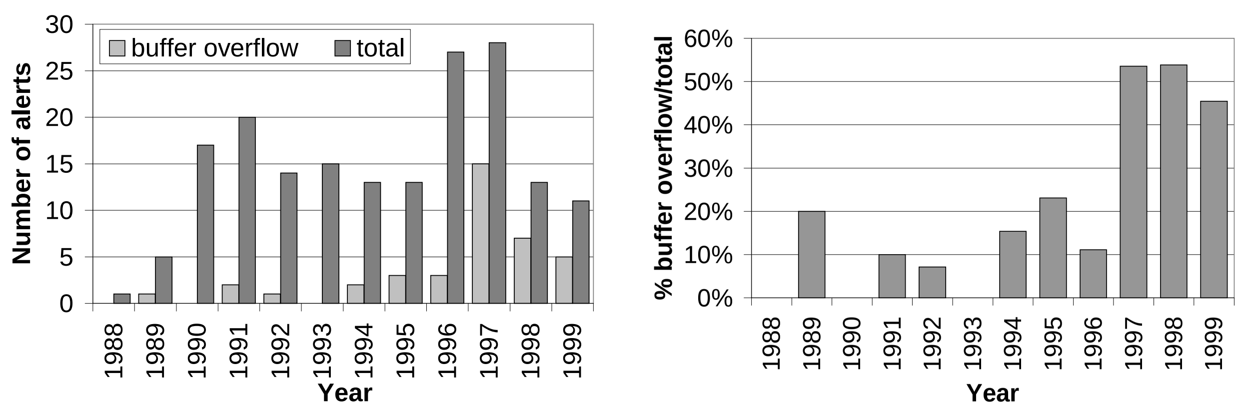

Figure 1: Number of Reported CERT Security Advisories Attributable to Buffer Overflow (Data from [24])

Perhaps, the most common form of attack takes advantage of the buffer overflow bug. Figure 1 shows the increase in the number of reported CERT [3] security advisories that are based on buffer overflow. In recent years, attacks that exploit buffer overflow bugs have accounted for approximately half of all reported CERT advisories. The buffer overflow bug may be due to errors in specifying function prototypes or in implementing functions. In either case, an inordinately large amount of data is written to the buffer, thus overflowing it and overwriting the memory immediately following the end of the buffer. The overflow injects additional code into an unsuspecting process and then hijacks control of that process to execute the injected code. The hijacking of control is usually accomplished by overwriting return addresses on the process stack or by overwriting function pointers in the process memory. In either case, an instruction that alters the control flow (such as a return, call, or jump instruction) may inadvertently transfer execution to the wrong address that points at the injected code instead of the intended code.

Table 1: Partial List of Unsafe Functions in the Standard C Library

Function prototype	Potential problem
`strcpy(char *dest, const char *src)`	May overflow the `dest` buffer.
`strcat(char *dest, const char *src)`	May overflow the `dest` buffer.
`getwd(char *buf)`	May overflow the `buf` buffer.
`gets(char *s)`	May overflow the `s` buffer.
`fscanf(FILE *stream, const char *format, ...)`	May overflow its arguments.
`scanf(const char *format, ...)`	May overflow its arguments.
`realpath(char *path, char resolved_path[])`	May overflow the `path` buffer.
`sprintf(char *str, const char *format, ...)`	May overflow the `str` buffer.

Programs written in C have always been plagued with buffer overflows. Two reasons contribute to this factor. First, the C programming language does not automatically bounds-check array and pointer references. Second, and more importantly, many of the functions provided by the standard C library, such as those listed in Table 1, are unsafe. Therefore, it is up to the programmers to check explicitly that the use of these functions cannot overflow buffers. However, programmers often omit these checks. Consequently, many programs are plagued with buffer overflows, which makes them vulnerable to security attacks.

Preventing buffer overflows is clearly desirable. If one did not have access to a C program's source code, the general problem of automatically bounds-checking array and pointer references is very difficult, if not impossible. So at first, it might seem natural to dismiss any attempts to perform automatic bounds checking at runtime when one does not have access to the source code. One of the contributions of this paper is to demonstrate that by leveraging some information that is available only at runtime, together with context-specific security knowledge, one can automatically prevent security attacks that exploit unsafe functions to overflow stack buffers. Such an exploit is illustrated in the following example.

2 Buffer Overflow Exploit

The most general form of security attack achieves two goals:

1. Inject the attack code, which is typically a small sequence of instructions that spawns a shell, into a running process.
2. Change the execution path of the running process to execute the attack code.

It is important to note that these two goals are mutually dependent on each other: injecting attack code without the ability to execute it is not a security vulnerability.

By far, the most popular form of buffer overflow exploitation is to attack buffers on the stack, referred to as the *stack smashing attack*. As is discussed below, the reason for this popularity is because overflowing stack buffers can achieve *both goals simultaneously*. Another form of buffer overflow attack known as the *heap smashing attack*, is to attack buffers residing on the heap (a similar attack involves buffers residing in data space). Heap smashing attacks are much harder to exploit, simply because it is difficult to change the execution path of a running process by overflowing heap buffers. For this reason, heap smashing attacks are far less prevalent.

```
#include <stdio.h>

char shellcode[] =
  "\xeb\x1f\x5e\x89\x76\x08\x31\xc0\x88\x46\x07\x89\x46\x0c\xb0\x0b"
  "\x89\xf3\x8d\x4e\x08\x8d\x56\x0c\xcd\x80\x31\xdb\x89\xd8\x40\xcd"
  "\x80\xe8\xdc\xff\xff\xff/bin/sh";

char large_string[128];
int i;
long *long_ptr;

int main() {
  char buffer[96];

  long_ptr = (long *)large_string;
  for (i=0; i<32; i++)
    *(long_ptr+i) = (int)buffer;
  for (i=0; i<(int)strlen(shellcode); i++)
    large_string[i] = shellcode[i];
  strcpy(buffer, large_string);
  return 0;
}
```

Figure 2: A Sample Program to Demonstrate a Stack Smashing Attack

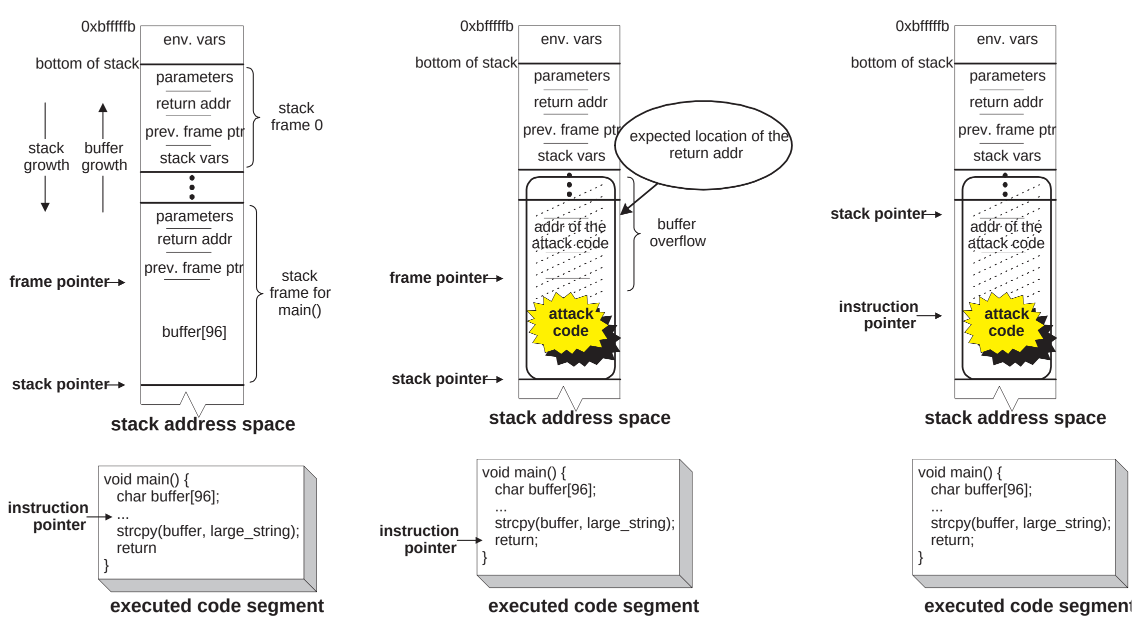

(a) before the attack (b) after injecting the attack code (c) executing the attack code

Figure 3: Buffer Overflow on Process Stack

A complete C program to demonstrate the stack smashing attack is shown in Figure 2. Figure 3 illustrates the address space of a process undergoing this attack. The process stack after executing the initialization code and entering the `main()` function (but before executing any of the instructions) is illustrated in Figure 3(a). Notice the structure of the top stack frame (i.e., the stack frame for `main()`). This stack frame contains, in order, the function parameters, the return address of the calling function, the previous frame pointer, and finally the stack variable `buffer`. Looking at the sample program in Figure 2, a sequence of instructions for spawning a shell is stored in a string variable called `shellcode` (lines 3-6). The two `for` loops in the `main` function prepare the attack code by writing two sequences of bytes to `large_string`: the `for` loop starting on line 16 writes the (future) starting address of the attack code; then the `for` loop starting on line 18 copies the attack code (excluding the terminating null character). The stack is smashed on line 20 by the `strcpy()` function. Figure 3(b) depicts the process' stack space after executing the `strcpy()` call. Notice how the unsafe use of `strcpy()` simultaneously achieves both requirements of the stack smashing attack: (1) it injects the attack code by writing it on the process' stack space, and (2) by overwriting the return address with the address of the attack code, it instruments the stack to alter the execution path. The attack completes once the `return` statement on line 21 is executed: the instruction pointer "jumps" and starts executing the attack code. This step is illustrated in Figure 3(c).

In a real security attack, the attack code would normally come from an environment variable, user input, or even worse, from a network connection. A successful attack on a privileged process would give the attacker an interactive shell with the user-ID of `root`!

3 Related Work

The Internet Worm that infected tens of thousands of hosts in 1988 was one of the first well-known buffer overflow attacks, although there are some anecdotal evidence that buffer overflow attacks date back to the 1960's [4]. In particular, the Internet Worm exploited a buffer overflow vulnerability of the finger daemon. The proportion of attacks based on buffer overflows is increasing each year—in recent years, buffer overflow attacks have become the most widely used type of security attack [24]. Among such attacks, the stack smashing attack is the most popular form [10, 22].

The majority of buffer overflow attacks, including the one exploited by the Internet Worm is based on the stack smashing attack. Detailed descriptions of stack smashing attacks are presented in [20, 22], and cook-book-like recipes are presented in [15, 16, 6].

Researchers in the areas of operating systems, static code analyzers and compilers, and run-time middleware systems have proposed solutions to circumvent stack smashing type of attacks. In most operating systems the stack region is marked as executable, which means that code located in the stack memory can be executed. Because this "feature" is used by stack smashing attacks, making the stack non-executable is a commonly proposed method for preventing overflow attacks. A kernel patch removing the stack execution permission has been made available [17]. This approach, however, has some drawbacks. First, patching and recompiling the kernel is not feasible for everyone. Second, *nested function calls* or *trampoline functions*, which are used extensively by LISP interpreters and Objective C compilers, and the most common implementation of signal handler returns on Unix (as well as Linux), rely on an executable stack to work properly. And finally, an alternative attack on stacks known as *return-into-libc*, which directs the program control into code located in shared libraries, cannot be prevented by making the stack non-executable [25]. Because of those reasons, Linus Torvalds has consistently refused to incorporate this change into the Linux kernel [23].

Snarskii has developed a custom implementation of the standard C library for FreeBSD [21]. Similar to libsafe, this library targets the set of unsafe functions, and inspects the process stack to detect buffer overflows that write across frame pointers. In contrast to libsafe, this is a custom implementation and replaces the standard C library.

Several commonly used tools, such as Lint [11], and those proposed in [8] use compile-time analysis to detect common programming errors. Existing compilers have also been augmented to perform bounds-checking [13]. These projects have demonstrated a limited success in preventing the general buffer overflow problem. Wagner *et al.* have recently proposed the use of compile-time range analysis to ensure the "safe" use of C library functions [24]. Similar to our libsafe method, this project specifically concentrates on the set of unsafe library functions. However, unlike our approach, this method requires source code, which is not always available, and may produce false positives: a correct program may produce warning or error messages.

StackGuard [5] is another compiler extension that instruments the generated code with stack-bounds checks. Specifically, on function entry, a *canary* is placed near the caller's return address on the stack. Before the function returns to the caller, the validity of this canary is checked and the program is terminated if a discrepancy is detected. This approach works on the assumption that if the return address is tampered with (due to buffer overflows), the canary will also be modified, thus causing validation of the canary to fail. With the exception of a few programs, this approach has shown to be effective. In contrast to libsafe, StackGuard introduces a noticeable run-time overhead. Furthermore, StackGuard requires source code access, and there are some programs, such as Netscape Navigator, Adobe Acrobat Reader, and Star Office, that it does not currently support.

Janus [9] is a run-time sand-boxing environment that confines each application to a set of predefined operations. It works on the principle that "an application can do little harm if its access to the underlying operating system is appropriately restricted." It relies on the operating system's debugging features, such as `trace` and `strace`, to observe and to confine a process to a sand-box. Similar to our work, this approach works with existing binary applications and does not require an application's source code. However, unlike our approach, Janus does not work with applications that legitimately need high privileges. For example, the Unix `login` process requires a high level of privilege to execute, but Janus is unable to selectively allow legitimate privileges while denying unauthorized privileges. This inherent limitation prevents Janus from being applied to high privileged applications, where secure execution is most critical.

4 Libsafe

This paper presents a novel method for performing detection and handling of buffer overflow attacks. In contrast to previous methods and without requiring source code, our novel method can transparently protect processes against stack smashing attacks, even on a system-wide basis. The method intercepts all calls to library functions that are known to be vulnerable. A substitute version of the corresponding function implements the original functionality, but in a manner that ensures that any buffer overflows are contained within the current stack frame.

The key idea is the ability to estimate a safe upper limit on the size of buffers automatically. This estimation cannot be performed at compile time because the size of the buffer may not be known at that time. Thus, the calculation of the buffer size must be made after the start of the function in which the buffer is accessed. Our method is able to determine the maximum buffer size by realizing that such local buffers cannot extend beyond the end of the current stack frame. This realization allows the substitute version of the function to limit buffer writes within the estimated buffer size. Thus, the return address from that function, which is located on the stack, cannot be overwritten and control of the process cannot be commandeered.

Table 2: List of Some Known Exploits That Are Detected

Program Name	Version	Description
xlockmore	3.10	Program to lock an X Window display
amd	6.0	Automatic remote file system mount daemon
imapd	3.6	IMAP mail server
elm	2.5 PL0pre8	ELM mail user agent
SuperProbe	2.11	Program to probe for and identify video hardware

We have implemented the previously described method on Linux as a dynamically loadable library called *libsafe*. Libsafe has demonstrated its ability to detect and prevent known security attacks on several commonly used applications, including those listed in Table 2.[1] Libsafe's key benefit, moreover, is its ability to prevent yet unknown attacks.

[1]The security attacks are available from Crv's Security Bugware Page (`http://oliver.efri.hr/~crv/`).

Table 3: Summary of Detection Technique Characteristics

	Instrumentation Techniques				
	None	Libsafe	StackGuard	Janus	Non-Executable Stack
Effectiveness (what types of errors are handled?)					
Kernel Errors	No	No	Yes	No	Yes
Specification Errors	No	Yes	Yes[a]	Maybe[b]	Maybe[c]
Implementation Errors	No	Maybe[d]	Yes[a]	Maybe[b]	Maybe[c]
User Code Errors	No	No	Yes	Maybe[b]	Maybe[c]
Other characteristics					
Performance Overhead	None	Very low	Medium	Medium	None
Disk Usage Overhead	None	Very low	Low	Very low	None
Source Code Needed	No	No	Yes	No	No
Ease of Use	—	Very easy	Easy[e]	Easy-Medium[f]	Easy-Medium[g]

[a] If libraries are instrumented.
[b] Cannot catch hijacked privileges that are similar to legitimate privileges.
[c] For certain types of exploits (see Section 3).
[d] If we know which functions have errors.
[e] Source code must be recompiled, and the compiler may also needed to be recompiled.
[f] Policies need to be written.
[g] Kernel may need to be patched and recompiled.

The characteristics of libsafe are shown in Table 3 along with the corresponding characteristics of two alternative methods, StackGuard and Janus, which were described earlier in Section 3. The first instrumentation technique labeled "None" is presented as a point of comparison and represents the original program with no modifications. The upper half of Table 3 describes the types of errors that each method is able to handle. Specification and implementation errors refer specifically to errors in standard library functions as described in the introductory section. Kernel errors and user code errors refer to implementation errors in kernel code and user code, respectively. The bottom half of the table describes other characteristics. The performance overhead includes only the run-time overhead. Time spent during configuration and compilation are not included. The disk usage overhead is the extra disk space required due to additional shared libraries, increased executable binary file sizes, and configuration files. The next to last row indicates whether source code is needed for that method. The ease of use considers the complexity and time requirement of human efforts needed for configuration and compilation.

5 Implementation

The fundamental observations forming the basis of the libsafe library are the following:

- Overflowing a stack variable—that is, injecting the attack code into a running process—does not necessarily lead to a successful stack smashing attack. The attack must also divert the execution sequence of a process to run the attack code.
- Although buffer overflows cannot be stopped in general, automatic and transparent run-time mechanisms can prevent the overflow from corrupting a return address and altering the control flow of a process.

Refer to Figure 3(a) for an example. The `strcpy()` function cannot determine the exact size of the destination variable `buffer`. At the time `strcpy()` is called, the frame pointer (i.e., `ebp` register in the

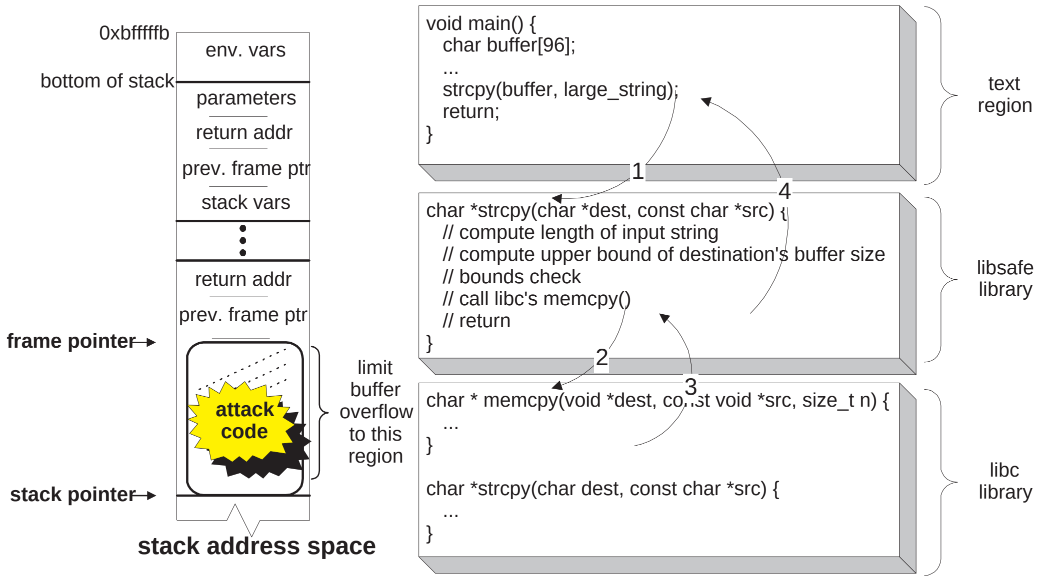

Figure 4: Libsafe Containment of Buffer Overflow

Intel Architecture) will be pointing to a memory location containing the previous frame's frame pointer. Furthermore, this memory address separates the stack variables (local to the current function) from the function arguments. Continuing with the example of Figure 3(a), the size of `buffer` and all other stack variables residing on the top frame cannot extend beyond the frame pointer—this is a safe upper limit. The size of variables residing on previous stack frames—below the top frame—can be bounded by traversing frame pointers to determine the stack frame locations and sizes for those variables. A correct C program should never explicitly modify any stored frame pointers, nor should it explicitly modify any return addresses (located next to the frame pointers). We use this knowledge to detect and limit stack buffer overflows. As a result, the attack executed by calling the `strcpy()` can be detected and terminated before the return address is corrupted (as in Figure 3(b)).

Libsafe implements the above technique. It is implemented as a dynamically loadable library that is preloaded with every process it needs to protect. The preloading injects the libsafe library between the program code and the dynamically loadable standard C library functions. The library can then intercept and bounds-check the arguments before allowing the standard C library functions to execute. In particular, it intercepts the unsafe functions listed in Table 1 to provide the following guarantees:

- Correct programs will execute correctly, i.e., no false positives.
- The frame pointers, and more importantly return addresses, can never be overwritten by an intercepted function. In most cases, an overflow that leads to overwriting the return address can be detected.

Figure 4 illustrates the memory of a process that has been linked with the libsafe library, and in particular, it shows the new implementation of `strcpy()` in the libsafe library. Once the program invokes `strcpy()`, the version implemented in the libsafe library gets executed—this is due to the order in which the libraries were loaded. The libsafe implementation of the `strcpy()` function first computes the length of the source string and the upper bound on the size of the destination buffer (as explained above). It then verifies that the length of the source string is less than the bound on the destination buffer. If the verification succeeds, then the `strcpy()` calls `memcpy()` (implemented in the standard C library) to perform the operation. However, if the verification fails, `strcpy()` creates a `syslog` entry and terminates the program. A similar approach is applied to the other unsafe functions in the standard C library.

The libsafe library has been implemented on Linux. It uses the preload feature of dynamically loadable ELF libraries to automatically and transparently load with processes it needs to protect. In essence, it can be used in one of two ways: (1) by defining the environment variable `LD_PRELOAD`, or (2) by listing the library in `/etc/ld.so.preload`. The former approach allows per-process control, where as the latter approach automatically loads the libsafe library machine-wide.

The libsafe library does not use any Linux specific feature of ELF; these ELF features are available for many other versions of Unix such as Solaris, and have been used for other purposes [1, 14]. Furthermore, an alternative technique with a similar feature can be used for Windows NT [2, 12].

We have installed the libsafe library on a Linux machine. The library is automatically loaded with every process and transparently protects each process from stack smashing attacks. The protected applications include daemon processes such as Apache HTTP server, sendmail, and NFS server, as well as those started by users such as XFree86 server, Enlightenment window manager, GNU Emacs, Netscape Navigator, and Adobe Acrobat Reader. We have used this machine for over a week and found the machine to be stable and running without a noticeable performance hit.

6 Performance

The libsafe library is effective in detecting and preventing stack smashing attacks. Extra code is needed to perform this detection, and that extra code incurs a performance overhead. In this section we quantify the performance overhead associated with use of the libsafe library. Section 6.1 describes the overheads associated with synthetic kernel programs to illustrate the range of possible overheads. Section 6.2 gives performance data for a selected set of actual applications.

All experiments were conducted on a 400 MHz Pentium II machine with 128 MB of memory running RedHat Linux version 6.0. Libsafe and all programs in Sections 6.1 and 6.2 were compiled (and optimized using -O2) with GCC compiler version 2.91.66.

6.1 Kernel Tests

The first time each libsafe function is activated, the initialization of that particular function makes a `dlsym()` call for each libc function that is called from this libsafe function. Because the libc function has the same name as the corresponding libc version, the `dlsym()` call is needed to obtain a pointer to the libc function. Each `dlsym()` call requires 1.26 μs. The interception and redirection of a C library function consists of an additional user-level function call, which approximately adds 0.04 μs of overhead.

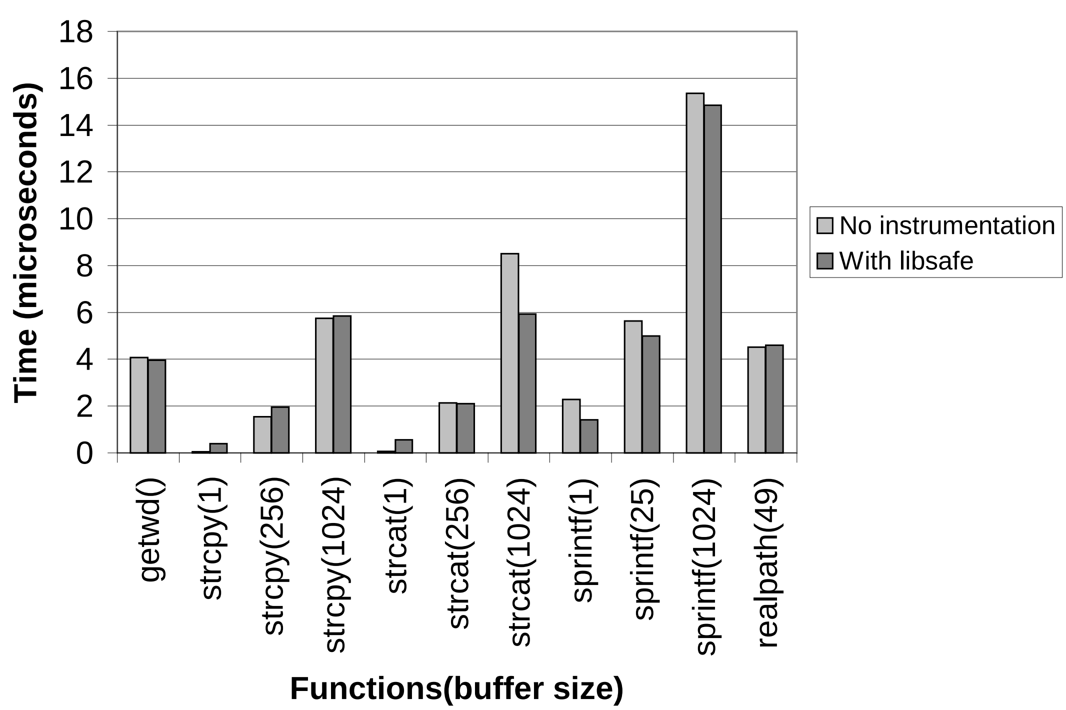

Figure 5: Performance of Libsafe Functions

To quantify the performance overhead of the libsafe library we measure the execution times of five unsafe C library functions and compare the results with our "safe" versions. The results are depicted in Figure 5. Reported times are "wall clock" elapsed times as reported by `gettimeofday()`. An interesting observation is that the libsafe versions of several functions outperform the original versions. This is a repeatable behavior, and we have observed consistent findings on different machines and operating system versions. This effect is due both to low-level optimizations and the fact that libsafe's implementation of most functions is different than those of C library. For example, consider the performance of `getwd()` and `sprintf()` functions. Our libsafe library replaces these functions with equivalent safe versions. In particular, `getwd()` is replaced with `getcwd()` and `sprintf()` is replaced with `snprintf()`; on Linux, the safe versions execute faster.

The figure also shows that the libsafe library can slow down the string operations `strcpy()` and `strcat()` by as much as 0.5 μs per function call. However, as the string size increases, the absolute overhead decreases because the execution time of the safe versions increases more slowly than that for the unsafe versions. In fact, the safe version of `strcat()` used with strings longer than 256 bytes is actually faster than the unsafe version! This is an example of how using a different implementation (e.g., using `memcpy()` to copy a string) can outperform the standard implementation for certain cases.

The slowdown effect of `strcpy()` is observed in the `realpath()` experiment. When a program calls `realpath()`, the libsafe library calls `realpath()` but stores the result in a buffer in its own memory region. It then uses `strcpy()` to copy the result to the final destination. As Figure 5 shows, the slowdown effect of `strcpy()` on `realpath()` is less than 0.05 μs.

6.2 Application Tests

We used four real-world applications to illustrate the performance overhead associated with libsafe. The applications are `quicksort` (a CPU-bound program), `imapd` (a network-bound program), `tar` (an I/O-bound program), and `xv` (a CPU and video-bound program). Figure 6 shows the execution time for each of these applications using (1) the original libc (i.e., without libsafe), (2) the libsafe method, and (3) StackGuard. The execution times are based on 100 runs and are given in seconds, with associated 95% confidence intervals. Reported times are "wall clock" elapsed times as reported by `/bin/time`.

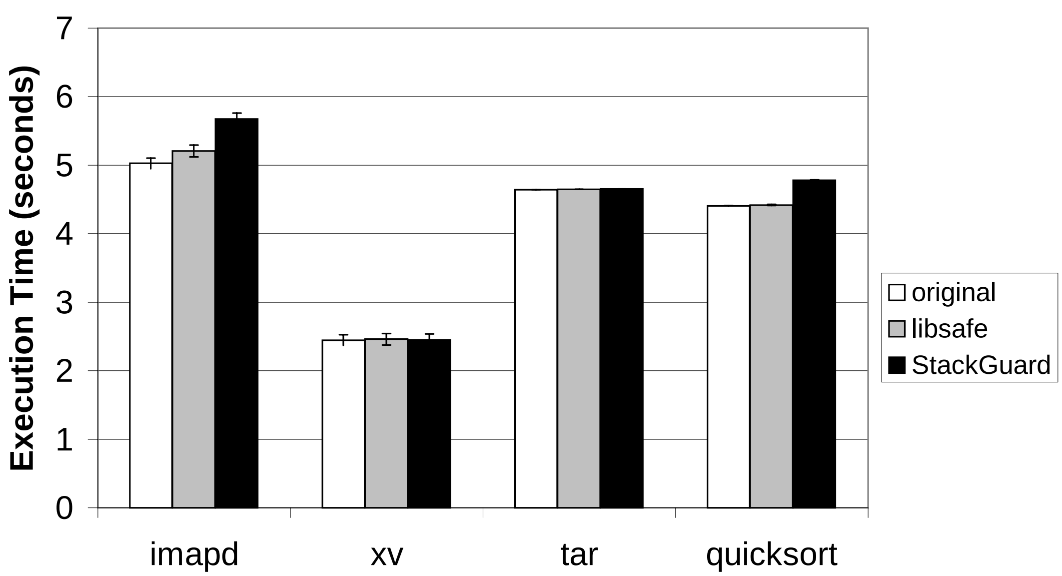

Figure 6: Mean Execution Times (With 95% confidence intervals) of Sample Applications

Figure 6 shows that the overheads associated with all detection methods are reasonable. Libsafe is the most efficient method because only the unsafe library functions are intercepted. The overall application test results are encouraging. We have installed and used libsafe on one of own machine, and in practice, we have found that this overhead is not noticeable.

7 Conclusions

We have described a new method for preventing stack smashing attacks that rely on corrupting the return address, and implemented this method in as a dynamically loaded library called libsafe. The libsafe library instruments a small set of library functions that are known to be vulnerable to buffer overflows.

An interesting finding is the performance of libsafe. We anticipated a low performance overhead at the onset of this project. We were happily surprised to find how little this overhead is in practice. Because of low-level optimizations and because libsafe's implementation of most functions is different than those of C library, for some applications we actually observed a speedup. This is encouraging since it indicates the viability of this approach. Furthermore, the elegance and simplicity of instrumenting the standard C library lead to a stable implementation.

We believe that the stability, minimal performance overhead, and ease of implementation (i.e., no modification or recompilation of source code) of libsafe makes it an attractive first line of defense against stack smashing attacks. We has demonstrated its effectiveness in testing it against several known buffer overflow attacks, but its real benefit, we believe, is its ability to prevent yet unknown attacks.

References

[1] Albert D. Alexandrov, Maximilian Ibel, Klaus E. Schauser, and Chris J. Scheiman. Extending the operating system at the user-level: the Ufo global file system. In *Proceedings of the 7th USENIX Annual Technical Conference*, 1997.

[2] Robert Balzer and Neil Goldman. Mediating connectors. In *Proceedings the 19th IEEE International Conference on Distributed Computing Systems Workshop*, 1999.

[3] CERT coordination center. `http://www.cert.org`.

[4] Crispin Cowan. `http://geek-girl.com/bugtraq/1999_1/0481.html`, 1999. Posting to Bugtraq Mailing List.

[5] Crispin Cowan, Calton Pu, Dave Maier, Heather Hinton, Peat Bakke, Steve Beattie, Aaron Grier, Perry Wagle, and Qian Zhang. StackGuard: automatic adaptive detection and prevention of buffer-overflow attacks. In *Proceedings of the 7th USENIX Security Conference*, 1998.

[6] dark spyrit aka Barnaby Jack. Win32 buffer overflows (location, exploitation and prevention). `http://www.insecure.org`.

[7] Mark W. Eichin and Jon A. Rochlis. With microscope and tweezers: An analysis of the internet virus of november 1988. In *Proceedings of the 1989 IEEE Computer Society Symposium on Security and Privacy (SSP '89)*, 1989.

[8] David Evans. Static detection of dynamic memory errors. *ACM SIGPLAN Notices*, 31(5):44–53, May 1996. Proceedings of the 1996 ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI).

[9] Ian Goldberg, David Wagner, Randi Thomas, and Eric A. Brewer. A secure environment for untrusted helper applications. In *Proceedings of the 6th USENIX Security Symposium*, 1996.

[10] Shawn Instenes. Stack smashing: What to do? *;login: the USENIX Association newsletter*, April 1997.

[11] Stephen C. Johnson. *Lint, a C program checker*. Bell Laboratories, Murray Hill, New Jersey, USA, December 1977. Computer Science Technical Report 65.

[12] Michael B. Jones. Interposition agents: Transparently interposing user code at the system interface. In *Proceedings of the 14th Symposium on Operating Systems Principles (SOSP)*, December 1993.

[13] Richard Jones. Bounds checking patches for gcc. `http://web.inter.NL.net/hcc/Haj.Ten.Brugge`.

[14] Alain Knaff. Zlibc - transparent access to compressed file. http://zlibc.linux.lu.

[15] Mudge. How to write buffer overflows. http://www.insecure.org/stf/mudge_buffer_overflow_tutorial.html, 1995.

[16] Aleph One. Smashing the stack for fun and profit. *Phrack Magazine*, 49(14), 1998.

[17] Openwall Project. Linux kernel patch from the openwall project. http://www.openwall.com/linux.

[18] Jon A. Rochlis and Mark W. Eichin. With microscope and tweezers: The worm from MIT's perspective. *Communications of the ACM*, June 1989.

[19] Donn Seeley. A tour of the worm. In *Proceedings 1989 Winter USENIX Technical Conference*, January 30 - February 3 1989.

[20] Nathan Smith. Stack smashing vulnerabilities in the UNIX operating system. http://millcomm.com/~nate/machines/security/stack-smashing/nate-buffer.%ps, 1997.

[21] Alexandre Snarskii. Increasing overall security.... ftp://ftp.lucky.net/pub/unix/local/libc-letter and http://www.lexa.ru:8100/snar/libparanoia, 1997.

[22] Evan Thomas. Attack class: Buffer overflows. *Hello World!*, 1999.

[23] Linus Torvalds. Posting to linux kernel mailing list. http://www.lwn.net/980806/a/linus-noexec.html, 1998.

[24] David Wagner, Jeffrey S. Foster, Eric A. Brewer, and Alexander Aiken. A first step towards automated detection of buffer overrun vulnerabilities. In *Proceedings 7th Network and Distributed System Security Symposium (to appear)*, February 2000.

[25] Rafel Wojtczuk. Defeating solar designer non-executable stack path. http://geek-girl.com/bugtraq, January 1998.