

## **Methods of Protecting the Stack Overflow Vulnerability**

Z. Muzaffar   A. Rasheed   K. Salah <sup>\*\*</sup>

*Department of Information and Computer Science*

*King Fahd University of Petroleum and Minerals*

*Dhahran 31261, Saudi Arabia*

*Email: {zeeshanm,aimanr,salah}@ccse.kfupm.edu.sa*

---

<sup>\*\*</sup> Corresponding Author

## Abstract

*With the need of sophisticated software systems, the size and complexity is also growing consequently increasing the number of vulnerabilities. Most of the vulnerabilities are related to security issues. The most dominant of such vulnerabilities are buffer overflows. Buffer overflow vulnerabilities dominate the area of remote network penetration vulnerabilities, where an anonymous network user tries to get control of host. In this paper we survey various forms of buffer overflow vulnerabilities and their solutions. Taxonomy of solutions is compiled based upon different classes of buffer overflow attacks. Further a comprehensive analysis and evaluation of each technique is carried out based on various criteria. Our analysis and evaluation can help security researchers and designer to choose among these existing solutions. In addition we propose potential enhancements to some existing solutions.*

**KEYWORDS:** Buffer Overflow Vulnerabilities, Network Security, and Qualitative Analysis.

## 1. Introduction

Most of the today's applications are designed by keeping in mind the importance of time and space efficiency, and quite often security is not trepidation. The security of application is now becoming a major concern due to the fact that as the size and complexity of applications are growing, the numbers of bugs are also increasing resulting into the exploitation and vulnerabilities of applications.

These vulnerabilities arose due to wide usage of C/C++ language in critical software including almost all the dominating operating systems of the present time. Actually these programs are optimized for high performance and offer little error checking capability, and hence open the doors for the attackers to exploit this deficiency. The common practice using these languages is to make excessive use of arrays, pointers and memory allocation/de-allocation, on which the programmers don't have any control during program execution. There are various C/C++ library routines that perform different sort of manipulations using these arrays and

pointes etc. For example, the *strcpy(char \* dst, const char \*src)* function copies a string pointed to by *src* to the string pointed to by *dst*. Now the problem is that this built-in C function doesn't perform the bound checking for the *dst* while copying from the *src* [12,16,20]. It's the responsibility of the programmers to ensure whether the destination buffer has ample memory to completely accommodate the source string. Unfortunately, these checks are often omitted in existing software (including well-known modern operating systems). Consequently, the missing checks could be exploited by the attackers to gain the control of a computer by launching buffer overflow attacks.

A buffer overflow occurs when a source writes beyond the boundaries of the destination buffer and hence overwrites the contents immediately adjacent to it. Generally this causes a memory corruption or memory fault of the program. However, it can also be exploited maliciously to alter the control flow of a program in order to break the security of the system [8,18]. Buffer overflow attacks can be generally divided into two sub-categories depending on the buffer allocation over stack and heap, thus known as Stack smashing attacks and Heap smashing attacks, respectively [20,24,26].

```
char shellcode[] =
"\xeb\x1f\x5e\x89\x76\x08\x31\xc0\x88\x46\x07\x89\x46\x0c\xb0\x
0b\x89\xf3\x8d\x4e\x08\x8d\x56\x0c\xcd\x80\x31\xdb\x89\xd8\x40
\xcd\x80\xe8\xdc\xff\xff\xbin/sh";
char large_string[128];
void main()
{ char buffer[96];
  int i;
  long *long_ptr = (long *) large_string;
  for (i = 0; i < 32; i++)
    *(long_ptr + i) = (int) buffer;
  for (i = 0; i < strlen(shellcode); i++)
    large_string[i] = shellcode[i];
  strcpy(buffer, large_string); }
```

**Figure 1. Stack overflow attack**

There are four attack targets that could be victimized to change the control flow [8]:

- i. *The return address, allocated on the stack*
- ii. *The old base pointer, allocated on the stack: Function pointers, allocated on the heap, in the BSS or DATA segment, or on the stack either as a local variable or as a parameter*

iii. *Longjmp buffers, allocated on the heap, in the BSS or DATA segment, or on the stack either as a local variable or as a parameter*

To cater such attacks researchers have come up with two approaches namely; static and dynamic discussed in section 3.

The rest of the paper is organized as follows; Section 2 discusses the basics of stack smashing attacks. Section 3 is about some of the renowned protection mechanisms against stack smashing attacks discussed in taxonomical manner. Section 4 contains the qualitative comparisons of selected techniques. Section 5 gives some suggestions for possible enhancements to already existing approaches and finally Section 6 concludes the paper.

## 2. Stack Smashing Attacks

When a function is called the return address of the calling function is stored on the stack together with the stack frame pointer. The naïve approach of launching the stack smashing attack is to overflow a buffer in such a way that the return address is overwritten and it points to the malicious code, so that when the called function returns the flow of program execution transfers to the malicious code rather than returning to the calling function [5,16,18,20]. This is the way an attacker can execute some malevolent code if it is cleverly placed in memory. This malicious code is usually a binary representation of a shell code with the help of which an attacker can spawn a new shell with the root privileges. This will give him full freedom of launching malicious activities, for example, installing some security breaching application, reading confidential information, or deleting essential information etc. Figure 1 shows a sample code for launching stack overflow attack [24]. Figure 2a and Figure 2b are the pictorial view of stack before and after the attack.

## 3. Protection Mechanism

The campaign for preventing stack-smashing attacks has given rise to two fundamental schools of thought. Each of these disciplines has proposed several solutions and most of them are claimed to be promising for tackling such attacks.

**Static Approaches.** The stack overflow vulnerabilities can be detected and/or prevented through an analysis of the program source code. The major advantage of such approaches is that security bugs can be eliminated before application is deployed.

Besides having their own pros and cons, the static approaches require an updated database of programming flaws to test for. Also they all require complete recompilation/modification of source code for a particular application. So all the vulnerable libraries in the existing operating systems must be recompiled and deployed to replace the running copies, unfortunately this is not that easy to accomplish. Therefore, these techniques are helpful for those applications that are being developed. Examples include ITS4, Integer range analysis, Source to source transformation Cyclone, CCured, Gemini etc.

**Dynamic Approaches.** These are the runtime preventive

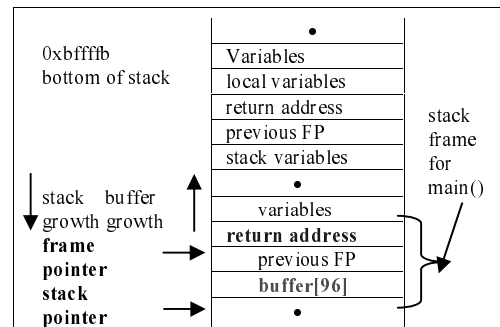


Figure 2a. Stack view before attack

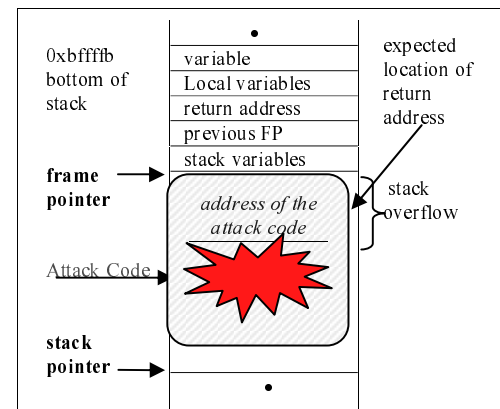


Figure 2b. Stack view after attack

measures for stack smashing attacks. This means that vulnerable programs become less harmful in dynamically protective environment. However, these programs can only tackle the known vulnerabilities because these approaches are based upon patching up the loopholes that are exploited by the attackers. Consequently these approaches are ineffective against the unknown future attacks. Some of the examples are, StackGuard, Stack Shield, Stack Ghost, Libsafe, Libverify, ProPolice, RAD, Purify, Point Guard, Valgrind etc.

Approaches stated above have given rise to many overlapping solutions. We can divide the solutions in three fundamental approaches, which are *Return Address Protection*, *Access Violation Probing*, and *Bounds Checking*.

### 3.1 Return Address Protection Approaches

In this approach, the Return Address is protected to prevent the change of control flow of a vulnerable program. This approach doesn't prevent the injection of malicious code.

**Canary Insertion, StackGuard.** This approach is devised mainly for encountering return address overwrites attacks [1,8,11,15,18,20]. In the stack smashing attack, the attacker has to write the malevolent contents to the *dest* buffer surpassing local variables and all the way to the return address. The core concept of encountering such attack is to put a canary between the old base pointer and return address, as shown in Figure 3. Upon returning from the function this value is validated, if it's tampered then a stack overflow is detected.

The potential ways to circumvent this detection method is to either leave the canary word intact while changing the return address to point to a malicious code, this is made possible by overwriting the return address using an abused pointer and not touching the canary word, or by overwriting it with its correct value and thus in a way not changing it.

To overcome these shortcomings the inventors have proposed two versions of canary words, *random canary* and *terminator canary* [18].

**Return Address as Canary, Libverify.** It is library for return address verification scheme like StackGuard. Unlike StackGuard, Libverify uses the actual function return address




Figure 3. The StackGuard stack frame

as the canary value for each function [8,16]. Moreover, Libverify introduces the verification code at the start of the process execution, rather than during compilation, through a binary re-write of the process memory and hence does not require the recompilation of the software. To accomplish this Libverify rearranges the code a little bit. First each function is fully copied to the heap (requires executable heap) where it can be altered. Then saving and verification of the return address is inserted into function by overwriting the first instruction of the original function with a call to *wrapper-entry* and return instruction with a call to *wrapper-exit*. The copying of functions to the heap is because of the Intel architecture unlike other platforms. Libverify does not protect the integrity of the canary stack. They propose protecting it with *mprotect()* as in the RAD but this will most probably impose a very serious performance penalty. Also, it has same drawback as in the case of Libsafe.

**Guard Value, Propolice.** ProPolice, a GCC compiler patch [8,28], borrows the main idea from StackGuard – it places a canary value which is given the name Guard to detect attacks on the stack. The novelty is that it protects the local stack allocated variables by relocating them in such a way that char buffers come before all the other variables and after the Guard value as shown in Figure 4.




Figure 4. The ProPolice stack frame

In this way apparently it seems that local char buffers can only be over flown to perturb each other, the old base pointer

and the return address, but by placing the guard value between these buffers and the old base pointer all the attacks outside the buffers segment will be detected. Once an attack is detected the process is terminated. The beauty of this approach is that unlike StackGuard, it can nicely tackle Emsi’s attack [27].

ProPolice when tested it was revealed that it is somewhat unstable in the sense that it skipped the protection over small character buffers even if the user had set the protector flag [8].

**Return Address Repository, RAD.** RAD is basically a simple GCC compiler patch that automatically adds protection code into the function prologues and epilogues of the programs compiled by it [1,15]. RAD protection mechanism doesn’t need source code modification. The key idea behind RAD is quite similar to return address protection of Libverify. Every time a function call is made and a new stack frame is created, RAD stores a copy of the new return address. When a function returns, the return address about to be de-referenced is first checked against its copy. There are two variants of RAD, MineZone RAD and Read-Only RAD. MineZone RAD is more efficient while Read-Only RAD is more secure.

**Global Ret Stack, StackShield.** It is a GCC compiler patch for three kinds of protections [8]; two are against the overwriting of return address and one against overwriting of function pointers. Return address overwriting is protected by implementing Ret Range Check and Global Ret Stack.

This preventive approach uses a separate global array of 32-bit entries in the stack for storing the return addresses. Whenever a function call is made, the return address that is pushed onto the stack is also copied into this global space. Even if an attacker succeeds in overwriting the return address, he will not be able to get any benefit out of it as when the function returns, its address is replaced with the global stored value without any comparison with the current address.

**Ret Range Check, StackShield.** This approach is similar to Global Ret Stack mentioned above. The only difference is that instead of replacing the address, it checks the validity of return address and if this check is not validated then the execution is halted [8].

**Return Address and Code Pointers Cryptography, PointGuard.** Point Guard is a generalization of StackGuard such that it places canaries adjacent to all code pointers (function pointers and longjmp buffers) [8,18]. This way it checks the integrity of these code pointers by checking the canaries, if the canaries are intact it means there is no attack otherwise it issues an intrusion alert and exits the program. This approach is used for debugging purposes because of its execution overhead.

### 3.2 Access Violation Probing Approaches

These techniques deal with statically and dynamically checking the memory access limits. Among the following approaches, the first two are static and the rest are dynamic.

**Type Safe Languages.** Unsafe languages like C and C++ do not perform array-bounds checking, which turns out to be a security-critical issue, particularly in handling strings. The risks increase even more dramatically when user-controlled data is on the program stack (i.e., is a local variable). Most of the vulnerabilities are due to lack of safety in languages [8,14]. The safe programming languages on the other hand offer security features making them lucrative for developing extensible environments on wide variety of platforms [7,18]. Such languages provide the mechanism of isolation for untrusted code.

**Code Segment Boundary Checking, StackShield.** Code segment boundary checking entails storing the boundary address of the code segment some where e.g., in a processor register or in data segment. A check is made before executing an instruction to assure that the instruction being executed is within the boundary of code segment by comparing the address of the instruction with the boundary address. Stack Shield makes it possible by declaring a global variable in the data segment, and its address is used as a boundary value [8,18]. Any reference by the function pointers beyond this value is taken as malicious activity and the process is terminated.

**Run-time Memory Access Checking.** Programs may incorrectly read or write memory in a variety of ways; these are called memory access errors. For example, the program may reference a block of memory which has been deallocated through a free() call for a heap block, or because a function returned a pointer to a local variable. Access errors may result in wild pointers in the program and can cause abnormal program behavior, including wrong outputs and segmentation violations. Some kinds of memory access errors can be very hard to track down. Research groups have developed tools to cater these problems, some of them are;

- i. Complete Memory Access Checking, Purify [5,11,18].
- ii. Cyclone [14].
- iii. Source Translator, CCured [14].
- iv. Invalid Memory Access Checking, Valgrind [32].

### 3.3 Bounds Checking Approaches

This technique is used to protect against the code injection attacks [1]. It provides prevention from all buffer overflow attacks, not just those that attempt to alter function activation records [22].

Performance of this technique is proportionate to the number of array and string operations, the greater are such operations the worse the performance. Usually its cost is substantial.

**Buffer Boundary Checking, Libsafe.** The key idea behind Libsafe is the ability to estimate a safe upper limit for the buffers on the stack automatically at the run-time and then ensure that vulnerable function is not allowed to write anything across this and hence protecting the return address from being overwritten [8,16]. Vulnerable functions are considered to be the ones in Table 1 below.

**Table 1. LibSafe protected functions**

Function	Vulnerability
strcpy(char * dest, const char *src)	May overflow 'dest'
strcat(char dest, const char *src)	May overflow 'dest'
getwd(char *buf)	May overflow 'buf'
gets(char *s)	May overflow 's'
scanf(const char *format, ...)	May overflow 'arguments'
Realpath(char*path, char res_path[])	May overflow 'resolve path'
sprintf(char *str, const char *format, ...)	May overflow 'str'

The old base pointer value is considered the boundary value. In this way a stack based buffer overflow can not over write the return address. This boundary is enforced by overloading the functions in the Table 1 while Libsafe guarantees that [16]:

- i. Correct programs will execute correctly, i.e., no false positives.
- ii. The frame pointers, and return addresses, can never be overwritten by an intercepted function.

**GCC Extension to Automatically Perform Array Bounds and Pointer Checking.** A gcc compiler patch [29] was developed by Jones and Kelly that performs complete array bounds checking for C programs. The representation of the pointers is not changed and the compiled programs are compatible with the gcc modules. They have derived a base pointer from each pointer expression, and check the attributes of that pointer to determine whether the expression is within bounds or not. The problem with this approach is its performance cost e.g. a pointer intensive program (3D matrix multiplication) experienced 30x slow down. This penalty is proportional to pointer usage which is quite common in privileged programs. The compiler is also not mature.

**Integer Range Analysis by Treating C Strings as an Abstract Type.** The approach has developed a system to statically detect buffer overflow in C [17]. It can be used effectively to find both known and unknown buffer overflow vulnerabilities in a version of 'sendmail'. The approach formulates the problem as an integer range analysis problem by treating C strings as an abstract type accessed through library functions and modeling pointers as integer ranges for allocated size and length. A consequence of modeling strings as an abstract data type is that buffer overflows involving non-character buffer cannot be detected.

**Source to Source Transformations.** It is a system that detects unsafe string operations in C programs [13]. This system performs a source-to-source transformation that instruments a program with additional variables that describe string attributes and contains assert statements that check for unsafe string operations. The instrumented program is then analyzed statically using integer analysis to determine possible assertion failures. This approach can handle many complex properties such as overlapping pointers. However, in the worst case the number of variables in the instrumented program is quadratic in terms of number of variables in the original program. To date, it has only been used on small example programs.

**Lexical Analysis, ITS4.** It is a lexical analysis tool that searches for security problem using a database of potentially dangerous constructs [10]. Lexical analysis techniques are fast and simple but their power is very limited since they do not take into account the syntax or semantics of the program.

### 3.4 Miscellaneous Approaches

The stack overflow protection schemes presented above are not the only ones. Other researchers have come up with some different ideas which are briefly described in the following discussion.

**Randomization Techniques.** These techniques rely on the randomization of memory layout and the application of cryptographic techniques using some randomly generated secrets.

*Memory Address Obfuscation.* This technique involves the randomization of memory layout, which actually does not make attacks impossible but it makes them highly improbable. This memory obfuscation vastly increases the security. The randomization can be done for [6]:

- i. *The base address of the memory regions*
- ii. *Permuting the order of variables and functions*
- iii. *The introduction of random length gaps*

*Randomization of Executable Code.* The main idea here is randomization of the code that is executed in a process [3,4]. The core of this scheme is to randomize the code in executable by encrypting it. This encryption is taking the XOR of executable code with a random 32-bit key. Similarly when the instruction is fetched, it is decrypted using 32-bit key to get the original code.

**Non Executable User Stack.** One of the methods to mitigate the buffer overflow attack is to make the stack non-executable. This type of defense method allows both code injection and change of Return Address thus allowing the change in control flow of the program to malicious code but it does not allow its execution completely [1], yet any attempt to transfer control to executable code placed in stack will result in system crash. Kernel patches are available for both Linux and Solaris [30,31] that make the stack segment of the programs address space non-executable. These approaches offer zero performance penalty and that programs are saved from re-compilation.

**Conversion of Stack Allocated Arrays into Pointers to Arrays, Gemini.** This tool transforms the stack allocated arrays of C to heap in order to prevent the return address from corruption as heap does not contain return addresses [2]. The real beauty of this approach is that it prevents all stack overflow exploits whereas the drawback is that it does not prevent the heap overflow. Moreover, it prohibits the attacker from injecting executable instructions such as shell instructions.

### 4. Qualitative Analysis

A qualitative comparison of different techniques based on various parameters is presented in Table 2. The columns

depict solutions based on taxonomy while parameters on the basis of which these solutions are evaluated are arranged in rows. These solutions employ both the static and dynamic approaches.

The code pointer protection and access violation probing involve the protection of return address, function pointers, base pointers and de-allocated memory locations. These schemes are usually platform dependent. All the static approaches here entail both the increase and the change in source code and due to these they can not be applied on the legacy systems. On the other hand, dynamic approaches do not need change in source code but may require recompilation; therefore it would be difficult to use them where source code for the legacy systems is not available. Based on the performance code pointer protection schemes can be used either as debugger or as normal security application whereas access violation probing schemes can be employed for debugging purposes only. Among all the techniques in these two categories only Cyclone and CCured are secured against all the attack targets.

The schemes addressing bounds checking are mostly static and usually platform independent. Although majority of these approaches have high performance penalty but Libsafe which is a dynamic solution offers very low performance degradation. Also Libsafe is secured against all the attack targets. Unlike static approaches in this category Libsafe is applicable to legacy systems upon recompilation.

Concluding the above analysis, it is found that Libsafe offers the least performance penalty than any other scheme. Therefore it can be used both for the development of security sensitive applications and legacy systems.

### 5. Proposed Solution

For last 15 years great amount of work for preventing stack smashing attack is being done. In most of the cases the solutions are extension of previous research e.g. using canary to protect return address. In the following discussion we propose some useful modifications to the existing solutions that would be helpful in boosting the security against the stack smashing attacks.

In order to protect function pointers as done by Stack Shield, one can slightly modify the PCB (Process Control Block) to hold the code segment boundary in the Process Identification Information block instead of storing the code segment boundary some where in the data segment. This is because code segment boundary for a particular process never changes that's why a process can hold this information upon its creation. Unlike Stack Shield where an attacker can overflow some pointer to overwrite the code segment boundary, this will provide a better solution because it is not vulnerable to malicious over write. This solution will require a slight change in the PCB data structure and the corresponding operating system routines.

The XORing of canary with the return address and then

**Table 2. Qualitative analysis of protection schemes with respect to different parameters**

Parameters	Solutions	Code Pointers Protection	Access Violation Probing	Bound Checking	Miscellaneous		
					Randomization Technique	Non-executable User Stack	Gemini
Require Changes in	Source Code	static approaches require changes; most of the dynamic approaches do not	static approaches require changes; dynamic approaches do not	static approaches require changes; dynamic approaches do not	No	No	Yes
	Compiler	Yes	static approaches require; dynamic approaches do not	No, except Libsafe	No	No	Yes
	Operating System	No	No	No	Yes	requires kernel patches	No
	Microprocessor	No	No, except Valgrind	No	No	No	No
	The Need for Source Code Recompilation	static approaches need; most of the dynamic approaches do not need	static approaches need; most of the dynamic approaches do not need	Yes	No	No	Yes
Platform Independence	Usually no	usually no	Yes	Yes	No	No	Yes
Code Size	increases in static approaches; remains same for dynamic approaches	increases in static approaches; remains same in dynamic approaches	increases	increases	remains same	remains same	increases
Impact on Performance	usually moderate degradation	high degradation	moderate degradation	moderate degradation	moderate degradation	low Degradation	low Degradation
Applicability on Legacy Systems	static approaches are not applicable; dynamic approaches are applicable	static approaches are not applicable; dynamic approaches are applicable	static approaches are not applicable; dynamic approaches are applicable	static approaches are not applicable; dynamic approaches are applicable	applicable	applicable for dynamic approaches	not applicable
Practical Applications	depending upon performance cost, usage varies from security application to debugging tool	it is used as debugging tool	it can be used as debugging tool as well as security application	it can be used as debugging tool as well as security application	it can be used as debugging tool as well as security application	it can be used as debugging tool	it is used as security application

storing the return address, canary and the XORed values by StackGuard will limit the number of times one can recursively call a particular function. Because at each function call a canary is generated and is stored some where in the data segment, which has limited size. In the case of excessive recursive or subsequent calls to functions a time will reach when this storage will start overflowing which means there must be some mechanism to handle this situation to avoid abnormal behavior. In our opinion the simplest solution is similar to what we have mentioned above i.e., assign the canary to process and not to functions. This canary can be stored in PCB and will not limit the number of recursive or subsequent calls; see Figure 5. Hence, whenever a function is called by the process we can XOR the return address with the canary and store the return address and the XORed value in the stack. This solution entails a slight change in compiler, PCB structure and the apposite operating system routines.




Figure 5. Modified process control block

Another interesting solution would be to put a memory mirror that contains a bit, initially reset, next to each memory location as depicted in Figure 6. This will allow the operating system to deny the access to those memory locations which contain return address and old base pointer. This mechanism is very simple, while preserving the return address and old base pointer to memory locations, upon the invocation of a function call, operating system will also set the pertinent memory bits. This will deny any malicious access to overwrite these memory locations. Upon returning from the function, the operating system will reset these memory bits to release them for later use. This technique will provide better protection than MemGuard[21], because there is no limit for recursive or subsequent function calls.




Figure 6. proposed memory architecture

## 6. Conclusion

One of the most common ways of breaching systems securities is to exploit buffer overflow vulnerabilities. Attackers have excessively exploited these vulnerabilities and caused various threats. According to a study carried out at University of California at Berkley, buffer overflows alone stand for about 60% of the vulnerabilities reported by CERT

in year 2002. Realizing the importance of severe damages caused by this security threat, researchers have done a great deal of study in order to prevent or detect it. We have studied and analyzed some of the well know approaches. The buffer overflow attacks can be generally classified into two types as stack and heap smashing attacks. In this study we have focused on stack smashing attacks and further we have categorized them in two basic classes, namely, Static and Dynamic. Moreover we have classified them according to protection mechanism. Also a detailed qualitative comparison of different techniques is presented. This will help security researchers and designers to choose among the techniques based upon the desired parameters and system specific criteria. We have also proposed some recommendations that, if incorporated, will strengthen the security and performance of some of the existing solutions. The performance analysis of proposed solutions in comparison to approaches given in the literature could be the future work.

## Acknowledgement

The authors acknowledge the support of King Fahd University of Petroleum and Minerals in the development of this work.

## References

- [1] Mark Shaneck, "An Overview of Buffer Overflow Vulnerabilities and Internet Worm", *CSCI 8980*, December 10, 2003.
- [2] Christopher Dahn and Spiros Mancoridis, "Using Program Transformation to Secure C Programs Against Buffer Overflows", *10th IEEE Working Conference on Reverse Engineering, Victoria, B.C., Canada*, November 13 - 17, 2003.
- [3] G. Kc, A. Keromytis, V. Prevelakis, "Countering Code-Injection Attacks with Instruction-Set Randomization," *In Proceedings of the 10th ACM Conference on Computer and Communication Security*, October 2003.
- [4] E. Barrantes, D. Ackley, S. Forrest, T. Palmer, D. Stefanovic, D. Zovi, "Randomized Instruction Set Emulation to Disrupt Binary Code Injection Attacks," *In Proceedings of the 10th ACM Conference on Computer and Communication Security*, October 2003.
- [5] Suan Yong and Susan Horwitz, "Protecting C Programs from Attacks via Invalid Pointer Dereferences", *The ESEC/FSE 2003 technical program on Safety and Security*, September 5, 2003.
- [6] S. Bhatkar, D. Du Warney, R. Sekar, "Address Obfuscation: an Efficient Approach to Combat a Broad Range of Memory Error Exploits", *In Proceedings of the 12th USENIX Security Symposium*, August 2003, pp 105-120.
- [7] Andrew Wright, "On Sapphire and Type-Safe Languages", *Proceedings, Texas Workshop on Security of Information Systems, Texas A&M University*, April 2 2003.



- [8] John Wilander and Mariam Kamkar, "A Comparison of Publicly Available Tools For Dynamic Buffer Overflow Prevention", *The 10th Annual Network and Distributed System Security Symposium, San Diego, California*, 6-7 February 2003.
- [9] B. Chess, "Improving Computer Security using Extended Static Checking", *In Proceedings of the 2002 IEEE Symposium on Security and Privacy*, May 2002
- [10] John Vigea, J. T. Bloch, Tada Woshi Kohno and Gary McGraw, "ITS4: Static Vulnerability Scanner for C/C++ Code", *Annual Computer Security Applications Conference*, December 2001.
- [11] Christof Fetzer and Zhen Xiao, "Detecting Heap Smashing Attacks Through Fault Containment Wrappers", *In Proceedings of the 20th IEEE Symposium on Reliable Distributed Systems*, October 2001.
- [12] Timothy K. Tsai, and Navjot Singh, "Libsafe: Protecting Critical Elements of Stacks", White Paper <http://www.research.avayalabs.com/project/libsafe>, August 29, 2001.
- [13] N. Dor, M. Rodeh, and S. Sagiv, "Cleanness Checking of String Manipulations in C programs via Integer Analysis", *In 8th International Symposium on Static Analysis (SAS)*, July 2001, pp 194–212.
- [14] D. Wagner and D. Dean, "Intrusion Detection via Static Analysis", *In Proceedings of the IEEE Symposium on Security and Privacy*, May 2001, pp 156-169.
- [15] T. cker Chiueh and F.H. Hsu, "RAD: A Compile Time Solution to Buffer Overflow Attacks", *In proceedings of the 21st International Conference on Distributed Computing Systems (ICDCS), Phoenix, Arizona, USA*, April 2001.
- [16] Arash Bratloo, Navjot Singh, and Timoyhy Tsai, "Transparent Runtime Defense Against Stack Smashing Attacks", *In proceedings of 2000 USENIX Annual Technical Conference, San Diego, California, USA*, June 18-23, 2000.
- [17] D. Wagner, J. Foster, E. Brewer, and A. Aiken, "A First Step towards Automated Detection of Buffer Overrun Vulnerabilities", *In Network and Distributed System Security Symposium, San Diego, CA*, February 2000.
- [18] C. Cowan, P. Wagle, C. Pu, S. Beattie, and J. Walpole, "Buffer Overflows: Attacks and Defenses for the Vulnerability of the Decade", *In Proceedings of the DARPA Information Survivability Conference and Expo (DISCEX), Hilton Head, South Carolina*, January 2000, pp 119–129.
- [19] A. Simon, A. King, "Analyzing String Buffers in C", *In International Conference on Algebraic Methodology and Software Technology*, 2000.
- [20] Crispin Cowan, Steve Beattie, Ryan Finnin Day, Calton Pu, Perry Wagle, and Erik Walthinsen, "Protecting Systems from Stack Smashing Attacks with StackGuard", *Presented at the Linux Expo* (<http://www.linuxexpo.org/~crispin/>), Raleigh, NC, May 18-22, 1999.
- [21] C. Cowan, C. Pu, D. Maier, H. Hinton, J. Walpole, P. Bakke, S. Beattie, A. Grier, P. Wagle, and Q. Zhang, "Stackguard: Automatic Adaptive Detection and Prevention of Buffer-Overflow Attacks", *In Proceedings of the 7th USENIX Security Symposium*, January 1998.
- [22] Jones, R. W. M. and Kelly, P. H. J. "Backwards-compatible bounds checking for arrays and pointers in C programs", *In Third International Workshop on Automated Debugging, Linkoping, Sweden*, May 1997.
- [23] Geogr C. Necula, "Proof-Carrying Code", Annual Symposium on Principles of Programming Languages, *Proceedings of the 24th ACM SIGPLAN-SIGACT, Paris, France*, 1997, pp 106 – 119.
- [24] "Aleph One, Smashing the Stack for Fun and Profit", *Phrack Magazine Volume 7, Issue 49*, November 1996.
- [25] *CERT Advisories*, <http://www.cert.org/advisories>, Feb 2003.
- [26] P. –A. Fayolle and V. Glaume, "A Buffer Overflow Study, Attacks & Defences", <http://www.enseirb.fr/~glaume/indexen.html>, March 2002.
- [27] <http://www.securiteam.com/securitynews/5TP0N0A2AW.html>, Aug 2000.
- [28] H. Etoh, "GCC Extension for Protecting Applications from Stack-Smashing Attacks", <http://www.tri.ibm.com/projects.security/ssp>, June 2000.
- [29] Richard Jones and Paul Kell, "Bounds Checking for C", <http://www-ala.doc.ic.ac.uk/phjk/BoundsChecking.html>, July 1995.
- [30] "Solar Designer, Non-Executable User Stack", <http://www.openwall.com/linux>
- [31] Casper Dik, "Non-Executable Stack for Solaris", Posting to [comp.security.unix](http://comp.security.unix), <http://x10.dejanews.com/getdoc>
- [32] J.Seward and N. Nethercote, Valgrind, "An Open Source Memory Debugger for X86-Linux", <http://developer.kde.org/~swardj/>
- [33] The *US-CERT Vulnerabilities Database*, <http://www.kb.cert.org/vuls/>