

Buffer overflow

In computer security and programming, a **buffer overflow**, or **buffer overrun**, is an anomaly where a program, while writing data to a buffer, overruns the buffer's boundary and overwrites adjacent memory locations.

Buffers are areas of memory set aside to hold data, often while moving it from one section of a program to another, or between programs. Buffer overflows can often be triggered by malformed inputs; if one assumes all inputs will be smaller than a certain size and the buffer is created to be that size, then an anomalous transaction that produces more data could cause it to write past the end of the buffer. If this overwrites adjacent data or executable code, this may result in erratic program behavior, including memory access errors, incorrect results, and crashes.

Exploiting the behavior of a buffer overflow is a well-known security exploit. On many systems, the memory layout of a program, or the system as a whole, is well defined. By sending in data designed to cause a buffer overflow, it is possible to write into areas known to hold executable code, and replace it with malicious code. Buffers are widespread in operating system (OS) code, so it is possible to make attacks that perform privilege escalation and gain unlimited access to the computer's resources. The famed Morris worm in 1988 used this as one of its attack techniques.

Programming languages commonly associated with buffer overflows include C and C++, which provide no built-in protection against accessing or overwriting data in any part of memory and do not automatically check that data written to an array (the built-in buffer type) is within the boundaries of that array. Bounds checking can prevent buffer overflows, but requires additional code and processing time. Modern operating systems use a variety of techniques to combat malicious buffer overflows, notably by randomizing the layout of memory, or deliberately leaving space between buffers and looking for actions that write into those areas ("canaries").

Contents

Technical description

- Example

Exploitation

- Stack-based exploitation
- Heap-based exploitation
- Barriers to exploitation
- Practicalities of exploitation
 - NOP sled technique
 - The jump to address stored in a register technique

Protective countermeasures

- Choice of programming language
- Use of safe libraries
- Buffer overflow protection
- Pointer protection
- Executable space protection
- Address space layout randomization
- Deep packet inspection
- Testing

History

- See also
- References
- External links

Technical description

A buffer overflow occurs when data written to a buffer also corrupts data values in memory addresses adjacent to the destination buffer due to insufficient bounds checking. This can occur when copying data from one buffer to another without first checking that the data fits within the destination buffer.

Example

In the following example expressed in C, a program has two variables which are adjacent in memory: an 8-byte-long string buffer, A, and a two-byte big-endian integer, B.

```
char A[8] = "";
unsigned short B = 1979;
```

Initially, A contains nothing but zero bytes, and B contains the number 1979.

variable name	A								B	
value	[null string]								1979	
hex value	00	00	00	00	00	00	00	00	07	BB

Now, the program attempts to store the null-terminated string "excessive" with ASCII encoding in the A buffer.

```
strcpy(A, "excessive");
```

"excessive" is 9 characters long and encodes to 10 bytes including the null terminator, but A can take only 8 bytes. By failing to check the length of the string, it also overwrites the value of B:

variable name	A									B	
value	'e'	'x'	'c'	'e'	's'	's'	'i'	'v'		25856	
hex	65	78	63	65	73	73	69	76	00	65	00

B's value has now been inadvertently replaced by a number formed from part of the character string. In this example "e" followed by a zero byte would become 25856.

Writing data past the end of allocated memory can sometimes be detected by the operating system to generate a segmentation fault error that terminates the process.

To prevent the buffer overflow from happening in this example, the call to strcpy could be replaced with strncpy, which takes the maximum capacity of A as an additional parameter and ensures that no more than this amount of data is written to A:

```
strncpy(A, "excessive", sizeof(A));
```

Note that the above code is not free from problems either; while a buffer overrun has been prevented this time, the `strcpy` library function does not null-terminate the destination buffer if the source string's length is greater than or equal to the size of the buffer (the third argument passed to the function), therefore A is, in this case, not null-terminated and cannot be treated as a valid C-style string.

Exploitation

The techniques to exploit a buffer overflow vulnerability vary by architecture, by operating system and by memory region. For example, exploitation on the heap (used for dynamically allocated memory), differs markedly from exploitation on the call stack.

Stack-based exploitation

A technically inclined user may exploit stack-based buffer overflows to manipulate the program to their advantage in one of several ways:

- By overwriting a local variable that is located near the vulnerable buffer on the stack, in order to change the behavior of the program
- By overwriting the return address in a stack frame. Once the function returns, execution will resume at the return address as specified by the attacker - usually a user-input filled buffer
- By overwriting a function pointer^[1] or exception handler, which is subsequently executed
- By overwriting a local variable (or pointer) of a different stack frame, which will be used by the function which owns that frame later.^[2]

If the address of the user-supplied data used to effect the stack buffer overflow is unpredictable, exploiting a stack buffer overflow to cause remote code execution becomes much more difficult. One technique that can be used to exploit such a buffer overflow is called "trampolining". In that technique, an attacker will find a pointer to the vulnerable stack buffer, and compute the location of their shellcode relative to that pointer. Then, they will use the overwrite to jump to an instruction already in memory which will make a second jump, this time relative to the pointer; that second jump will branch execution into the shellcode. Suitable instructions are often present in large code. The Metasploit Project, for example, maintains a database of suitable opcodes, though it lists only those found in the Windows operating system.^[3]

Heap-based exploitation

A buffer overflow occurring in the heap data area is referred to as a heap overflow and is exploitable in a manner different from that of stack-based overflows. Memory on the heap is dynamically allocated by the application at runtime and typically contains program data. Exploitation is performed by corrupting this data in specific ways to cause the application to overwrite internal structures such as linked list pointers. The canonical heap overflow technique overwrites dynamic memory allocation linkage (such as malloc meta data) and uses the resulting pointer exchange to overwrite a program function pointer.

Microsoft's GDI+ vulnerability in handling JPEGs is an example of the danger a heap overflow can present.^[4]

Barriers to exploitation

Manipulation of the buffer, which occurs before it is read or executed, may lead to the failure of an exploitation attempt. These manipulations can mitigate the threat of exploitation, but may not make it impossible. Manipulations could include conversion to upper or lower case, removal of metacharacters and filtering out of non-alphanumeric

strings. However, techniques exist to bypass these filters and manipulations; alphanumeric code, polymorphic code, self-modifying code and return-to-libc attacks. The same methods can be used to avoid detection by intrusion detection systems. In some cases, including where code is converted into unicode,^[5] the threat of the vulnerability has been misrepresented by the disclosers as only Denial of Service when in fact the remote execution of arbitrary code is possible.

Practicalities of exploitation

In real-world exploits there are a variety of challenges which need to be overcome for exploits to operate reliably. These factors include null bytes in addresses, variability in the location of shellcode, differences between environments and various counter-measures in operation.

NOP sled technique

A NOP-sled is the oldest and most widely known technique for successfully exploiting a stack buffer overflow.^[6] It solves the problem of finding the exact address of the buffer by effectively increasing the size of the target area. To do this, much larger sections of the stack are corrupted with the no-op machine instruction. At the end of the attacker-supplied data, after the no-op instructions, the attacker places an instruction to perform a relative jump to the top of the buffer where the shellcode is located. This collection of no-ops is referred to as the "NOP-sled" because if the return address is overwritten with any address within the no-op region of the buffer, the execution will "slide" down the no-ops until it is redirected to the actual malicious code by the jump at the end. This technique requires the attacker to guess where on the stack the NOP-sled is instead of the comparatively small shellcode.^[7]

Because of the popularity of this technique, many vendors of intrusion prevention systems will search for this pattern of no-op machine instructions in an attempt to detect shellcode in use. It is important to note that a NOP-sled does not necessarily contain only traditional no-op machine instructions; any instruction that does not corrupt the machine state to a point where the shellcode will not run can be used in place of the hardware assisted no-op. As a result, it has become common practice for exploit writers to compose the no-op sled with randomly chosen instructions which will have no real effect on the shellcode execution.^[8]

While this method greatly improves the chances that an attack will be successful, it is not without problems. Exploits using this technique still must rely on some amount of luck that they will guess offsets on the stack that are within the NOP-sled region.^[9] An incorrect guess will usually result in the target program crashing and could alert the system administrator to the attacker's activities. Another problem is that the NOP-sled requires a much larger amount of memory in which to hold a NOP-sled large enough to be of any use. This can be a problem when the allocated size of the affected buffer is too small and the current depth of the stack is shallow (i.e. there is not much space from the end of the current stack frame to the start of the stack). Despite its problems, the NOP-sled is often the only method that will work for a given platform, environment, or situation; as such it is still an important technique.

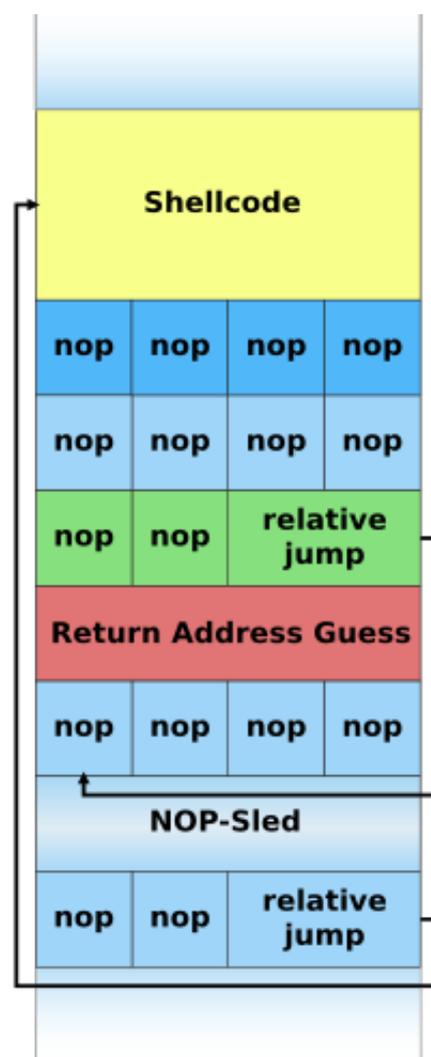
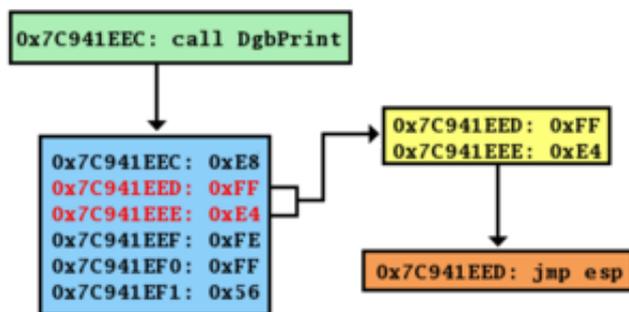


Illustration of a NOP-sled payload on the stack.

The jump to address stored in a register technique

The "jump to register" technique allows for reliable exploitation of stack buffer overflows without the need for extra room for a NOP-sled and without having to guess stack offsets. The strategy is to overwrite the return pointer with something that will cause the program to jump to a known pointer stored within a register which points to the controlled buffer and thus the shellcode. For example, if register A contains a pointer to the start of a buffer then any jump or call taking that register as an operand can be used to gain control of the flow of execution.^[10]



An instruction from ntdll.dll to call the `DbgPrint()` routine contains the i386 machine opcode for `jmp esp`.

the `jmp esp` instruction, and will then jump to the top of the stack and execute the attacker's code.^[12]

When this technique is possible the severity of the vulnerability increases considerably. This is because exploitation will work reliably enough to automate an attack with a virtual guarantee of success when it is run. For this reason, this is the technique most commonly used in Internet worms that exploit stack buffer overflow vulnerabilities.^[13]

This method also allows shellcode to be placed after the overwritten return address on the Windows platform. Since executables are mostly based at address `0x00400000` and x86 is a Little Endian architecture, the last byte of the return address must be a null, which terminates the buffer copy and nothing is written beyond that. This limits the size of the shellcode to the size of the buffer, which may be overly restrictive. DLLs are located in high memory (above `0x01000000`) and so have addresses containing no null bytes, so this method can remove null bytes (or other disallowed characters) from the overwritten return address. Used in this way, the method is often referred to as "DLL Trampolining".

Protective countermeasures

Various techniques have been used to detect or prevent buffer overflows, with various tradeoffs. The most reliable way to avoid or prevent buffer overflows is to use automatic protection at the language level. This sort of protection, however, cannot be applied to legacy code, and often technical, business, or cultural constraints call for a vulnerable language. The following sections describe the choices and implementations available.

Choice of programming language

Assembly and C/C++ are popular programming languages that are vulnerable to buffer overflow, in part because they allow direct access to memory and are not strongly typed.^[14] C provides no built-in protection against accessing or overwriting data in any part of memory; more specifically, it does not check that data written to a buffer is within the boundaries of that buffer. The standard C++ libraries provide many ways of safely buffering data, and C++'s Standard Template Library (STL) provides containers that can optionally perform bounds checking if the programmer explicitly

calls for checks while accessing data. For example, a vector's member function `at()` performs a bounds check and throws an `out_of_range` exception if the bounds check fails.^[15] However, C++ behaves just like C if the bounds check is not explicitly called. Techniques to avoid buffer overflows also exist for C.

Languages that are strongly typed and don't allow direct memory access, such as COBOL, Java, Python, and others, prevent buffer overflow from occurring in most cases.^[14] Many programming languages other than C/C++ provide runtime checking and in some cases even compile-time checking which might send a warning or raise an exception when C or C++ would overwrite data and continue to execute further instructions until erroneous results are obtained which might or might not cause the program to crash. Examples of such languages include Ada, Eiffel, Lisp, Modula-2, Smalltalk, OCaml and such C-derivatives as Cyclone, Rust and D. The Java and .NET Framework bytecode environments also require bounds checking on all arrays. Nearly every interpreted language will protect against buffer overflows, signaling a well-defined error condition. Often where a language provides enough type information to do bounds checking an option is provided to enable or disable it. Static code analysis can remove many dynamic bound and type checks, but poor implementations and awkward cases can significantly decrease performance. Software engineers must carefully consider the tradeoffs of safety versus performance costs when deciding which language and compiler setting to use.

Use of safe libraries

The problem of buffer overflows is common in the C and C++ languages because they expose low level representational details of buffers as containers for data types. Buffer overflows must thus be avoided by maintaining a high degree of correctness in code which performs buffer management. It has also long been recommended to avoid standard library functions which are not bounds checked, such as `gets`, `scanf` and `strcpy`. The Morris worm exploited a `gets` call in `fingerd`.^[16]

Well-written and tested abstract data type libraries which centralize and automatically perform buffer management, including bounds checking, can reduce the occurrence and impact of buffer overflows. The two main building-block data types in these languages in which buffer overflows commonly occur are strings and arrays; thus, libraries preventing buffer overflows in these data types can provide the vast majority of the necessary coverage. Still, failure to use these safe libraries correctly can result in buffer overflows and other vulnerabilities; and naturally, any bug in the library itself is a potential vulnerability. "Safe" library implementations include "The Better String Library",^[17] `Vstr`^[18] and `Erwin`.^[19] The OpenBSD operating system's C library provides the `strncpy` and `strlcat` functions, but these are more limited than full safe library implementations.

In September 2007, Technical Report 24731, prepared by the C standards committee, was published;^[20] it specifies a set of functions which are based on the standard C library's string and I/O functions, with additional buffer-size parameters. However, the efficacy of these functions for the purpose of reducing buffer overflows is disputable; it requires programmer intervention on a per function call basis that is equivalent to intervention that could make the analogous older standard library functions buffer overflow safe.^[21]

Buffer overflow protection

Buffer overflow protection is used to detect the most common buffer overflows by checking that the stack has not been altered when a function returns. If it has been altered, the program exits with a segmentation fault. Three such systems are `Libsafe`,^[22] and the StackGuard^[23] and ProPolice^[24] `gcc` patches.

Microsoft's implementation of Data Execution Prevention (DEP) mode explicitly protects the pointer to the Structured Exception Handler (SEH) from being overwritten.^[25]

Stronger stack protection is possible by splitting the stack in two: one for data and one for function returns. This split is present in the Forth language, though it was not a security-based design decision. Regardless, this is not a complete solution to buffer overflows, as sensitive data other than the return address may still be overwritten.

Pointer protection

Buffer overflows work by manipulating pointers (including stored addresses). PointGuard was proposed as a compiler-extension to prevent attackers from being able to reliably manipulate pointers and addresses.^[26] The approach works by having the compiler add code to automatically XOR-encode pointers before and after they are used. Because the attacker (theoretically) does not know what value will be used to encode/decode the pointer, he cannot predict what it will point to if he overwrites it with a new value. PointGuard was never released, but Microsoft implemented a similar approach beginning in Windows XP SP2 and Windows Server 2003 SP1.^[27] Rather than implement pointer protection as an automatic feature, Microsoft added an API routine that can be called at the discretion of the programmer. This allows for better performance (because it is not used all of the time), but places the burden on the programmer to know when it is necessary.

Because XOR is linear, an attacker may be able to manipulate an encoded pointer by overwriting only the lower bytes of an address. This can allow an attack to succeed if the attacker is able to attempt the exploit multiple times or is able to complete an attack by causing a pointer to point to one of several locations (such as any location within a NOP sled).^[28] Microsoft added a random rotation to their encoding scheme to address this weakness to partial overwrites.^[29]

Executable space protection

Executable space protection is an approach to buffer overflow protection which prevents execution of code on the stack or the heap. An attacker may use buffer overflows to insert arbitrary code into the memory of a program, but with executable space protection, any attempt to execute that code will cause an exception.

Some CPUs support a feature called NX ("No eXecute") or XD ("eXecute Disabled") bit, which in conjunction with software, can be used to mark pages of data (such as those containing the stack and the heap) as readable and writable but not executable.

Some Unix operating systems (e.g. OpenBSD, macOS) ship with executable space protection (e.g. W^X). Some optional packages include:

- PaX^[30]
- Exec Shield^[31]
- Openwall^[32]

Newer variants of Microsoft Windows also support executable space protection, called Data Execution Prevention.^[33] Proprietary add-ons include:

- BufferShield^[34]
- StackDefender^[35]

Executable space protection does not generally protect against return-to-libc attacks, or any other attack which does not rely on the execution of the attackers code. However, on 64-bit systems using ASLR, as described below, executable space protection makes it far more difficult to execute such attacks.

Address space layout randomization

Address space layout randomization (ASLR) is a computer security feature which involves arranging the positions of key data areas, usually including the base of the executable and position of libraries, heap, and stack, randomly in a process' address space.

Randomization of the virtual memory addresses at which functions and variables can be found can make exploitation of a buffer overflow more difficult, but not impossible. It also forces the attacker to tailor the exploitation attempt to the individual system, which foils the attempts of internet worms.^[36] A similar but less effective method is to rebase processes and libraries in the virtual address space.

Deep packet inspection

The use of deep packet inspection (DPI) can detect, at the network perimeter, very basic remote attempts to exploit buffer overflows by use of attack signatures and heuristics. These are able to block packets which have the signature of a known attack, or if a long series of No-Operation instructions (known as a NOP-sled) is detected, these were once used when the location of the exploit's payload is slightly variable.

Packet scanning is not an effective method since it can only prevent known attacks and there are many ways that a NOP-sled can be encoded. Shellcode used by attackers can be made alphanumeric, metamorphic, or self-modifying to evade detection by heuristic packet scanners and intrusion detection systems.

Testing

Checking for buffer overflows and patching the bugs that cause them naturally helps prevent buffer overflows. One common automated technique for discovering them is fuzzing.^[37] Edge case testing can also uncover buffer overflows, as can static analysis.^[38] Once a potential buffer overflow is detected, it must be patched; this makes the testing approach useful for software that is in development, but less useful for legacy software that is no longer maintained or supported.

History

Buffer overflows were understood and partially publicly documented as early as 1972, when the Computer Security Technology Planning Study laid out the technique: "The code performing this function does not check the source and destination addresses properly, permitting portions of the monitor to be overlaid by the user. This can be used to inject code into the monitor that will permit the user to seize control of the machine." (Page 61)^[39] Today, the monitor would be referred to as the kernel.

The earliest documented hostile exploitation of a buffer overflow was in 1988. It was one of several exploits used by the Morris worm to propagate itself over the Internet. The program exploited was a service on Unix called finger.^[40] Later, in 1995, Thomas Lopatic independently rediscovered the buffer overflow and published his findings on the Bugtraq security mailing list.^[41] A year later, in 1996, Elias Levy (also known as Aleph One) published in Phrack magazine the paper "Smashing the Stack for Fun and Profit",^[42] a step-by-step introduction to exploiting stack-based buffer overflow vulnerabilities.

Since then, at least two major internet worms have exploited buffer overflows to compromise a large number of systems. In 2001, the Code Red worm exploited a buffer overflow in Microsoft's Internet Information Services (IIS) 5.0.^[43] and in 2003 the SQL Slammer worm compromised machines running Microsoft SQL Server 2000.^[44]

In 2003, buffer overflows present in licensed Xbox games have been exploited to allow unlicensed software, including homebrew games, to run on the console without the need for hardware modifications, known as modchips.^[45] The PS2 Independence Exploit also used a buffer overflow to achieve the same for the PlayStation 2. The Twilight hack accomplished the same with the Wii, using a buffer overflow in The Legend of Zelda: Twilight Princess.

See also

- Billion laughs
- Buffer over-read
- Computer security
- End-of-file
- Heap overflow
- Ping of death
- Port scanner
- Return-to-libc attack
- Security-focused operating system
- Self-modifying code
- Shellcode
- Stack buffer overflow
- Uncontrolled format string

References

1. "CORE-2007-0219: OpenBSD's IPv6 mbufs remote kernel buffer overflow" (<http://www.securityfocus.com/archive/1/462728/30/150/threaded>). Retrieved 2007-05-15.
2. "Modern Overflow Targets" (<http://packetstormsecurity.com/files/download/121751/ModernOverflowTargets.pdf>) (PDF). Retrieved 2013-07-05.
3. "The Metasploit Opcode Database" (<https://web.archive.org/web/20070512195939/http://www.metasploit.com/users/opcode/msfopcode.cgi>). Archived from the original (<http://metasploit.com/users/opcode/msfopcode.cgi>) on 12 May 2007. Retrieved 2007-05-15.
4. "Microsoft Technet Security Bulletin MS04-028" (<http://www.microsoft.com/technet/security/bulletin/MS04-028.mspx>). Retrieved 2007-05-15.
5. "Creating Arbitrary Shellcode In Unicode Expanded Strings" (<https://web.archive.org/web/20060105041036/http://www.net-security.org/dl/articles/unicodebo.pdf>) (PDF). Archived from the original (<http://www.net-security.org/dl/articles/unicodebo.pdf>) (PDF) on 2006-01-05. Retrieved 2007-05-15.
6. Vangelis (2004-12-08). "Stack-based Overflow Exploit: Introduction to Classical and Advanced Overflow Technique" (<https://web.archive.org/web/20070818115455/http://www.neworder.box.sk/newsread.php?newsid=12476>). Wowhacker via Neworder. Archived from the original (<http://www.neworder.box.sk/newsread.php?newsid=12476>) (text) on August 18, 2007.
7. Balaban, Murat. "Buffer Overflows Demystified" (<http://www.enderunix.org/docs/en/bof-eng.txt>) (text). Enderunix.org.
8. Akritidis, P.; Evangelos P. Markatos; M. Polychronakis; Kostas D. Anagnostakis (2005). "STRIDE: Polymorphic Sled Detection through Instruction Sequence Analysis." (<http://dcs.ics.forth.gr/Activities/papers/stride-IFIP-SEC05.pdf>) (PDF). *Proceedings of the 20th IFIP International Information Security Conference (IFIP/SEC 2005)*. IFIP International Information Security Conference. Retrieved 2012-03-04.
9. Klein, Christian (September 2004). "Buffer Overflow" (<https://web.archive.org/web/20070928011639/http://c0re.23.nu/~chris/presentations/overflow2005.pdf>) (PDF). Archived from the original (<http://c0re.23.nu/~chris/presentations/overflow2005.pdf>) (PDF) on 2007-09-28.
10. Shah, Saumil (2006). "Writing Metasploit Plugins: from vulnerability to exploit" (<http://conference.hitb.org/hitbsecconf2006kl/materials/DAY%201%20-%20Saumil%20Shah%20-%20Writing%20Metasploit%20Plugins.pdf>) (PDF).

Hack In The Box. Kuala Lumpur. Retrieved 2012-03-04.

11. *Intel 64 and IA-32 Architectures Software Developer's Manual Volume 2A: Instruction Set Reference, A-M* (<https://web.archive.org/web/20071129123212/http://developer.intel.com/design/processor/manuals/253666.pdf>) (PDF). Intel Corporation. May 2007. pp. 3–508. Archived from the original (<http://developer.intel.com/design/processor/manuals/253666.pdf>) (PDF) on 2007-11-29.
12. Alvarez, Sergio (2004-09-05). "Win32 Stack BufferOverflow Real Life Vuln-Dev Process" (http://packetstormsecurity.org/papers/Win2000/Intro_to_Win32_Exploits.pdf) (PDF). IT Security Consulting. Retrieved 2012-03-04.
13. Ukai, Yuji; Soeder, Derek; Permech, Ryan (2004). "Environment Dependencies in Windows Exploitation" (<https://www.blackhat.com/presentations/bh-asia-04/bh-jp-04-ukai-eng.ppt>). *BlackHat Japan*. Japan: eEye Digital Security. Retrieved 2012-03-04.
14. https://www.owasp.org/index.php/Buffer_Overflows Buffer Overflows article on OWASP
15. "vector::at - C++ Reference" (<http://www.cplusplus.com/reference/vector/vector/at/>). Cplusplus.com. Retrieved 2014-03-27.
16. <http://wiretap.area.com/Gopher/Library/Techdoc/Virus/inetvir.823>
17. "The Better String Library" (<http://bstring.sf.net/>).
18. "The Vstr Homepage" (<http://www.and.org/vstr/>). Retrieved 2007-05-15.
19. "The Erwin Homepage" (<http://www.theiling.de/projects/erwin.html>). Retrieved 2007-05-15.
20. International Organization for Standardization (2007). "Information technology — Programming languages, their environments and system software interfaces — Extensions to the C library — Part 1: Bounds-checking interfaces" (<https://www.iso.org/obp/ui/#iso:std:iso-iec:tr:24731:-1:ed-2:v1:en:sec:4>). *ISO Online Browsing Platform*.
21. "CERT Secure Coding Initiative" (<https://www.securecoding.cert.org/confluence/x/QwY>). Retrieved 2007-07-30.
22. "Libsafe at FSF.org" (<http://directory.fsf.org/libsafe.html>). Retrieved 2007-05-20.
23. "StackGuard: Automatic Adaptive Detection and Prevention of Buffer-Overflow Attacks by Cowan et al" (https://www.usenix.org/publications/library/proceedings/sec98/full_papers/cowan/cowan.pdf) (PDF). Retrieved 2007-05-20.
24. "ProPolice at X.ORG" (<https://web.archive.org/web/20070212032750/http://wiki.x.org/wiki/ProPolice>). Archived from the original (<http://wiki.x.org/wiki/ProPolice>) on 12 February 2007. Retrieved 2007-05-20.
25. "Bypassing Windows Hardware-enforced Data Execution Prevention" (<http://www.uninformed.org/?v=2&a=4&t=txt>). Retrieved 2007-05-20.
26. PointGuard: Protecting Pointers From Buffer Overflow Vulnerabilities (http://www.usenix.org/events/sec03/tech/full_papers/cowan/cowan_html/index.html)
27. Protecting Against Pointer Subterfuge (Kinda!) (http://blogs.msdn.com/michael_howard/archive/2006/01/30/520200.aspx)
28. Defeating Compiler-Level Buffer Overflow Protection (<http://www.usenix.org/publications/login/2005-06/pdfs/alexander0506.pdf>)
29. Protecting against Pointer Subterfuge (Redux) (http://blogs.msdn.com/michael_howard/archive/2006/08/16/702707.aspx)
30. "PaX: Homepage of the PaX team" (<http://pax.grsecurity.net>). Retrieved 2007-06-03.
31. "KernelTrap.Org" (<https://archive.is/20120529183334/http://kerneltrap.org/node/644>). Archived from the original (<http://kerneltrap.org/node/644>) on 2012-05-29. Retrieved 2007-06-03.
32. "Openwall Linux kernel patch 2.4.34-ow1" (<https://web.archive.org/web/20120219111512/http://linux.softpedia.com/get/System/Operating-Systems/Kernels/Openwall-Linux-kernel-patch-16454.shtml>). Archived from the original (<http://linux.softpedia.com/get/System/Operating-Systems/Kernels/Openwall-Linux-kernel-patch-16454.shtml>) on 2012-02-19. Retrieved 2007-06-03.
33. "Microsoft Technet: Data Execution Prevention" (<http://technet2.microsoft.com/WindowsServer/en/Library/b0de1052-4101-44c3-a294-4da1bd1ef2271033.mspx?mfr=true>).
34. "BufferShield: Prevention of Buffer Overflow Exploitation for Windows" (http://www.sys-manage.com/english/products/products_BufferShield.html). Retrieved 2007-06-03.
35. "NGSec Stack Defender" (<https://web.archive.org/web/20070513235539/http://www.ngsec.com/ngproducts/stack>)

- defender/). Archived from the original (<http://www.ngsec.com/ngproducts/stackdefender/>) on 2007-05-13. Retrieved 2007-06-03.
36. "PaX at GRSecurity.net" (<http://pax.grsecurity.net/docs/aslr.txt>). Retrieved 2007-06-03.
 37. "The Exploitant - Security info and tutorials" (<http://raykoid666.wordpress.com>). Retrieved 2009-11-29.
 38. Larochelle, David; Evans, David (13 August 2001). "Statically Detecting Likely Buffer Overflow Vulnerabilities" (https://www.usenix.org/legacy/events/sec01/full_papers/larochelle/larochelle_html/). *USENIX Security Symposium*. **32**.
 39. "Computer Security Technology Planning Study" (<http://csrc.nist.gov/publications/history/ande72.pdf>) (PDF). Retrieved 2007-11-02.
 40. "'A Tour of The Worm" by Donn Seeley, University of Utah" (<https://web.archive.org/web/20070520233435/http://world.std.com/~frani/worm.html>). Archived from the original (<http://world.std.com/~frani/worm.html>) on 2007-05-20. Retrieved 2007-06-03.
 41. "Bugtraq security mailing list archive" (https://web.archive.org/web/20070901222723/http://www.security-express.com/archives/bugtraq/1995_1/0403.html). Archived from the original (http://www.security-express.com/archives/bugtraq/1995_1/0403.html) on 2007-09-01. Retrieved 2007-06-03.
 42. "'Smashing the Stack for Fun and Profit" by Aleph One" (<http://www.phrack.com/issues.html?issue=49&id=14>). Retrieved 2012-09-05.
 43. "eEye Digital Security" (<http://research.eeye.com/html/advisories/published/AL20010717.html>). Retrieved 2007-06-03.
 44. "Microsoft Technet Security Bulletin MS02-039" (<http://www.microsoft.com/technet/security/bulletin/ms02-039.msp>). Retrieved 2007-06-03.
 45. "Hacker breaks Xbox protection without mod-chip" (https://web.archive.org/web/20070927210513/http://www.gamesindustry.biz/content_page.php?aid=1461). Archived from the original (http://www.gamesindustry.biz/content_page.php?aid=1461) on 2007-09-27. Retrieved 2007-06-03.

External links

- "Discovering and exploiting a remote buffer overflow vulnerability in an FTP server" (<http://raykoid666.wordpress.com/2009/11/28/remote-buffer-overflow-from-vulnerability-to-exploit-part-1/>) by Raykoid666
- "Smashing the Stack for Fun and Profit" (<http://phrack.org/issues/49/14.html#article>) by Aleph One
- An Overview and Example of the Buffer-Overflow Exploit. pps. 16-21. (http://iac.dtic.mil/iatac/download/Vol7_No4.pdf)
- CERT Secure Coding Standards (<https://www.securecoding.cert.org/>)
- CERT Secure Coding Initiative (<http://www.cert.org/secure-coding>)
- Secure Coding in C and C++ (<http://www.cert.org/books/secure-coding>)
- SANS: inside the buffer overflow attack (http://www.sans.org/reading_room/whitepapers/securecode/386.php)
- "Advances in adjacent memory overflows" (<https://web.archive.org/web/20130126024851/http://www.awarenetwork.org/etc/alpha/?x=5>) by Nomenclura
- A Comparison of Buffer Overflow Prevention Implementations and Weaknesses (<https://www.blackhat.com/presentations/bh-usa-04/bh-us-04-silberman/bh-us-04-silberman-paper.pdf>)
- More Security Whitepapers about Buffer Overflows (<https://web.archive.org/web/20090817230359/http://doc.bughunter.net/buffer-overflow/>)
- Chapter 12: Writing Exploits III (https://web.archive.org/web/20071129123212/http://www.syngress.com/book_catalog/327_SSPC/sample.pdf) from *Sockets, Shellcode, Porting & Coding: Reverse Engineering Exploits and Tool Coding for Security Professionals* by James C. Foster (ISBN 1-59749-005-9). Detailed explanation of how to use Metasploit to develop a buffer overflow exploit from scratch.
- Computer Security Technology Planning Study (<http://csrc.nist.gov/publications/history/ande72.pdf>), James P. Anderson, ESD-TR-73-51, ESD/AFSC, Hanscom AFB, Bedford, MA 01731 (October 1972) [NTIS AD-758 206]
- "Buffer Overflows: Anatomy of an Exploit" (<https://www.exploit-db.com/docs/18346.pdf>) by Nevermore
- Secure Programming with GCC and Glibc (<https://cansecwest.com/csw08/csw08-holtmann.pdf>) (2008), by Marcel Holtmann

Retrieved from "https://en.wikipedia.org/w/index.php?title=Buffer_overflow&oldid=813105473"

This page was last edited on 1 December 2017, at 20:14.

Text is available under the [Creative Commons Attribution-ShareAlike License](#); additional terms may apply. By using this site, you agree to the [Terms of Use](#) and [Privacy Policy](#). Wikipedia® is a registered trademark of the [Wikimedia Foundation, Inc.](#), a non-profit organization.