


# Writing Secure Privileged Programs

*linux.conf.au 2010*

**Wellington, New Zealand; 22 Jan 2010**



Michael Kerrisk, *jambit* GmbH

<http://www.kernel.org/doc/man-pages/>  
<http://www.man7.org/>

## Overview

## Introduction

- Avoiding creation of security holes by:
  - Using safe C coding practices
  - Understanding and using Linux/UNIX OS-level concepts to improve security
- Excludes topics like:
  - Authentication techniques
  - Cryptography & Random-number generation
  - Linux Security Modules
  - Virtualization
  - Cross-site scripting
  - SQL injection
  - etc.

## Topics

- Process credentials
- Gaining privileges
- Why worry?
- Guidelines for secure programming
- Useful reading

# Process Credentials

## Process credentials

- Every process has credentials:
  - User IDs (UIDs)
  - Group IDs (GIDs)
  - Supplementary group IDs
- Credentials determine:
  - Ownership of process
  - File access permissions
  - Privileges to perform other operations
- Set for login shell on startup
- Inherited by child of *fork()*
- Some credentials can change during *exec()*

## User IDs

- Real UID (RUID)
  - Who owns process
- Effective UID (EUID)
  - File access permissions
  - Privileges for other operations
    - EUID == 0 vs EUID != 0
- Saved set-user-ID (SUID)
  - (Described later)
- (For unprivileged program, all IDs have same value)

## Group IDs

- Real GID (RGID)
  - Which group owns process
- Effective GID (EGID)
  - File access permissions
- Saved set-group-ID (SGID)
  - (Described later)
- (For unprivileged program, all IDs have same value)
- Supplementary group IDs
  - Additional groups used to check file access permissions
  - Derived from `/etc/group` file

## File access permissions

- Determined by EUID, EGID, & supplementary GIDs
- Each file has permissions: *owner group other*  

R	W	X	R	W	X	R	W	X
---	---	---	---	---	---	---	---	---
- EUID matches file *owner*?  
→ granted permissions for file *owner*
- EGID or any *supp. GID* matches file *group*?  
→ granted permissions for file *group*
- Otherwise:  
→ granted permissions for *other*
- *User perms can be < group perms; & group < other*

# Gaining Privileges

## Gaining privileges

- Two ways to obtain privileges of user or group:
  - Run program in process belonging to user or group
  - Execute set-user-ID or set-group-ID program
    - Usual way to give privilege to nonprivileged users
    - Abbreviations: `setuid` program & `setgid` program

## Setuid programs

- Like other executables, except that set-user-ID permission bit is enabled:

```
$ su
Password:
# ls -l prog
-rwxr-xr-x  1 root  root  302585 Jan 22 10:05 prog
# chmod u+s prog
# ls -l prog
-rwsr-xr-x  1 root  root  302585 Jan 22 10:05 prog
```
- When executed, process EUID is set to owner of file
- Setuid-root program == setuid program owned by *root*

## Setgid programs

- Analogous to setuid program

```
$ su
Password:
# ls -l prog
-rwxr-xr-x 1 root  root  302585 Jan 22 10:05 prog
# chmod g+s prog
# ls -l prog
-rwxr-sr-x 1 root  root  302585 Jan 22 10:05 prog
```

## Credential changes during exec()

- What happens to process IDs during `exec()`?
- RUID: unchanged
- EUID: is set-user-ID permission bit enabled for executable file?
  - Yes → EUID set to file owner
  - No → EUID unchanged
- SUID: copied from EUID (*after preceding step*)
- (Analogous changes to GIDs for setgid program)
- Supplementary GIDs: unchanged

## Credential changes during exec() (cont.)

- Example:
  - Login as "mtk"

```
$ whoami  
mtk
```

*RUID == EUID == SUID == <mtk>*

```
$ ls -l prog  
-rwsr-xr-x 1 root root 302585 Jan 22 10:05 prog  
$ ./prog # Create new process that execs "prog"
```

*RUID == <mtk>; EUID == SUID == <root (0)>*

# Why Worry?



## What's the problem?

- Privileged program grants rights of another user or group
- If subverted, security is compromised
- Especially dangerous for EUID of 0 and for N/W services
- Many ways to create bugs that lead to subversion
- Some guidelines...

## Guidelines

Click to add title

## Guideline: Avoid writing setuid-root programs

### **Setuid-root? Just say no!**

- If there's a way to avoid setuid-root, use it
  - (Maybe you don't really need privilege at all)
- Limits potential damage if program is compromised
- Two useful techniques:
  - Privilege separation
  - Use an ID other than *root*

## Avoiding setuid-root: Privilege separation

- Isolate functionality requiring *root* privileges into a separate process running as *root*
- Request operations via IPC, or info passed across `exec()`
- **Make inputs and functionality of program as limited as possible!**
  - Less flexibility == fewer chances to compromise

## Example of privilege separation

- Example:
  - `grantpt(3)` library function
  - Forks child process that execs a setuid-*root* program, `pt_chown`
  - Changes ownership and permissions of pseudo-tty slave corresponding to master specified via open file descriptor

## Avoiding setuid-root: Use an ID other than root

- Suppose we have a program that updates a file that shouldn't be updated by normal users
- *Bad*: make file writable only by *root*, and use setuid-root program
- *Better*: create new, dedicated *group* ID, make file writable by that group, and use setgid program
  - Damage if compromised is greatly limited
  - Examples:
    - *wall(1)*, *write(1)* (*tty* group)
    - many games (*games* group)

## Always check return status

- Almost every system call and library function returns a status indicating success or failure
- Someday, the call you thought could never fail, **will**
  - Some system calls can fail even for *root*
    - e.g.,:
      - open a file for writing on a read-only file system;
      - fork()* fails if process table is full

Click to add title

## Guideline: Check return statuses

### Always check return status

- Almost every system call and library function returns a status indicating success or failure
- Someday, the call you thought could never fail, **will**
  - Some system calls can fail even for *root*
    - e.g.,:
      - open a file for writing on a read-only file system;
      - fork()* fails if process table is full
- **Always check the return status**

Click to add title

Guideline: If the unexpected occurs, fail safely

## Handling unexpected errors

- What if an “unexpected” error occurs?
- Trying to “fix” things usually requires assumptions that may not be valid (i.e., safe) in all cases
- **When the unexpected occurs, log a message and give up:**
  - Locally executed program: terminate
  - Network server: drop client request
- *Fail safely*

Click to add title

## Guideline: Operate with least privilege

### Operate with least privilege

- A setuid program doesn't need privileged EUID all the time
- If compromise occurs while program is unprivileged, damage is limited
- → **Operate with “least privilege”**:
  - Drop privilege (immediately!) at start of execution
  - Raise privilege temporarily when needed
  - Drop privilege permanently when it will never again be required
  - (Techniques rely on saved set-\*IDs)

## Saved set-user-ID

- When setuid program is executed:
  - EUID of process is made same as file owner
  - EUID is copied to SUID
- e.g., after exec of setuid-root program by *mtk*:
  - RUID: mtk (unprivileged ID)
  - EUID: 0
  - SUID: 0 (privileged ID)
- (SGID is analogous for setgid programs)

## Changing privileges in a setuid program

- Raising/dropping privileges == switching EUID between:
  - RUID (unprivileged)
  - SUID (privileged)
- Permanently dropping privileges == setting EUID and SUID to RUID
- Changes via system calls



## Changing process credentials

System call	IDs changed	Notes
<i>setuid(u)</i> <i>setgid(g)</i>	effective	If EUID == 0, <i>real</i> and <i>saved</i> are also changed to same value; semantics vary across systems
<i>seteuid(e)</i> <i>setegid(e)</i>	effective	
<i>setreuid(r, e)</i> <i>setregid(r, e)</i>	real, saved	Also changes saved UID, if real ID is changed
<i>setresuid(r, e, s)</i> <i>setresgid(r, e, s)</i>	real, effective, saved	Nonstandard
<i>setgroups(n, list)</i>	supp. GIDs	

- General rules:
  - EUID == 0: arbitrary changes to IDs
  - EUID != 0: change an ID to be same as another of the IDs
  - For some calls, -1 argument value means “no change”

## Dropping and raising privilege

```
uid_t orig_euid;

orig_euid = geteuid();           /* Save privileged EUID */
                                /* (same as value in SUID) */

if (seteuid(getuid()) == -1)    /* Drop privileges */
    errExit("seteuid");       /* (Switch to RUID) */

/* Do unprivileged work */

if (seteuid(orig_euid) == -1)  /* Raise privileges */
    errExit("seteuid");       /* (Switch back to SUID) */

/* Do privileged work */

if (seteuid(getuid()) == -1)   /* Drop privileges */
    errExit("seteuid");

/* Do unprivileged work */
```

## Dropping privileges permanently

- Dropping UID 0 in *setuid-root* program:

```
if (setuid(getuid()) == -1) /* Sets RUID, EUID, SUID */  
    errExit("setuid");
```

- **But!** Doesn't work if EUID != 0
  - i.e., *setuid-non-root* program; or *setuid-root* program with privilege currently dropped
  - *setuid()* changes only EUID
    - (And call returns success...)

## Dropping privileges permanently (*cont.*)

- If EUID != 0, either:
  - *setuid-root* program:
    - ***seteuid(orig\_euid)***); /\* Regain privileged EUID \*/
    - setuid(getuid)***); /\* Drop all privileged UIDs \*/
  - *setuid-non-root* program:
    - ***setreuid(getuid(), getuid())***); /\* Changes all UIDs \*/
  - In any (Linux) program:
    - ***setresuid(getuid(), getuid(), getuid())***

## Problems with changing credentials

- Too many system calls – confusing!
- Some calls aren't available on some systems
  - `setres[ug]id()` (but: nicest interface!)
- Differing semantics when `EUID==0` and `EUID!=0`
  - `set[ug]id()`
- Differing semantics across systems
  - `set[ug]id()`
- Kernel bugs or unusual scenarios mean calls may unexpectedly fail (perhaps without error!)
  - <http://userweb.kernel.org/~morgan/sendmail-capabilities-war-story.html>
- *Easy to get it wrong!*

## Safely changing process credentials

- Read the documentation!
  - <http://www.kernel.org/doc/man-pages/>
  - `credentials(7)`
- Check return status from `set*id()` calls
- Verify that IDs have actually changed
  - `getres[ug]id()` [Linux]; or `/proc` [other systems]
- Write/employ a portable package to do the above
  - See [Tsafrir et al., 2008]

Click to add title

Guideline: Be careful when executing another program

## Executing programs

- Drop privileges permanently before `exec()`
  - (See earlier techniques)
  - `setuid(getuid())` is sufficient
    - **successful** `exec()` copies EUID to SUID
- Never `exec()` a shell with EUID 0:
  - Shells are too complex to avoid all security loopholes
  - Likewise other interpreters, such as `awk`...
  - Avoid `system()`, `popen()`
  - Avoid `setuid` scripts
    - (Not even permitted on Linux)

## Executing programs (*cont.*)

- Close unneeded file descriptors
  - Privileged programs can open files that are not be accessible to others
  - Leaving descriptors open across `exec()` is a security leak
    - Close-on-exec flag may be useful (see `fcntl(2)`)

Click to add title

Guideline: Avoid exposing sensitive information

## Avoid exposing sensitive information


- Sometimes, info in memory can land on disk
- Erase sensitive info from memory as soon as no longer needed
- Lock pages into memory (*mlock()*) if having data written to disk (swap area) is a concern
- Disable core dumps
  - Set `RLIMIT_CORE` limit to 0; see *setrlimit(2)*

Click to add title

Guideline: Be careful of signals

## Be aware of signals

- Users can send arbitrary signals to a program at **any point time**



- If necessary, block, ignore, or catch signals to prevent security problems
- Keep signal handler design simple
  - e.g., just set global flag checked by main program
  - Minimizes bugs that may occur because of races

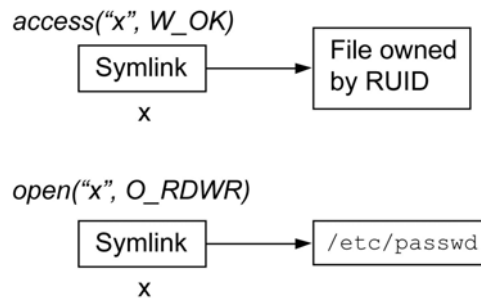
## Race conditions

- Time-of-check, time-of-use (TOCTOU) race condition:
  - 1) Program checks details of run-time environment
  - 2) The user manages to change details of RTE
    - e.g., [change file permissions](#); [change target of symlink](#)
  - 3) Program continues, based on false assumptions

→ [Security breach](#)

## Race conditions (cont.)

- Classic example:
  - **Check**: Setuid-root program uses `access(2)` to determine that a file is writable by real user (RUID)
  - **Use**: Program opens file for writing
- Between check and use, things change:



## Race conditions (cont.)

- Chances of exploit can be greatly increased by:
  - Using stop/continue signals to widen window between TOC and TOU
  - Repeated execution + scripts to deliver signals rapidly
- Write code so as to avoid TOCTOU races; e.g.:
  - Program (temporarily) drops privileged ID
  - Program opens file for writing and checks return status



Click to add title

## Guideline: Be careful when performing file operations and file I/O

### File operations and file I/O

- If creating a file, ensure that it is never vulnerable (even briefly) to malicious manipulation
  - Ensure file is never publicly writable (*umask(2)* may help)
  - Ownership of new files taken from EUID
    - **Don't** create file and *chown(2)*
    - **Do** ensure EUID is correctly set before creating file
- Perform checks on file descriptors, not pathnames
  - For example:
    - **Don't** use *stat()* and then *open()* (TOCTOU race!)
    - **Do** use *open()* and then *fstat()*

## File operations and file I/O (*cont.*)

- To ensure that *you* are creator of a file, use `open() O_EXCL`
- Avoid creating files in publicly writable directories (e.g., `/tmp`)
  - Can be manipulated or removed by other users
  - If you must do so, use random filename (`mkstemp(3)`)

Click to add title

Guideline: Don't trust user inputs!

## Don't trust user inputs!

- Never trust input from users:
  - Interactive input
  - Command-line arguments
  - User-supplied files
  - Email
  - IPC channels
  - CGI inputs
  - Network packets
  - etc.

## Don't trust user inputs! (*cont.*)

- Validate and sanitize all inputs
  - Are numbers inside acceptable limits?
  - Are strings of acceptable length?
  - Are characters in string valid?
  - etc.

## Don't trust user inputs! (cont.)

- Classic example:

```
char cmd[CMDLEN], pat[PATLEN];
fgets(pat, PATLEN, stdin);
snprintf(cmd, CMDLEN, "ls %s", pat);
system(cmd);
```

- Suppose user supplies following input to *snprintf()*:

```
x; rm /etc/passwd
==> system("ls x; rm /etc/passwd");
```

- In this example:
  - Check that characters are in set `[[a-z_A-Z?*~]]`; or
  - Escape shell metacharacters with `\`

Click to add title

Guideline: Don't trust  
environment variables

## Don't trust environment variables!

- Values should be checked (like other user input)
- EVs affect operation of many programs & libraries
  - e.g., PATH affects shell (and *system()*, *popen()*, *execvp()*, *execip()*)
    - Manipulation may cause unexpected program to be executed
    - Ensure PATH value is safe, or (better) use absolute pathnames
- Safest approach:
  - Erase entire environment
  - Restore selected variables with known-safe values

Click to add title

Guideline: Don't trust the Run-time Environment

## Don't trust the run time environment!

- Do you expect standard input, output, and error to be open?
  - What about library functions?
- What happens if you run out of disk space?
- What happens if resource limits are set very low?
  - CPU time, file size, stack size, number of open files
- What if *fork()* fails because there are too many processes on system?

## Don't trust the run time environment! (cont.)

- Do you check status of *malloc()* calls?
- Do you expect signal mask to be empty?
  - *sigprocmask(2)*, *signal(7)*
- Are you assuming that initial umask is okay?
  - What if umask is 0700?
- **Attackers may try to subvert program by forcing unexpected or low-resource scenarios**

Click to add title

Guideline: Beware of buffer overruns!!

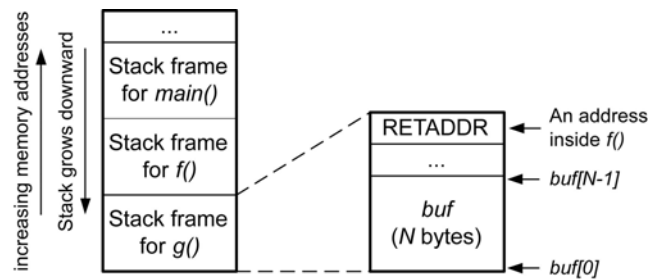
## Buffer overruns

- A.K.A. stack smashing
- Subvert program by making it run injected code
- Extremely common flaw
  - See CERT ([www.cert.org](http://www.cert.org)), Bugtraq ([www.securityfocus.com](http://www.securityfocus.com)), LWN.net
- Examples
  - 2001 Code Red worm (MS IIS web server)
  - 2003 SQL Slammer (MS SQL Server)
  - 1988 Morris worm (*fingerd (gets())*, UNIX)

## Buffer overruns (cont.)


- Idea is to exploit this type of situation:

```
void g() {  
    char buf[N];  
    /* Code allowing user input for buf */  
}  
  
void f() { g(); }  
  
main() { f(); }
```



## Buffer overruns (cont.)

- Suppose we can:
  - provide user input into *buf*, and
  - write beyond end of *buf* (no bounds checking)
- Exploit is achieved by this transformation:





## Buffer overruns (*cont.*)

- Modern OSes/hardware use techniques to make stack smashing harder:
  - Address-space randomization
  - Nonexecutable stacks
- But, can be circumvented with more effort
  - “Return to libc” (see [Anley et al., 2007])
  - Repeated execution driven by scripts
  - See also [Erickson, 2008], [Aleph One, 1996]

## Avoiding buffer overruns

- **Always bounds check user input**
- Never use *gets(3)*!
  - Use *fgets(strbufp, len, stream)*
- Use *scanf()*, *sprintf()*, *strcpy()*, *strcat()* with caution
  - Guard use with boundary checking code, or
  - Use *snprintf()*, *strncpy()*, *strncat()*; but
    - Check for truncated result string
    - **NB** If result string is too long, *strncpy()* does not include NULL terminator!

## Format-string attacks

- e.g., `printf(argv[1]);`
- By including “%n” specifier in string, we can write arbitrary values at specific address
  - %n == write # of characters so far output to address given in arg. list
- Can achieve similar exploit to buffer overruns, but requires more work
- See [Anley et al., 2007]
- **Don't permit user input as part of format string!**

## Heap overflows

- Dynamic memory allocated and freed via *malloc* package
- All allocations from same area of memory (“the heap”)
- Similar concept to stack smashing: overrun a buffer, in order to write data into “sensitive” buffer elsewhere on heap
  - e.g., buffer might contain name of file to open for writing
- Avoid in same way as for stack buffer overruns

Click to add title

## Guideline: Be prepared for denial-of-service attacks

### Denial of (network) service attacks

- What happens if client or server doesn't reply to your message?
  - Use timeouts
- Be prepared for overload attacks
  - What happens if traffic is 100x expected?
  - (Network attacks may be distributed; source addresses may be spoofed)
- If traffic volume exceeds expectations:
  - Degrade gracefully: throttle load (drop *some* requests)
  - Log details of situation ([but throttle logging too!](#))

## Denial of service attacks (*cont.*)

- Can too many outstanding requests cause data structures (e.g., arrays) to overflow?
- Beware of algorithmic-complexity attacks
  - E.g., does right sequence of inputs turn your binary search tree or hash table into a linked list?
  - Use data structures that avoid these problems
  - [Crosby & Wallach, 2003]

Click to add title

Guideline: Confine the  
program / consider using  
capabilities

## Confine the program

- Use *chroot(2)* to restrict process to subset of file-system tree
  - But: *setuid-root* programs can break out of *chroot* jails
- Consider using capabilities

## Capabilities

- Divide all-or-nothing power of *root* into distinct units (34, as at Linux 2.6.33):
  - CAP\_AUDIT\_CONTROL, CAP\_AUDIT\_WRITE, CAP\_CHOWN, CAP\_DAC\_OVERRIDE, CAP\_DAC\_READ\_SEARCH, CAP\_FOWNER, CAP\_FSETID, CAP\_IPC\_LOCK, CAP\_IPC\_OWNER, CAP\_KILL, CAP\_LEASE, CAP\_LINUX\_IMMUTABLE, CAP\_MAC\_ADMIN, CAP\_MAC\_OVERRIDE, CAP\_MKNOD, CAP\_NET\_ADMIN, CAP\_NET\_BIND\_SERVICE, CAP\_NET\_BROADCAST, CAP\_NET\_RAW, CAP\_SETFCAP, CAP\_SETGID, CAP\_SETPCAP, CAP\_SETUID, CAP\_SYS\_ADMIN, CAP\_SYS\_BOOT, CAP\_SYS\_CHROOT, CAP\_SYS\_MODULE, CAP\_SYS\_NICE, CAP\_SYS\_PACCT, CAP\_SYS\_PTRACE, CAP\_SYS\_RAWIO, CAP\_SYS\_RESOURCE, CAP\_SYS\_TIME, CAP\_SYS\_TTY\_CONFIG

## Capabilities (*cont.*)

- Instead of having UID 0, process can have selected capabilities, without having other powers of superuser
  - Takes “operate with least privilege” to finer granularity
  - Can have privileged program that can't access files owned by *root*
- Linux-specific...

## Capabilities (*cont.*)

- Partial implementation since Linux 2.2
- File capabilities added in Linux 2.6.24
  - Capabilities can be associated with executable file
    - *setcap(8)* and *getcap(8)*
  - When file is executed, process gains capabilities (analogous to *setuid* program)
- Process has permitted and effective capability sets
  - Analogous to SUID and EUID in *setuid* programs
  - Use *libcap* API to raise/drop effective capabilities
- Further info:
  - *capabilities(7)*, [Hallyn, 2007], [Kerrisk, 2010]

## Summary

- Avoid writing *setuid-root* programs
- Check return status from every call
- Fail safely
- Operate with least privilege at all times
- Drop privileges permanently when no longer needed
- Drop privilege before execing another program
- Avoid exposing sensitive information
- Be aware of signals
- Avoid TOCTOU races
- Be careful with file operations and file I/O
- Don't trust: user inputs; environment variables; run time environment
- Beware of buffer overruns
- Be prepared for denial of service attacks
- Consider using capabilities

# Useful Reading

## Useful reading

*A very small sample of a very wide range of publications on security*

- Aleph One. 1996. Smashing the Stack for Fun and Profit  
<http://www.phrack.com/issues.html?issue=49&id=14#article>
- Anley, C., Heasman, J., Lindner, F., and Richarte, G. 2007. *The Shellcoder's Handbook: Discovering and Exploiting Security Holes*. Wiley.
- Bishop, M. Various papers at <http://nob.cs.ucdavis.edu/~bishop/secprog>.
- Bishop, M. 2003. *Computer Security: Art and Science*. Addison-Wesley.
- Bishop, M. 2005. *Introduction to Computer Security*. Addison-Wesley.
- Chen, H., Wagner, D., and Dean, D. 2002. "Setuid Demystified," *Proceedings of the 11th USENIX Security Symposium*.  
<http://www.cs.berkeley.edu/~daw/papers/setuid-usenix02.pdf>
- Crosby, S. A., and Wallach, D. S. 2003. "Denial of Service via Algorithmic Complexity Attacks," *Proceedings of the 12th USENIX Security Symposium*.  
[http://www.cs.rice.edu/~scrosby/hash/CrosbyWallach\\_UenixSec2003.pdf](http://www.cs.rice.edu/~scrosby/hash/CrosbyWallach_UenixSec2003.pdf)
- Drepper, U. 2009. "Defensive Programming for Red Hat Enterprise Linux"  
<http://people.redhat.com/drepper/defprogramming.pdf>

## Useful reading

- Erickson, J. M. 2008. *Hacking: The Art of Exploitation (2e)*. No Starch Press.
- Garfinkel, S., et al.. 2003. *Practical Unix and Internet Security (3e)*. O'Reilly.
- Hallyn, S. A. 2007. "POSIX file capabilities: Parceling the power of root."  
<http://www.ibm.com/developerworks/library/l-posixcap.html>
- Kerrisk, M., et al. *capabilities(7)* manual page
- Kerrisk, M. 2010. *The Linux Programming Interface*. No Starch Press.
- Peikari, C., and Chuvakin, A. 2004. *Security Warrior*. O'Reilly.
- Tsafrir, D., da Silva, D., and Wagner, D. "The Murky Issue of Changing Process Identity: Revising 'Setuid Demystified'," *login: The USENIX Magazine*, June 2008.  
<http://www.usenix.org/publications/login/2008-06/pdfs/tsafrir.pdf>
- Viega, J., and McGraw, G. 2002. *Building Secure Software*. Addison-Wesley.
- Wheeler, D., *Secure Programming for Linux and Unix HOWTO*  
<http://www.dwheeler.com/secure-programs/>



# Thanks!

[http://userweb.kernel.org/~mtk/papers/lca2010/  
writing\\_secure\\_privileged\\_programs.pdf](http://userweb.kernel.org/~mtk/papers/lca2010/writing_secure_privileged_programs.pdf)

Michael Kerrisk  
jambit GmbH

*The Linux Programming Interface*  
No Starch Press, 2010 (soon)  
<http://blog.man7.org/>