

CSci 5271
Introduction to Computer Security
Defensive programming and OS security
(combined lecture)

Stephen McCamant
University of Minnesota, Computer Science & Engineering

Outline

Bernstein's perspective (cont'd)
Secure use of the OS
Techniques for privilege separation
Announcements intermission
OS security: protection and isolation
OS security: authentication
Basics of access control

Eliminating bugs

- Enforce explicit data flow
- Simplify integer semantics
- Avoid parsing
- Generalize from errors to inputs

Eliminating code

- Identify common functions
- Automatically handle errors
- Reuse network tools
- Reuse access controls
- Reuse the filesystem

The "qmail security guarantee"

- \$500, later \$1000 offered for security bug
- Never paid out
- Issues proposed:
 - Memory exhaustion DoS
 - Overflow of signed integer indexes
- Defensiveness does not encourage more submissions

qmail today

- Originally had terms that prohibited modified redistribution
 - Now true public domain
- Latest release from Bernstein: 1998; netqmail: 2007
- Does not have large market share
- All MTAs, even Sendmail, are more secure now

Outline

Bernstein's perspective (cont'd)

Secure use of the OS

Techniques for privilege separation

Announcements intermission

OS security: protection and isolation

OS security: authentication

Basics of access control

Avoid special privileges

- Require users to have appropriate permissions
 - Rather than putting trust in programs
- Anti-pattern 1: `setuid/setgid` program
- Anti-pattern 2: privileged daemon
- But, sometimes unavoidable (e.g., email)

One slide on `setuid/setgid`

- Unix users and process have a user id number (UID) as well as one or more group IDs
- Normally, process has the IDs of the user who starts it
- A `setuid` program instead takes the UID of the program binary

Don't use shells or Tcl

- ...in security-sensitive applications
- String interpretation and re-parsing are very hard to do safely
- Eternal Unix code bug: path names with spaces

Prefer file descriptors

- Maintain references to files by keeping them open and using file descriptors, rather than by name
- References same contents despite file system changes
- Use `openat`, etc., variants to use FD instead of directory paths

Prefer absolute paths

- Use full paths (starting with `/`) for programs and files
- `$PATH` under local user control
- Initial working directory under local user control
 - But FD-like, so can be used in place of `openat` if missing

Prefer fully trusted paths

- Each directory component in a path must be write protected
- Read-only file in read-only directory can be changed if a parent directory is modified

Don't separate check from use

- Avoid pattern of e.g., `access` then `open`
- Instead, just handle failure of `open`
 - You have to do this anyway
- Multiple references allow races
 - And `access` also has a history of bugs

Be careful with temporary files

- Create files exclusively with tight permissions and never reopen them
 - See detailed recommendations in Wheeler
- Not quite good enough: reopen and check matching device and inode
 - Fails with sufficiently patient attack

Give up privileges

- Using appropriate combinations of `set*id` functions
 - Alas, details differ between Unix variants
- Best: give up permanently
- Second best: give up temporarily
- Detailed recommendations: *Setuid Demystified* (USENIX'02)

Whitelist environment variables

- Can change the behavior of called program in unexpected ways
- Decide which ones are necessary
 - As few as possible
- Save these, remove any others

Outline

Bernstein's perspective (cont'd)

Secure use of the OS

Techniques for privilege separation

Announcements intermission

OS security: protection and isolation

OS security: authentication

Basics of access control

Restricted languages

- ▣ Main application: code provided by untrusted parties
- ▣ Packet filters in the kernel
- ▣ JavaScript in web browsers
 - Also Java, Flash ActionScript, etc.

SFI

- ▣ Software-based Fault Isolation
- ▣ Instruction-level rewriting like (but predates) CFI
- ▣ Limit memory stores and sometimes loads
- ▣ Can't jump out except to designated points
- ▣ E.g., Google Native Client

Separate processes

- ▣ OS (and hardware) isolate one process from another
- ▣ Pay overhead for creation and communication
- ▣ System call interface allows many possibilities for mischief

System-call interposition

- ▣ Trusted process examines syscalls made by untrusted
- ▣ Implement via `ptrace` (like `strace`, `gdb`) or via kernel change
- ▣ Easy policy: deny

Interposition challenges

- ▣ Argument values can change in memory (TOCTTOU)
- ▣ OS objects can change (TOCTTOU)
- ▣ How to get canonical object identifiers?
- ▣ Interposer must accurately model kernel behavior
- ▣ Details: Garfinkel (NDSS'03)

Separate users

- ▣ Reuse OS facilities for access control
- ▣ Unit of trust: program or application
- ▣ Older example: gmail
- ▣ Newer example: Android
- ▣ Limitation: lots of things available to any user

chroot

- Unix system call to change root directory
- Restrict/virtualize file system access
- Only available to root
- Does not isolate other namespaces

OS-enabled containers

- One kernel, but virtualizes all namespaces
- FreeBSD jails, Linux LXC, Solaris zones, etc.
- Quite robust, but the full, fixed, kernel is in the TCB

(System) virtual machines

- Presents hardware-like interface to an untrusted kernel
- Strong isolation, full administrative complexity
- I/O interface looks like a network, etc.

Virtual machine designs

- (Type 1) hypervisor: 'superkernel' underneath VMs
- Hosted: regular OS underneath VMs
- Paravirtualization: modify kernels in VMs for ease of virtualization

Virtual machine technologies

- Hardware based: fastest, now common
- Partial translation: e.g., original VMware
- Full emulation: e.g. QEMU proper
 - Slowest, but can be a different CPU architecture

Modern example: Chrom(ium)

- Separates "browser kernel" from less-trusted "rendering engine"
 - Pragmatic, keeps high-risk components together
- Experimented with various Windows and Linux sandboxing techniques
- Blocked 70% of historic vulnerabilities, not all new ones
- <http://seclab.stanford.edu/websec/chromium/>

Outline

Bernstein's perspective (cont'd)

Secure use of the OS

Techniques for privilege separation

Announcements intermission

OS security: protection and isolation

OS security: authentication

Basics of access control

Rescheduled office hour

- My usual Monday 10-11am office hour will be rescheduled due to travel
- Substitute time 3-4pm Wednesday 2/27, usual place 4-225E Keller

Outline

Bernstein's perspective (cont'd)

Secure use of the OS

Techniques for privilege separation

Announcements intermission

OS security: protection and isolation

OS security: authentication

Basics of access control

OS security topics

- Resource protection
- Process isolation
- User authentication
- Access control

Protection and isolation

- Resource protection: prevent processes from accessing hardware
- Process isolation: prevent processes from interfering with each other
- Design: by default processes can do neither
- Must request access from operating system


Reference monitor

- Complete mediation: all accesses are checked
- Tamperproof: the monitor is itself protected from modification
- Small enough to be thoroughly verified

Hardware basis: memory protection

- Historic: segments
- Modern: paging and page protection
 - Memory divided into pages (e.g. 4k)
 - Every process has own virtual to physical page table
 - Pages also have R/W/X permissions

Linux 32-bit example



Hardware basis: supervisor bit

- Supervisor (kernel) mode: all instructions available
- User mode: no hardware or VM control instructions
- Only way to switch to kernel mode is specified entry point
- Also generalizes to multiple "rings"

Outline

- Bernstein's perspective (cont'd)
- Secure use of the OS
- Techniques for privilege separation
- Announcements intermission
- OS security: protection and isolation
- OS security: authentication
- Basics of access control

Authentication factors

- Something you know (password, PIN)
- Something you have (e.g., smart card)
- Something you are (biometrics)
- CAPTCHAs, time and location, ...
- Multi-factor authentication

Passwords: love to hate

- Many problems for users, sysadmins, researchers
- But familiar and near-zero cost of entry
- User-chosen passwords proliferate for low-stakes web site authentication

Password entropy

- Model password choice as probabilistic process
- If uniform, $\log_2 |S|$
- Controls difficulty of guessing attacks
- Hard to estimate for user-chosen passwords
 - Length is an imperfect proxy

Password hashing

- Idea: don't store password or equivalent information
- Password 'encryption' is a long-standing misnomer
 - E.g., Unix `crypt(3)`
- Presumably hard-to-invert function h
- Store only $h(p)$

Dictionary attacks

- Online: send guesses to server
- Offline: attacker can check guesses internally
- Specialized password lists more effective than literal dictionaries
 - Also generation algorithms ($s \rightarrow \$$, etc.)
- ~25% of passwords consistently vulnerable

Better password hashing

- Generate random salt s , store $(s, h(s, p))$
 - Block pre-computed tables and equality inferences
 - Salt must also have enough entropy
- Deliberately expensive hash function
 - AKA password-based key derivation function (PBKDF)
 - Requirement for time and/or space

Password usability

- User compliance can be a major challenge
 - Often caused by unrealistic demands
- Distributed random passwords usually unrealistic
- Password aging: not too frequently
- Never have a fixed default password in a product

Backup authentication

- Desire: unassisted recovery from forgotten password
- Fall back to other presumed-authentic channel
 - Email, cell phone
- Harder to forget (but less secret) shared information
 - Mother's maiden name, first pet's name
- Brittle: ask Sarah Palin or Mat Honan

Centralized authentication

- Enterprise-wide (e.g., UMN ID)
- Anderson: Microsoft Passport
- Today: Facebook Connect, Google ID
- May or may not be single-sign-on (SSO)


Biometric authentication

- Authenticate by a physical body attribute
- + Hard to lose
- Hard to reset
- Inherently statistical
- Variation among people

Example biometrics

- (Handwritten) signatures
- Fingerprints, hand geometry
- Face and voice recognition
- Iris codes

Error rates: ROC curve



Outline

Bernstein's perspective (cont'd)
Secure use of the OS
Techniques for privilege separation
Announcements intermission
OS security: protection and isolation
OS security: authentication
Basics of access control

Mechanism and policy

- Decision-making aspect of OS
- Should subject S (user or process) be allowed to access object (e.g., file) O ?
- Complex, since admin must specify what should happen

Access control matrix

	grades.txt	/dev/hda	/usr/bin/bcvi
Alice	r	rw	rx
Bob	rw	-	rx
Carol	r	-	rx

Slicing the matrix

- $O(nm)$ matrix impractical to store, much less administer
- Columns: access control list (ACL)
 - Convenient to store with object
 - E.g., Unix file permissions
- Rows: capabilities
 - Convenient to store by subject
 - E.g., Unix file descriptors

Groups/roles

- Simplify by factoring out commonality
- Before: users have permissions
- After: users have roles, roles have permissions
- Simple example: Unix groups
- Complex versions called role-based access control (RBAC)