

Operating Systems

20. Protection

Paul Krzyzanowski
Rutgers University
Spring 2015

April 13, 2015 © 2014-2015 Paul Krzyzanowski 1

Protection & Security

- **Security**
 - Prevention of unauthorized access to a system
 - Prevent malicious or accidental access
 - “access” may be:
 - user login, a process accessing things it shouldn’t, physical access
 - The access operations may be reading, destruction, or alteration
- **Protection**
 - The mechanism that provides and enforces controlled access of resources to processes
 - A protection mechanism *enforces* security policies

2

Principle of Least Privilege

At each abstraction layer, every element (user, process, function) should be able to access **only** the resources necessary to perform its task

- Even if an element is compromised, the scope of damage is limited
- Consider:
 - **Good:** You cannot kill another user’s process
 - **Good:** You cannot open the /etc/hosts file for writing
 - **Good:** Private member functions & local variables in functions limit scope
 - **Violation:** a compromised print daemon allows someone to add users
 - **Violation:** a process can write a file even though there is no need to
 - **Violation:** admin privileges set by default for any user account
- Least privilege is often difficult to define & enforce

3

Privilege Separation

Divide a program into multiple parts: high & low privilege components

- **Example on POSIX systems**
 - Each process has a *real* and *effective* user ID
 - Privileges are evaluated based on the effective user ID
 - Normally, `uid == euid`
 - An executable file may be tagged with a *setuid bit*
 - `chmod +sx filename`
 - When run: `uid = user’s ID`
`euid = file owner’s ID (without setuid, runs with user’s ID)`
 - Separating a program
 1. Run a *setuid* program
 2. Create a communication link to self (*pipe, socket, shared memory*)
 3. `fork`
 4. One of the processes will call `seteuid(getuid())` to lower its privilege

4

Security Goals

- **Authentication**
 - Ensure that users, machines, programs, and resources are properly identified
- **Integrity**
 - Verify that data has not been compromised: deleted, modified, added
- **Confidentiality**
 - Prevent unauthorized access to data
- **Availability**
 - Ensure that the system is accessible

5

The Operating System

The OS provides processes with access to resources

Resource	OS component
Processor(s)	Process scheduler
Memory	Memory Management + MMU
Peripheral devices	Device drivers & buffer cache
Logical persistent data	File systems
Communication networks	Sockets

- Resource access attempts go through the OS
- OS decides whether access should be granted
 - Rules that guide the decision = *policy*

6

Domains of protection

- Processes interact with **objects**
 - Objects include:
 - hardware (CPU, memory, I/O devices)
 - software: files, processes, semaphores, messages, signals
- A process should be allowed to access only objects that it is authorized to access
 - A process operates in a **protection domain**
 - Protection domain defines the objects the process may access and how it may access them

Modeling Protection: Access Matrix

Rows: domains
Columns: objects

Each entry represents an access right of a domain on an object

		objects		
		F ₀	F ₁	Printer
domains of protection	D ₀	read	read-write	print
	D ₁	read-write-execute	read	
	D ₂	read-execute		
	D ₃		read	print
	D ₄			print

Access Matrix: Domain Transfers

Switching from one domain to another is a configurable policy

A process in D₀ can switch to running in domain D₁

		objects							
		F ₀	F ₁	Printer	D ₀	D ₁	D ₂	D ₃	D ₄
domains of protection	D ₀	read	read-write	print	-	switch	switch		
	D ₁	read-write-execute	read			-			
	D ₂	read-execute				switch	-		
	D ₃		read	print					
	D ₄			print					

Access Matrix: Additional operations

Copy: allow delegation of rights

- Copy a specific access right on an object from one domain to another
- Rights may specify either a copy or a transfer of rights

A process executing in D₁ can give a read right on F₁ to another domain

		objects							
		F ₀	F ₁	Printer	D ₀	D ₁	D ₂	D ₃	D ₄
domains of protection	D ₀	read	read-write	print	-	switch	sv		
	D ₁	read-write-execute	read*						
	D ₂	read-execute				switch	-		
	D ₃		read	print					
	D ₄			print					

Access Matrix: Additional operations

Owner: allow new rights to be added or removed

- An object may be identified as being *owned* by the domain
- Owner can add and remove any right in any **column** of the object

A process executing in D₀ can give a read right on F₀ to domain D₃ and remove the execute right from D₁

		objects							
		F ₀	F ₁	Printer	D ₀	D ₁	D ₂	D ₃	D ₄
domains of protection	D ₀	read owner	read-	print	-	switch	sv		
	D ₁	read-write-execute	read*						
	D ₂	read-execute				switch	-		
	D ₃		read	print					
	D ₄			print					

Access Matrix: Additional operations

Control: change entries in a row

- If access(i, j) includes a **control** right, then a process executing in Domain i can change access rights for Domain j

A process executing in D₁ can modify any rights in domain D₄

		objects							
		F ₀	F ₁	Printer	D ₀	D ₁	D ₂	D ₃	D ₄
domains of protection	D ₀	read owner	read-write	print	-	switch	switch		
	D ₁	read-write-execute	read*						control
	D ₂	read-execute				switch	-		
	D ₃		read	print					
	D ₄			print					

Implementing an access matrix

- A single table is usually impractical
 - Big size: # domains (users) x # objects (files)
 - Objects may come and go frequently
- Access Control List**
 - Associate a column of the table with each object

Implementing an access matrix

- Access Control List**
 - Associate a column of the table with each object

	F ₀	F ₁	Printer	D ₀	D ₁	D ₂	D ₃	D ₄
D ₀	read owner	read- write						
D ₁	read- write- execute	read*			-			
D ₂	read- execute				switch	-		
D ₃		read	print					
D ₄			print					

Example: Limited ACLs in POSIX systems

Problem: an ACL takes up a varying amount of space (possibly a lot!)

- Won't fit in an inode

UNIX Compromise:

- A file defines access rights for three domains:
 - the **owner**, the **group**, and **everyone** else
- Permissions
 - Read, write, execute, directory search
 - Set user ID on execution
 - Set group ID on execution
- Default permissions set by the **umask** system call
- chown** system call changes the object's owner
- chmod** system call changes the object's permissions

Example: Full ACLs in POSIX systems

- What if we really want a full ACL?**
- Extended attributes:** stored outside of the inode
 - Hold an ACL
 - And other name:value attributes
- Enumerated list of permissions on users and groups
 - Operations on all objects:
 - delete, readattr, writeattr, readextattr, writeextattr, readsecurity, writesecurity, chown*
 - Operations on directories
 - list, search, add_file, add_subdirectory, delete_child*
 - Operations on files
 - read, write, append, execute*
 - Inheritance controls

Implementing an access matrix

Capability List

- Associate a row of the table with each domain

	F ₀	F ₁	Printer	D ₀	D ₁	D ₂	D ₃	D ₄
D ₀	read owner	read- write	print	-	switch	switch		
D ₁	read- write- execute	read*			-			
D ₂	read- execute				switch	-		
D ₃		read	print					
D ₄			print					

Capability Lists

- List of objects together with the operations allowed on the objects
- Each item in the list is a **capability**: the operations allowed on a specific object
- A process presents the capability along with a request
 - Possessing the capability means that access is allowed
- A process cannot modify its capability list

Access Control Models: MAC vs. DAC

- **DAC: Discretionary Access Control**
 - A subject (domain) can pass information onto **any** other subject
 - In some cases, access rights may be transferred
 - *Most systems use this*
- **MAC: Mandatory Access Control**
 - Policy is centrally controlled
 - Users cannot override the policy

19

Multi-level Access Control

- Typical MAC implementations use a **Multi-Level Secure (MLS)** access model
- **Bell-LaPadula model**
 - Identifies the ability to access and communicate data
 - Objects are classified into a hierarchy of sensitivity levels
 - Unclassified, Confidential, Secret, Top Secret
 - Each user is assigned a clearance
 - “**No read up; no write down**”
 - Cannot read from a higher clearance level
 - Cannot write to a lower clearance level
- Works well for government information
- Does not translate well to civilian life



20

The End

4/13/2015

© 2014-2015 Paul Krzyzanowski

21