Protection and OS security

Definitions

Computer security:

Techniques for computing in the presence of adversaries.

Cryptography:

Techniques for communicating in the presence of adversaries.

Definitions

Protection:

The subset of computer security which deals with controlling access to information stored locally.

Traditionally excludes issues arising from networks.

Traditionally associated with operating systems, and in particular with centralized mechanisms.
Note that the typical role of an OS is to manage multiplexed access to resources; protection is about access control, but not more general resource management issues.

Protection

Three categories of security goals:

- Confidentiality: preventing unauthorized release of information
- Integrity: preventing unauthorized modification of information
- Availability: preventing denial of service attacks

Protection is about providing all of these in the setting of a single machine.

Protection

The protection problem can be approximately factored into three sub-problems:

- Access control: enforcement of restrictions on who may access what
- Policy selection: a human-level statement of goals for the security subsystem
- Policy specification: human-computer interface

Access control

The access control problem is to decide whether a given request should be granted.

Can be conceptually divided into two steps

- Authentication: verifying the identity of an entity making a request.
- Authorization: verifying whether the policy allows the request.

Policy

Policy:

- A specification of the desired security properties of the system
- Roughly speaking, the goals which security mechanisms are intended to achieve
- “Human intent”
- Thus, policy cannot be automated; it is a fundamentally human-centered decision about “what ought to be”
- A good rule of thumb for formulating computer security policy is that it should roughly correspond to real-world policies
  - For example, at a bank, two officers may be required to both authorize the opening of a bank vault

  - Further constraints might be placed (i.e. that the vault may only be opened between the hours of 9:00–3:00)
  - If documents are moved from the vault and stored on bank computers, the computer security policy should probably uphold those same properties

Subjects and objects

A useful way of thinking about access control is in terms of *subjects* and *objects*.

A *subject* is an entity making a request. (An executing program, a user sitting at a terminal, etc.) Also called a *principal*. In Unix, this is just a user-id.

An *object* is an item to which access is requested. (A byte in memory, a file on disk, a printer, etc.)

Subjects and objects, cont.

Then, the policy specification can be translated into an *access control matrix*.

The rows of the matrix are labelled by the subjects, and the columns by the objects. The entry at row S and column O specifies whether S may access O.

In this language, we see that authentication is the problem of identifying the subject making the request, and checking authorization just comes down to checking a matrix entry.

Access control, cont.

Usually, “all-or-nothing” access control is insufficient; we want more fine-grained controls.

In this case, we may establish a set of *permissions*. (For example: read, write, execute, etc.) Then, the matrix entry at location (S, O) is filled with the set of permissions which S may acquire on O.

This provides a way to view the system which is convenient for analysis. For instance, suppose we want to find whether there is any way which a subject S may send a message to another subject S'. Then we may simply search for all objects O which S has write permission to, and check whether S' has read access to them.

Mandatory access control

Discretionary access control: owner of each object gets to set the permissions on that object.

Mandatory access control (or MAC): system administrator gets some control over the access controls on every object. Usually found in military systems (e.g. anything marked “classified” may not be sent to a printer in an unclassified facility). We’ll discuss this in detail next week.

Objects, cont.

What are objects in real life?

Some examples: pages of memory, files on a storage system, I/O devices, communications channels to other subjects.

Access control lists

One typical way to implement access control is via *access control lists*, or ACLs.

In this approach, each object is tagged with an ACL. The ACL identifies a list of subjects S which may access the object. For each subject S, the ACL lists the permissions S has on the object.

An example:

```
allow read,write Alice,Jim
allow read John,Jim,Bob
allow delete Alice
deny read Eric
```


Capabilities

Capability: An unforgeable "ticket" which presents proof that the presenter is authorized to access the specified object.

Example:

$$(O, \text{read})$$

means the holder of this capability has read permissions on object O.

Permissions may be further restricted, e.g. by including an expiration date, etc.

Capabilities, cont.

Example implementations:

- Some old computer architectures supported capabilities in hardware. The memory system provided the unforgeability properties by tagging each memory word; only special privileged processes (e.g. the OS kernel) could create new capabilities.
- More commonly, enforced in the OS: the OS keeps track of all your capabilities for you, and manipulating them requires a system call.
- Pointer solutions: a capability is a pointer to the object in question. The address space is so large that one cannot guess the pointer by pure trial and error.
- Cryptographic solutions: a capability is a ticket encrypted under a key only the OS knows.
- Safe languages: a type-safe language gives capabilities a different type and enforces strict rules on the use of that type.

Comparing ACLs and capabilities

Access control decisions:

- ACLs: made at time of each access request (late binding).
- Capabilities: made once when capability is granted (early binding).
- Thus, capabilities potentially allow much higher performance by caching access control decisions.

Passing privileges to another subject:

- ACLs: usually not possible without explicit provisions for delegation.
- Capabilities: usually may be copied and passed to other entities at will.

Comparing ACLs and capabilities, cont.

Revocation of priviliges:

- ACLs: easy; just change the ACL list.
- Capabilities: usually hard; once a capability is granted, it usually cannot be revoked. (Expiration dates help somewhat.)

It's clear that all of the interesting tradeoffs come in the *dynamic* behavior of the system.

ACLs and capabilities

A useful implementation technique:

> When an ACL is too slow, inflexible, or unwieldy, grant access to a capability by checking an ACL and then grant access to objects according to the capability.

This is known as a hybrid approach.

For example, in Unix, when you `open()` a file, it checks the file's ACL. The `open()` returns a file descriptor, which is just a capability to access the file; the file descriptor may be passed around at will.

Address spaces

For a nice example, let's look at how memory protection works in Unix.

Element of protection is a memory page.

Enforcement is provided by the combined effort of the OS kernel and the hardware memory management subsystem.

The main idea is to use virtual memory for protection: authorization is checked by OS when you load a page of memory into your virtual address space, and hardware ensures that memory accesses are restricted to valid pages.

You can think of the set of the virtual memory map loaded into the hardware as capabilities granting access to physical memory; the OS kernel enforces an ACL on physical pages of memory.

Dynamics

Enforcing access control on a static system with an unchanging access control policy is “easy.” The problem becomes much harder when you introduce *dynamics*, e.g. the ability to modify policy while the system is running.

A useful permission (like read, write, etc.) in this context is the permission to modify the ACL for the object in question.

But note how difficult it becomes to analyze the system. For instance, suppose we want to find whether there is any way which a subject S may affect an object O in any way. Then it is not enough to just check whether S has write permission on O; one must also check whether S has permission to modify the permissions for O, or whether S can modify anything which can modify O, etc.

Groups

Often it is useful to group a list of subjects together and specify a common set of permissions for the group. This is intended to simplify the specification (and management)

Examples: graduate-students, cs261-students, developers.

Then an ACL for this course’s syllabus might look like

```
allow  read,write   David Wagner
allow  read         cs261-students, graduate-students
```

Warning: be careful with dynamics of group membership!

Authentication

Authentication is about identifying the person who initiated a given request.

A few categories:

- Something you know: passwords, cryptographic keys
- Something you have: hardware tokens
- Something you are: biometrics
  - e.g. fingerprints, iris scans, hand-geometry, key-typing rates, etc.

Strongest approach is to use several of these in combination.

Password authentication

Many systems just require the user to enter a secret *password* known to the system.

Disadvantage: the sys-admin knows your password.
Fix: when the account is created, apply a non-invertible one-way hash function to the secret password and store the result. Then a login request may be validated by hashing the password provided and comparing it to the value stored earlier.

Passwords, cont.

Weaknesses:

- Guessability: most people pick poor passwords.
  - The Internet Worm (1988) included automatic password-guessing code.
- Eavesdropping: passwords are usually sent in the clear over the network, where they may be intercepted.
  - In 1995, CERT issued a security advisory noting that intruders had broken into many ISPs and backbone routers and intercepted tens of thousands of account passwords.

    CERT's advice? Every Internet user should change their password.

These weaknesses have been known for decades.

Secure authentication

For authentication across a network, passwords are obsolete technology.

The only known solution is cryptography; wait a few weeks.

Enforcement

How is access control enforced?

- Hardware
- Operating system kernel
- Trusted subsystems

In practice, hardware provides very simple mechanisms, which are used to bootstrap an OS kernel. Then the security policy is refined by trusted programs run under the OS.

Hardware architectures

The concept of privileged execution states is critical.

- One simple concept: two modes of execution, *supervisor* mode (or *privileged* mode) and *unprivileged* mode (or *user* mode).
- May only access some features (memory mapping management, I/O, etc.) in supervisor mode.
- OS kernel runs in supervisor mode; normal programs run in unprivileged mode.
- May only switch to *privileged* mode by way of predefined gates (e.g.: a system call) which contain special checks and guards.

A refinement of this approach is *rings*, where there are multiple levels of privilege.

Upgrading privileges

How does a call to a privileged-mode process work?

- Trap to operating system
- Operating system checks arguments
  - Be careful with call-by-reference arguments and return addresses!
- OS executes privileged system call code
- Return to user-mode

A shift of topics

Military security,
the Orange Book,
what worked,
and what didn't

Military security

Military computer security types traditionally worry a lot about

- Confidentiality: the military keeps lots of high-grade secrets.
- High assurance: lives are on the line!
- Insider threats: with millions of insiders, you start to get worried about spies. (Walker, etc.)

The Orange Book: a specification for designing and certifying military-use trusted systems.
Intended to encourage commercial developers to build systems that the military could trust.

Military security, cont.

To understand some of the Orange Book, it helps to understand how military security folks are coming from.

Background:

- Hierarchical authority structures.
- Existing precedents from paper document handling:
  classification markings, codewords, need-to-know, specialized procedures, etc.
- Serious concerns about personnel security.
- Lots of pre-existing laws and regulations.
- Dedicated and highly secure communications lines.

Orange Book concepts

Some extraordinarily useful concepts from the Orange Book:

- Reference monitors: a mechanism that implements access control, which is
  - tamper-proof,
  - simple enough to trust,
  - and ensures complete mediation
- Trusted computing base (TCB): the smallest subset of the system which must be trusted if it is to be secure
- Assurance: how do you know the TCB does what it claims?
- Multi-level security, mandatory access control, and covert channels.
- Object reuse limitations: objects must be zeroized before reuse.
- Heavy audit logs.
- Trusted path.

Trusted computing base

- often a useful question to ask of any system: can you point to its TCB?
- e.g., in Java, the TCB includes the JVM, the bytecode verifier, the security manager, the classloader, . . .
  - not a good sign for Java's security.

Object reuse

Some spectacular failures:

- MS Word doesn't zero disk blocks before writing to them, so sending a Word document to someone else may leak confidential data. Oops!
- Kerberos 4 also has object reuse vulnerabilities: in some cases, it will send a packet containing data from an unsanitized buffer.

Trusted path

- Trusted path means when you see a "login:" prompt, knowing it's really the trusted login program and not a Trojan horse surrogate.
  - e.g. in NT, pressing CTRL-ALT-DEL gives you a guaranteed-good login prompt
- More generally, "trusted path" is about how you bootstrap trust.

Mandatory access control

- Motivation: subjects are not entirely trusted.
- Intended to prevent attacks by insiders (both untrusted users and untrusted code).
- Based on pre-existing classification procedures:

  No person may have access to classified information unless: (a) that person has been determined to be trustworthy, i.e., granted a personnel security clearance – MANDATORY, and (b) access is necessary for the performance of official duties, i.e., determined to have a need-to-know – DISCRETIONARY.
- All data is tagged with a classification level (UNCLASSIFIED, SECRET, TOP SECRET, . . .) and optionally a list of codewords (NOFORN, ULTRA, . . .).
- Codewords allow for compartmentalization (need-to-know).
- Information cannot flow "down" without proper authorization (de-classification by an authorized security officer).

The lattice model

Mathematically, the MAC policy above lends itself to a lattice model.

- Let C be the set of classifications, W the set of codewords.
- A label l is $l = (l_c, l_w)$ with $l_c \in C$, $l_w \subset W$.
- You get the partial ordering $\leq$ on tags by

$$l \leq l' \qquad \Leftrightarrow \qquad l_c \leq l_{c'} \text{ and } l_w \subset l_{w'}$$

- Informal policy: you may not read "up" or write "down"
- Formal rules for access by a subject with tag l to an object with tag l':

  read allowed only if $l \geq l'$; write allowed only if $l \leq l'$.
- Only trusted software may violate these rules.

MAC and integrity

A serious problem with mandatory access control:

- The policy above assures only confidentiality, not integrity.
- Note that it allows an unclassified user to tamper with even the most sensitive data!
- Can obviously ensure integrity with a lattice model, too
  - but the rules are the dual: you may not write up or read down
- And if you try for both mandatory confidentiality and integrity, the access rules degenerate to the trivial condition that $l = l'$, which is too inflexible.
- Moral: you can't have your cake and eat it too.

What goes up may never come down. . .

- In practice, the tendency is for objects to migrate up the lattice, and soon everything takes on the highest classification level possible.
- A few simple thought experiments:
  - Imagine an Excel spreadsheet containing only unclassified data. You add a single classified field somewhere near the end. Then you try to print the first page, which contains only unclassified fields, to an unclassified printer. Will the system let you?
  - When sending messages from low to high, how do you handle ACKs?
  - What's the security label on the program counter register?

What goes up may never come down. . . (cont)

- You end up needing to build “trusted” data pumps which can de-classify data.
  - requires either human intervention or solving a nearly AI-complete problem
  - and thus is a very difficult problem
- This is yet another serious problem with MAC.

Covert channels

A covert channel is a stealthy way to send information from high to low:

- e.g. by CPU load: to send a “1” (“0”), spin-loop (sleep) for a second; to receive, measure CPU load for a second.
- All sorts of crazy covert channels: monitor disk head timings, amount of free memory, etc.
- Theoretically, the *non-interference* property is sufficient to prevent covert channels.
- In practice, you end up isolating all programs with different security labels, and eliminating all resource sharing/multiplexing from your OS.
  - Once you’re done, does it really deserve to be called an OS anymore?
- In real life, eliminating covert channels is mind-bogglingly difficult.

This is yet another serious problem with MAC.

Assurance

- Arguably the most important criterion of all!
  - After all, a false sense of security is worse than no security at all
- Note that the Orange Book ties assurance into the certification levels; its more recent European cousin makes assurance and features orthogonal.

Problems with the Orange Book

It doesn’t fit the commercial world, and in some ways it is just plain silly.

Problems with the Orange Book

- Certification takes many years.

  In a market living on Internet time scales where time-to-market dominates, this is unthinkably deadly.
- Certification is linked to hardware configuration.

  You end up with obsolete hardware, too.
- Any fix, patch, version rev., or change in hardware configuration requires a re-certification.

  Oops, there goes another year. . .
- If you want to buy a B1-system, you end up with some 386 running 10-year old software, and it'll cost you $100k for the privilege.

Problems with the Orange Book

- The security policy is all wrong.
- The Orange Book is heavily concerned about confidentiality and insider threats.
- But most commercial enterprises trust all insiders.
  - Anyhow, defending against employees leaking information by covert channels is silly, when they can easily walk out with floppies in their briefcase.
- And availability and integrity controls are more important than confidentiality in many cases.

Problems with the Orange Book

- The threat model is all wrong.
- In a B1 firewall, you turn off all web access:
  - just in case some spy on a secret network does something totally goofy like going to the web page

    `http://www.kgb.gov.su/the/attack/begins/at/midnight`
  - so you end up telling your users to *go home* and use an AOL account.
- Anyhow, worrying about covert channels is just plain silly:
  - nobody in their right mind is going to try to steal information from disk head timings when they can have Aldridge Ames for $100k!

Credits: Marcus Ranum

Problems with the Orange Book

Interestingly, the Orange Book applies only to non-networked computers. This is very explicit.

Once you plug that C2-certified NT box into a network, it loses its certification!

Results?

Software glitches leave Navy Smart Ship dead in the water

July 13, 1998: In September 1997, the Yorktown suffered a systems failure during maneuvers off the coast of Cape Charles, VA., apparently as a result of the failure to prevent a divide by zero in a Windows NT application. The zero seems to have been an erroneous data item that was manually entered. Atlantic Fleet officials said the ship was dead in the water for about 2 hours and 45 minutes.

To first order, noone buys multilevel certified systems anymore; everyone is using a Windoze 95 or NT box.

Lessons

- Worrying about covert channels is the wrong mindset.
  - Much safer: simply never let untrusted programs get access to confidential information.
- MAC probably doesn't help much against untrusted insiders.
  - But it can help stop Trojan horses.

Networks

So what happens when you introduce networks anyway?

Answer: the whole game changes.

Networks and centralized control

- Networks eliminate all hope of centralized control.
  - unless you magically ensure that your network is secure, and only trusted machines are connected to it
  - and even if you can, it is totally fragile.
- This dooms all hope for multi-level systems or mandatory access control.
  - Do you label each packet with its classification level?

Networks and the threat model

In practice, just about the only reasonable threat model is to assume that the network is totally under the control of the adversary.

Why?

- Hard to ensure physical security of network links.
- Multiplexed access means the wire may go right into the enemy's doorstep.
- Every piece of infrastructure (routers, nameservers, . . .) introduces single points of failure.
- When every computer on the net has access to every other, scaling is your enemy.
  - This makes it especially critical that all public services be perfect.

Networks

Networks require new tools:

- Techniques for dealing with scale.
- Techniques for secure communications, and for managing trust.

One answer to the latter is cryptography. More later. . .