

# Lecture 26

# December 1, 2025

ECS 235A, Computer and Information Security

# Administrative Stuff

- Friday, December 5 is the last day of class
- Please get all homework and extra credit in by then
  - I must submit grades 3 days after the final exam date

***No work, including homework, extra credit, or projects, can be accepted after the final exam date, December 9, 2025***

# Integrity Mechanisms

- The above also works with Biba, as it is mathematical dual of Bell-LaPadula
- All constraints are simply duals of confidentiality-based ones presented above

# Example 1

For information flow of assignment statement:

$$y := f(x_1, \dots, x_n)$$

the relation  $\text{glb}\{x_1, \dots, x_n\} \geq y$  must hold

- Why? Because information flows from  $x_1, \dots, x_n$  to  $y$ , and under Biba, information must flow from a higher (or equal) class to a lower one

# Example 2

For information flow of conditional statement:

**if**  $f(x_1, \dots, x_n)$  **then**  $S_1$ ; **else**  $S_2$ ; **end**;

then the following must hold:

- $S_1, S_2$  must satisfy integrity constraints
- $\text{glb}\{x_1, \dots, x_n\} \geq \text{lub}\{y \mid y \text{ target of assignment in } S_1, S_2\}$

# Example Information Flow Control Systems

- Privacy and Android Cell Phones
  - Analyzes data being sent from the phone
- Firewalls

# Privacy and Android Cell Phones

- Many commercial apps use advertising libraries to monitor clicks, fetch ads, display them
  - So they send information, ostensibly to help tailor advertising to you
- Many apps ask to have full access to phone, data
  - This is because of complexity of permission structure of Android system
- Ads displayed with privileges of app
  - And if they use Javascript, that executes with those privileges
  - So if it has full access privilege, it can send contact lists, other information to others
- Information flow problem as information is flowing from phone to external party

# Analyzing Android Flows

- Android based on Linux
  - App executables in bytecode format (Dalvik executables, or DEX) and run in Dalvik VM
  - Apps event driven
  - Apps use system libraries to do many of their functions
  - Binder subsystem controls interprocess communication
- Analysis uses 2 security levels, *untainted* and *tainted*
  - No categories, and *tainted* < *untainted*

# TaintDroid: Checking Information Flows

- All objects tagged *tainted* or *untainted*
  - Interpreters, Binder augmented to handle tags
- Android native libraries trusted
  - Those communicating externally are *taint sinks*
- When untrusted app invokes a taint sink library, taint tag of data is recorded
- Taint tags assigned to external variables, library return values
  - These are assigned based on knowledge of what native code does
- Files have single taint tag, updated when file is written
- Database queries retrieve information, so tag determined by database query responder

# TaintDroid: Checking Information Flows

- Information from phone sensor may be sensitive; if so, *tainted*
  - TaintDroid determines this from characteristics of information
- Experiment 1 (2010): selected 30 popular apps out of a set of 358 that required permission to access Internet, phone location, camera, or microphone; also could access cell phone information
  - 105 network connections accessed *tainted* data
  - 2 sent phone identification information to a server
  - 9 sent device identifiers to third parties, and 2 didn't tell user
  - 15 sent location information to third parties, none told user
  - No false positives

# TaintDroid: Checking Information Flows

- Experiment 2 (2012): revisited 18 out of the 30 apps (others did not run on current version of Android)
  - 3 still sent location information to third parties
  - 8 sent device identification information to third parties without consent
    - 3 of these did so in 2010 experiment
    - 5 were new
  - 2 new flows that could reveal *tainted* data
  - No false positives

# Firewalls

- Host that mediates access to a network
  - Allows, disallows accesses based on configuration and type of access
- Example: block Conficker worm
  - Conficker connects to botnet, which can use system for many purposes
    - Spreads through a vulnerability in a particular network service
  - Firewall analyze packets using that service remotely, and look for Conficker and its variants
    - If found, packets discarded, and other actions may be taken
  - Conficker also generates list of host names, tried to contact botnets at those hosts
    - As set of domains known, firewall can also block outbound traffic to those hosts

# Filtering Firewalls

- Access control based on attributes of packets and packet headers
  - Such as destination address, port numbers, options, etc.
  - Also called a *packet filtering firewall*
  - Does not control access based on content
  - Examples: routers, other infrastructure systems

# Proxy

- Intermediate agent or server acting on behalf of endpoint without allowing a direct connection between the two endpoints
  - So each endpoint talks to proxy, thinking it is talking to other endpoint
  - Proxy decides whether to forward messages, and whether to alter them

# Proxy Firewall

- Access control done with proxies
  - Usually bases access control on content as well as source, destination addresses, etc.
  - Also called an *applications level* or *application level firewall*
- Example: virus checking in electronic mail
  - Incoming mail goes to proxy firewall
  - Proxy firewall receives mail, scans it
  - If no virus, mail forwarded to destination
  - If virus, mail rejected or disinfected before forwarding

# Example

- Want to scan incoming email for malware
- Firewall acts as recipient, gets packets making up message and reassembles the message
  - It then scans the message for malware
  - If none, message forwarded
  - If some found, mail is discarded (or some other appropriate action)
- As email reassembled at firewall by a mail agent acting on behalf of mail agent at destination, it's a proxy firewall (application layer firewall)

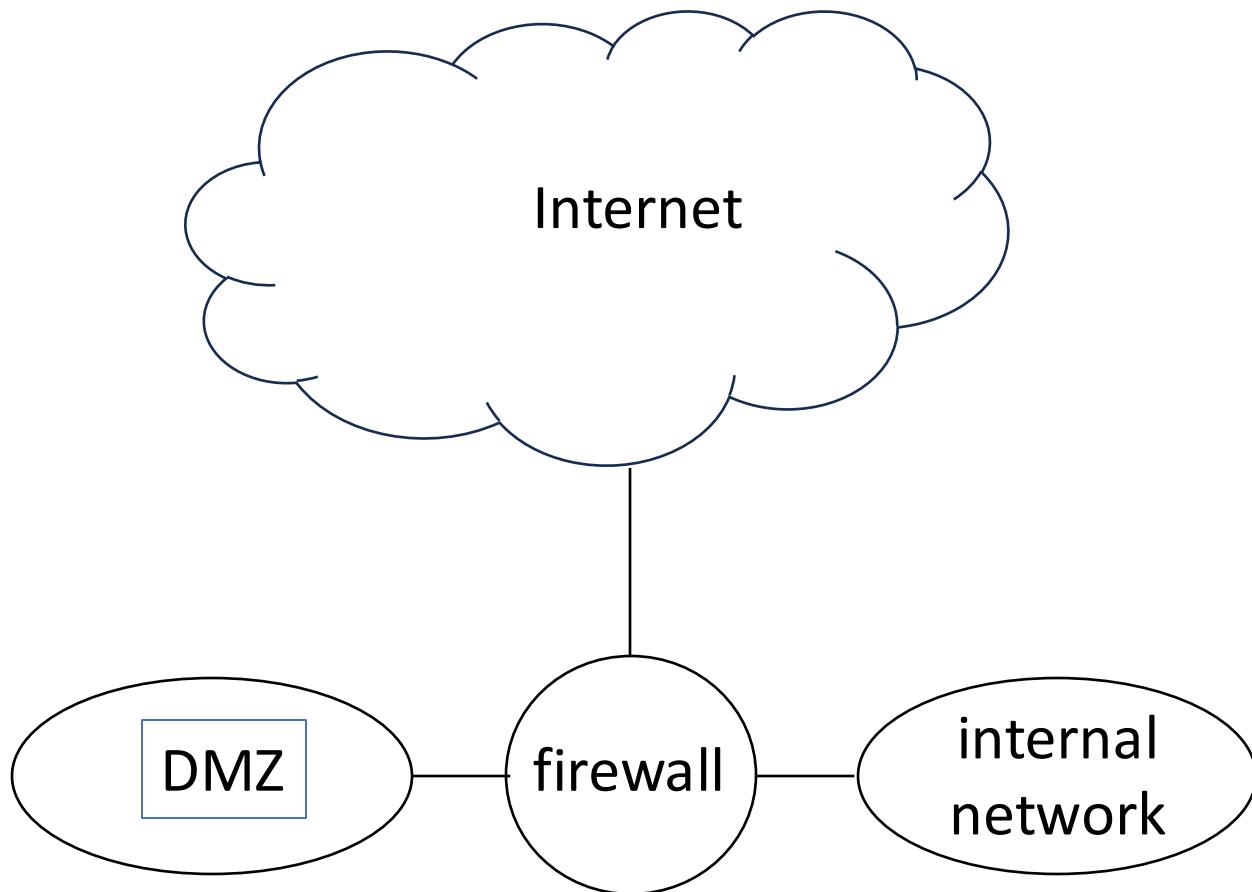
# Stateful Firewall

- Keeps track of the state of each connection
- Similar to a proxy firewall
  - No proxies involved, but this can examine contents of connections
  - Analyzes each packet, keeps track of state
  - When state indicates an attack, connection blocked or some other appropriate action taken

# Network Organization: DMZ

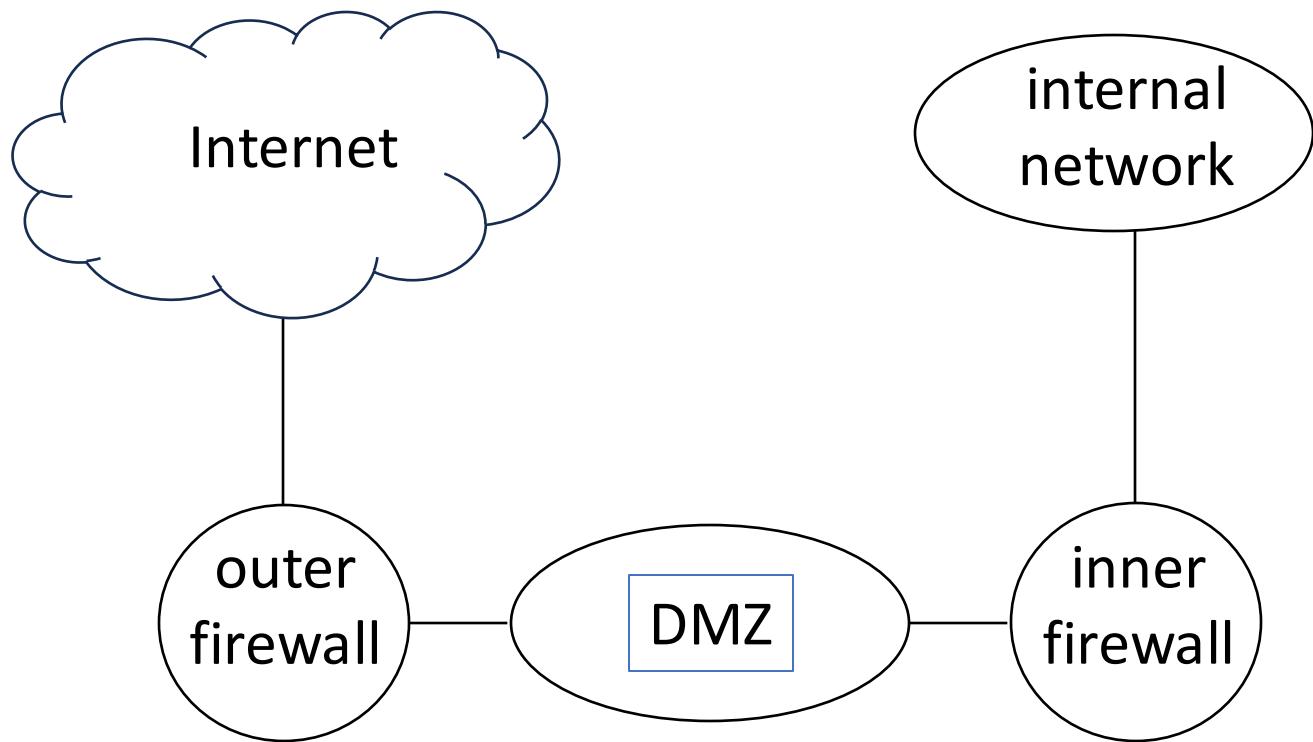
- DMZ is portion of network separating a purely internal network from external network
  - Comes from "DeMilitarized Zone", a term meaning an area between conflicting entities in which no hostilities can occur
- Usually put systems that need to connect to the Internet here
- Firewall separates DMZ from purely internal network
- Firewall controls what information is allowed to flow through it
  - Control is bidirectional; it controls flow in both directions

# One Setup of DMZ



One dual-homed firewall that routes messages to internal network or DMZ as appropriate

# Another Setup of DMZ



Two firewalls, one (outer firewall) connected to the Internet, the other (inner firewall) connected to internal network, and the DMZ is between the firewalls

# Identity

- *Principal*: a unique entity
- *Identity*: specifies a principal
- *Authentication*: binding of a principal to a representation of identity internal to the system
  - All access, resource allocation decisions assume binding is correct

# Files and Objects

- Identity depends on system containing object
- Different names for one object
  - Human use, *eg.* file name
  - Process use, *eg.* file descriptor or handle
  - Kernel use, *eg.* file allocation table entry, inode

# More Names

- Different names for one context
  - Human: aliases, relative vs. absolute path names
  - Kernel: deleting a file identified by name can mean two things:
    - Delete the object that the name identifies
    - Delete the name given, and do not delete actual object until *all* names have been deleted
- Semantics of names may differ

# Example: Names and Descriptors

- Interpretation of UNIX file name
  - Kernel maps name into an inode using iterative procedure
  - Same name can refer to different objects at different times without being deallocated
    - Causes race conditions
- Interpretation of UNIX file descriptor
  - Refers to a specific inode
  - Refers to same inode from creation to deallocation

# Direct vs. Indirect Alias File Names

- Direct alias: name identifying specific entry in inode table or file allocation table (FAT)
  - Kernel maps name to directory entry that contains metadata of the file and addresses of data blocks make up the file contents
  - In UNIX/Linux/\*BSD, etc. sometimes called a *hard link*
- Indirect alias: name identifying another file
  - Contents of indirect alias file are interpreted as the name of another file
  - Kernel then iterates process to obtain contents
  - In UNIX/Linux/\*BSD, etc. called a *symbolic link*

# Example: Different Systems

- Object name must encode location or pointer to location
  - *SSH* style: *host:object*
  - URLs: *protocol://host/object*
- Need not name actual object
  - *SSH* style may name pointer (link) to actual object
  - URL may forward to another host

# Users

- Exact representation tied to system
- Example: UNIX/\*BSD/Linux systems
  - Login name: used to log in to system
    - Logging usually uses this name
  - User identification number (UID): unique integer assigned to user
    - Kernel uses UID to identify users
    - One UID per login name, but multiple login names may have a common UID

# Multiple Identities

- UNIX/\*BSD/Linux systems again
  - Real UID: user identity at login, but changeable
  - Effective UID: user identity used for access control
    - Setuid changes effective UID
  - Saved UID: UID before last change of UID
    - Used to implement least privilege
    - Work with privileges, drop them, reclaim them later
  - Audit/Login UID: user identity used to track original UID
    - Cannot be altered; used to tie actions to login identity

# Groups

- Used to share access privileges
- First model: alias for set of principals
  - Processes assigned to groups
  - Processes stay in those groups for their lifetime
- Second model: principals can change groups
  - Rights due to old group discarded; rights due to new group added

# Roles

- Group with membership tied to function
  - Rights given are consistent with rights needed to perform function
- Uses second model of groups
- Example: DG/UX
  - User *root* does not have administration functionality
  - System administrator privileges are in *sysadmin* role
  - Network administration privileges are in *netadmin* role
  - Users can assume either role as needed

# Naming and Certificates

- Certificates issued to a principal
  - Principal uniquely identified to avoid confusion
- Problem: names may be ambiguous
  - Does the name “Matt Bishop” refer to:
    - The author of this book?
    - A programmer in Australia?
    - A stock car driver in Muncie, Indiana?
    - Someone else who was named “Matt Bishop”

# Disambiguating Identity

- Include ancillary information in names
  - Enough to identify principal uniquely
  - X.509v4 Distinguished Names do this
- Example: X.509v4 Distinguished Names
  - /O=University of California/OU=Davis campus/OU=Department of Computer Science/CN=Matt Bishop/  
refers to the Matt Bishop (*CN* is *common name*) in the Department of Computer Science (*OU* is *organizational unit*) on the Davis Campus of the University of California (*O* is *organization*)

# CAs and Policies

- Matt Bishop wants a certificate from Certs-from-Us
  - How does Certs-from-Us know this is “Matt Bishop”?
    - CA’s *authentication policy* says what type and strength of authentication is needed to identify Matt Bishop to satisfy the CA that this is, in fact, Matt Bishop
  - Will Certs-from-Us issue this “Matt Bishop” a certificate once he is suitably authenticated?
    - CA’s *issuance policy* says to which principals the CA will issue certificates

# Example: Verisign CAs

- Class 1 CA issued certificates to individuals
  - Authenticated principal by email address
    - Idea: certificate used for sending, receiving email with various security services at that address
- Class 2 CA issued certificates to individuals
  - Authenticated by verifying user-supplied real name and address through an online database
    - Idea: certificate used for online purchasing

# Example: Verisign CAs

- Class 3 CA issued certificates to individuals
  - Authenticated by background check from investigative service
    - Idea: higher level of assurance of identity than Class 1 and Class 2 CAs
- Class 4 CA issued certificates to web servers
  - Same authentication policy as Class 3 CA
    - Idea: consumers using these sites had high degree of assurance the web site was not spoofed

# Registration Authority

- Third party delegated by CA the authority to check data to be put into certificate
  - This includes identity
- RA determines whether CA's requirements are met
- If so, then it informs CA to issue certificates

# Internet Certification Hierarchy

- Tree structured arrangement of CAs
  - Root is *Internet Policy Registration Authority*, or IPRA
    - Sets policies all subordinate CAs must follow
    - Certifies subordinate CAs (called *policy certification authorities*, or PCAs), each of which has own authentication, issuance policies
    - Does not issue certificates to individuals or organizations other than subordinate CAs
  - PCAs issue certificates to ordinary CAs
    - Does not issue certificates to individuals or organizations other than subordinate CAs
  - CAs issue certificates to organizations or individuals

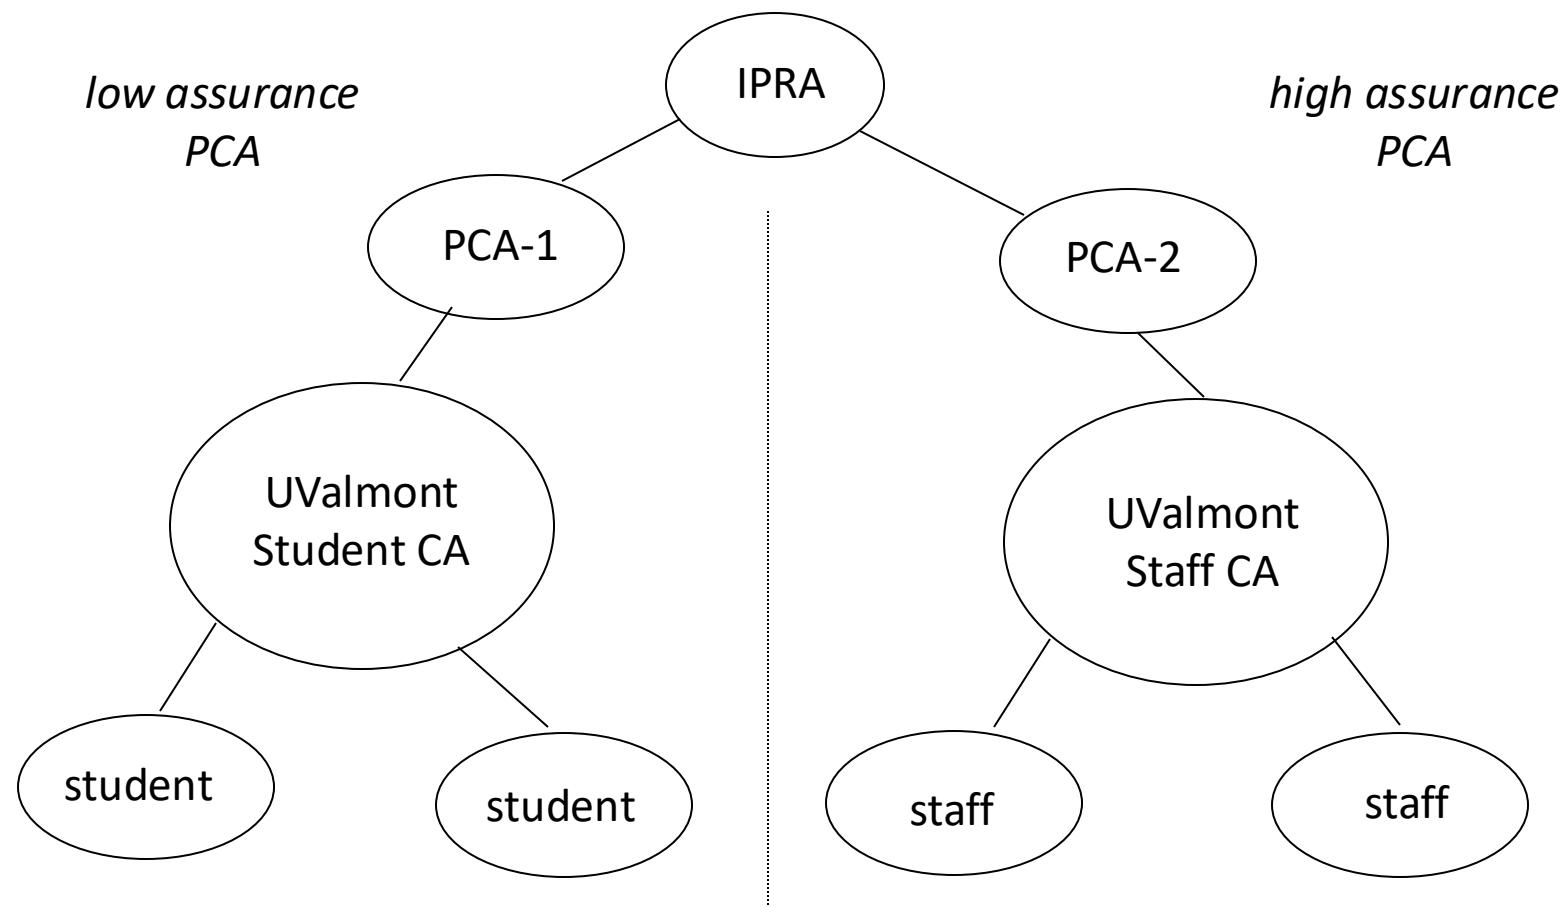
# Example

- University of Valmont issues certificates to students, staff
  - Students must present valid reg cards (considered low assurance)
  - Staff must present proof of employment and fingerprints, which are compared to those taken when staff member hired (considered high assurance)

# UValmont and PCAs

- First PCA: requires subordinate CAs to make good-faith effort to verify identities of principals to whom it issues certificates
  - Student authentication requirements meet this
- Second PCA: requires use of biometrics to verify identity
  - Student authentication requirements do not meet this
  - Staff authentication requirements do meet this
- UValmont establishes two CAs, one under each PCA above

# UValmont and Certification Hierarchy



# Certificate Differences

- Student, staff certificates signed using different private keys (for different CAs)
  - Student's signed by key corresponding to low assurance certificate signed by first PCA
  - Staff's signed by key corresponding to high assurance certificate signed by second PCA
- To see what policy used to authenticate:
  - Determine CA signing certificate, check its policy
  - Also go to PCA that signed CA's certificate
    - CAs are restricted by PCA's policy, but CA can restrict itself further