

# ECS 235B Module 44

## Building Systems with Assurance

# Threats and Goals

- *Threat* is a danger that can lead to undesirable consequences
- *Vulnerability* is a weakness allowing a threat to occur
- Each identified threat requires countermeasure
  - Unauthorized people using system mitigated by requiring identification and authentication
- Often single countermeasure addresses multiple threats

# Architecture

- Where do security enforcement mechanisms go?
  - Focus of control on operations or data?
    - Operating system: typically on data
    - Applications: typically on operations
  - Centralized or distributed enforcement mechanisms?
    - Centralized: called by routines
    - Distributed: spread across several routines

# Layered Architecture

- Security mechanisms at any layer
  - Example: 4 layers in architecture
    - *Application layer*: user tasks
    - *Services layer*: services in support of applications
    - *Operating system layer*: the kernel
    - *Hardware layer*: firmware and hardware proper
- Where to put security services?
  - Early decision: which layer to put security service in

# Security Services in Layers

- Choose best layer
  - User actions: probably at applications layer
  - Erasing data in freed disk blocks: OS layer
- Determine supporting services at lower layers
  - Security mechanism at application layer needs support in all 3 lower layers
- May not be possible
  - Application may require new service at OS layer; but OS layer services may be set up and no new ones can be added

# Security: Built In or Add On?

- Think of security as you do performance
  - You don't build a system, then add in performance later
    - Can “tweak” system to improve performance a little
    - Much more effective to change fundamental algorithms, design
- You need to design it in
  - Otherwise, system lacks fundamental and structural concepts for high assurance

# Reference Validation Mechanism

- *Reference monitor* is access control concept of an abstract machine that mediates all accesses to objects by subjects
- *Reference validation mechanism* (RVM) is an implementation of the reference monitor concept.
  - Tamperproof
  - Complete (always invoked and can never be bypassed)
  - Simple (small enough to be subject to analysis and testing, the completeness of which can be assured)
    - Last engenders trust by providing evidence of correctness

# Examples

- *Security kernel* combines hardware and software to implement reference monitor
- *Trusted computing base (TCB)* consists of all protection mechanisms within a system responsible for enforcing security policy
  - Includes hardware and software
  - Generalizes notion of security kernel



# Adding On Security

- Key to problem: analysis and testing
- Designing in mechanisms allow assurance at all levels
  - Too many features adds complexity, complicates analysis
- Adding in mechanisms makes assurance hard
  - Gap in abstraction from requirements to design may prevent complete requirements testing
  - May be spread throughout system (analysis hard)
  - Assurance may be limited to test results

# Example

- 2 AT&T products with same goal of adding mandatory controls to UNIX system
  - SV/MLS: add MAC to UNIX System V Release 3.2
  - SVR4.1ES: re-architect UNIX system to support MAC

# Comparison

- Architecting of System
  - SV/MLS: used existing kernel modular structure; no implementation of least privilege
  - SVR4.1ES: restructured kernel to make it highly modular and incorporated least privilege

# Comparison

- File Attributes (*inodes*)
  - SV/MLS added separate table for MAC labels, DAC permissions
    - UNIX inodes have no space for labels; pointer to table added
    - Problem: 2 accesses needed to check permissions
    - Problem: possible inconsistency when permissions changed
    - Corrupted table causes corrupted permissions
  - SVR4.1ES defined new inode structure
    - Included MAC labels, DAC attributes
    - Only 1 access needed to check permissions

# Requirements Assurance

- *Specification* describes of characteristics of computer system or program
- *Security specification* specifies desired security properties
- Must be clear, complete, unambiguous
  - Something like “meets C2 security requirements” not good: what *are* those requirements (actually, 34 of them!)

# Example

- “Users of the system must be identified and authenticated” is ambiguous
  - Type of ID required—driver’s license, token?
  - What is to be authenticated—user, representation of identity, system?
  - Who is to do the authentication—system, guard?
- “Users of the system must be identified to the system and must have that identification authenticated by the system” is less ambiguous
  - Under what conditions must the user be identified to the system—at login, time of day, or something else?

# Example

- “Users of the system must be identified to the system and must have that identification authenticated by the system before the system performs any functions on behalf of that identity”
  - Type of identification is user name
  - User identification (name) to be authenticated
  - System to authenticate
  - Authentication to be done at login (before system performs any action on behalf of user)

# Methods of Definition

- Extract applicable requirements from existing security standards
  - Tend to be semiformal
- Combine results of threat analysis with components of existing policies to create a new policy
- Map the system to existing model
  - If model appropriate, creating a mapping from model to system may be cheaper than requirements analysis



# Example

- System X: UNIX system with MAC based on Bell-LaPadula Model
  - Mapping constructed in series of stages
  - Auditing also required

# Example Stage 1

- Map elements, state variables of BLP to entities in System X
  - Subject set  $S$  in BLP  $\rightarrow$  set of processes in System X
  - Object set  $O$  in BLP  $\rightarrow$  set of inode objects, IPC objects, mail messages, processes as destinations, passive entities in System X
  - Right set  $P$  in BLP  $\rightarrow$  set of rights of system functions in System X
    - Functions that create entities, write entities, have write  $\underline{w}$
    - Functions that read entities have right  $\underline{r}$
    - Functions that execute, search entities have right  $\underline{r}$
  - Access set  $b$  in BLP  $\rightarrow$  types of access
    - Subjects can use rights  $\underline{r}$ ,  $\underline{w}$ ,  $\underline{a}$  to access inode objects

# Example Stage 1

- Access control matrix  $a$  for current state in BLP  $\rightarrow$  current state of mandatory and discretionary controls in System X
- Functions  $f_s, f_o,$  and  $f_c$  in BLP  $\rightarrow$  three functions in System X
  - $f(s)$  is the maximum security level of subject  $s$
  - $current-level(s)$  is current security level of subject  $s$
  - $f(o)$  is the security level of object  $o$
- Hierarchy  $H$  in BLP  $\rightarrow$  differently for different objects in System X
  - Inode objects are hierarchical trees represented by the file system hierarchy
  - Other object types map to discrete points in the hierarchy

# Example Stage 2

- Define BLP properties in language of System X and show each property is consistent with BLP
  - MAC property of BLP  $\rightarrow$  user having over an object:
    - read access iff user's clearance dominates object's classification
    - write access over an object iff object's classification dominates user's clearance.
  - DAC property of BLP  $\rightarrow$  user having access to object iff owner of object has explicitly granted that user access to object

# Example Stage 2

- Label inheritance, user level changes specific to System X
  - Security level of newly created object inherited from creating subject
  - Security level of initial process at user login, security level of initial process after user level change, bounded by security level range defined for that user and for the terminal
  - Security level of newly spawned process inherited from parent, except for first process after a user level change
  - When user's level raised, child process does not inherit write access to objects opened by parent
  - When user's level lowered, all processes, accesses associated with higher privilege terminated

# Example Stage 2

- Reclassification property of System X
  - Specially trusted users allowed to downgrade objects they own within constraints of user's authorizations.
- System X property of owner/group transfer allows ownership or group membership of process to be transferred to another user or group
- Status property is property of System X
  - Restricts visibility of status information available to users when they use standard System X set of commands

# Example Stage 3

- Designers define System X rules by mapping System X system calls, commands, and functions to BLP rules
  - Simple security condition, \*-property, and discretionary security property interpreted for each type of access
  - From these interpretations, designers can extract specific requirements for specific accesses to particular types of objects.

# Example Stage 4

- Designers show System X rules preserve security properties
  - Show that the rules enforce the properties directly; or
  - Map the rules directly to a BLP rule or a sequence of BLP rules
    - 9 rules about current access
    - 5 rules about functions and security levels
    - 8 access permission rules
    - 8 more rules about subjects and objects
  - Designers must show that each rule is consistent with actions of System X.



# Justifying Requirements

- Show policy complete and consistent
- Example: ITSEC suitability analysis
  - Map threats to requirements and assumptions
  - Describe how references address threat

# Example: System Y Evaluation

- Threat T1: A person not authorized to use the system gains access to the system and its facilities by impersonating an authorized user.
  - Requirement IA1: A user is permitted to begin a user session only if the user presents a valid unique identifier to the system and if the claimed identity of the user is authenticated by the system by authenticating the supplied password.
  - Requirement IA2: Before the first user/system interaction in a session, successful identification and authentication of the user take place.

# System Y Assumptions

- Assumption A1: The product must be configured such that only the approved group of users has physical access to the system.
- Assumption A2: Only authorized users may physically remove from the system the media on which authentication data is stored.
- Assumption A3: Users must not disclose their passwords to other individuals.
- Assumption A4: Passwords generated by the administrator shall be distributed in a secure manner.

# System Y Mapping

<b>Threat</b>	<b>Security Target Reference</b>
T1	IA1, IA2, A1, A2, A3, A4

# System Y Justifications

1. Referenced requirements and assumptions guard against unauthorized access.
  - Assumption A1 restricts physical access to the system to those authorized to use it.
  - Requirement IA1 requires all users to supply a valid identity and confirming password.
  - Requirement IA2 ensures that requirement IA1 cannot be bypassed.

# System Y Justifications

2. Referenced assumptions prevent unauthorized users from gaining access by using valid user's identity and password
  - Assumption A3 ensures that users keep passwords secret
  - Assumption A4 prevents unauthorized users from intercepting new passwords when those passwords are distributed to users
  - Assumption A2 prevents unauthorized access to authentication information stored on removable media.

These justifications provide an informal basis for asserting that, if the assumptions hold and the requirements are met, the threat is adequately handled.