



INF523: Assurance in Cyberspace Applied to Information Security


Covert Channel Analysis (continued)

Prof. Clifford Neuman

Lecture 11
6 November 2020



“Assurance Waterfall”





Non-interference Definition

- Intuitively:
 - Low-level user’s “view” of the system should not be affected by anything that a high-level user does
- More formally:
 - Suppose L is a subject in the system
 - Now suppose you:
 1. run the system normally, interleaving the operations of all users
 2. run the system again after deleting all operations requested by subjects which should not be able to pass information to (interfere with) L
 - From L ’s point of view, there should be no visible difference
 - The system is “non-interference secure” if this is true of every subject in the system

Non-interference Implementation



-
- Non-interference is another policy, more abstract than BLP
 - The enforcement mechanisms may be anything, including the BLP rules
 - The more system state you add to the definition of “view”, can catch covert channels that uses that state

Limitations of Non-interference



- Non-interference is very difficult to achieve for realistic systems
- It requires identifying within the view function all potential channels of information
- Realistic systems have many such channels
- Modeling must be at very low level to capture many such channels
- Dealing with timing channels is possible, but difficult
- Very few systems are completely deterministic
- Some “interferences” are benign, e.g., encrypted files

TCSEC Bandwidth Guidelines



- Low bandwidths represent a lower risk
- Rate of one hundred (100) bps is considered "high"
 - not appropriate to call a computer system "secure"
- Rate $<$ one (1) bps acceptable in most environments
- Audit any rate $>$ one (1) bit in ten (10) seconds
- Trade-off system performance and CC bandwidth
 - Provide information for system developer to assess



Measuring Capacity

- Intuitively, difference between unmodulated, modulated channel
- E.g.,
 - Normal uncertainty in channel is 8 bits
 - Attacker modulates channel to send information, reducing uncertainty to 5 bits
 - Covert channel capacity is 3 bits
 - Modulation in effect fixes those bits

Mitigation of Covert Channels



- Problem: channels work by varying use of shared resources
- One solution:
 - Require processes to say what resources they need before running
 - Provide access to them in a way that no other process can access them
- Cumbersome!
 - Includes running (CPU covert channel)
 - Resources stay allocated for lifetime of process



Alternate Approach

- Obscure amount of resources being used
 - Receiver cannot distinguish between what the sender is using and what is added
- How? Two ways:
 - Devote uniform resources to each process
 - Inject randomness into allocation, use of resources



Uniformity or Randomness

- **Uniformity:** Subjects always use same amount of resources
 - Variation of isolation
 - Process can't tell if second process using resource
- **Example: KVM/370 covert channel via CPU usage**
 - Give each VM a time slice of fixed duration
 - Do not allow VM to surrender its CPU time
 - Can no longer send 0 or 1 by modulating CPU usage
- **Randomness:** Make noise dominate channel
 - Does not close it, but makes it useless



Randomness

- Example: MLS database
 - Probability of transaction being aborted by user other than sender, receiver approaches 1 -> very high noise
 - How to do this: have participants abort transactions randomly

Problem: Loss of Efficiency



- Fixed allocation constrains use and wastes resources
- Randomness wastes resources
- Policy question: Is the inefficiency preferable to the covert channel?



Example

-
- Goal: limit covert timing channels on VAX/VMM
 - “Fuzzy time” reduces accuracy of system clocks by generating random clock ticks
 - Random interrupts take any desired distribution
 - System clock updates only after each timer interrupt
 - Kernel rounds time to nearest 0.1 sec before giving it to VM
 - Means it cannot be more accurate than timing of interrupts



Example

- I/O operations have random delays
- Kernel distinguishes 2 kinds of time:
 - *Event time* (when I/O event occurs)
 - *Notification time* (when VM told I/O event occurred)
 - Random delay between these prevents VM from figuring out when event actually occurred)
 - Delay can be randomly distributed as desired (in security kernel, it's 1–19ms)
 - Added enough noise to make covert timing channels hard to exploit



Improvement

- Modify scheduler to run processes in increasing order of security level
 - Now we're worried about "reads up", so ...
- Countermeasures needed only when transition from *dominating* VM to *dominated* VM
 - Add random intervals between quanta for these transitions

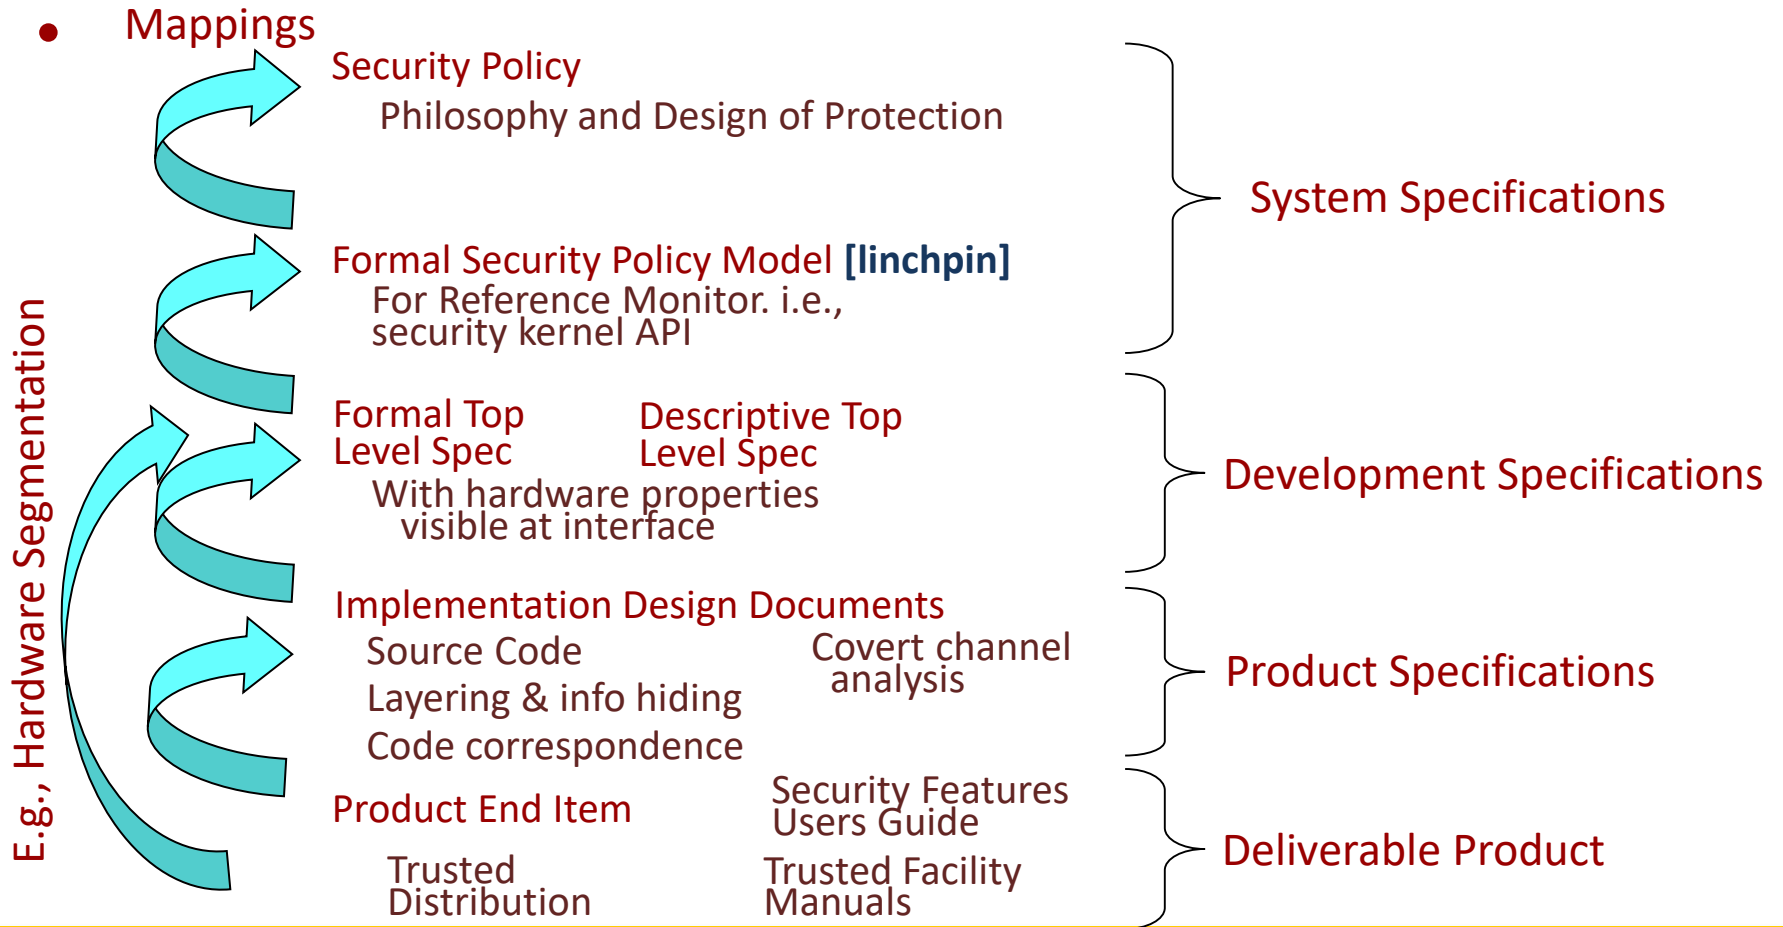


INF523: Case Studies and Security Kernels

Professor Clifford Neuman

Lecture 11 CONT
6 November 2020

Systematic Kernel Engineering Process






Only Proven Solution: Security Kernel

“The only way we know . . . to build highly secure software systems of any practical interest is the kernel approach.”

-- ARPA Review Group, 1970s (Butler Lampson, Draper Prize recipient)

Secure System
Engineering



INF 525 Focus:
Verifiably
Secure
Platform

Truly a paradigm shift: no Class A1 security patches for kernel in years of use

Secure System Case Studies



CASE STUDIES

- GARNETS MLS File System Architecture
- NFS Artifice Demonstration Properties
- MLS Cloud NFS-Based Storage Design
- POSSIBLY:
 - Crypto Seal Guard Demonstration Concepts
 - Crypto Seals in RECON Guard for CIA

Kernel Implementation Strategies



- New operating system
 - Simple mapping of O/S features to SK features
 - Distinctive is lack of backward compatibility
- Compatible operating system (emulation)
- Emulate insecure operating system (ISOS)
 - Typically emulator runs in each process
 - Renders O/S calls into kernel calls
- Identical operating system (virtual machine)
 - Provides isolation, but not sharing, of memory
 - Kernel is virtual machine monitor (VMM)
 - Principal “objects” are virtual disks, not files
 - Subjects – kernel users and VMs

Designing a Security Kernel



- Most used highly secure technique
 - Not easy to build a security kernel
- SK is reference validation mechanism (RVM)
 - Defined as H/W and S/W that implements RM
- Most RMs implement multilevel security (MLS)
- Non-security-relevant functions managed by O/S
- Subject must invoke RM to reference **any** object
 - Basis for *completeness*
- Must have domain control, e.g., protection rings
 - Basis for *isolation* to block subversion
- SK software engineered for RM *verifiability*

Security Analysis of Trusted Systems



- Need independent 3rd party evaluation/analysis
- TCSEC/TNI security kernel evaluation factors
 - System architecture
 - Design specification & verification
 - Sensitivity label management
 - External interfaces
 - Human interfaces
 - Trusted system design properties
 - Security analysis
 - System use and management
 - Trusted system development and delivery

Secure System Design & Development

- Architectural considerations (Gasser chapter 11)
 - Applies to development of security kernels
 - As well as to their applications
- Operating-system layering
 - Promote structured design for kernel assurance
 - Operating systems services on kernel
- Asynchronous attacks and argument validation
- Protected subsystems
 - Static process, on-demand process, multiple domains
- Secure file systems
 - Alternate naming structures; unique identifiers

Secure System Design & Development

- Architectural considerations (Gasser chapter 11)
 - Applies to development of security kernels
 - As well as to their applications
- Operating-system layering
 - Promote structured design for kernel assurance

Designing a Security Kernel



- Most used highly secure technique
 - Not easy to build a security kernel
- SK is reference validation mechanism (RVM)
 - Defined as H/W and S/W that implements RM
- Most RMs implement multilevel security (MLS)
- Non-security-relevant functions managed by O/S
- Subject must invoke RM to reference **any** object
 - Basis for *completeness*
- Must have domain control, e.g., protection rings
 - Basis for *isolation* to block subversion
- SK software engineered for RM *verifiability*

GEMSOS Security Kernel Layering

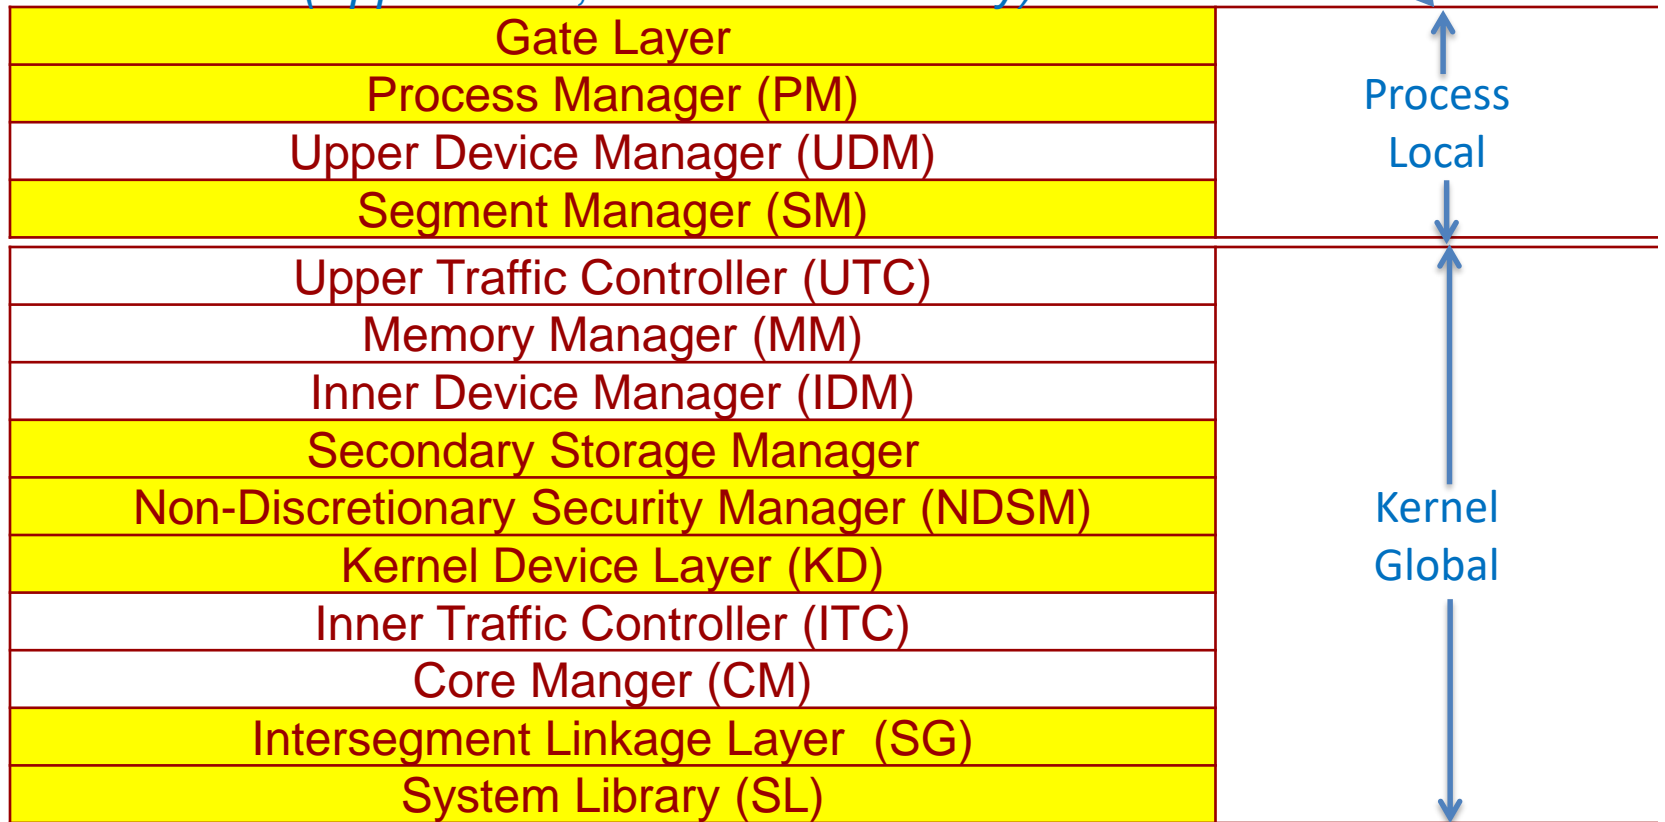


- Segment make known is illustrative example
 - All modules in call chain are strictly layered
- Top gate layer is kernel API – implements FSPM
 - Receives call and completes entry to Ring 0
 - Parameters copied from outer ring to Ring 0
 - Entry point call to next layer
- Process-local modules at the top
 - Their “information hiding” data bases are per process
 - Code is shared with same PLSN by all processes
- Kernel-global modules
 - Kernel API “effects” reflect data all processes share

GEMSOS Make Known Example



(Applications, Kernel Gate Library)



(Hardware)

Secure System Design & Development



- Architectural considerations (Gasser chapter 11)
 - Applies to development of security kernels
 - As well as to their applications
- Operating-system layering
 - Promote structured design for kernel assurance
 - Operating systems services on kernel

Operating System Layering Strategies

- New operating system: layer on security kernel
 - Simple mapping of O/S features to SK features
 - Distinctive is lack of backward compatibility
- Compatible operating system (emulation)
 - Emulate insecure operating system (ISOS)
 - Typically emulator runs in each process
 - Renders O/S calls into kernel calls
- Identical operating system (virtual machine)
 - Provides isolation, but not sharing, of memory
 - Kernel is virtual machine monitor (VMM)
 - Principal “objects” are virtual disks, not files
 - Subjects – kernel users and VMs

Secure System Design & Development

- Architectural considerations (Gasser chapter 11)
 - Applies to development of security kernels
 - As well as to their applications
- Operating-system layering
 - Promote structured design for kernel assurance
 - Operating systems services on kernel
- Asynchronous attacks and argument validation

Cross-Domain Asynchronous Attacks



- Hardware support for cross-domain validation
 - Pointer validation is particularly challenging
- Multiprocessor and multiprogramming issues
 - Multiple processes can access pointer argument
 - Time of check/time of use (TOC/TOU) problem
 - Safest to copy parameters to new domain before use
- OS must prevent changes to validation data
 - There are no generic solutions
 - May require appropriate locks inside OS
 - May require total atomic copy of data
 - Kernel support for this is valuable aid
- Examine I/O operations: are also asynchronous

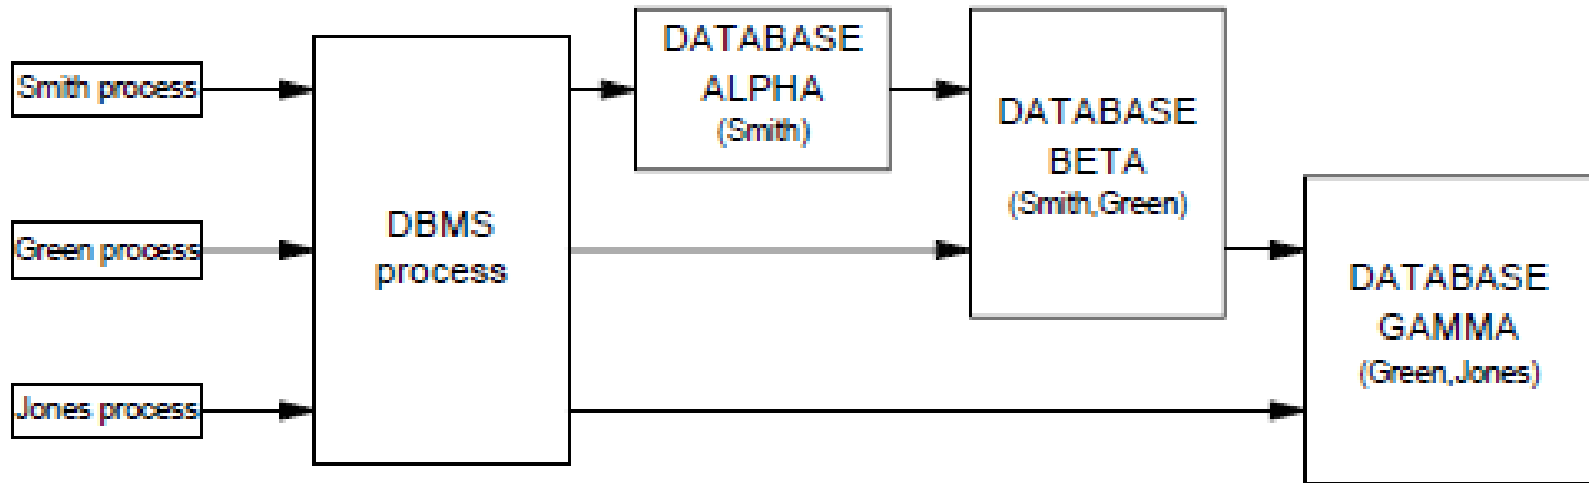
Synchronization for a Trusted System

- Useful operating system needs synchronization
 - Is usual and customary service applications expect
- Need synchronization across access classes
 - RVM must insure information flow meets MAC policy
- Mutexes and semaphores imply shared object
 - Read and write make not secure across access levels
- Use alternative NOT based on mutual exclusion
- Two kinds of objects are defined for computation
 - **Eventcount** to signal and observe progress
 - Primitives **advance(E)**, **read(E)**, **await(E, v)**
 - **Sequencer** to assign an order to events occurring

Secure System Design & Development

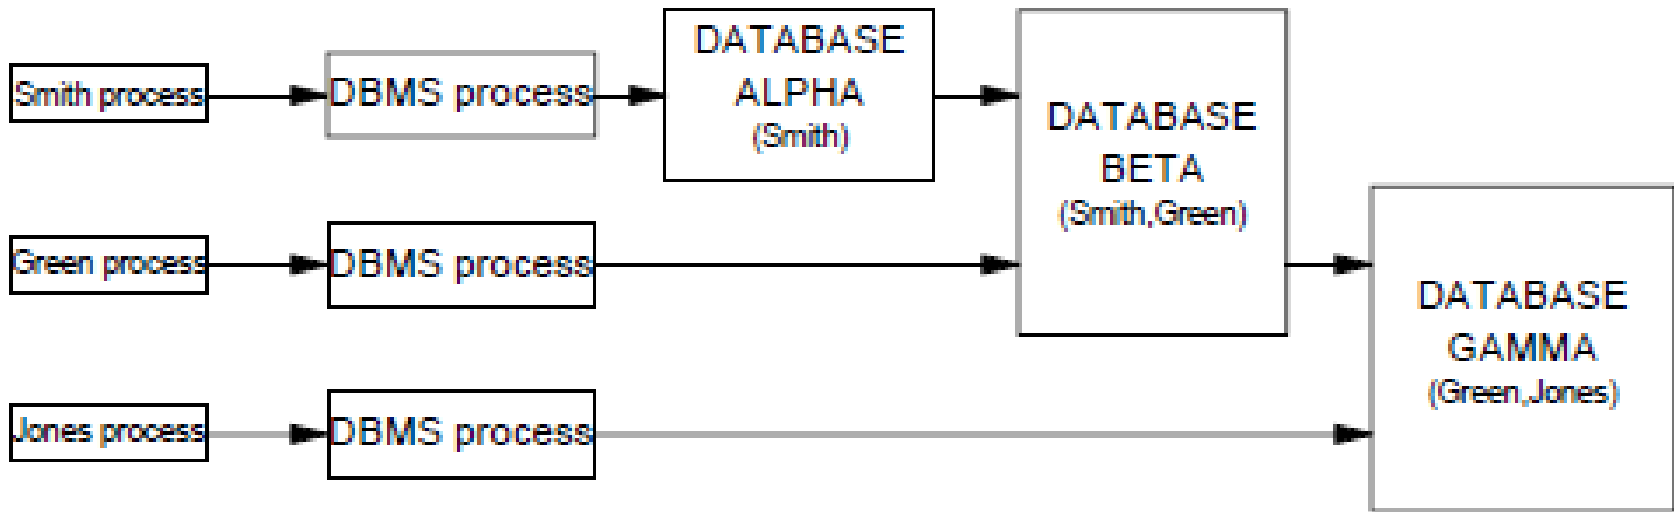
- Architectural considerations (Gasser chapter 11)
 - Applies to development of security kernels
 - As well as to their applications
- Operating-system layering
 - Promote structured design for kernel assurance
 - Operating systems services on kernel
- Asynchronous attacks and argument validation
- Protected subsystems
 - Static process, on-demand process, multiple domains

Protected Subsystem in Active Process



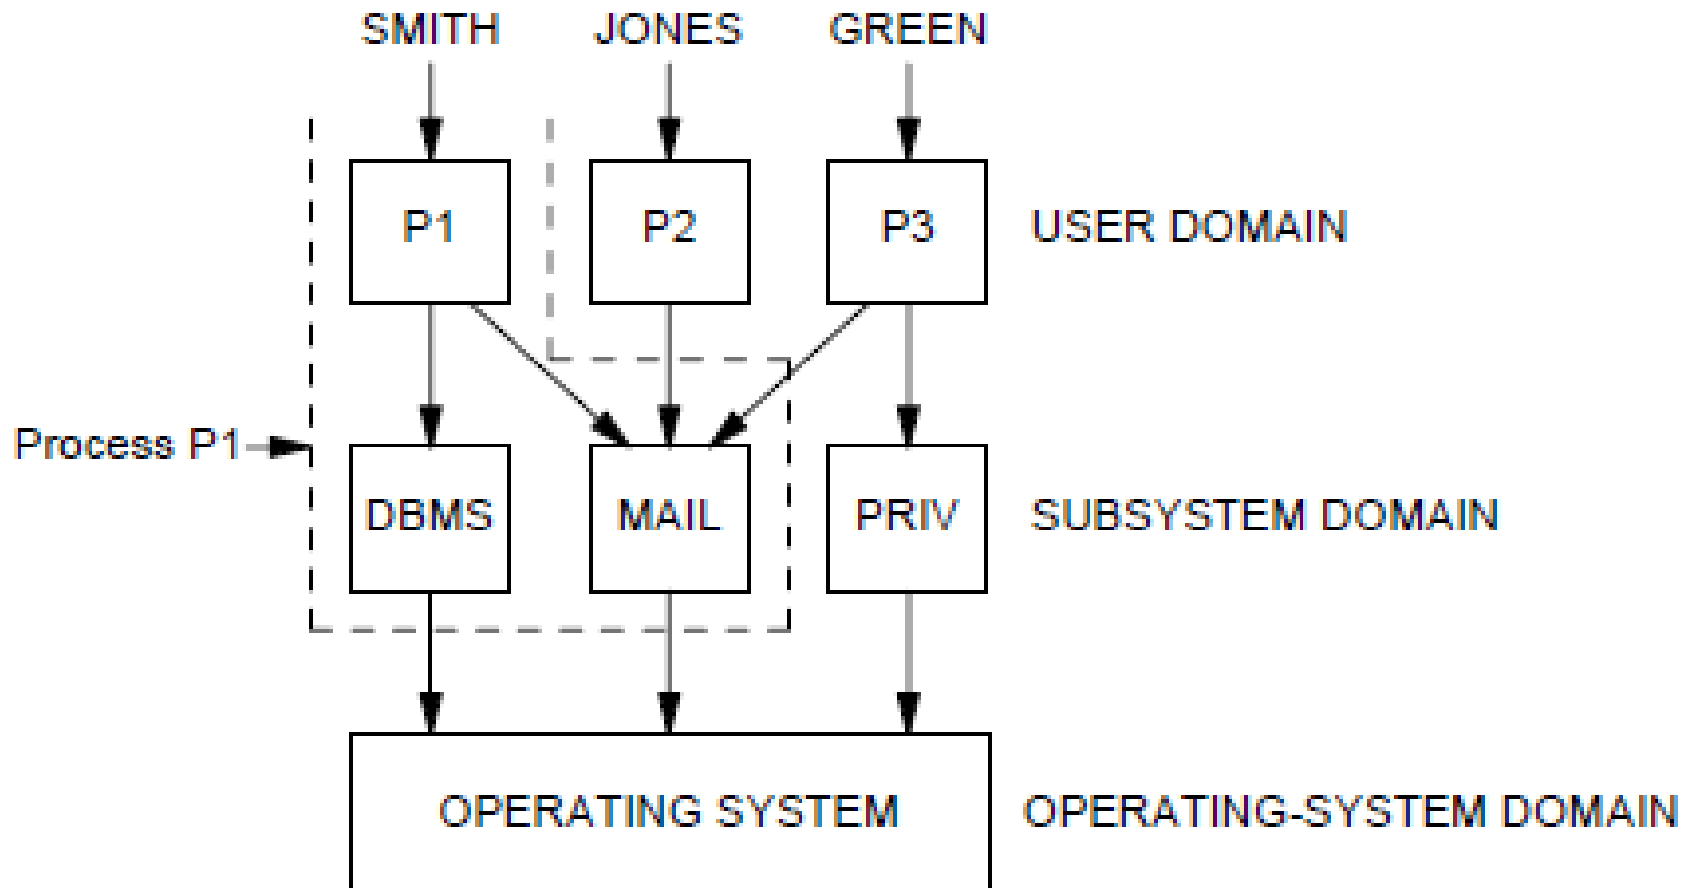
- **DBMS runs as a process**
 - Inherently no different from a user process
 - Normal OS access controls limit access to DBMS
- **DBMS must control individual users' access**
 - To different files and different portions of files

Protected Subsystem on Request



- Subsystem activated as a separate process
 - Each time it is needed
- While retaining its own identity
 - Separate from that of the invoking process

Mutually Suspicious Subsystems



Management of SK Rings and Labels



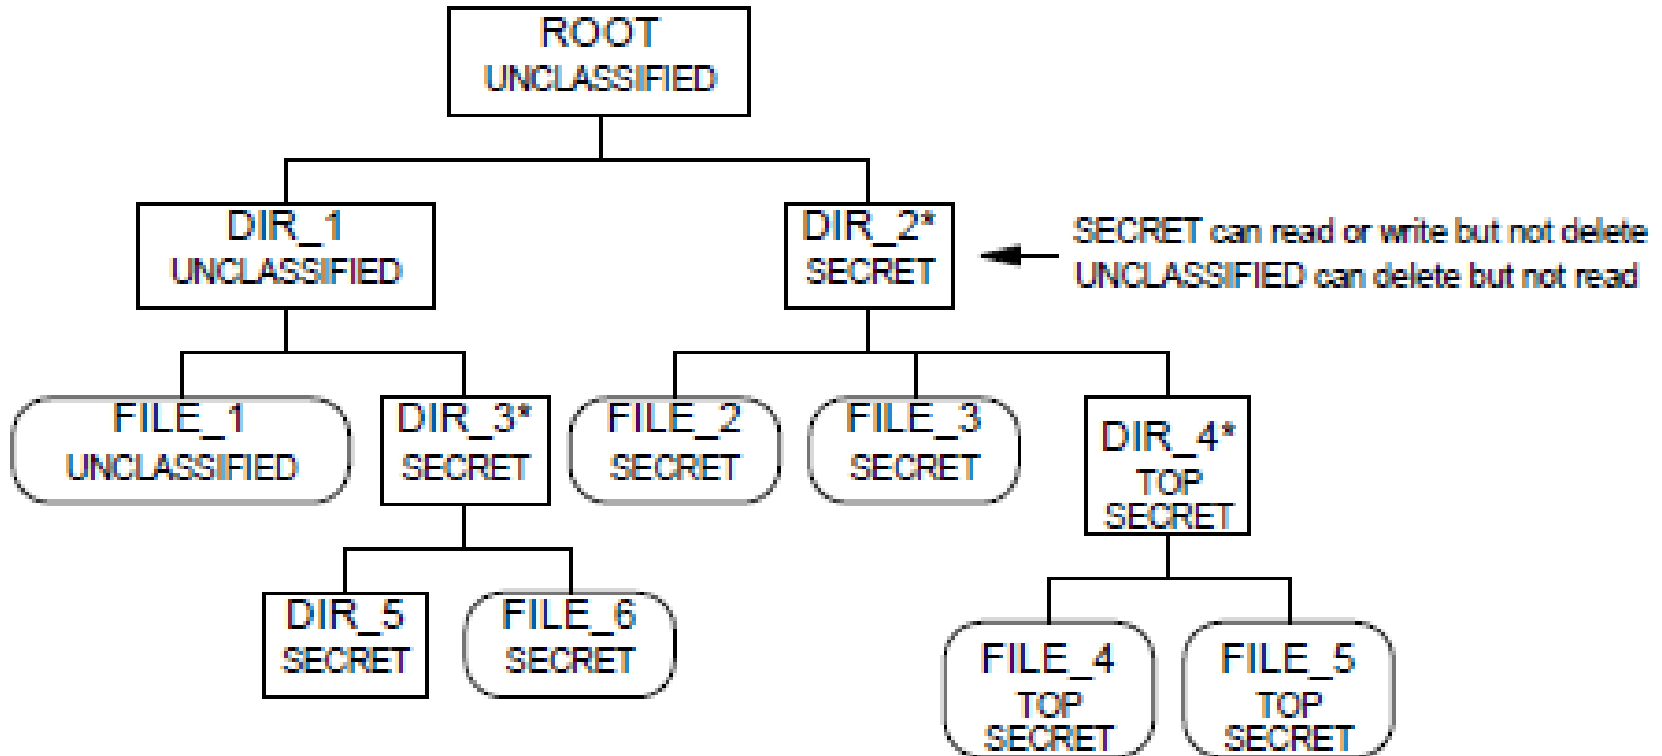
- For a system, privileged services bound to ring
 - Static binding for given system architecture
 - Kernel is permanently bound to PL0
- Ring bracket (RB) associates object with domain
 - RB encoded in 3 ring numbers (RB1, RB2, RB3)
- General trusted system has at least 3 domains
 - Kernel, operating system, applications
- Non-discretionary is mandatory policy, i.e., MAC
- Each can be represented by access class label
 - Labels can be compared by dominance relation
 - Combine confidentiality, integrity, dominance domain

Secure System Design & Development

- Architectural considerations (Gasser chapter 11)
 - Applies to development of security kernels
 - As well as to their applications
- Operating-system layering
 - Promote structured design for kernel assurance
 - Operating systems services on kernel
- Asynchronous attacks and argument validation
- Protected subsystems
 - Static process, on-demand process, multiple domains
- Secure file systems
 - Alternate naming structures; unique identifiers



MLS Hierarchical File System



* Upgraded Directories

Security Kernel Objects



- Minimization of kernel
 - “Economy of mechanism”
 - “Significantly more complicated OS outside kernel”
 - Implies kernel cannot be compatible with insecure O/S
- All subjects need system-wide name for objects
 - Each subject must be able to identify shared object
 - “Flat” naming is classic covert channel example
- Object hierarchy naming
 - BLP hierarchy with “compatibility” meets need
 - Biba “inverse compatibility” for integrity needed
- Least common mechanism drives reuse
 - OS creates “directories” out of objects from kernel

Security Kernel Support for Segments

- Segmented instruction execution
 - Enforced for code is executing outside the kernel
- All memory addresses a pair of integers [s, i]
 - "s" is called the segment number;
 - "i" the index within the segment.
- Segment number is process local name (PLSN)
- Systems have similar kernel API to add segment
 - Kernel is invoked to “make known” a new segment
- Descriptor table defines process address space
 - Is a list of all the segments CPU can address
 - Must include code segment and stack segment

Secure System Architecture Summary

- Architectural considerations (Gasser chapter 11)
 - Applies to development of security kernels
 - As well as to their applications
- Operating-system layering
 - Promote structured design for kernel assurance
 - Operating systems services on kernel
- Asynchronous attacks and argument validation
- Protected subsystems
 - Static process, on-demand process, multiple domains
- Secure file systems
 - Alternate naming structures; unique identifiers



INF523

Introduction to MLS File System: GARNETS Case Study

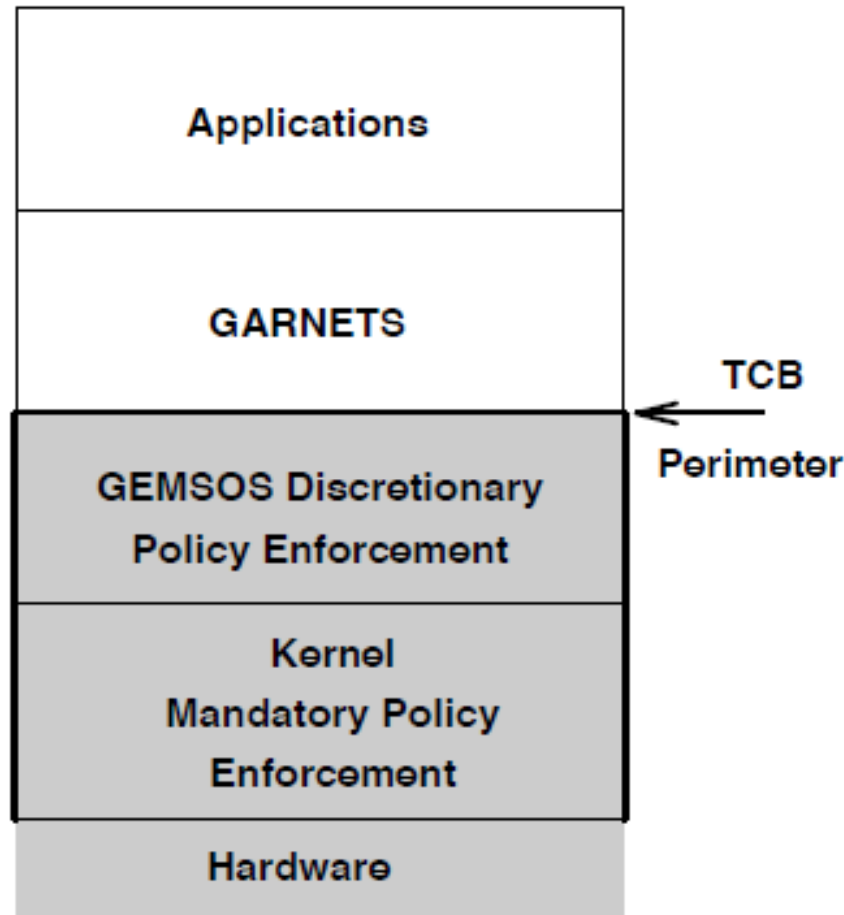
Professor Clifford Neuman

Lecture 12

GARNETS Example on Security Kernel

- Case study of broad application on top of TCB
- Security kernel TCB minimization severely limits
 - Can present only a primitive interface
- Lacks typical OS rich variety of functions
 - Argument that high assurance is “unusable”
- MAC enforcement constrains untrusted subjects
 - Argue renders application development “impossible”
- Analysis of GARNETS operating system
 - Gemini Application Resource and Network Support
 - Uses only TCB mechanism to provide interface
 - Interface is “friendly” and flexible

Standard GARNETS Architecture



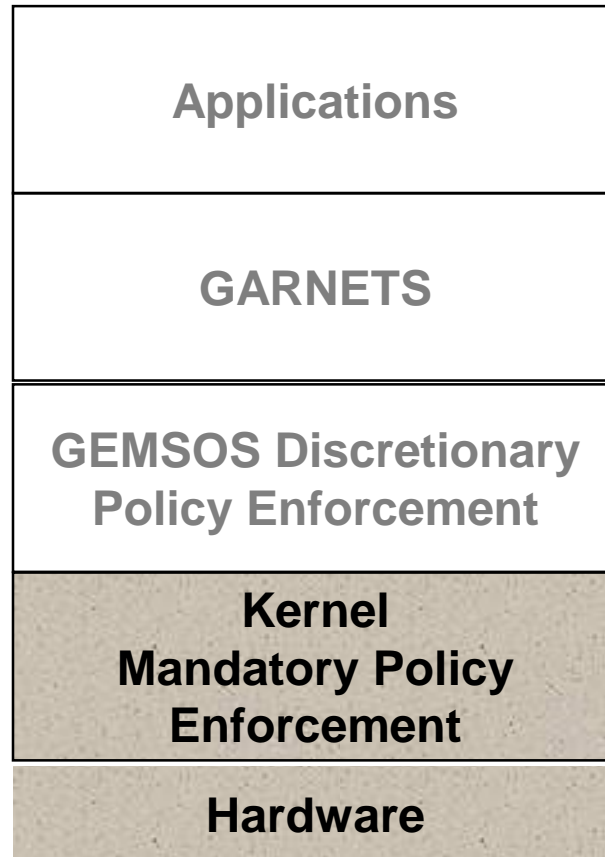


Design Objectives

1. General purpose file system interfaces
 - Permit application libraries to be ported to interface
2. Both MAC and DAC exclusively from TCB
 - DAC subject on top of kernel provides “strong DAC”
3. File system is multilevel
 - Managed by single-level subjects
4. All file system operations are atomic
5. No read locks are used
 - Applications with can “read down” subject to DAC
6. Application access only GARNETS file system
7. GARNETS itself designed to meet Class B2



Overview of GARNETS Architecture






Kernel Exports Segments

- Segmented instruction execution
 - Enforced for code executing outside the kernel
- CPU memory addresses are pair of integers [s, i]
 - "s" is called the segment number;
 - "i" the index within the segment.
- Segment number is process local name (PLSN)
- TCB has API similar to kernel to add segment
 - Kernel is invoked to “make known” a new segment
- Descriptor table defines process address space
 - Is a list of all the segments CPU can address
 - DAC subject needs code segment and stack segment

Domains in GARNETS Architecture



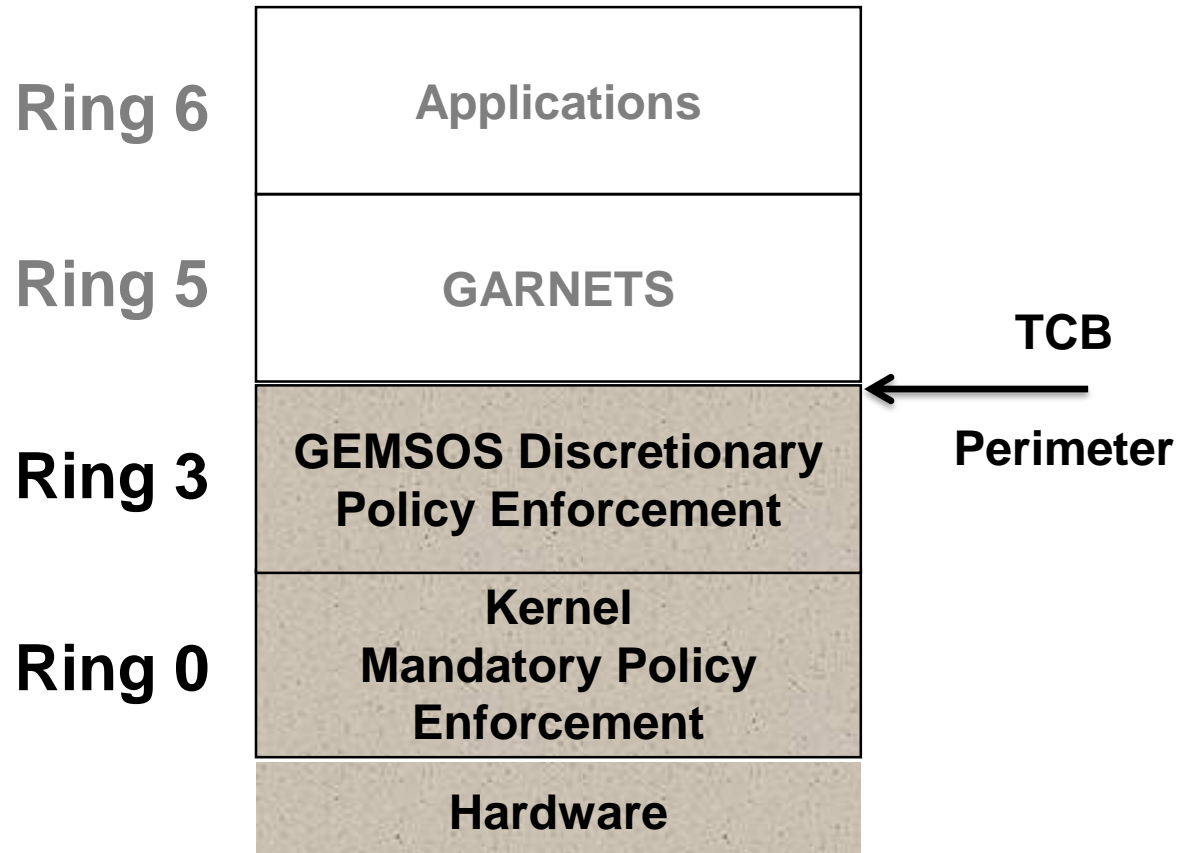
Kernel Creates (Domains) Rings



- For a system, privileged services bound to ring
 - Static binding for given system architecture
 - Kernel is permanently bound to PL0
- Ring bracket (RB) associates object with domain
 - RB encoded in 3 ring numbers (RB1, RB2, RB3)
- General trusted system has at least 3 domains
 - Kernel, operating system, applications
- Non-discretionary is mandatory policy, i.e., MAC
- Each subject represented by access class labels
 - Labels can be compared by dominance relation
 - Combine confidentiality, integrity, dominance domain



DAC TCB in GARNETS Architecture

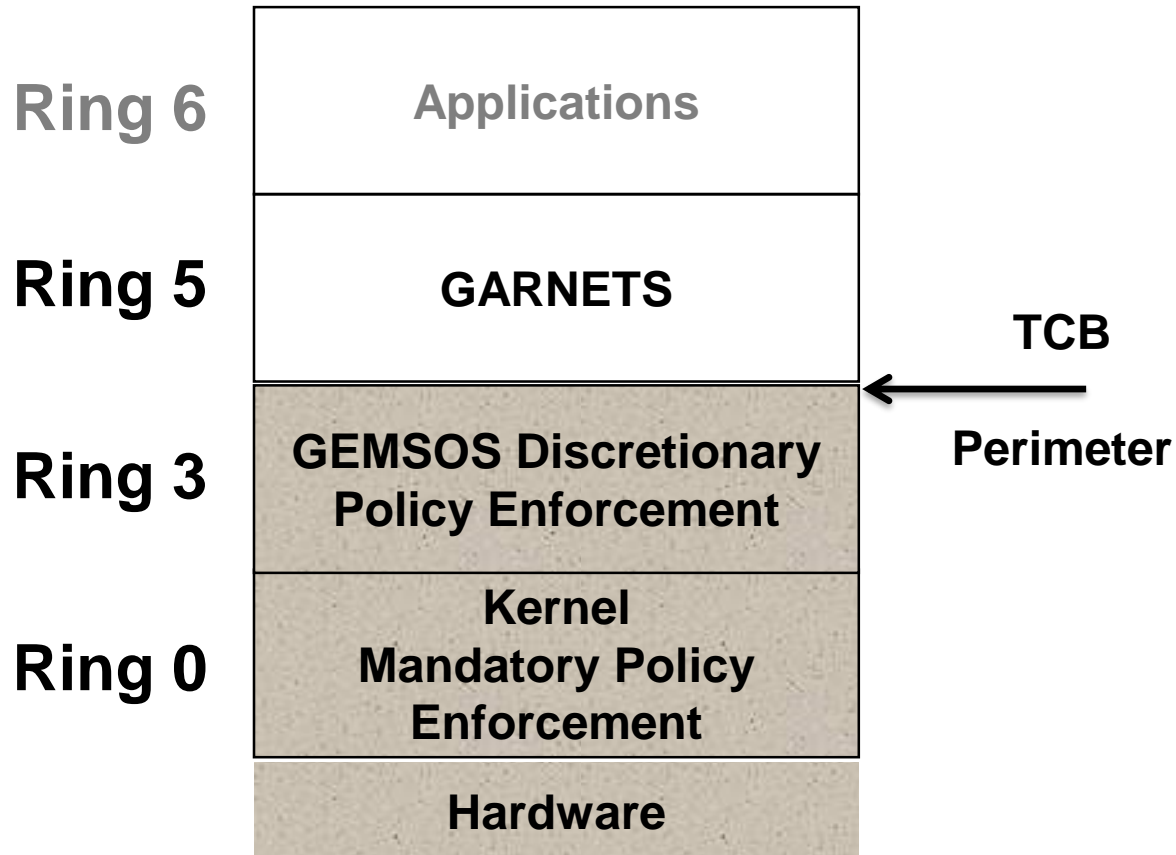


DAC TCB Exports DACLS & Msegs



- Segments
 - Fundamental storage object
 - Loci of mandatory control
 - Processes simultaneously and independently share
- DAC Access Control Lists (DACLS)
 - A segment containing limited number of ACLs
 - Interpretively accessed object exported by TCB
 - Building block for GARNETS directories
- Multisegments (“msegs”) exported by DAC TCB
 - Collection of zero or more segments
 - Segments are accessible only as elements of msegs
 - All its segments hierarchically related to single base

GARNETS in the Architecture



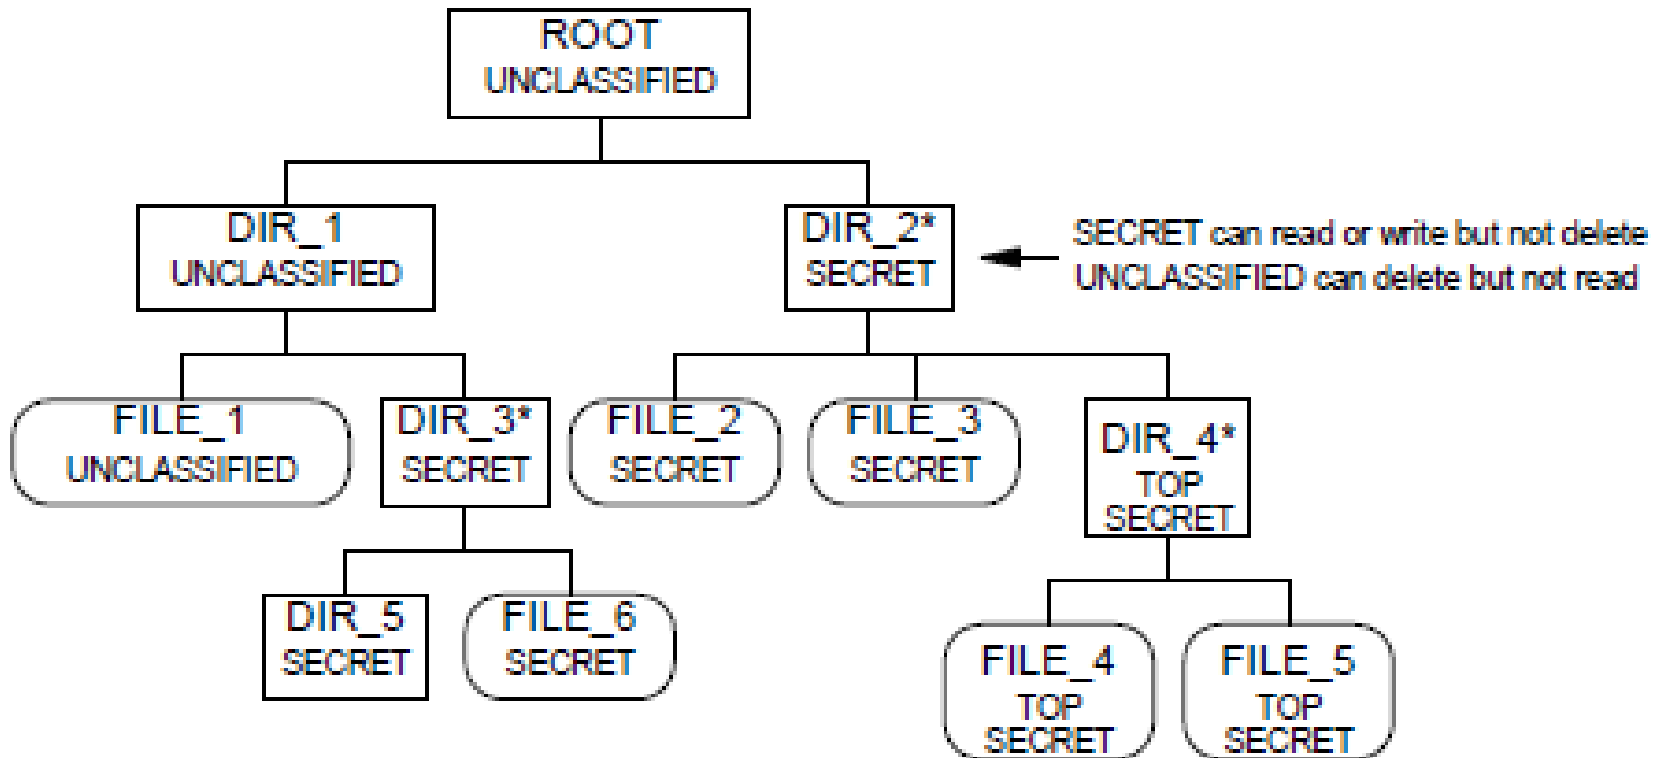
GARNETS Creates Files & Directories



- Directories
 - Rely only on TCB (not GARNETS) for access control
 - DAC access to an object based only on its ACL
- Named multisegments, distinct from files
 - Are namable directory entries
 - Its segments directly accessed via hardware
- Files interpretively access by applications
- File system built from three parallel trees
 - Directory tree of DACLs with mseg for dynamic data
 - File tree of DACLs which host file mseg for each file
 - Huge mseg with segments that mirror directory tree



Gasser: MLS Hierarchical File System



* Upgraded Directories

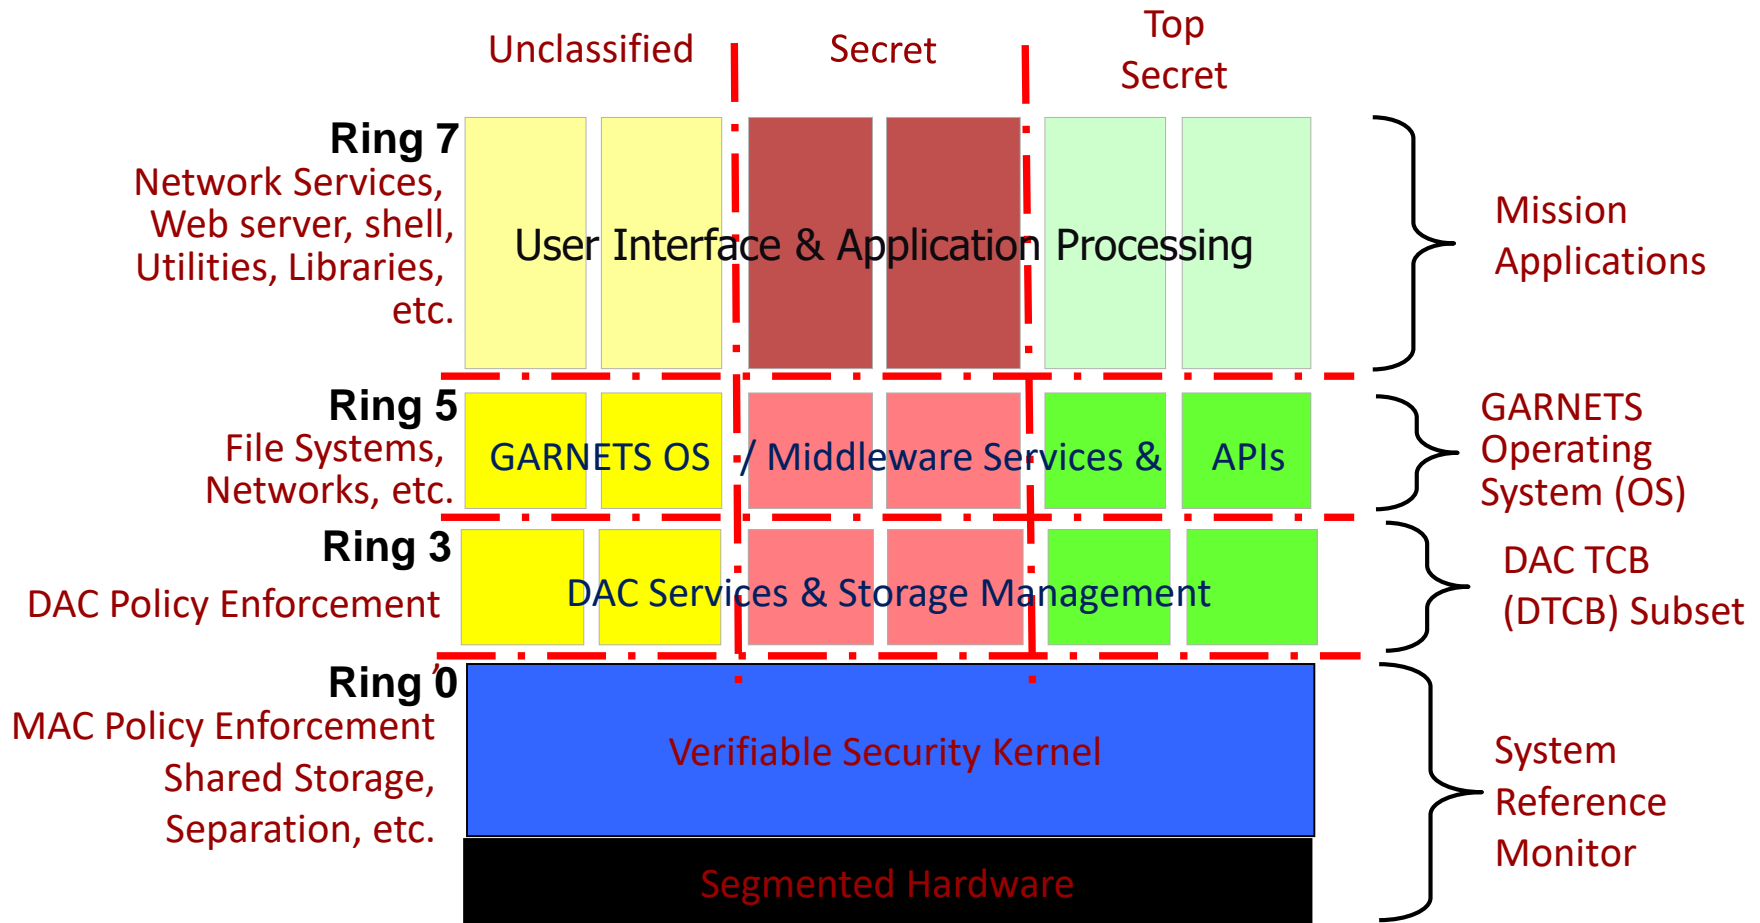


Directory Properties

- Single-level directories
 - Contains information all at one access class
 - Subdirectory creation information is in parent
 - Names and creation dates
 - Visible to parent-level subjects
- Upgraded directories
 - Kernel forces compatibility property to be met
 - Dynamic information in upgraded directory itself
 - Time of last modification and directory contents
 - Visible only at the upgraded access class



TCB Subsets for GARNETS



Components of GARNETS Directory



-
- GM – Huge mseg gives directory tree roadmap

GARNETS Directory Structure






Components of GARNETS Directory

- GM – Huge mseg gives directory tree roadmap
- DTD – Directory Tree DACL
 - Control access to tree from which directories are built
 - ACLs for directory entries

Directory Tree DACL Component

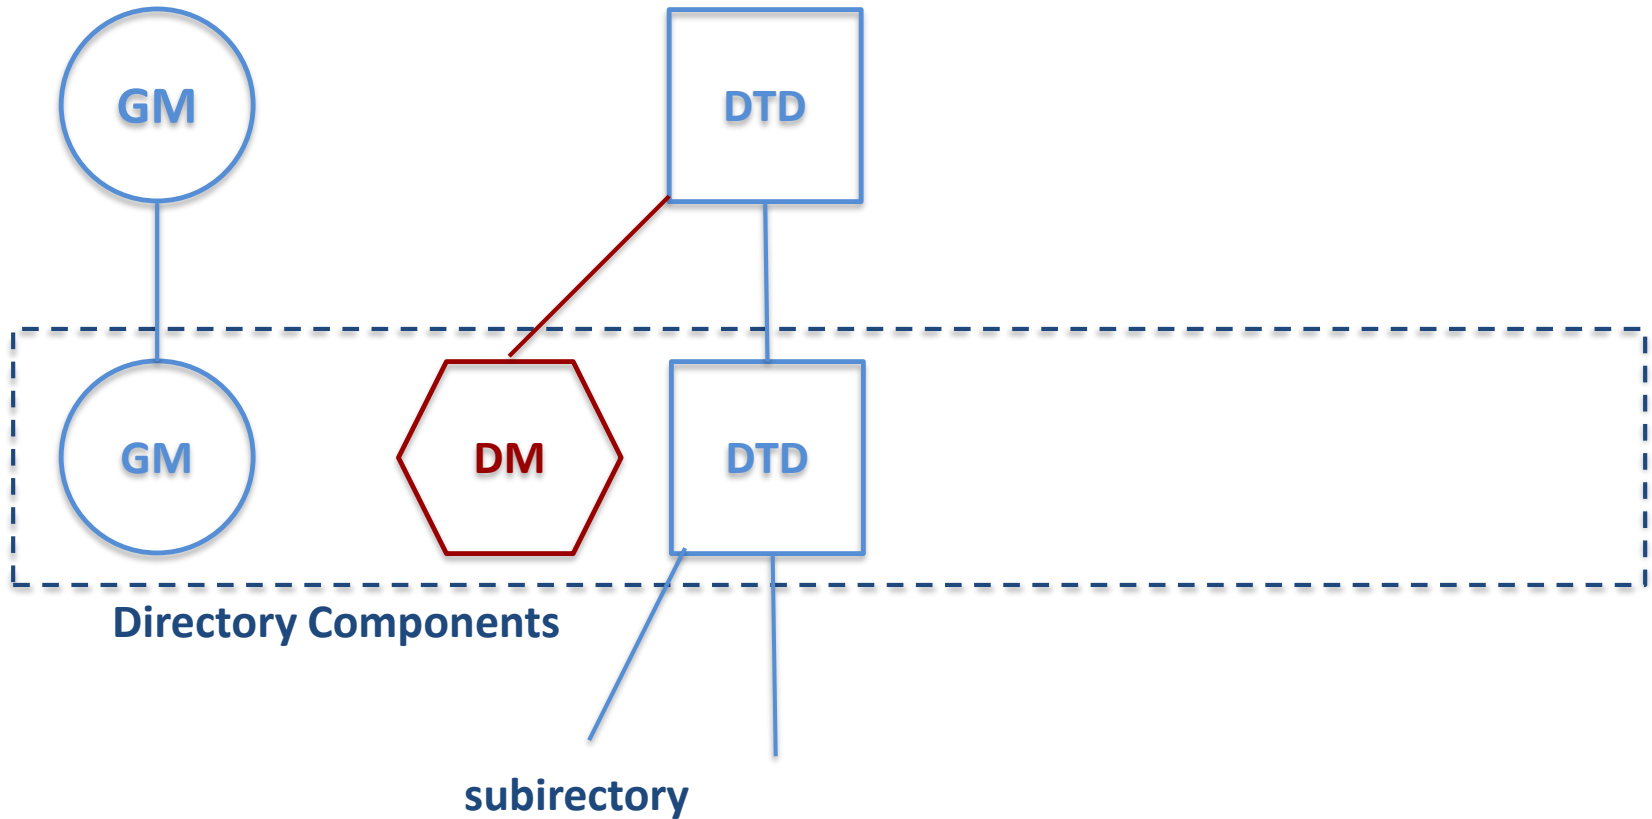




Components of GARNETS Directory

- GM – Huge mseg gives directory tree roadmap
- DTD – Directory Tree DACL
 - Control access to tree from which directories are built
 - ACLs for directory entries
- DM – Directory Multisegment
 - Dynamic data for directory entries

Directory Multisegment Component






Components of GARNETS Directory

- GM – Huge mseg gives directory tree roadmap
- DTD – Directory Tree DACL
 - Control access to tree from which directories are built
 - ACLs for directory entries
- DM – Directory Multisegment
 - Dynamic data for directory entries
- **FTD – File Tree DACL**
 - Used to extend the tree



File Tree DACL Component



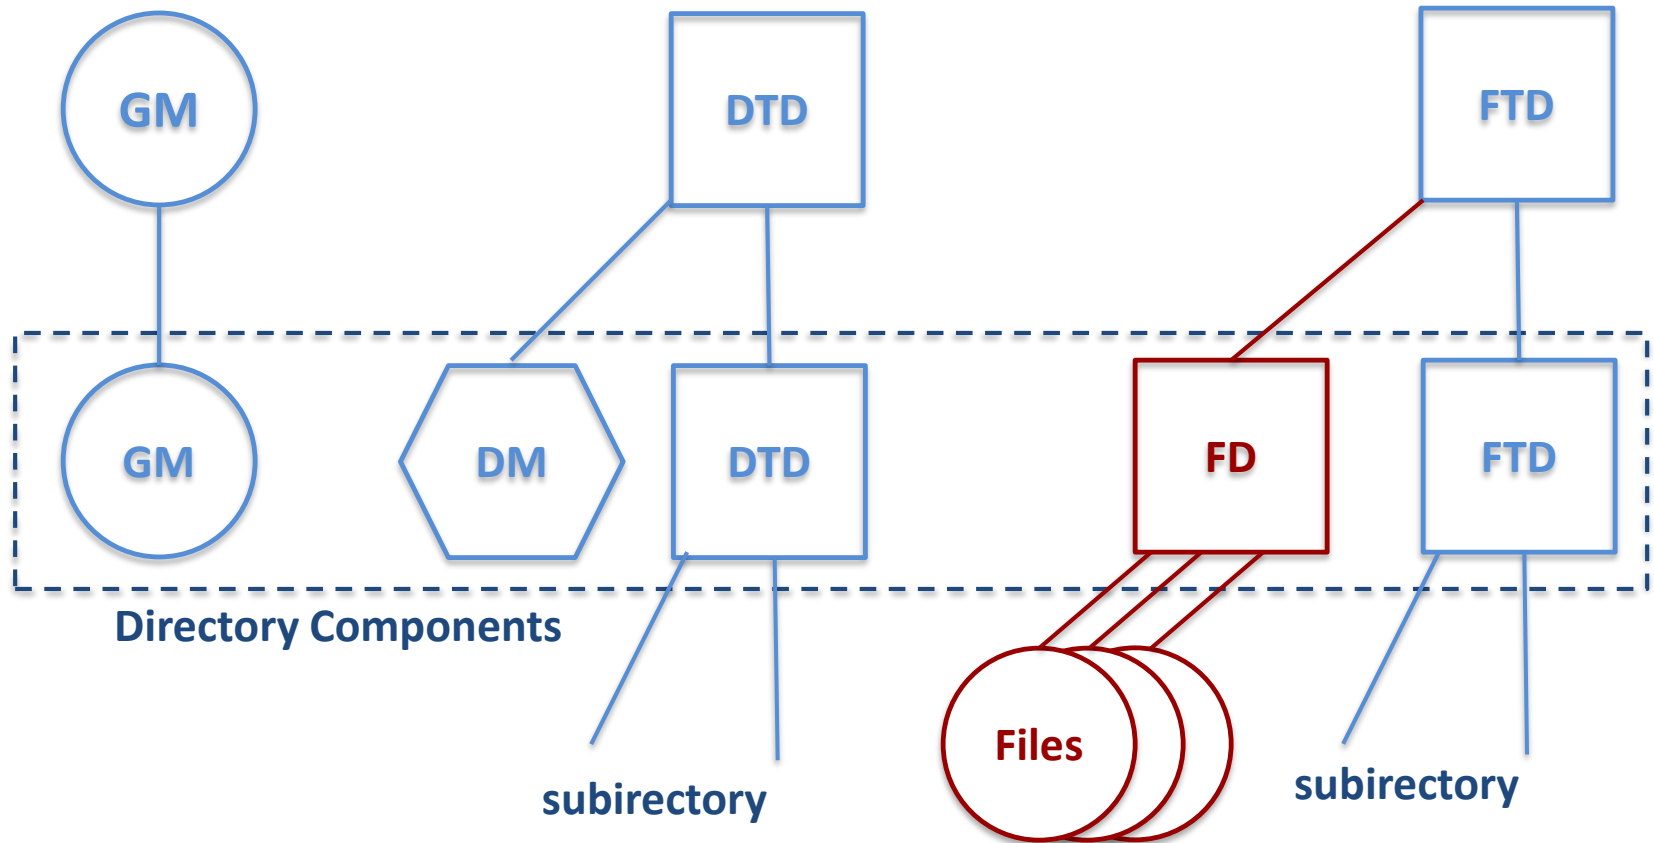


Components of GARNETS Directory

- GM – Huge mseg gives directory tree roadmap
- DTD – Directory Tree DACL
 - Control access to tree from which directories are built
 - ACLs for directory entries
- DM – Directory Multisegment
 - Dynamic data for directory entries
- FTD – File Tree DACL
 - Used to extend the tree
- FD – File DACL
 - ACLs for file entries in the directory



File DACL Component






Components of GARNETS Directory

- GM – Huge mseg gives directory tree roadmap
- DTD – Directory Tree DACL
 - Control access to tree from which directories are built
 - ACLs for directory entries
- DM – Directory Multisegment
 - Dynamic data for directory entries
- FTD – File Tree DACL
 - Used to extend the tree
- FD – File DACL
 - ACLs for file entries in the directory



GARNETS Directory Structure

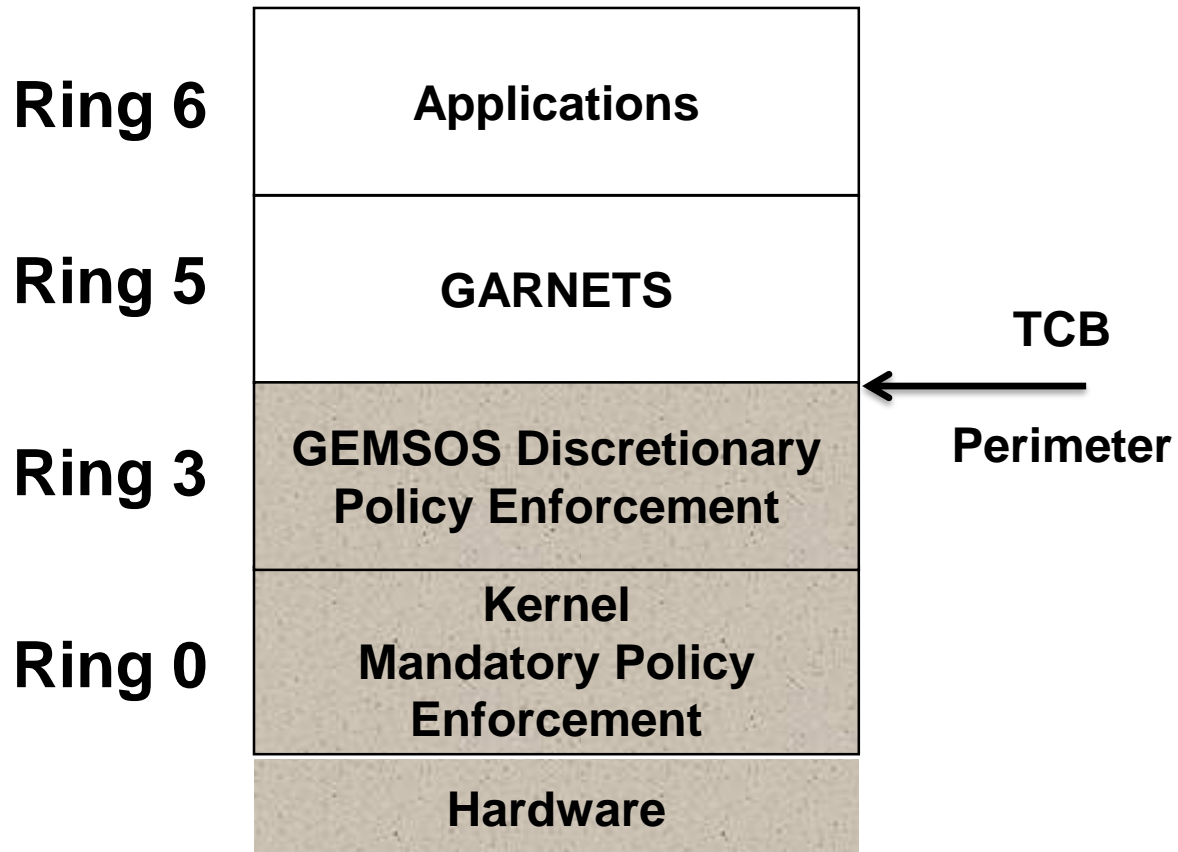


Management of Upgraded Directories



- Initialization is exported to GARNETS interface
 - Must be done by subject at upgraded access class
 - Cannot be done at access class of parent directory
 - For uniformity is same for both normal and upgraded
- Implication for deletion of upgraded directories
 - To meet kernel restriction will require trusted subject
 - Need not have entire range of system's access classes
 - Range encompasses parent and upgraded
- To limit deletion, GARNETS limits creation
 - Users are required to have a “special authorization”
 - Some environments might prohibit user creation

Review of GARNETS Architecture



Named Multisegments (Msegs)



- Msegs are namable directory entries
 - A hierarchically structured collection of segments
 - Single base segment ACL, for all segments in mseg
 - Each segment has an explicit access class
- In contrast to files, not interpretively accessed
 - Segments accessed directly by available hardware
 - Segments are included in process address space
- Msegs used to contain GARNETS internal data
 - Must be protected from less-privileged subjects
 - Uses GEMSOS ring mechanisms to insure integrity
 - Uses DACLs from DTCB to protect internal data



Benefits of Named Msegs

- Avoid unnecessary buffering
 - No per-process file buffer for data
 - Code is not copied from file into executable segments
 - Code is store in right size segments, executed directly
 - Code is not modifiable, so many process can execute
- Promotes sharing of executables
 - Reduced use of real random-access memory
 - Reduced swapping increases performance
- Direct hardware access reduces context switch
- Application can use for databases and libraries
- Highly efficient IPC and synchronization



Single Level Files

- Interpretively accessed as GARNETS interface
 - Applications make calls for file operations
 - Inside GARNETS are created from segments
- GARNETS manages each individual file
 - One ACL is associated with each file
 - Maintains attributes, e.g., time of last modification
 - Time of last read only updated for same level subjects
- Design rejected multilevel files
 - Careful application design eliminates most needs
 - Would create incoherent interface
 - Not clear how to avoid multilevel (trusted) subjects



The Gizillion Problem

- Problem of very large number of access classes
 - Must be addressed for flexible untrusted applications
 - Potential many access classes in underlying TCB
 - GEMSOS: two sets of 16 levels and 96 categories
- TCB minimization limits complex data structures
 - Objective is avoiding elaborate constructs
 - GEMSOS provides **one** base object per access class
 - Each access class must construct its own data
- Handle previously unencountered access class
 - GARNETS subject must create data for applications
 - At minimum OS creates application stack at level

Alternatives for New Access Classes



- GARNETS administrator creates data structures
 - Users requests directory at the new access class
 - New upgraded directory below system low root
 - File system data structures at each new access class
 - Administrator requires access to full range of classes
 - Depends on timely response by administrators
- Create all possible access classes a priori
 - A base directory is always available when needed
 - BUT is untenable to created a gizillion bases
- Trusted process to automate creation process
 - Designers would fail to meet “untrusted” objective
 - Far too complex to meet high assurance

GARNETS Gizillion Problem Strategy

- DTCB creates DACL at first occurrence of class
 - DTCB has function to locate that DACL
 - Has “access class to path” algorithm to base segment
- At first occurrence GARNETS builds a directory
 - Per Access Class (PAC) directory at new class
 - Location for application “home” directory
- GARNETS can then support non-TCB subjects
- TCB tools install GARNETS bootstrapping code
 - Code not located in GARNETS file system.
 - In separate data structures at predefined location



File System Object Naming

- Alias names for objects supported by GARNETS
 - All names must be deleted before object deleted
- Symbolic links are path to target object
 - TCB prevents creation of hard links
 - Can have links to files and named megs
 - Can have links to directories and other links
- Existence of intervening links invisible on access
 - Cycles controlled by number of links traversed in path
- Per Access Class (PAC) links
 - Link has a field for the access class
 - GARNETS finds PAC directory for access class

Leveraging Gizillion Solutions



- Supports use of single-level volumes
 - Single level file systems distributed on volumes
 - Symbolic links permit binding to multilevel structures
- Volumes transparent to application data access
 - Volume access class range implies covert channels
 - Volume range simplifies physical control of media
- GARNETS supports working directories
 - Simplifies naming of subordinate objects
 - Multiple working directory employed for volumes

GARNETS Self-protection



- GARNETS uses rings properly to be effective
 - Applications operate in a less privileged domain
 - Interpretively access objects protected, e.g., files
 - Internal data structures protected from applications
- GARNETS ring brackets
 - Some directories are dedicated to use by GARNETS
 - Range of rings of subjects that will be granted access
 - Apply to all objects in a directory
 - Permanently set when directory is created

Consistency and Concurrency



- File consistency
 - Used to address discontinuities in operation
 - Permit fine-grained robustness selection
- File System Concurrency Control
 - Doesn't insure total ordering of file system operations
 - Each file system object has a version number
- Leverages TCB primitive for atomic updates
 - Avoids conflict with real-time properties
 - Strict two-phase commit for directory components
 - Kernel API can atomically update doubly threaded list

Summary of MLS Implications




- GARNETS file system on high assurance TCB
 - Represents complex general-purpose application
 - Untrusted implementation is MLS context
 - Sufficiently flexible for broad spectrum of uses
- File system managed by single-level subjects
 - Leverage symbolic links
 - Solution to gizillion problem
 - Employable in single-level volume configurations
 - Permits upgraded directories and multilevel msecs
- GARNET protects itself and its data structures
 - Exploits rings and DACLs from high assurance TCB



INF523

Subversion Case Study NFS

Professor Clifford Neuman





Network File Service (NFS) Security

- Case study of NFS subversion demonstration
 - Running example by US Navy masters student
 - Emory A. Anderson, III, for Prof Cynthia Irvine (NPS)
 - Shown to Richard Clarke, “first cybersecurity czar”
- First, consider security implications for system
 - How deeply rooted are adverse consequences
- Second, explore applicability to other systems
 - Address whether attack approach is limited to NFS
 - Briefly examine Anderson SSL subversion design
- Follow on – Later NFS case study of mitigation
 - Compare to Anderson recommended solution

Likely Tool of Professional Attacker



- Subversion is technique of choice [And 1.D]
 - Professional distinguished from amateur
- A primary objective is avoiding detection
 - Amateur often motivated by desire for notoriety
- Professional often well-funded
 - Resources to research and test in closed environment
 - Amateur tends to numerous attempts on live target
 - Flawless execution reduces risk of detection
- Coordinated attacks are mark of a professional
- Professional will invest and be patient to use
 - Subverter is likely different than attacker

Demonstration of Subversion

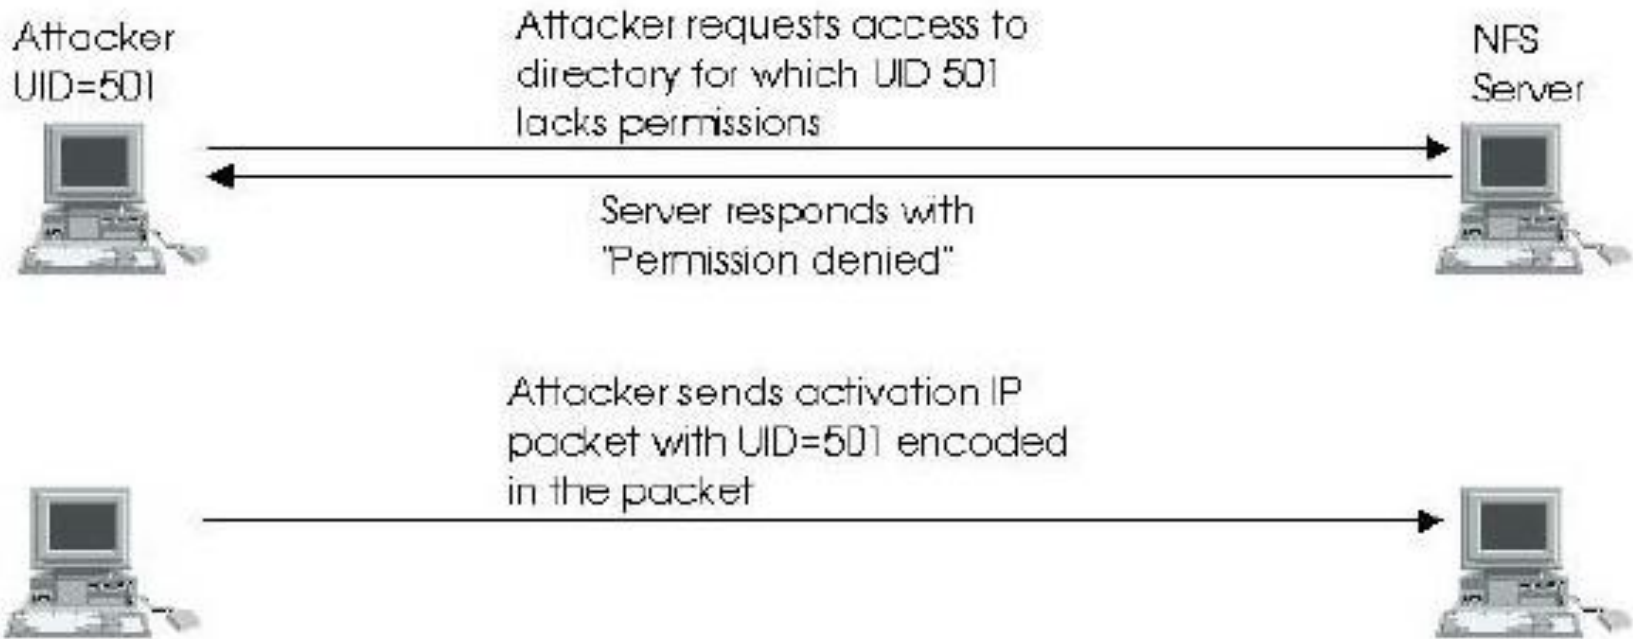


- Obfuscation of artifice not given serious attention
 - Would be of utmost importance to professional attack
- Subversion can occur multiple points in lifecycle
- Selected distribution phase for demonstration
 - Driven by limited resources and access of student
 - Facilitated by NFS on open source Linux system
 - Representative of attacker mirror site opportunities
- Closed source not daunting for professional
 - May involve reverse engineering application
 - Might create a binary patch to insert in binaries
 - Entirely within anticipated professional resources

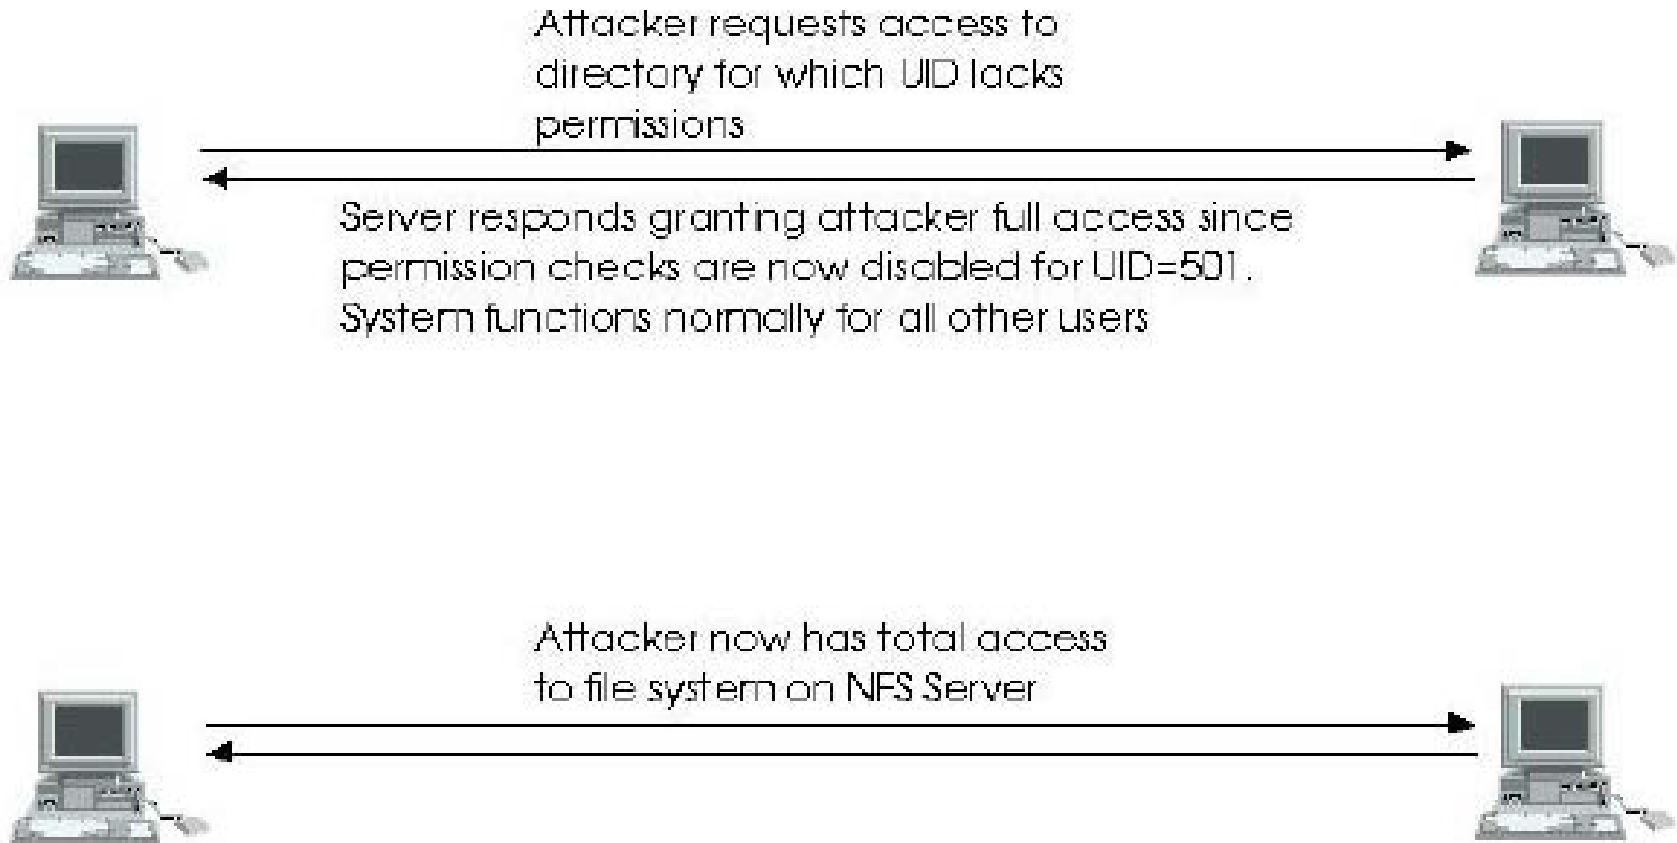
Choice of NFS as Suitable Application

- For impact, need readily apparent significance
 - NFS is application familiar to typical IT user
 - Users understand notion of need to protect data
- Activation needs to be straightforward
 - Network interface chosen for ease of explanation
 - Internet technology is widely used
- Choose to have remote activation
 - Representative of low risk for attacker
 - Also supports local activation, e.g., via loopback
 - Trigger is a malformed Internet packet
- **Study of subversion method benefits student**

Case System and Activate the Artifice



Attacker Uses Artifice for NFS Access





End Session by Deactivating Artifice



Attacker sends deactivation IP packet. Server returns to normal operation for all users



Design Properties of NFS Artifice



- Purpose of artifice to bypass file permissions
 - Bypass check for a specified user at will
 - Then re-enable the normal system operation
- Exhibits all the characteristics of subversion
 - Exception was no attempt for hide or obfuscate
- Artifice is small – eleven C statements
 - Small in relation to millions LOC in Linux kernel
 - Unlikely to be notice by those in Linux development
- Can be activated and deactivated
 - Further complicates attempts to discover existence
- Does not depend on activities of a system user




Artifice Functions


- Composed of two parts in two unrelated areas
- Subvert a portion of kernel that receives packets
 - Recognize a packet with distinguishing characteristics
 - Activation based on trigger known only to subverter
 - Extends normal check for packet checksum
- Activation recorded in global variable in kernel
- Subverts Linux file system permission checks
 - Check global kernel variable to see if activated
 - Grants attacker access to any file in the system
 - Bypass behavior limited to specified user ID
 - System functions normally for all other users



Artifice Activation



Subverted File Permission Checks





Separate Design of SSL Subversion

- Secure Sockets Layer (SSL) widespread use
 - Secure communications between client and server
 - Client and server negotiate session keys
 - Encrypt traffic using symmetric encryption algorithm
- Options available to attacker for subversion
 - Duplicate all communications and send to attacker
 - Weaken key generation mechanism – limit entropy
 - Simply send the session keys out to the attacker
- Advantages of exfiltrating session keys
 - Attacker is passive and maintains anonymity.
 - Subverting either client or server gives total access

NFS Subversion Technical Conclusions



- Practice for showing security inadequate at best
 - Penetration tests and add-on third party products
 - Layered defenses and security patches irrational
- Bad defense more dangerous than poor security
 - Leads to flawed belief system has adequate security
 - Can increase risk by more dependence on defense
- Have technology to provide appropriate security
 - Evaluation criteria tried and tested
 - These approaches have fallen into disfavor
- The need to address subversion is increasing
 - Threat sources multiplying and reliance increasing

NFS Subversion System Decisions




- Must address subversion for justification of trust
 - Irresponsible not to consider when deploying systems
 - Otherwise flawed belief system security is adequate
- Nurture a vast industry with add-on applications
 - Huge drain on resources for little or no assurance
- Objective of demonstration to raise awareness
 - Enable decision maker to understand the problem
 - Need to understand motive, means and opportunity
 - Consider subversion practicality and consequences
 - Make decision makers aware of proven technology
 - Verifiable protection technology applied successfully
- **Security professionals have a responsibility**

Recall Study Goals for NFS Subversion

- First, consider security implications for system
 - How deeply rooted are adverse consequences
- Second, explore applicability to other systems
 - Address whether attack approach is limited to NFS
 - Briefly examine Anderson SSL subversion design
- Next – NFS case study of mitigation
 - Compare to Anderson recommended solution
- **What else can be learned from the demo?**

Notional Cloud Storage Security



MLS File Sharing Server for Cloud



- Cloud storage service
 - Specific type of cloud computing
 - Managed resource is storage
- Needs security as good as enterprise
 - Typically replaces services of enterprise environment
 - Many of the same vulnerability as self-managed
 - Additional vulnerabilities specific to the cloud
- Current solutions are completely ineffective
 - Essential problem is construct of shared infrastructure
 - Built on low-assurance commodity technology
- Highly vulnerable to software subversion

Present-day Vulnerability Examples



Apple pushes security and privacy credentials after iCloud hack

Dark cloud: Study finds security risks in virtualization

Hackers penetrated Nasdaq computer network

Security Requirements of cloud



- Three primary cloud security requirements
 - Controlled sharing of information
 - Cloud isolation
 - High Assurance

Trap Door Subversion Vulnerability



- Malicious code in platform
 - Software, e.g., operating system, drivers, tools
 - Hardware/firmware, e.g., BIOS in PROM
 - Artifice can be embedded any time during lifecycle
 - Adversary chooses time of activation
- Can be remotely activated/deactivated
 - Unique “key” or trigger known only to attacker
 - Needs no (even unwitting) victim use or cooperation
- Efficacy and Effectiveness Demonstrated
 - Exploitable by malicious applications, e.g., Trojans
 - Long-term, high potential future benefit to adversary
 - Testing not at all a practical way to detect

Alternatives for Controlled Sharing



- Three ways controlled sharing can be facilitated:
- Massive copies of data from all lower levels
 - High assurance one-way flow of information
 - Light diode interface uses physics for high assurance
- File Caching (Local Active Copy)
 - Retain at high level only actually used lower data
 - No way to securely make requests for lower data
 - Security requires manual intervention
- High assurance segmented virtual memory




Massive Copies Approach

Top Secret

Secret

Unclassified



- **Highly inefficient!**
- **Does not scale!**

Computer Utility: Predecessor to Cloud

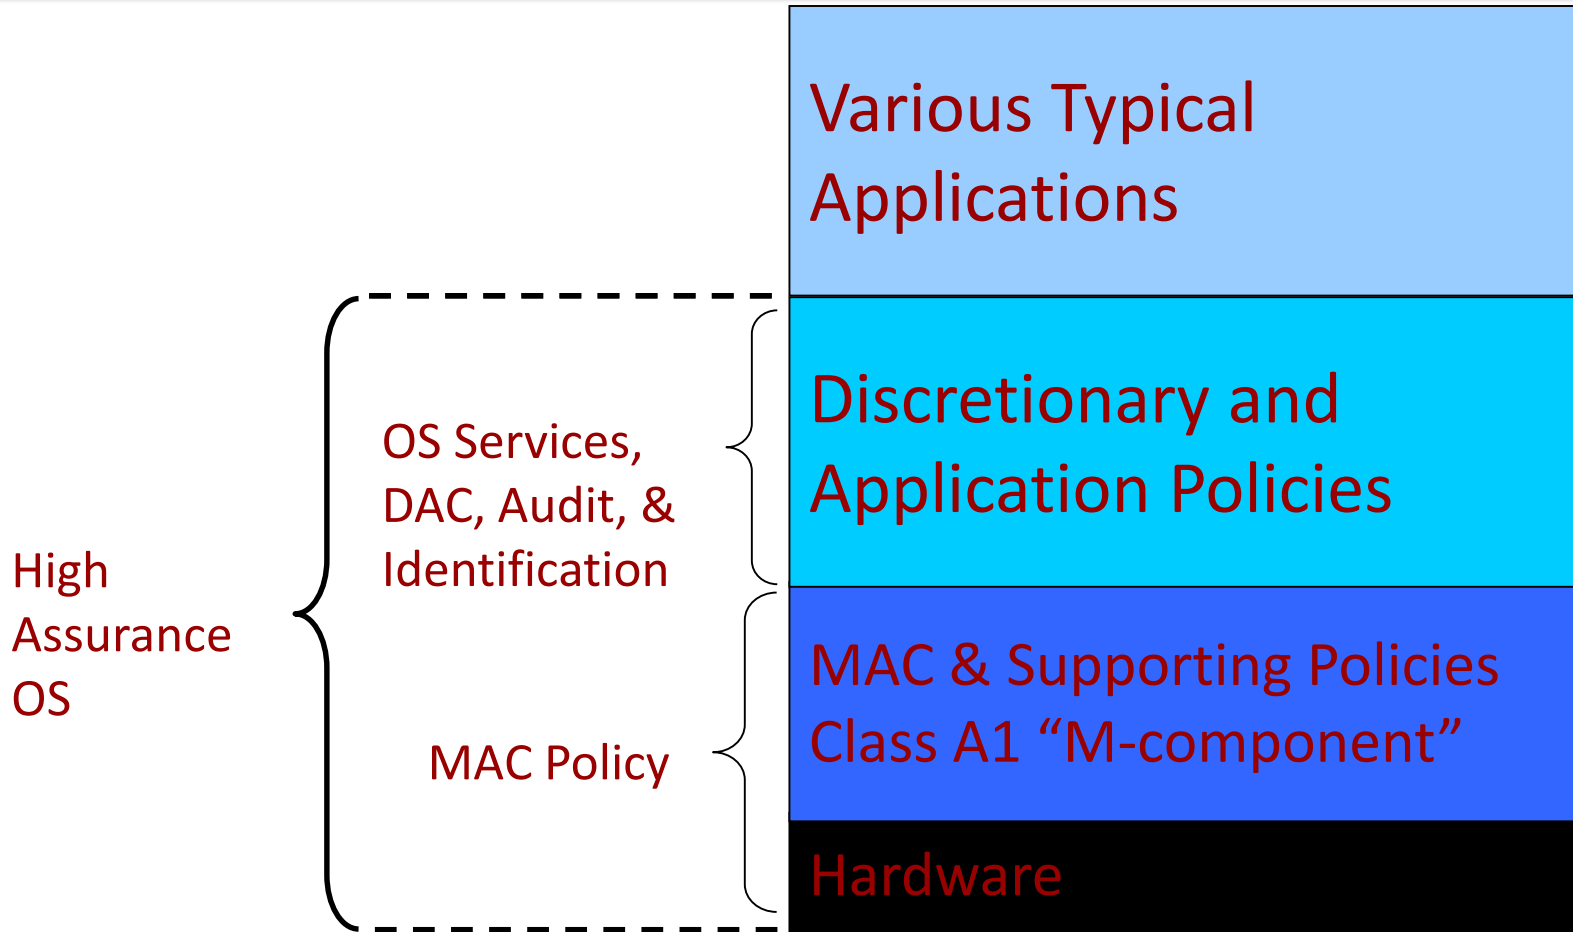
- Computer Utility made security a priority
- Anticipated present day cloud environment
 - Controlled sharing was an integral requirement
 - Incorporated MAC policy
 - Evident from commercial Multics product.
 - Consistent with high assurance
 - Evident from the BLP Multics interpretation.
- Didn't gain widespread acceptance
 - Before Internet enabled today's cloud computing

Basis to Consolidate Networks




- Foundation: Trusted Computing Base (TCB)
 - The totality of protection mechanisms
 - within a computer system
 - including hardware, firmware, and software
 - Responsible for enforcing a security policy
- The security perimeter is TCB boundary
 - Software within the TCB is trustworthy
 - Software outside the TCB is untrusted
- Mature technology -- 30 years experience
 - Derived from and used for commercial products
- Key choice is assurance: low to very high

TCB Subsets – “DAC on MAC”





Use Platform for Controlled Sharing



Motivation to Address Cloud Security



- Cloud storage flexible, cost-effective
- How to implement in multi-level environment?
 - Duplicate for each level? Loses advantages.
- Tempting target for attackers
 - Can increase privilege
- Want high-assurance, MLS solution
- Want cross-domain sharing (CDS)
 - Share resources are core cloud value proposition
 - Controlled sharing of up to date information
- Solution: MLS cloud network file service
 - Based on high-assurance, evaluated Class A1 TCB

Cloud Storage Security Challenges

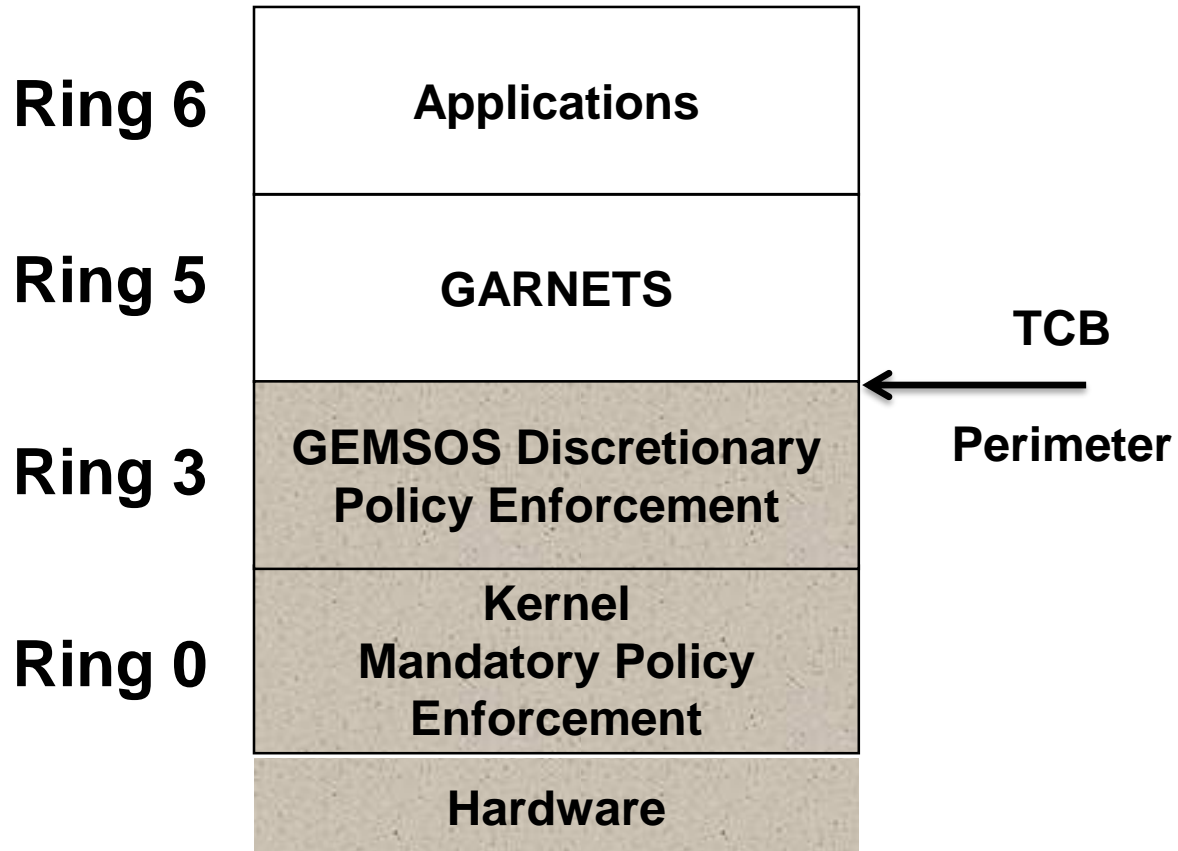


- Migrate storage from local security domain
 - Lose protection of “air gap” between domains
 - Dramatically expands opportunity for adversaries
- Common security approaches can’t work
 - Power of insidious software subversion attack tools
 - Exploding “cloud computing” amplifies impact & risk
- Proven reference validation mechanism (RVM)
 - Systematically codified as TCSEC “Class A1”
- Need architecture that leverages trusted RVM
 - Use verifiable (i.e., Class A1) trusted computing base
 - Want high cloud compatibility, e.g., standard NFS

Current Cloud Technology is Vulnerable

- Typically on commodity low-assurance platforms
- NetApp StorageGRID Systems
 - On Linux servers or VMware hypervisors
- OnApp Storage cloud
 - On various hypervisors, including Xen
- Amazon Elastic Compute Cloud
 - On Xen hypervisor to isolate guest virtual machines
- Xen is representative example of low-assurance
 - So-called TCB includes untrustworthy Linux, drivers
 - Largely unconstrained opportunities for subversion
- NSA said system on partition kernel too complex

Review of GARNETS Architecture




Build on MLS File System Foundation

- GARNETS MLS File is untrusted application
 - In a layer, in a separate ring, above the TCB
 - Creates file management that uses TCB objects
- MLS file system acts like standard file system
 - Spans multiple domains protected by the TCB
 - Creates one logical name space for all domains
- Run GARNETS instances at multiple domains
 - Means no massive copies needed for sharing
 - Can read both directories and files of lower domains
- For cloud storage need to add network interface
 - Multiple network interfaces for multiple domains



Target Secure Cloud Storage

- Operates like standard network file storage
- BUT, verifiable security for
 - MAC separation of security domains

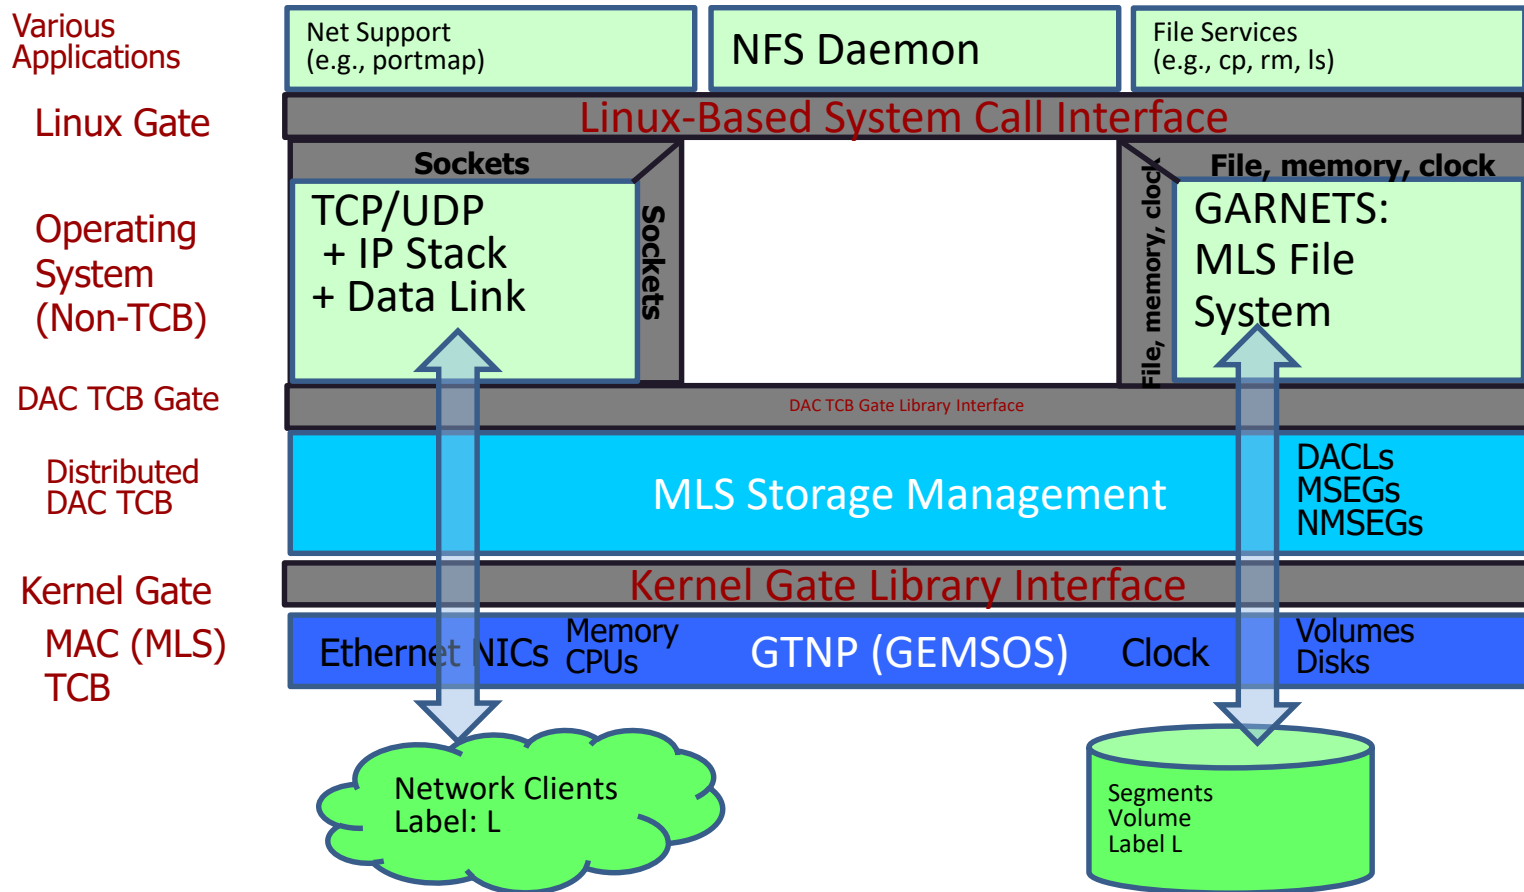


Run NFS on Top of MLS File System



- **System** challenge is a mostly compatible NFS
 - Accept the few intrinsic limitations, e.g., time last read
- Ease of porting NFS wants familiar OS interface
 - Requires creating “compatible OS” [GAS 10.7.2]
 - Demonstration chose POSIX style interface
 - Intended to support Linux source code compatibility
- In contrast, **GARNETS is not compatible**
 - Have to create a Linux system call interface
- Clients access files on server through NSF calls
 - Clients are untrusted
 - Any client using standard NSF protocol can access

NSF Low Domain S/W Instantiation



For MLS Run Multiple NSF Instances



- Need a separate instance for each level
 - Is the only way NFS can be untrusted software
- Clouds often use virtual machine monitor (VMM)
 - Have noted low-assurance virtual machine problem
- Secure NFS demo leverages GEMSOS VMM
 - FER from NSA describes trusted MLS virtualization
 - Is NOT Type I virtualization, i.e., cannot run binary
 - Each NFS instance runs in its own virtual machine
- MLS from TCB cannot be bypassed by VMM
 - Can be configured for typical isolation
 - In contrast to most VMMs, has controlled sharing

FER References GTNP VMM



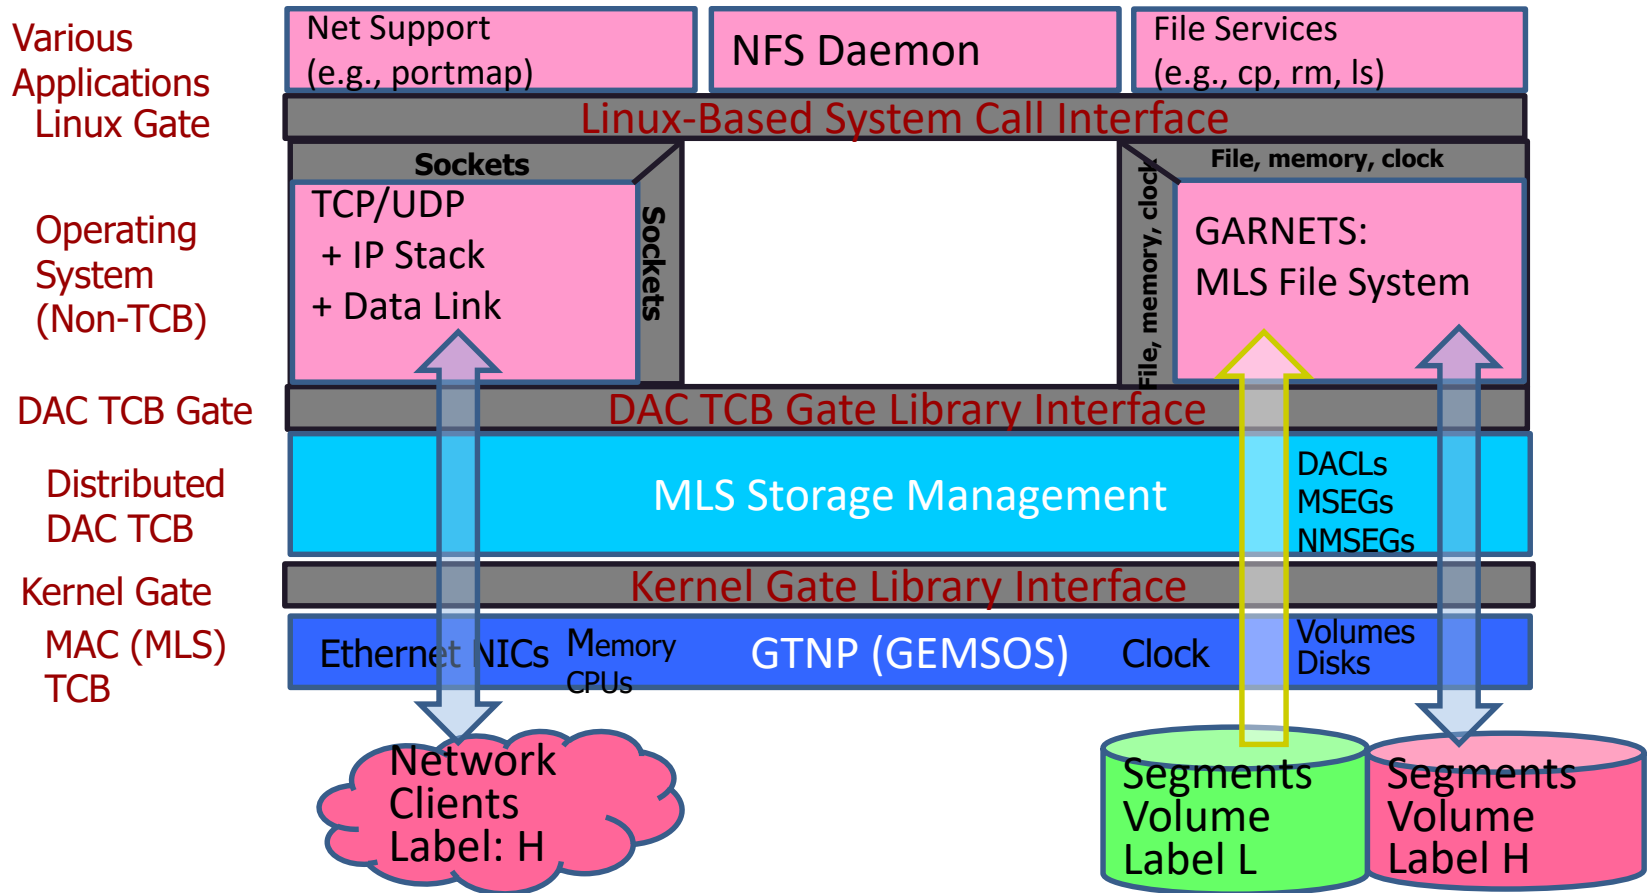
- Sec 2.3, para 2:
 - “The GTNP is intended to support "internal subjects"
 - Virtual machines per Section 1.3.2.2.2 of the TNI.
- Sec 2.3.1, page 12:
 - Implementing A-components as virtual machines
 - Layer between M-NTCB and untrusted VM subjects
- Sec 2.3.1, page 12:
 - VM on top of VMM provided by M-component
- Sect 2.3.1.2 Virtual Machine Interface,
 - Way to compose other components with GTNP
- Sec 4.2.1.8.2.1, para 2: VM supports users

More FER References GTNP VMM



- Sec 4.2.2, para 1:
 - Gate with hardware provide a VMM base
- Sec 9.2, page 181:
 - Rings support VMs the NTCB MAC partition.
- Sect10.3, subpara 2:
 - Multilevel device built as a VM on GNTN VMM
- Sect 10.6, para 1:
 - Implement other network partitions as VM
- Appendix C, EPL Entry, page C-2:
 - GTNP supports a virtual machine monitor interface

NSF High Domain S/W Instantiation



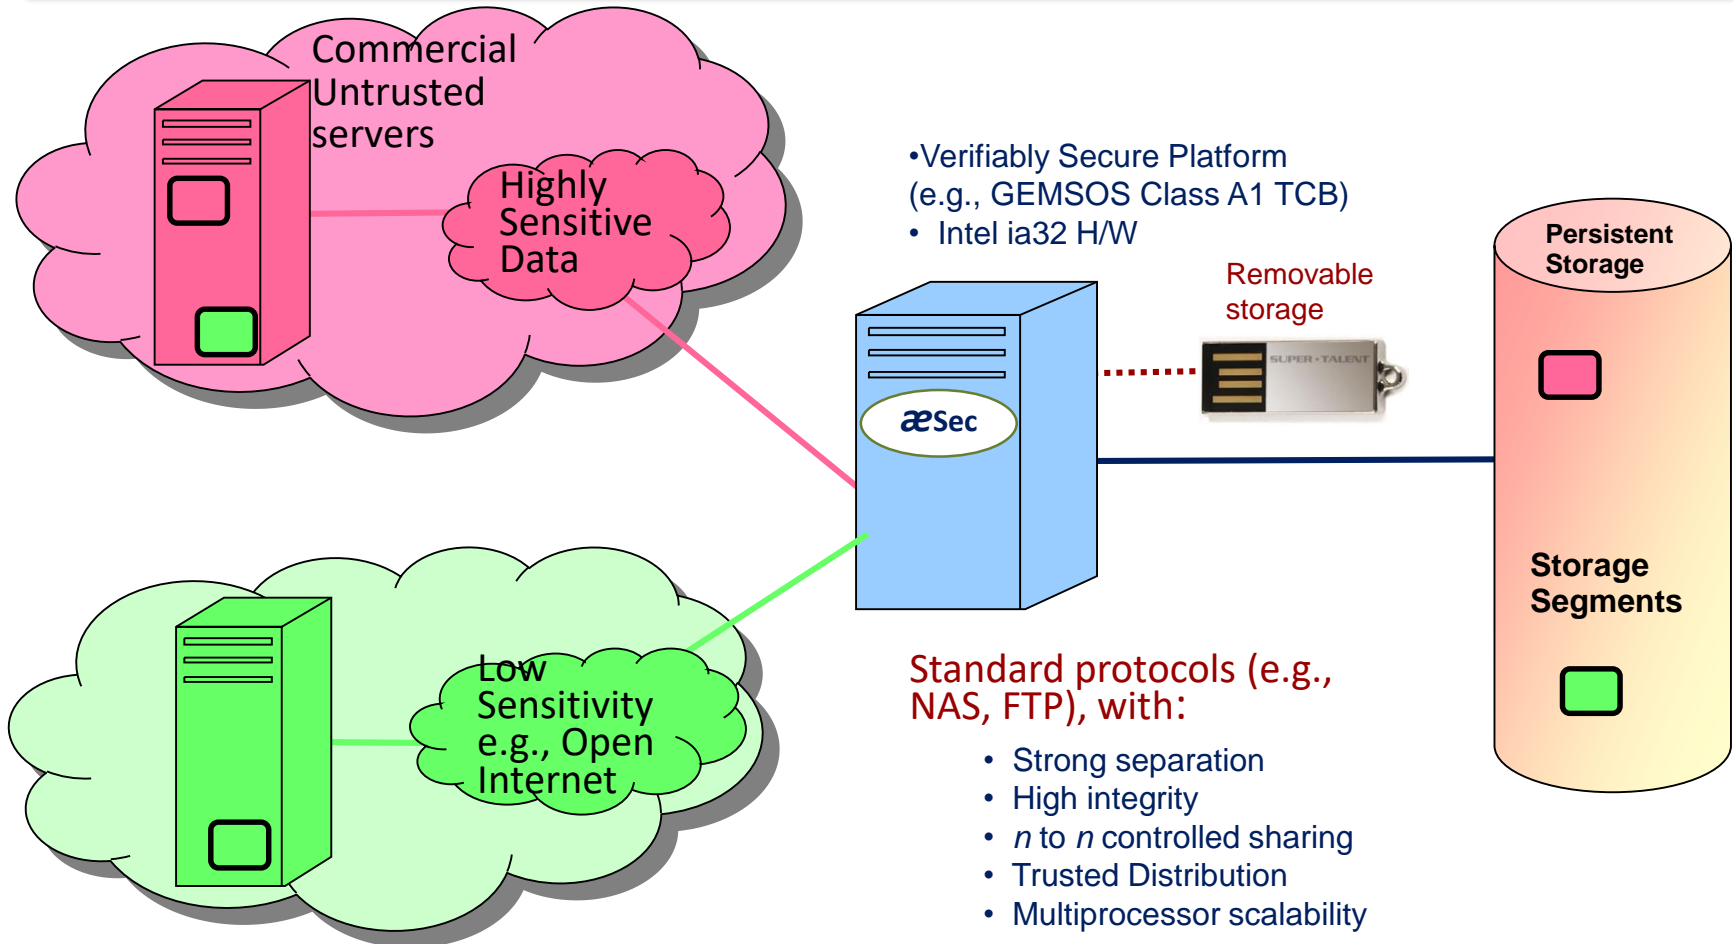
Can Support Extended Cloud Services



- Usually have metadata server for cloud services
 - Maps object names to actual locations in the cloud
- Single level servers in the network can access
 - May use FTP to transfer files to/from MLS file server
 - Applications NAS protocols can access it files
- **BEWARE** of sloppy system security engineering
 - Can't initiate information transfers **between** domains
 - TNI provide the network engineering framework
- Separate NIC per domain does not scale well
 - Next look at how to use a single hardware interface
 - Need equivalent of a TCSEC/TNI multilevel device



Security for Untrusted Servers



Cloud Controlled Sharing Summary



- Key is Mandatory Access Control (MAC)
 - Control isolation & sharing between security domains
- Defining properties: global and persistent
 - Control information flow (confidentiality)
 - Prevents malicious information exfiltration
 - Control contaminated modification (integrity)
 - Sound mathematical foundation
- Implement with distinct access class labels
 - Label domain of information with access class
 - User has authorized access classes, i.e., domains
- Supports MAC, viz., multilevel security (MLS)



INF527: Secure System Engineering

MLS Cloud Storage

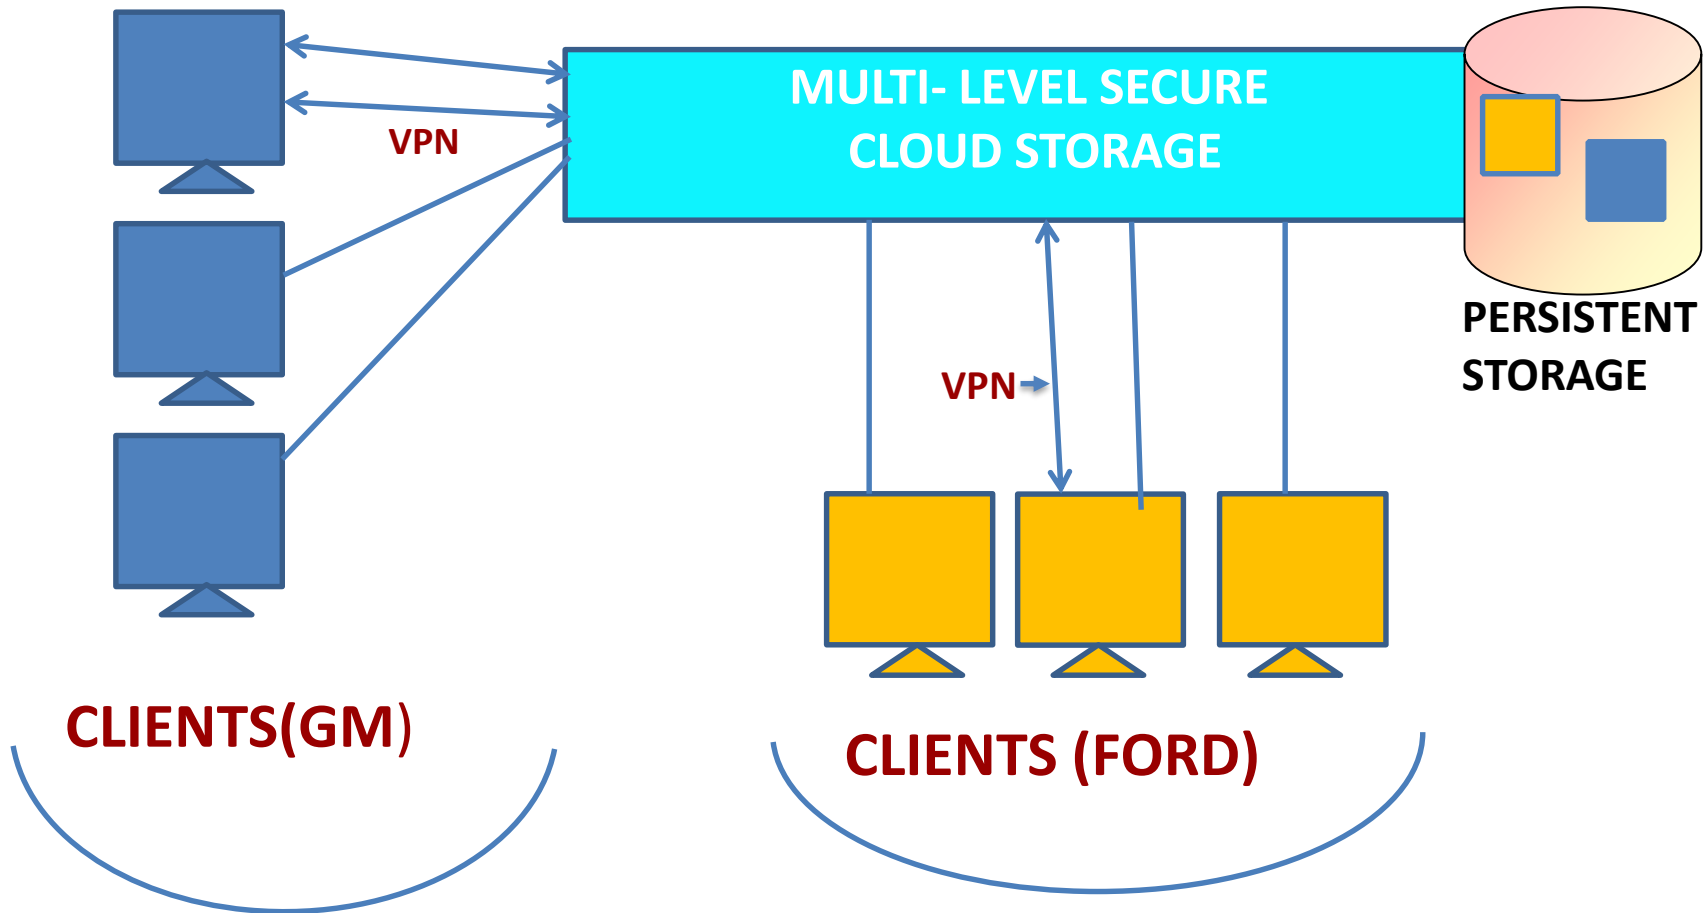
Professor Clifford Neuman

Lecture 13 continued

Recall Study Goals for NFS Subversion

- First, consider security implications for system
 - How deeply rooted are adverse consequences
- Second, explore applicability to other systems
 - Address whether attack approach is limited to NFS
 - Briefly examine Anderson SSL subversion design
- Next – NFS case study of mitigation
 - Compare to Anderson recommended solution
- **What else can be learned from the demo?**

Notional Cloud Storage Security



MLS File Sharing Server for Cloud



- Cloud storage service
 - Specific type of cloud computing
 - Managed resource is storage
- Needs security as good as enterprise
 - Typically replaces services of enterprise environment
 - Many of the same vulnerability as self-managed
 - Additional vulnerabilities specific to the cloud
- Current solutions are completely ineffective
 - Essential problem is construct of shared infrastructure
 - Built on low-assurance commodity technology
- Highly vulnerable to software subversion

Present-day Vulnerability Examples



Apple pushes security and privacy credentials after iCloud hack

Dark cloud: Study finds security risks in virtualization

Hackers penetrated Nasdaq computer network

Security Requirements of cloud



- Three primary cloud security requirements
 - Controlled sharing of information
 - Cloud isolation
 - High Assurance

Trap Door Subversion Vulnerability



- Malicious code in platform
 - Software, e.g., operating system, drivers, tools
 - Hardware/firmware, e.g., BIOS in PROM
 - Artifice can be embedded any time during lifecycle
 - Adversary chooses time of activation
- Can be remotely activated/deactivated
 - Unique “key” or trigger known only to attacker
 - Needs no (even unwitting) victim use or cooperation
- Efficacy and Effectiveness Demonstrated
 - Exploitable by malicious applications, e.g., Trojans
 - Long-term, high potential future benefit to adversary
 - Testing not at all a practical way to detect

Alternatives for Controlled Sharing



- Three ways controlled sharing can be facilitated:
- Massive copies of data from all lower levels
 - High assurance one-way flow of information
 - Light diode interface uses physics for high assurance
- File Caching (Local Active Copy)
 - Retain at high level only actually used lower data
 - No way to securely make requests for lower data
 - Security requires manual intervention
- High assurance segmented virtual memory




Massive Copies Approach

Top Secret

Secret

Unclassified



- **Highly inefficient!**
- **Does not scale!**

Computer Utility: Predecessor to Cloud


- Computer Utility made security a priority
- Anticipated present day cloud environment
 - Controlled sharing was an integral requirement
 - Incorporated MAC policy
 - Evident from commercial Multics product.
 - Consistent with high assurance
 - Evident from the BLP Multics interpretation.
- Didn't gain widespread acceptance
 - Before Internet enabled today's cloud computing

Basis to Consolidate Networks




- Foundation: Trusted Computing Base (TCB)
 - The totality of protection mechanisms
 - within a computer system
 - including hardware, firmware, and software
 - Responsible for enforcing a security policy
- The security perimeter is TCB boundary
 - Software within the TCB is trustworthy
 - Software outside the TCB is untrusted
- Mature technology -- 30 years experience
 - Derived from and used for commercial products
- Key choice is assurance: low to very high

TCB Subsets – “DAC on MAC”





Use Platform for Controlled Sharing



Motivation to Address Cloud Security



- Cloud storage flexible, cost-effective
- How to implement in multi-level environment?
 - Duplicate for each level? Loses advantages.
- Tempting target for attackers
 - Can increase privilege
- Want high-assurance, MLS solution
- Want cross-domain sharing (CDS)
 - Share resources are core cloud value proposition
 - Controlled sharing of up to date information
- Solution: MLS cloud network file service
 - Based on high-assurance, evaluated Class A1 TCB

Cloud Storage Security Challenges




- Migrate storage from local security domain
 - Lose protection of “air gap” between domains
 - Dramatically expands opportunity for adversaries
- Common security approaches can’t work
 - Power of insidious software subversion attack tools
 - Exploding “cloud computing” amplifies impact & risk
- Proven reference validation mechanism (RVM)
 - Systematically codified as TCSEC “Class A1”
- Need architecture that leverages trusted RVM
 - Use verifiable (i.e., Class A1) trusted computing base
 - Want high cloud compatibility, e.g., standard NFS

Current Cloud Technology is Vulnerable

- Typically on commodity low-assurance platforms
- NetApp StorageGRID Systems
 - On Linux servers or VMware hypervisors
- OnApp Storage cloud
 - On various hypervisors, including Xen
- Amazon Elastic Compute Cloud
 - On Xen hypervisor to isolate guest virtual machines
- Xen is representative example of low-assurance
 - So-called TCB includes untrustworthy Linux, drivers
 - Largely unconstrained opportunities for subversion
- NSA said system on partition kernel too complex

Review of GARNETS Architecture




Build on MLS File System Foundation

- GARNETS MLS File is untrusted application
 - In a layer, in a separate ring, above the TCB
 - Creates file management that uses TCB objects
- MLS file system acts like standard file system
 - Spans multiple domains protected by the TCB
 - Creates one logical name space for all domains
- Run GARNETS instances at multiple domains
 - Means no massive copies needed for sharing
 - Can read both directories and files of lower domains
- For cloud storage need to add network interface
 - Multiple network interfaces for multiple domains



Target Secure Cloud Storage

- Operates like standard network file storage
- BUT, verifiable security for
 - MAC separation of security domains




Run NFS on Top of MLS File System



- **System** challenge is a mostly compatible NFS
 - Accept the few intrinsic limitations, e.g., time last read
- Ease of porting NFS wants familiar OS interface
 - Requires creating “compatible OS” [GAS 10.7.2]
 - Demonstration chose POSIX style interface
 - Intended to support Linux source code compatibility
- In contrast, **GARNETS is not compatible**
 - Have to create a Linux system call interface
- Clients access files on server through NSF calls
 - Clients are untrusted
 - Any client using standard NSF protocol can access

NSF Low Domain S/W Instantiation




For MLS Run Multiple NSF Instances



- Need a separate instance for each level
 - Is the only way NFS can be untrusted software
- Clouds often use virtual machine monitor (VMM)
 - Have noted low-assurance virtual machine problem
- Secure NFS demo leverages GEMSOS VMM
 - FER from NSA describes trusted MLS virtualization
 - Is NOT Type I virtualization, i.e., cannot run binary
 - Each NFS instance runs in its own virtual machine
- MLS from TCB cannot be bypassed by VMM
 - Can be configured for typical isolation
 - In contrast to most VMMs, has controlled sharing

NSF High Domain S/W Instantiation



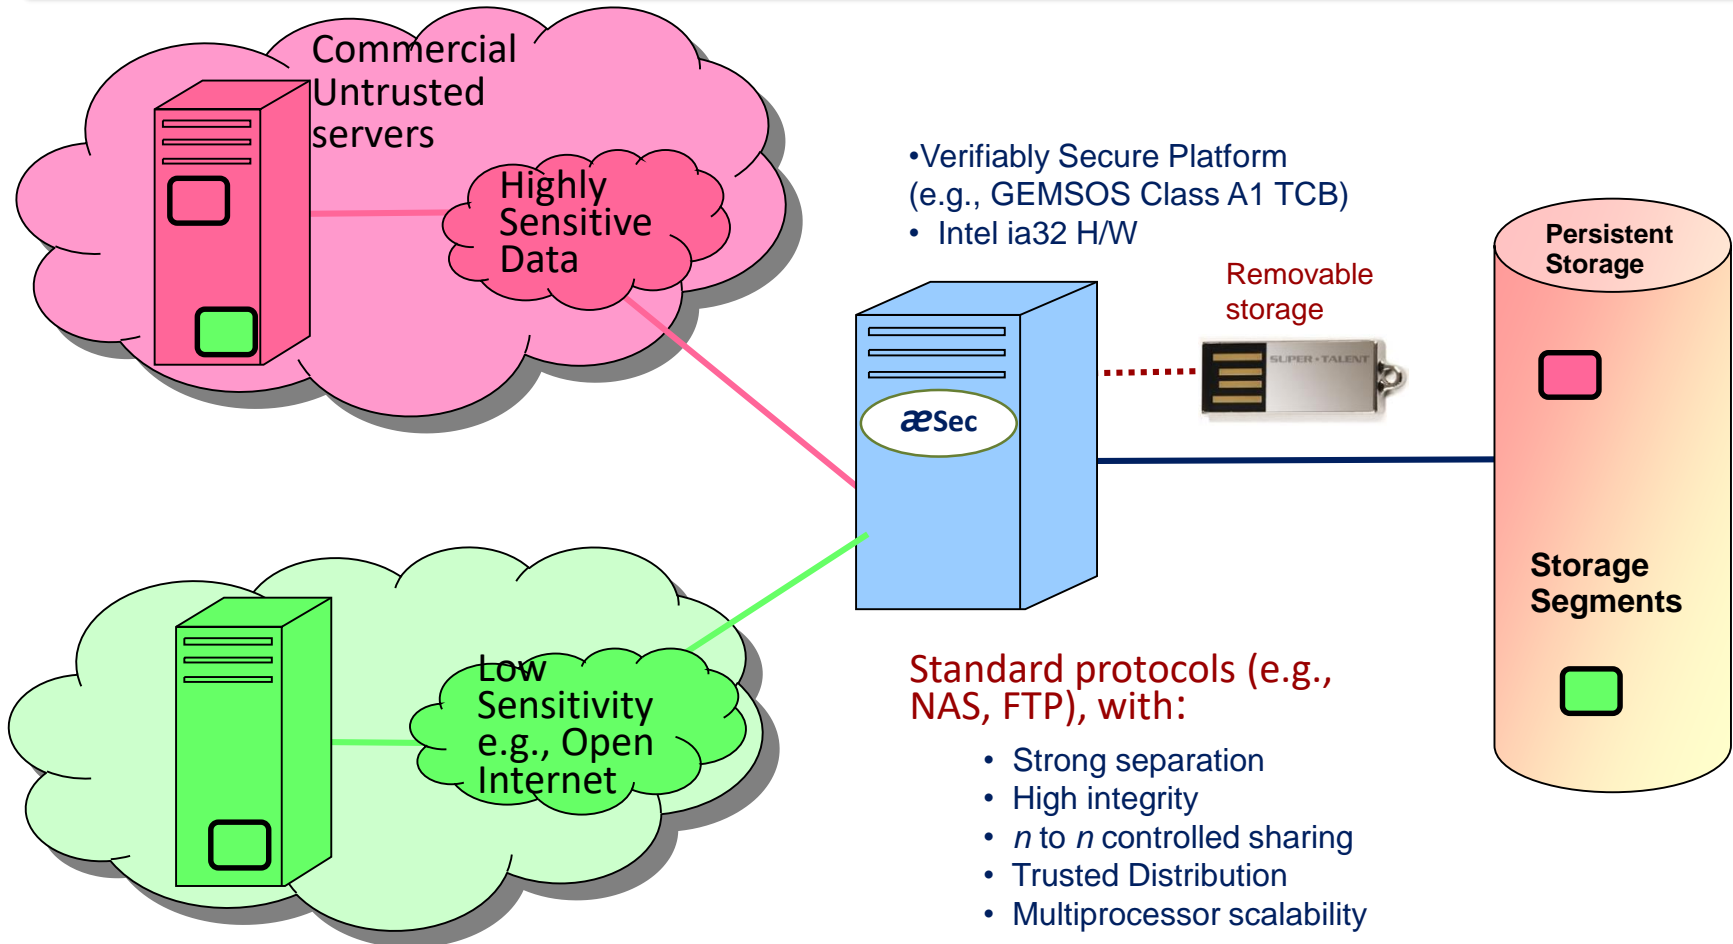
Can Support Extended Cloud Services



- Usually have metadata server for cloud services
 - Maps object names to actual locations in the cloud
- Single level servers in the network can access
 - May use FTP to transfer files to/from MLS file server
 - Applications NAS protocols can access it files
- **BEWARE** of sloppy system security engineering
 - Can't initiate information transfers **between** domains
 - TNI provide the network engineering framework
- Separate NIC per domain does not scale well
 - Next look at how to use a single hardware interface
 - Need equivalent of a TCSEC/TNI multilevel device



Security for Untrusted Servers



Cloud Controlled Sharing Summary



- Key is Mandatory Access Control (MAC)
 - Control isolation & sharing between security domains
- Defining properties: global and persistent
 - Control information flow (confidentiality)
 - Prevents malicious information exfiltration
 - Control contaminated modification (integrity)
 - Sound mathematical foundation
- Implement with distinct access class labels
 - Label domain of information with access class
 - User has authorized access classes, i.e., domains
- Supports MAC, viz., multilevel security (MLS)



INF523 SUPPLEMENTAL MATERIAL (if time in semester)

Introduction to Crypto Seal Guards
Case Study

Professor Clifford Neuman

Supplemental

Crypto Seal Guard Technology History

- Concept: label cryptographically sealed to data
- Conceived ~1980 for AF Korean Air Intelligence
- GEMSOS uses to meet TCSEC “Label Integrity”
 - Gemini Trusted Network Processor (GTNP) (1995)
 - Stored data (disk, tape) in Class A1 Evaluation
- GEMSOS uses for “Trusted Distribution”
 - Authoritative distribution media crypto sealed
 - Only sealed TCB software can be installed and run
- POC applied to packets exchanged by guards
 - Each guard is MLS – both a high and low interface

GEMSOS Support for Crypto Seals



- GEMSOS used crypto seals to meet Class A1
 - To meet Class A1 Label Integrity requirements
 - Integral to Trusted Recovery & Trusted Distribution
- GEMSOS publishes security services via APIs:
 - Data Sealing Device (and Cryptographic Services)
 - Key Management
 - Trusted Recovery & Distribution
- GemSeal uses GEMSOS APIs for crypto seals
 - Previously evaluated, stable, public interfaces
 - Minimal new trusted code
 - Generate seal
 - Validate integrity/authenticity of sealed packet & label

Overview of Seals for Shared Networks


- Proof of Concept (POC) demonstration done
 - Crypto seal release guards
 - Preproduction Class A1 MLS platform
- Access low network across system high network
 - Controlled interface protects system high data
 - Vertical data fusion with reduced footprint
- Benefits of crypto seal release guards
 - Swift implementation for MLS systems
 - Available core enabling technology for MLS
 - Rapid path to certification and accreditation (C&A)
 - Supports entire range of security domains
 - Mature deployed NSA Class A1 TCB and RAMP plan

Constraints to Access Lower Networks



High
Network

Multi-Level
Secure
Connection



Low
Network

- Any low connection = Multi-Level
 - Must be Multi-Level Secure (MLS)
 - Low/Medium assurance ineffective
 - Doesn't protect against subversion
 - Vulnerabilities unknown (unknowable)
- Isolation obstructs missions
 - Vertical data fusion
 - Tactical situational awareness
 - Timely access to open source data
 - Efficient utilization of resources

GemSeal POC Uses MLS Technology

- Class A1 TCB - GEMSOS™ security kernel
- Class A1 Ratings Maintenance Plan (RAMP)
- MLS aware crypto seal release guard
 - Gemini call it the GemSeal™ concept
- Technology Benefits
 - Minimize new trusted code development
 - Extensible to gamut of MLS capable systems
- High assurance resists subversion
 - Verifies absence of malicious code
 - Effective application of science
 - Key enabler for demanding accreditation, e. g., PL-5



How Guard Seals a Packet

- Packet switched network design, e.g. Internet
- Concept involves multiple guards
 - POC has one or more “workstation” guards
 - POC has one or more “sensor” guards
 - Connected via a common system-high network
- Each guard has both high and low interfaces
- Sealing packets – forwarding from low to high
 - Bind source interface (low) label to each packet
 - Generate cryptographic seal of packet data + label
 - “Low-Sealed” packets include packet data + seal
 - “Low-Sealed” packets via high network interface

How Guard Releases a Packet




- Releasing packets – delivering from high to low
 - Release ONLY packets with seal-validated labels
 - Seal and label are removed before being released
- Only released to interfaces matching labels
 - Allows low data to traverse & exit high network
 - Concept supports multiple release guards
- Assures integrity of BOTH data AND label
 - Packet data is not altered
 - Source sensitivity label is authentic for this packet



AF Crypto Seal POC Demonstration

Crypto seal release guards connect low resources across the high network but protect high data.





Summary of AF POC Demonstration

- Sensor (video) stream + command and control
 - Low sensor to low workstation connectivity
 - Uses existing high network infrastructure
 - Delivers access to low devices
 - For users lacking low network infrastructure
 - From controlled interface
- High network data is protected and unchanged
 - Guard validates low-sealed packets before release
 - Unsealed high packets cannot exit via guard

Summary of POC Configuration



- Two untrusted workstations with browsers
 - One (“Low”) connected to “workstation guard”
 - One (“High”) connected to high network
- One web server
 - Connected to low-side of the “sensor guard”
- A “high” Ethernet LAN
 - Connected to high-side of both guards
 - Also connected to second system high workstation
- The demonstration shows that
 - “Low” workstation can browse the “Low” web server
 - “High” workstation has no access to “Low” web server

Prior Evaluation Aids Accreditors



- Simplify job with reusable accreditation results
 - Certify or assess the platform once
 - Focus on system-specific additions & configurations
- GTNP provides evaluable TCB platform
 - Previously evaluated Class A1 for TNI M-Component
 - Class A1 RAMP in place and already proved useful
- Outside of the GTNP trusted computing base
 - Most of the application software will be untrusted
 - Only cryptographic seal operations need be trusted
 - Generate seals & release packets with validated seals
- Customer's certification and accreditation needs
 - The verifiably secure MLS TCB and

POC to Deployable System Summary

- Don't have to evaluate platform first
 - RAMP is already proven
 - No formal specification changes anticipated
- First: Evaluate and accredit the parts separately
 - Platform (very stable, accredit new hardware ports)
 - Crypto Seal implementation (as a trusted application)
 - Guard applications themselves evaluated separately
 - Supporting policies - audit, DAC, etc.
 - Untrusted application pieces, including network stack
 - Each protected by security kernel
- Last: refresh platform evaluation + accreditation
 - Because already successfully evaluated & accredited

Introduction to RECON Guard Security

- Review a classic and seminal paper
 - Cite: J. P. Anderson, "On the Feasibility of Connecting RECON to an External Network," Technical Report, J. P. Anderson Co., March 1981
 - Often cited for both databases and communications
- RECON is on-line interactive data base
 - Citations for both raw and finished intelligence reports
 - Also overnight batch and canned query capability
 - User may specify which file(s) to search
- Sponsor's security concerns are twofold
 - Subject to penetration from external network
 - Spillage of sensitive information from internal failure



Data Security Protection

- The data base contains two kinds of records
 - Those which can be widely distributed
 - Those whose distribution is restricted
 - Compartmented
 - Proprietary
 - Originator-controlled
- Operative aspects of the security problem
 - Commodity mainframe operating system
 - Must be prudently assumed that trapdoors exist
 - In some or much of application or operating systems
 - May be activated from externally connected users

Previously Considered Approaches




- Put two kinds of records in separate systems
 - Make entries not deemed "special" accessible
 - Protected the sponsor's assets from penetration
 - Rejected because of the cost of duplicate facilities
- Multilevel secure operating system
 - In principle, would go far to defeat direct attacks
 - Could defeat placing trapdoors and Trojan Horses
 - Produce totally incompatible (with anything!) systems
 - Very expensive
- Filters added to RECON software to limit access
 - Nothing to control internal or external penetration

Guard Authenticate Releasability



- Is akin to the problem of "sanitizing" SCI
 - For release to activities without proper clearances
- Permit arbitrary queries by all users
 - Route query result of uncleared users to sanitizer
 - Sanitization officer would manually examine output
- Sanitization officer approach works in principle
 - Not practical solution because of excessive delays
 - Delays cascade to produce large response times
- Adapted as proposal to solve RECON problem
 - Adopt the idea of “sanitization” in a GUARD station
 - Automate the identification of releasable citations

RECON Guard Technical Approach

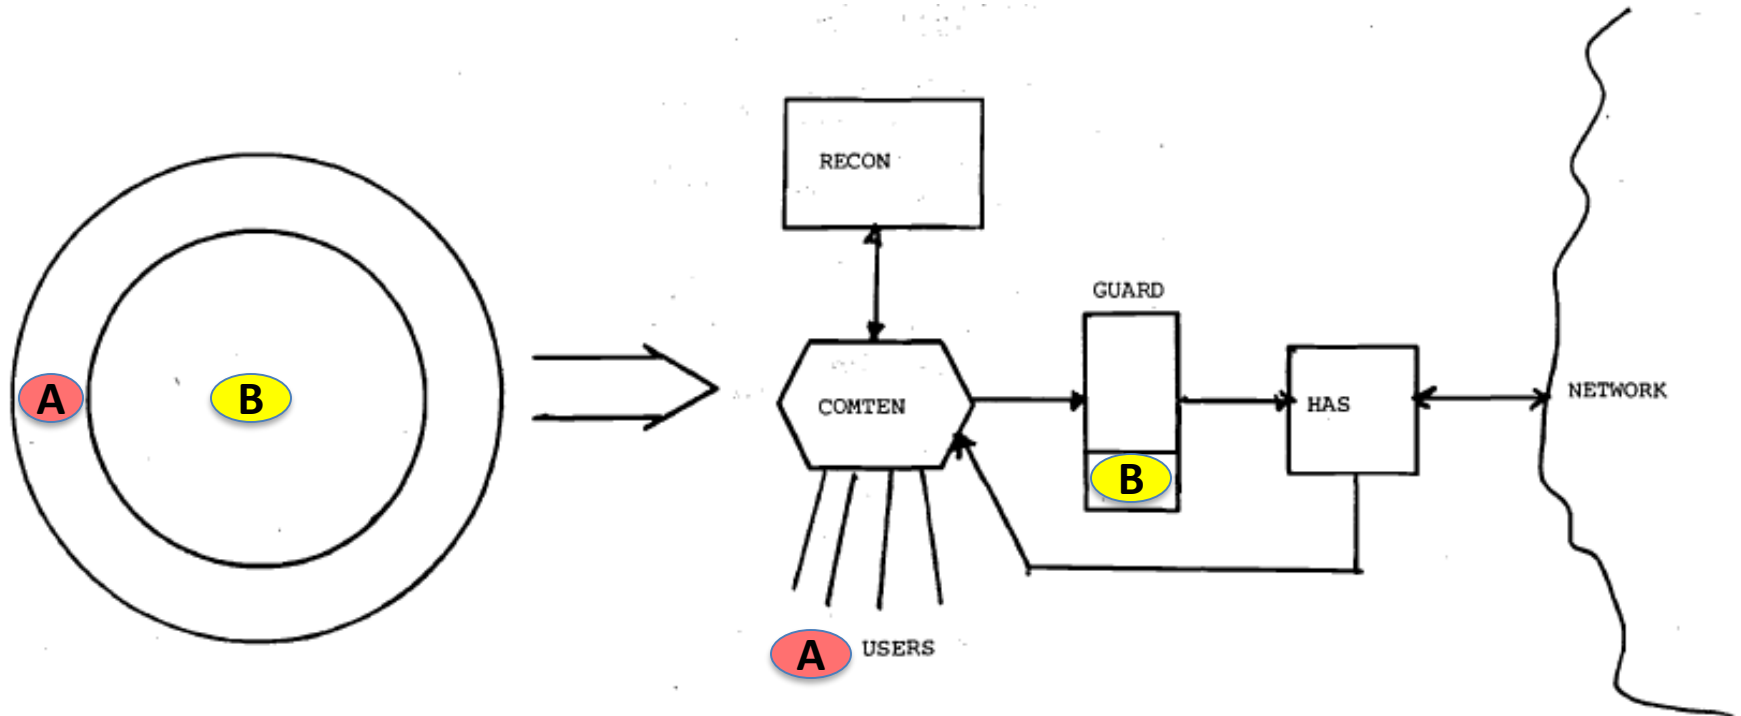


RECON Guard Concept of Operation



- Consider all citations in one of two cases
 - Releasable even if not approved for "special" citations
 - Releasable only to approved individuals
- Each RECON entry designated by **originator**
 - Whether (or not) it is releasable to external users
- Create cryptographic checksum for releasable
 - Computed as the data enters the system
 - A function of the entire record
 - Computed by a special authentication device
 - Checksum is appended to the record and stays with it

Representation of Basic Capability



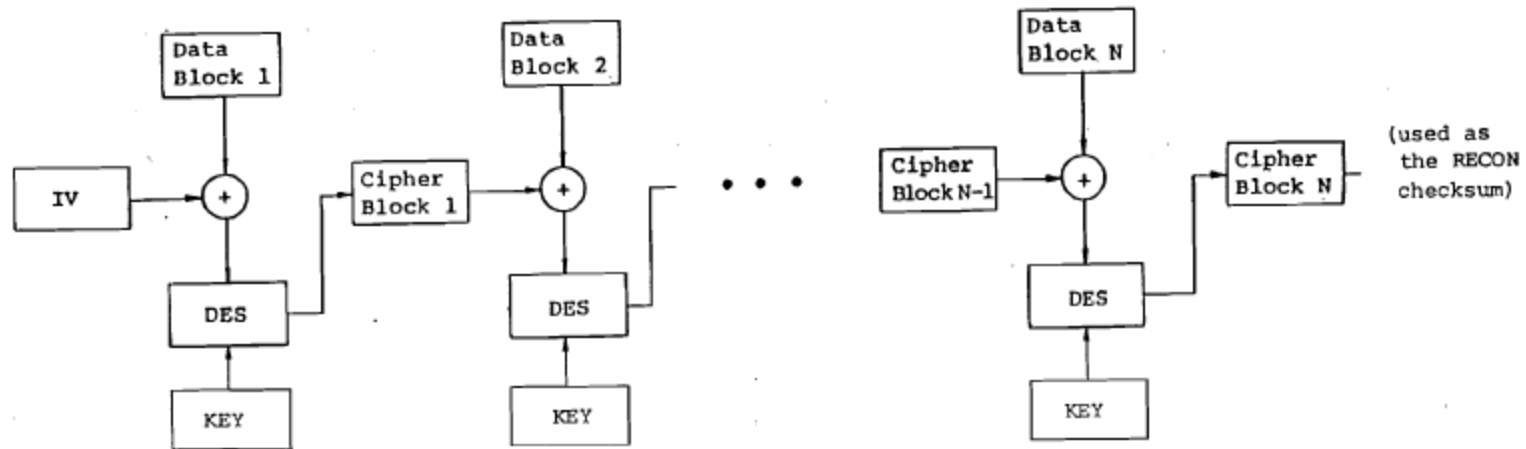
Cryptographic Checksums Properties



- Principle need is assuring checksum not forged
 - Good modern crypto algorithm
 - Perform checksum functions outside RECON hosts
 - Separate entities to create and do Guard functions
- Secret key is known only to checksum devices
 - Key is never available within RECON system
 - Hardwired on board with crypto processor
 - Only method to forge is random guess (brute force)
- Key used for block-chained encipherment
 - Excellent error or tampering detection
 - Initial variable (IV) is used as half of the “secret”
 - A security “kernel” in devices control their operation



Process to Create Crypto Checksum



- \oplus Exclusive OR
- The secret key(s) are the Initial Variable (IV) and the KEY.

BLOCK CHAINING

Security Properties of Guard



- No spill from RECON failure or compromise
- No manipulation of RECON will cause release
- Will “fail safe” if checksum detached from data
- Not protecting against manipulation of data
- Not preventing denial of service
- Guard system itself defends against its failure
 - Advanced design techniques, e.g., formal specs
 - Programs placed in read-only memory
 - Permits RECON to test guard message w/ loop back