

# Host-Based Intrusion Detection

---

Vitaly Shmatikov

# Reading Assignment

---

- ◆ Wagner and Dean. “Intrusion Detection via Static Analysis” (Oakland 2001).
- ◆ Feng et al. “Formalizing Sensitivity in Static Analysis for Intrusion Detection” (Oakland 2004).
- ◆ Garfinkel. “Traps and Pitfalls: Practical Problems in System Call Interposition Based Security Tools” (NDSS 2003).

# After All Else Fails

---

## ◆ Intrusion prevention

- Find buffer overflows and remove them
- Use firewall to filter out malicious network traffic

## ◆ **Intrusion detection** is what you do after prevention has failed

- Detect attack in progress
  - Network traffic patterns, suspicious system calls, etc.
- Discover telltale system modifications

# What Should Be Detected?

---

- ◆ Attempted and successful break-ins
- ◆ Attacks by legitimate users
  - For example, illegitimate use of root privileges
  - Unauthorized access to resources and data
- ◆ Trojan horses
- ◆ Viruses and worms
- ◆ Denial of service attacks

# Where Are IDS Deployed?

---

## ◆ Host-based

- Monitor activity on a single host
- Advantage: better visibility into behavior of individual applications running on the host

## ◆ Network-based (NIDS)

- Often placed on a router or firewall
- Monitor traffic, examine packet headers and payloads
- Advantage: single NIDS can protect many hosts and look for global patterns

# Intrusion Detection Techniques

---

## ◆ Misuse detection

- Use attack “signatures” (need a model of the attack)
  - Sequences of system calls, patterns of network traffic, etc.
- Must know in advance what attacker will do (how?)
- Can only detect known attacks

## ◆ Anomaly detection

- Using a model of normal system behavior, try to detect deviations and abnormalities
  - E.g., raise an alarm when a statistically rare event(s) occurs
- Can potentially detect unknown attacks

## ◆ Which is harder to do?

# Misuse Detection (Signature-Based)

---

- ◆ Set of rules defining a behavioral signature likely to be associated with attack of a certain type
  - Example: buffer overflow
    - A setuid program spawns a shell with certain arguments
    - A network packet has lots of NOPs in it
    - Very long argument to a string function
  - Example: denial of service via SYN flooding
    - Large number of SYN packets without ACKs coming back  
...or is this simply a poor network connection?
- ◆ Attack signatures are usually very specific and may miss variants of known attacks
  - Why not make signatures more general?

# Extracting Misuse Signatures

---

- ◆ Use **invariant characteristics** of known attacks
  - Bodies of known viruses and worms, port numbers of applications with known buffer overflows, RET addresses of overflow exploits
  - Hard to handle mutations
    - Polymorphic viruses: each copy has a different body
- ◆ Big research challenge: fast, automatic extraction of signatures of new attacks
  - Honeypots are useful - try to attract malicious activity, be an early target



# Anomaly Detection

---

- ◆ Define a **profile** describing “normal” behavior
  - Works best for “small”, well-defined systems (single program rather than huge multi-user OS)
- ◆ Profile may be statistical
  - Build it manually (this is hard)
  - Use machine learning and data mining techniques
    - Log system activities for a while, then “train” IDS to recognize normal and abnormal patterns
  - Risk: attacker trains IDS to accept his activity as normal
    - Daily low-volume port scan may train IDS to accept port scans
- ◆ IDS flags deviations from the “normal” profile

# Level of Monitoring

---

- ◆ Which types of events to monitor?
  - OS system calls
  - Command line
  - Network data (e.g., from routers and firewalls)
  - Processes
  - Keystrokes
  - File and device accesses
  - Memory accesses
- ◆ Auditing / monitoring should be scalable

# STAT and USTAT

[Ilgun, Porras]

- ◆ Intrusion signature = sequence of system states
  - State machine describing the intrusion must be specified by an expert
    - Initial state: what system looks like before attack
    - Compromised state: what system looks like after attack
    - Intermediate states and transitions
- ◆ State transition analysis is then used to detect when system matches known intrusion pattern


# USTAT Example

```
user% ln someSetuidRootScript -x
user% -x
root%
```

Opens interactive subshell

Creates symbolic link with same permissions and ownership as target

## State diagram of this attack



1. Fileset #1 != empty
2. Files are setuid

1. name(File1) == -\*
2. typeof(File1) == link
3. owner(link\_to(File1)) != user
4. name(link\_to(File1)) exists\_in Fileset #1

1. access(user,euid) == root

# Statistical Anomaly Detection

---

- ◆ Compute statistics of certain system activities
- ◆ Report an alert if statistics outside range
- ◆ Example: **IDES** (Denning, mid-1980s)
  - For each user, store daily count of certain activities
    - For example, fraction of hours spent reading email
  - Maintain list of counts for several days
  - Report anomaly if count is outside weighted norm

Problem: the most unpredictable user is the most important

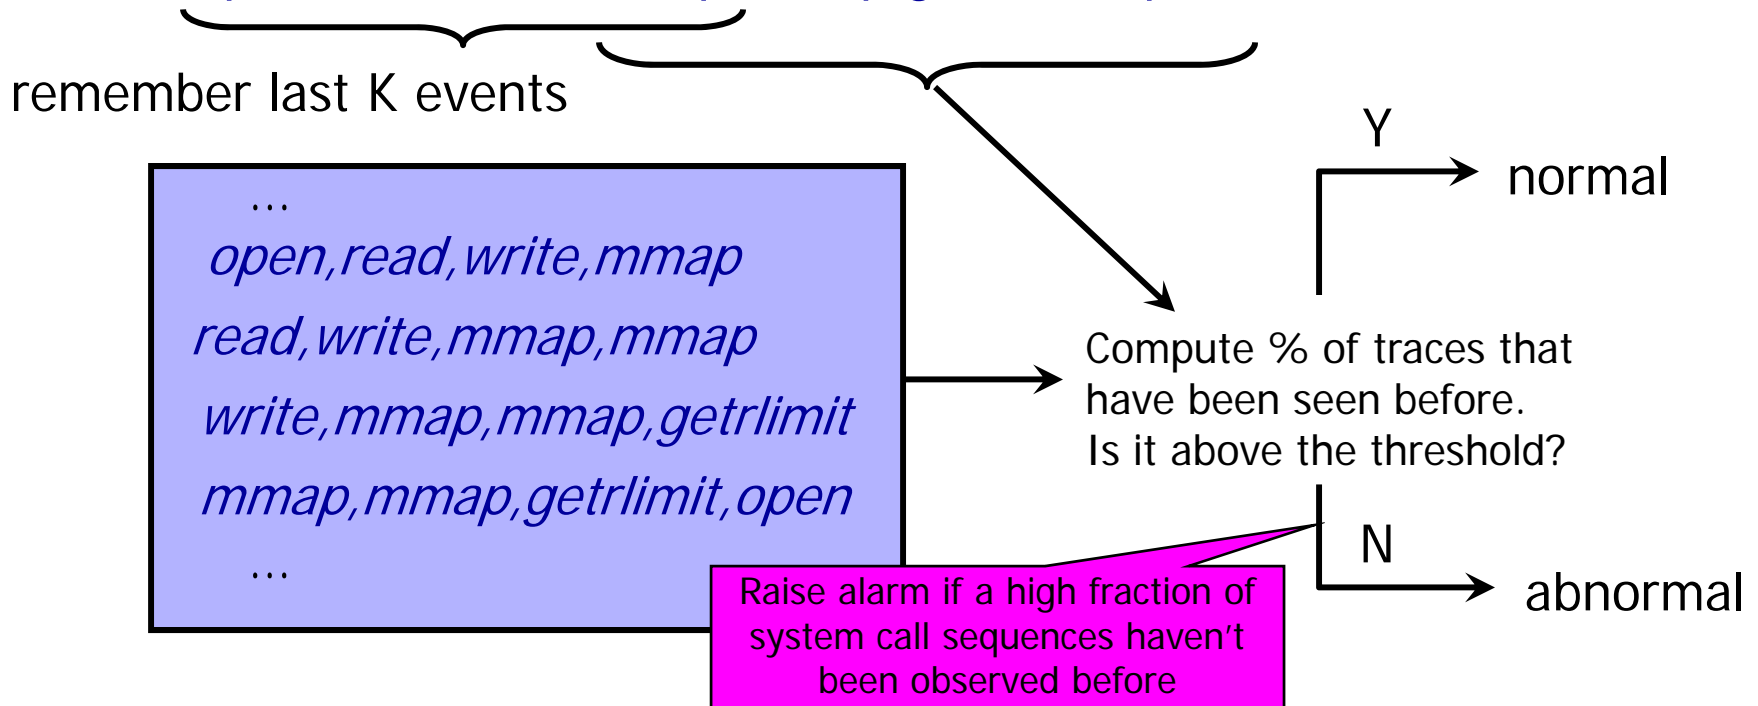
# "Self-Immunology" Approach

[Forrest]

## ◆ Normal profile: short sequences of system calls

- Use strace on UNIX

... *open, read, write, mmap, mmap, getrlimit, open, close* ...



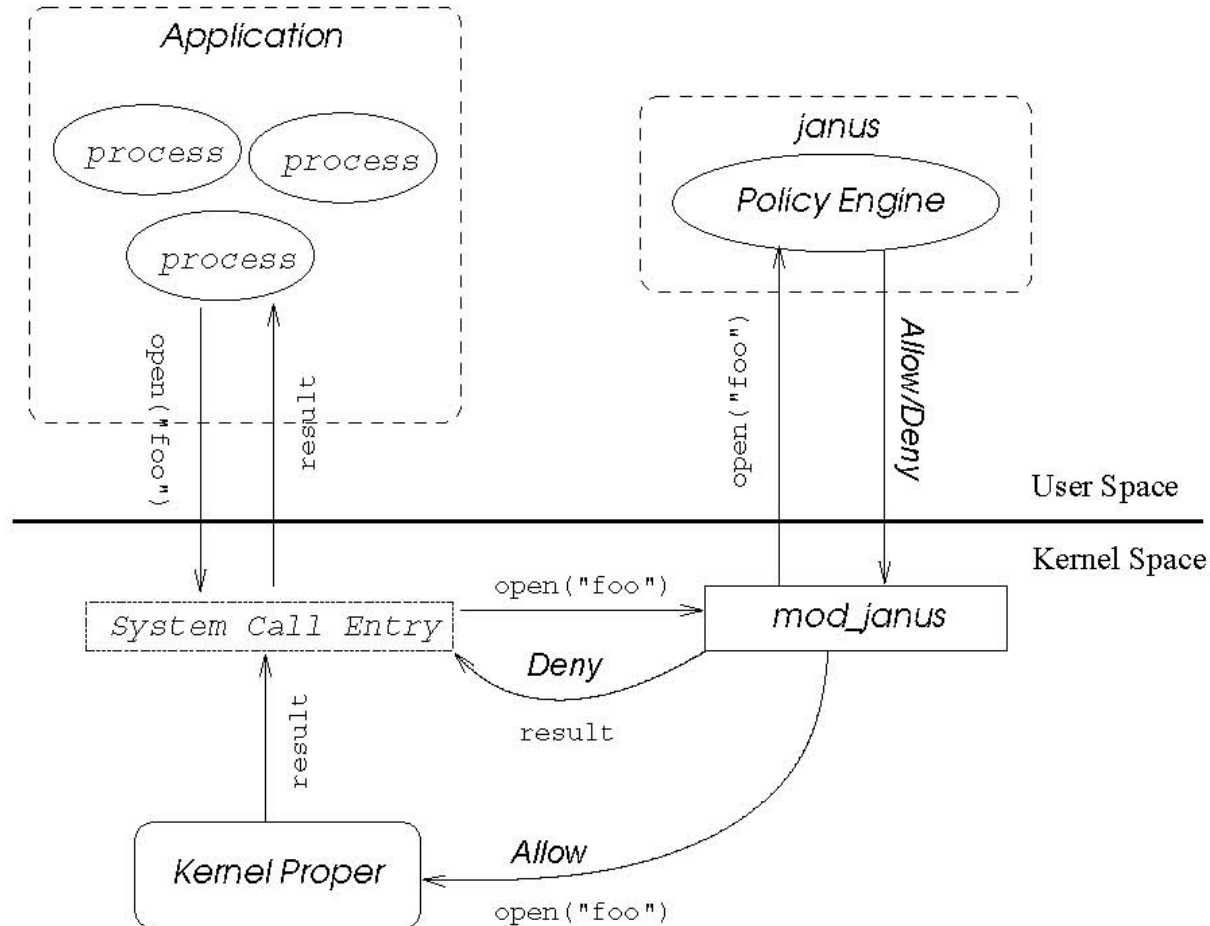
# System Call Interposition

---

- ◆ Observation: all sensitive system resources are accessed via OS system call interface
  - Files, sockets, etc.
- ◆ Idea: monitor all system calls and block those that violate security policy
  - Inline reference monitors
  - Language-level: Java runtime environment inspects stack of the function attempting to access a sensitive resource to check whether it is permitted to do so
  - Common OS-level approach: **system call wrapper**
    - Want to do this without modifying OS kernel (why?)

# Janus

[Berkeley project, 1996]



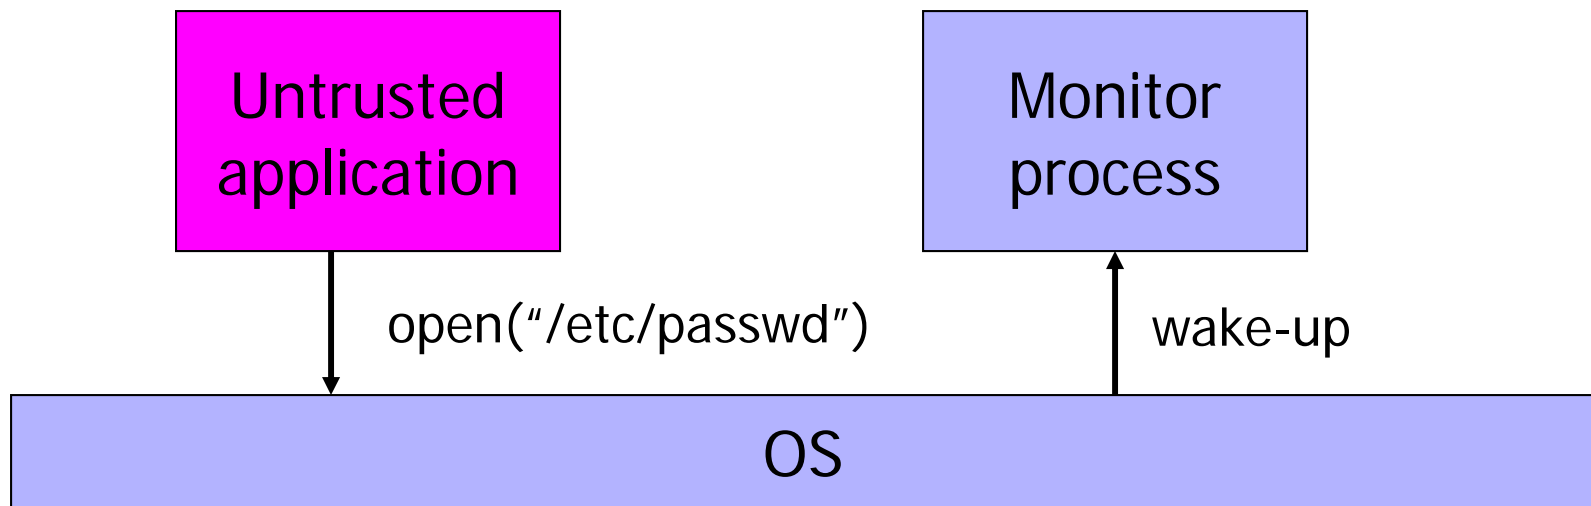


# Policy Design

---

- ◆ Designing a good system call policy is not easy
- ◆ When should a system call be permitted and when should it be denied?
- ◆ Example: ghostscript
  - Needs to open X windows
  - Needs to make X windows calls
  - But what if ghostscript reads characters you type in another X window?

# Trapping System Calls: ptrace()



- ◆ `ptrace()` – can register a callback that will be called whenever process makes a system call
  - Coarse: trace all calls or none
  - If traced process forks, must fork the monitor, too

Note: Janus used `ptrace` initially, later discarded...

# Problems and Pitfalls

[Garfinkel]

- ◆ Incorrectly mirroring OS state
- ◆ Overlooking indirect paths to resources
  - Inter-process sockets, core dumps
- ◆ Race conditions (TOCTTOU)
  - Symbolic links, relative paths, shared thread meta-data
- ◆ Unintended consequences of denying OS calls
  - Process dropped privileges using setuid but didn't check value returned by setuid... and monitor denied the call
- ◆ Bugs in reference monitors and safety checks
  - What if runtime environment has a buffer overflow?

# Incorrectly Mirroring OS State

[Garfinkel]

Policy: “process can bind TCP sockets on port 80,  
but cannot bind UDP sockets”

6 = socket(UDP, ...)

Monitor: “6 is UDP socket”

7 = socket(TCP, ...)

Monitor: “7 is TCP socket”

close(7)

dup2(6,7)

Monitor's state now inconsistent with OS

bind(7, ...)

Monitor: “7 is TCP socket, Ok to bind”

Oops!

# TOCTTOU in Syscall Interposition

---

- ◆ User-level program makes a system call
  - Direct arguments in stack variables or registers
  - Indirect arguments are passed as pointers
- ◆ Wrapper enforces some security condition
  - Arguments are copied into kernel memory (why?) and analyzed and/or substituted by the syscall wrapper
- ◆ **What if arguments change right here?**
- ◆ If permitted by the wrapper, the call proceeds
  - Arguments are copied into kernel memory
  - Kernel executes the call


# Exploiting TOCTTOU Conditions

[Watson]

- ◆ Page fault on an indirect syscall argument
  - Force wrapper to wait on disk I/O and use a concurrent process to replace already checked arguments
  - Example: `rename()` – see Watson's paper
    - Page out the target path of `rename()` to disk
    - Wrapper checks the source path, then waits for target path
    - Concurrent attack process replaces the source path
- ◆ Voluntary thread sleeps
  - Example: `TCP connect()` – see Watson's paper
    - Kernel copies in the arguments
    - Thread calling `connect()` waits for a TCP ACK
    - Concurrent attack process replaces the arguments

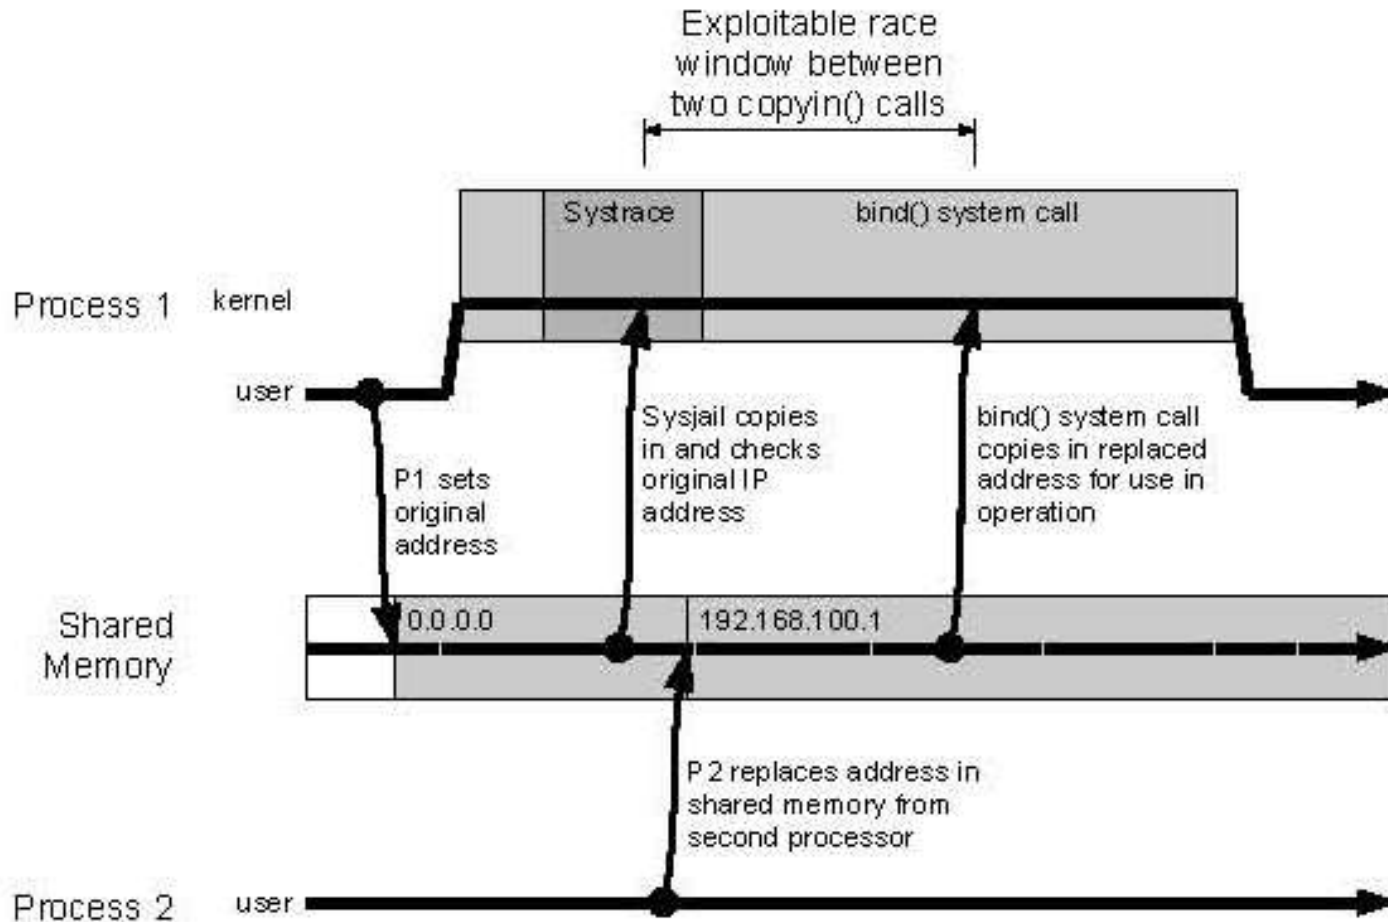
# TOCTTOU via a Page Fault

[Watson]



# TOCTTOU on Sysjail

[Watson]





# Mitigating TOCTTOU

---

- ◆ Make pages with syscall arguments read-only
  - Tricky implementation issues
  - Prevents concurrent access to data on the same page
- ◆ Avoid shared memory between user process, syscall wrapper and the kernel
  - Argument caches used by both wrapper and kernel
  - Message passing instead of argument copying
    - Why does this help?
- ◆ System transactions
- ◆ Integrate security checks into the kernel
  - Requires OS modification!

# Interposition + Static Analysis

---

Assumption: attack requires making system calls

1. Analyze the program to determine its expected behavior
2. Monitor actual behavior
3. Flag an intrusion if there is a deviation from the expected behavior
  - System call trace of the application is constrained to be consistent with the source or binary code
  - Main advantage: a conservative model of expected behavior will have zero false positives

# Runtime Monitoring

---

- ◆ One approach: run slave copy of application
  - Replication is hard; lots of non-determinism in code
    - Random number generation, process scheduling, interaction with outside environment
  - Slave is exposed to the same risks as master
    - Any security flaw in the master is also present in the slave
  - Virtual machines make problem easier!
- ◆ Another approach: simulate control flow of the monitored application

# Trivial “Bag-O’Calls” Model

---

- ◆ Determine the set  $S$  of all system calls that an application can potentially make
  - Lose all information about relative call order
- ◆ At runtime, check for each call whether it belongs to this set
- ◆ Problem: large number of false negatives
  - Attacker can use any system call from  $S$
- ◆ Problem:  $|S|$  very big for large applications

# Callgraph Model


[Wagner and Dean]

- ◆ Build a **control-flow graph** of the application by static analysis of its source or binary code
- ◆ Result: **non-deterministic finite-state automaton (NFA)** over the set of system calls
  - Each vertex executes at most one system call
  - Edges are system calls or empty transitions
  - Implicit transition to special “Wrong” state for all system calls other than the ones in original code
  - All other states are accepting
- ◆ System call automaton is conservative
  - **No false positives!**

# NFA Example

[Wagner and Dean]


```
f(int x) {  
  x ? getuid() : geteuid();  
  x++;  
}  
g() {  
  fd = open("foo", O_RDONLY);  
  f(0); close(fd); f(1);  
  exit(0);  
}
```



- No false positives
- Monitoring is  $O(|V|)$  per system call
- Problem: attacker can exploit impossible paths
  - The model has no information about stack state!

# Another NFA Example


[Giffin]



```
void
mysetuid (uid_t uid)
{
    setuid(uid);
    log("Set UID", 7);
}
```

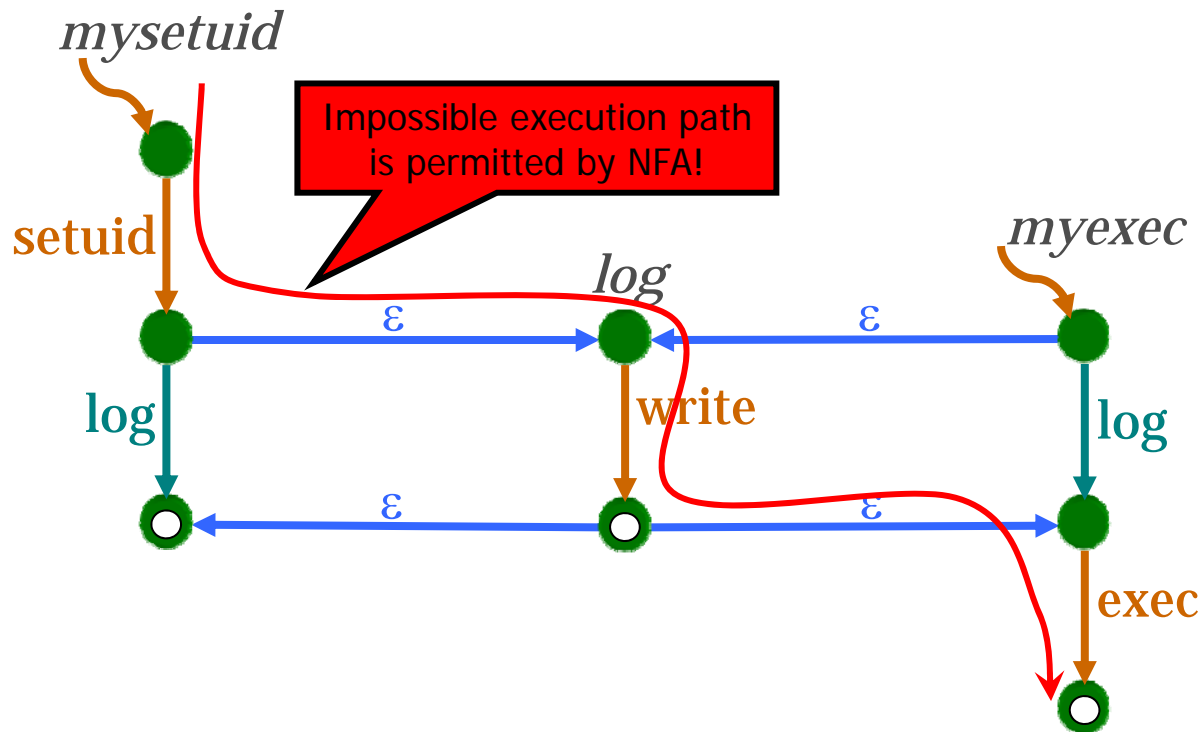


```
void
log (char *msg,
     int len)
{
    write(fd, msg, len);
}
```



```
void
myexec (char *src)
{
    log("Execing", 7);
    exec("/bin/ls");
}
```

# NFA Permits Impossible Paths





# NFA: Modeling Tradeoffs

◆ A good model should be...

- **Accurate:** closely models expected execution
  - Need context sensitivity!
- **Fast:** runtime verification is cheap

	<i>Inaccurate</i>	<i>Accurate</i>
<i>Slow</i>		
<i>Fast</i>	<b>NFA</b>	

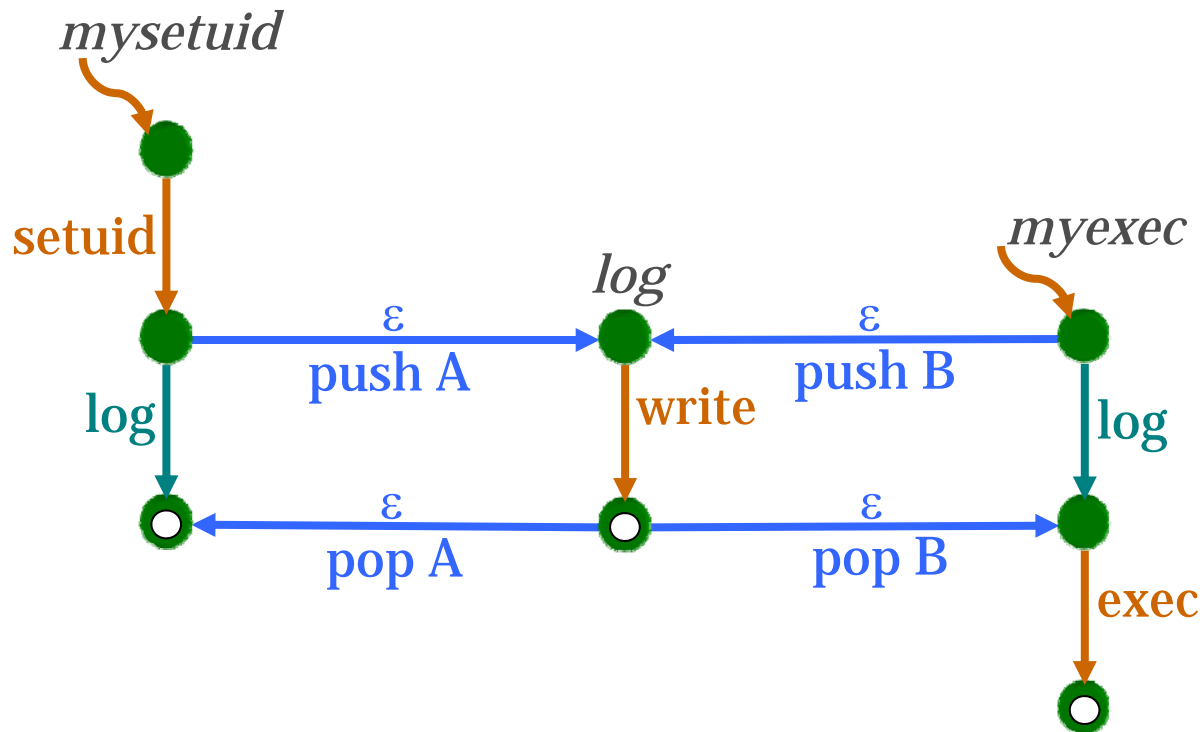
# Abstract Stack Model

---

- ◆ NFA is not precise, loses stack information
- ◆ Alternative: model application as a **context-free language** over the set of system calls
  - Build non-deterministic pushdown automaton (PDA)
  - Each symbol on the PDA stack corresponds to single stack frame in the actual call stack
  - All valid call sequences accepted by PDA; enter “Wrong” state when an impossible call is made

# PDA Example

[Giffin]



# Another PDA Example

[Wagner and Dean]

```
f(int x) {
  x ? getuid() : geteuid();
  x++;
}
g() {
  fd = open("foo", O_RDONLY);
  f(0); close(fd); f(1);
  exit(0);
}
```

```
Entry(f) ::= getuid() Exit(f)
           | geteuid() Exit(f)
Exit(f)   ::= ε
Entry(g)  ::= open() v
           v   ::= Entry(f) v'
           v'  ::= close() w
           w   ::= Entry(f) w'
           w'  ::= exit() Exit(g)
Exit(g)   ::= ε
```

```
while (true)
  case pop() of
    Entry(f) ⇒ push(Exit(f)); push(getuid())
    Entry(f) ⇒ push(Exit(f)); push(geteuid())
    Exit(f)  ⇒ no-op
    Entry(g) ⇒ push(v); push(open())
    v        ⇒ push(v'); push(Entry(f))
    v'       ⇒ push(w); push(close())
    w        ⇒ push(w'); push(Entry(f))
    w'       ⇒ push(Exit(g)); push(exit())
    Exit(g)  ⇒ no-op
    a ∈ Σ    ⇒ read and consume a from the input
    otherwise ⇒ enter the error state, Wrong
```

# PDA: Modeling Tradeoffs

- ◆ Non-deterministic PDA has high cost
  - Forward reachability algorithm is cubic in automaton size
  - Unusable for online checking

	<i>Inaccurate</i>	<i>Accurate</i>
<i>Slow</i>		PDA
<i>Fast</i>	NFA	

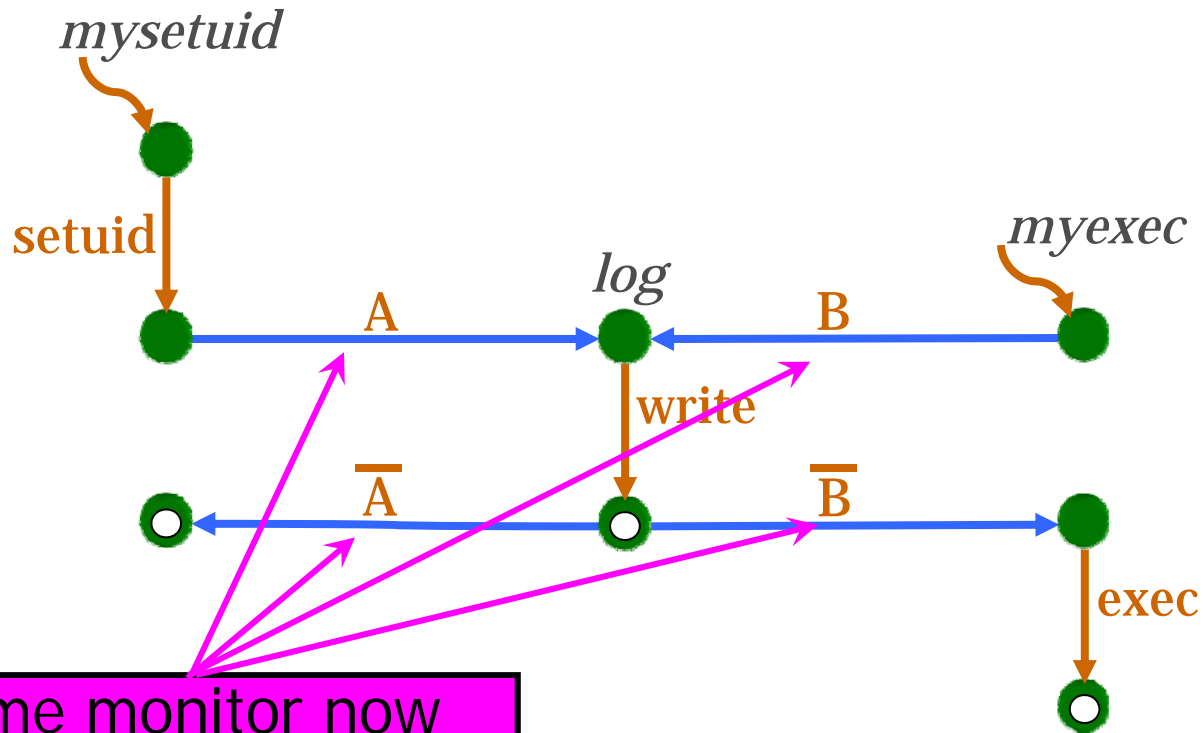
# Dyck Model

[Giffin et al.]

- ◆ Idea: make stack updates (i.e., function calls) explicit symbols in the automaton alphabet
  - Result: stack-deterministic PDA
- ◆ At each moment, the monitor knows where the monitored application is in its call stack
  - Only one valid stack configuration at any given time
- ◆ How does monitor learn about function calls?
  - Use binary rewriting to instrument the code to issue special “null” system calls to notify the monitor
    - Potential high cost of introducing many new system calls
  - Can't rely on instrumentation if application is corrupted

# Example of Dyck Model

[Giffin]



Runtime monitor now  
"sees" these transitions

# CFG Extraction Issues

---

## ◆ Function pointers

- Every pointer could refer to any function whose address is taken

## ◆ Signals

- Pre- and post-guard extra paths due to signal handlers

## ◆ `setjmp()` and `longjmp()`

- At runtime, maintain list of all call stacks possible at a `setjmp()`
- At `longjmp()` append this list to current state



# System Call Processing Complexity

<i>Model</i>	<i>Time &amp; Space Complexity</i>
NFA	$O(n)$
PDA	$O(nm^2)$
Dyck	$O(n)$

$n$  is state count

$m$  is transition count

# Dyck: Runtime Overheads

Execution times in seconds

Program	Unverified execution	Verified against Dyck	Increase
procmail	0.5	0.8	56%
gzip	4.4	4.4	1%
eject	5.1	5.2	2%
fdformat	112.4	112.4	0%
cat	18.4	19.9	8%


- ◆ Many tricks to improve performance
  - Use static analysis to eliminate unnecessary null system calls
  - Dynamic “squelching” of null calls

# Persistent Interposition Attacks

[Parampalli et al.]

- ◆ Observation: malicious behavior need not involve system call anomalies
- ◆ Hide malicious code inside a server
  - Inject via a memory corruption attack
  - Hook into a normal execution path (how?)
- ◆ Malicious code communicates with its master by “piggybacking” on normal network I/O
  - No anomalous system calls
  - No anomalous arguments to any calls except those that read and write

# Virtual Machine Monitors




Software layer between hardware and OS  
virtualizes and manages hardware resources

# History of Virtual Machines

---

- ◆ IBM VM/370 – A VMM for IBM mainframe
  - Multiple OS environments on expensive hardware
  - Desirable when few machines around
- ◆ Popular research idea in 1960s and 1970s
  - Entire conferences on virtual machine monitors
  - Hardware/VMM/OS designed together
- ◆ Interest died out in the 1980s and 1990s
  - Hardware got cheap
  - OS became more more powerful (e.g., multi-user)

# VMware



VMware virtual machine is an application execution environment with its own operating system

# VM Terminology

---

- ◆ **VMM** (Virtual Machine Monitor) – software that creates VMs
- ◆ **Host** – system running the VMM
- ◆ **Guest** – “monitored host” running inside a VM

# Isolation at Multiple Levels

---

## ◆ Data security

- Each VM is managed independently
  - Different OS, disks (files, registry), MAC address (IP address)
  - Data sharing is not possible; mandatory I/O interposition

## ◆ Fault isolation

- Crashes are contained within a VM

## ◆ Performance


- Can guarantee performance levels for individual VMs on VMWare ESX server

## ◆ No assumptions required for software inside a VM (important for security!)



# Observation by Host System

- ◆ “See without being seen” advantage
  - Very difficult within a computer, possible on host
- ◆ Observation points
  - Networking (through vmnet), physical memory, disk I/O and any other I/O

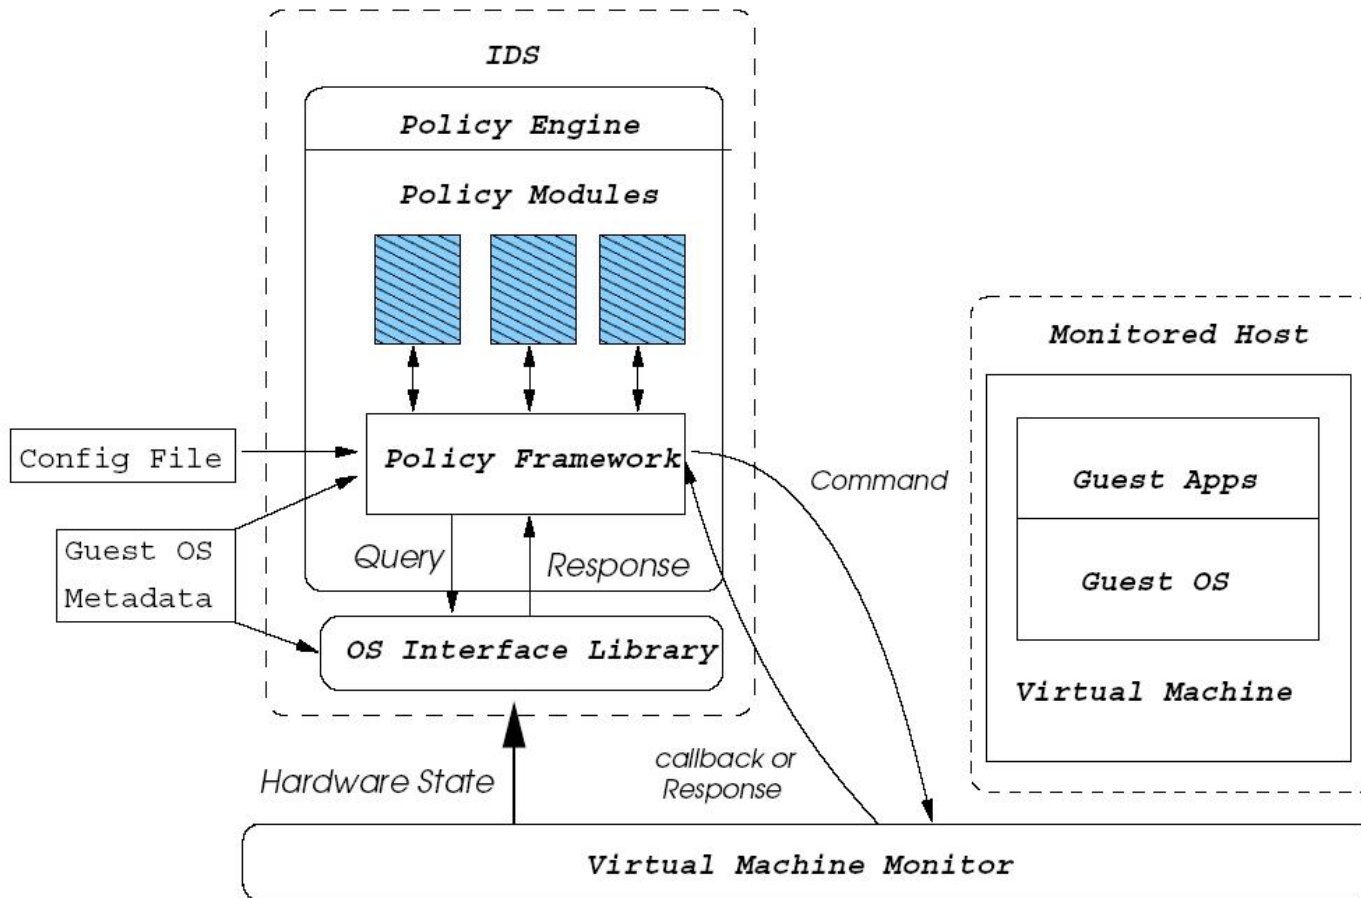


# Virtual Machine-Based IDS

---

- ◆ Run the monitored host within a **virtual machine (VM) sandbox** on a different host
- ◆ Run the intrusion detection system (IDS) outside the VM
- ◆ Allow the IDS to pause the VM and inspect the hardware state of host
- ◆ Policy modules determine if the state is good or bad, and how to respond

# VMI IDS Architecture



# Components of VMI IDS

---

## ◆ OS interface library

- Interprets hardware state into OS-level events
  - For example, list of all processes
- Hard to implement and tied to a particular guest OS

## ◆ Policy modules

- Determine if the OS has been compromised and what action to take
- Many detection techniques can be implemented as policy modules and thus used to prevent intrusions

# Livewire

[Garfinkel and Rosenblum]

- ◆ VMware + interposition/inspection hooks
- ◆ Six sample policy modules
  - Most related to existing host-based intrusion detection techniques
- ◆ Uses **crash** as OS interface library
  - Linux crash dump examination tool
  - Same goal: interpret the host's raw memory in terms of OS-level events

# Livewire Policy Modules

---

- ◆ User program integrity detector
  - Periodically hashes unchanging sections of running programs, compares to those of known good originals
- ◆ Signature detector
  - Periodically scans guest memory for substrings belonging to known malware
    - Finds malware in unexpected places, like filesystem cache
- ◆ Lie detector
  - Detects inconsistencies between hardware state and what is reported by user-level programs (ls, netstat, ...)
- ◆ Raw socket detector

# Enforcing Confinement Policies

---

- ◆ **Event-driven modules** run in response to a change in hardware state
- ◆ Memory access enforcer
  - Prevents sensitive portions of the kernel from being modified
- ◆ NIC access enforcer
  - Prevents the guest's network interface card (NIC) from entering promiscuous mode or having a non-authorized MAC address

# Other Advantages

---

- ◆ Resistant to “swap-out” malware
  - Sophisticated malware might detect an IDS running on the infected host and remove itself from memory when a scan is performed
  - Difficult to detect a scanner when the scanner is running outside your VM
- ◆ Can save entire system state for forensics
- ◆ Fail closed



# VM Issues

---

- ◆ Guest can often tell that he is running inside a VM
  - Timing: measure time required for disk access
    - VM may try to run clock slower to prevent this attack...
    - ...but slow clock may break an application like music player
- ◆ Better visibility into guest  $\Rightarrow$  worse performance
- ◆ OS interface library is complex, can be fooled
- ◆ **Policy problem**
  - Even with perfect visibility into the monitored system, how to differentiate good and bad behavior?
  - Hard to express and enforce fine-grained policies
    - “Do not write any SSNs from payroll into Web database”