

# Selecting and Improving System Call Models for Anomaly Detection

(or, 30 minutes before CIPHER 5's CTF results)

Alessandro Frossi

Federico Maggi

Gian Luigi Rizzo

Stefano Zanero

July 10, 2009

Topic of this talk

# System Call Based Anomaly Detection

Detecting intrusions using system call flows w/ data models

# System Call Based Anomaly Detection

Detecting intrusions using system call flows w/ data models

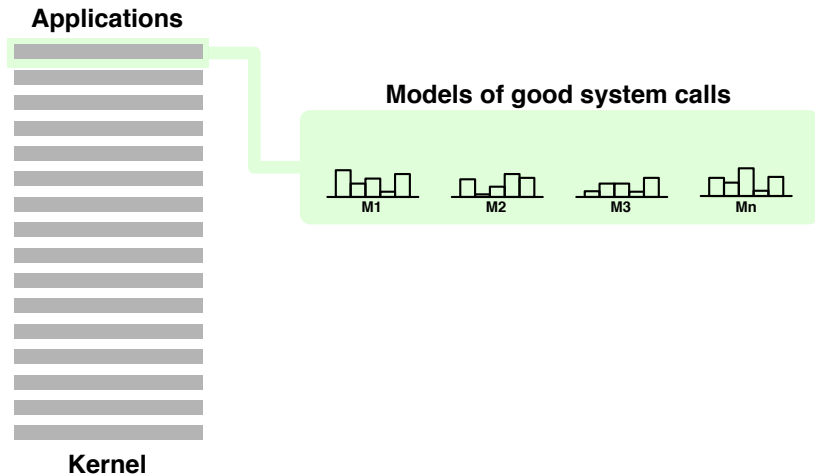


1. run applications into processes
2. intercept system calls
3. create models of good system calls
4. flag deviations to detect anomalies

**Let's take a look at a simple, generic example.**

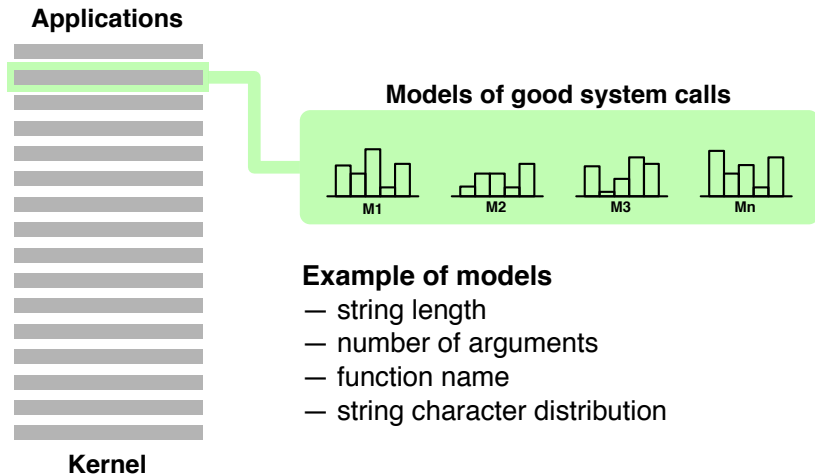
# System Call Based Anomaly Detection

A set of models is created based on certain features of the system calls



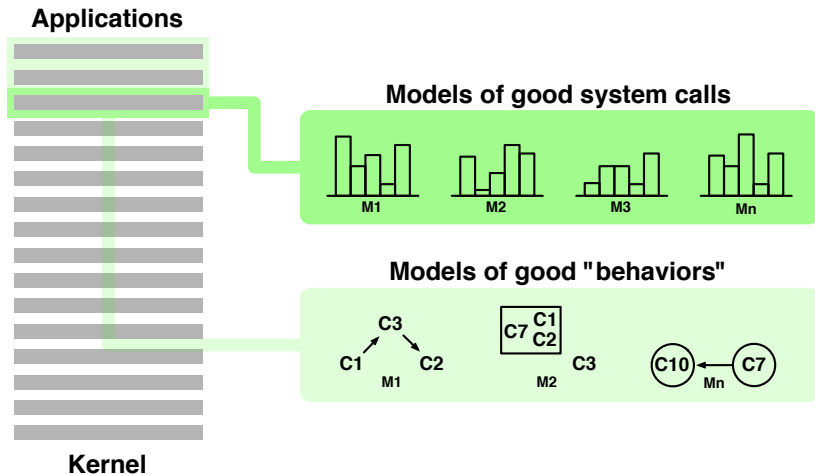
# System Call Based Anomaly Detection

Models estimate feature values observed in “good” system calls



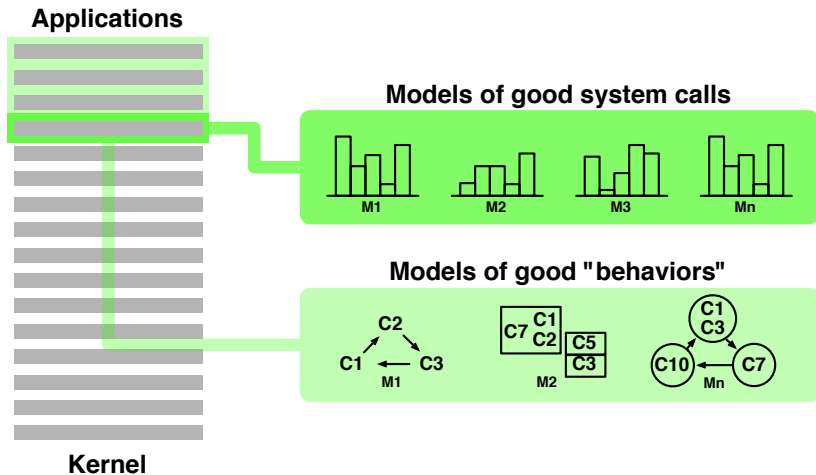
# System Call Based Anomaly Detection

Estimations become more accurate as more system calls are analyzed



# System Call Based Anomaly Detection

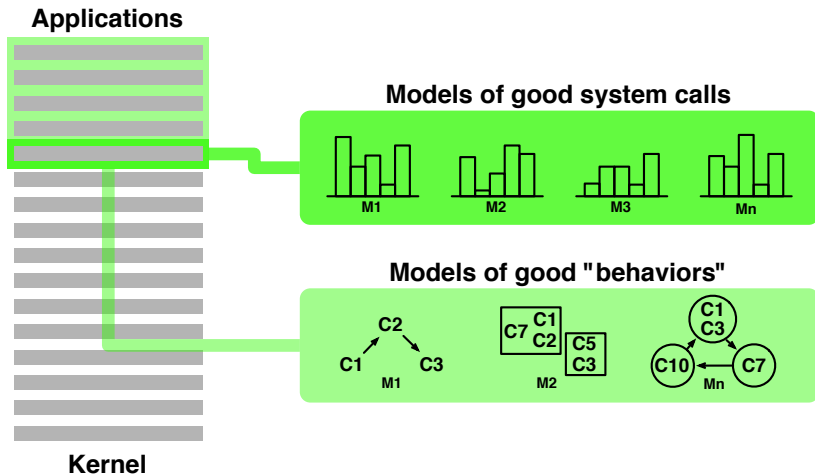
Also, models based on sets of system calls can be constructed





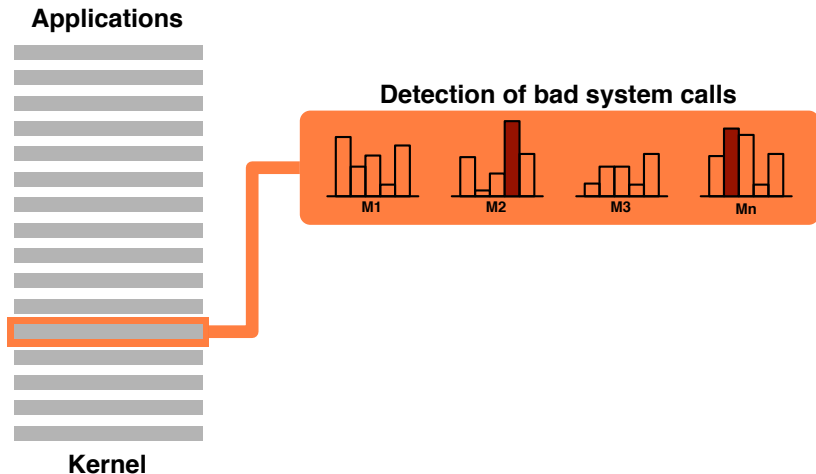
# System Call Based Anomaly Detection

Knowledge about system calls' context is learned



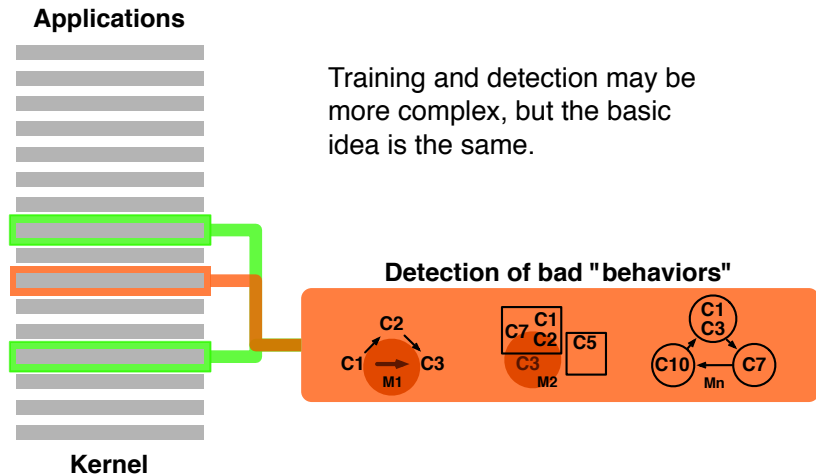
# System Call Based Anomaly Detection

In detection mode, the same models can be used to spot malicious system calls...



# System Call Based Anomaly Detection

...or malicious execution contexts



## The systems we analyzed

# Different Approaches: Deterministic vs. Stochastic

We analyzed two anomaly detectors based on different approaches

# Different Approaches: Deterministic vs. Stochastic

We analyzed two anomaly detectors based on different approaches

## **FSA-DF** [IEEE S&P 2006]

- ▶ Deterministic
- ▶ Control-flow: FSA
- ▶ Data-flow: unary/binary relations

# Different Approaches: Deterministic vs. Stochastic

We analyzed two anomaly detectors based on different approaches

## **FSA-DF** [IEEE S&P 2006]

- ▶ Deterministic
- ▶ Control-flow: FSA
- ▶ Data-flow: unary/binary relations

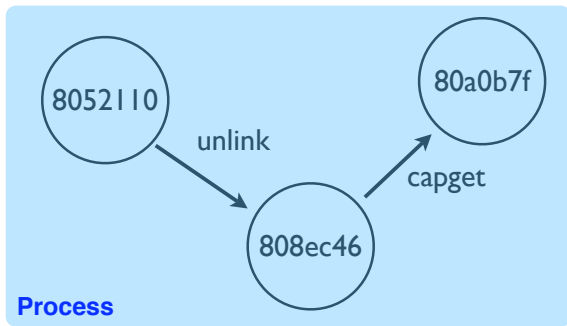
## **S<sup>2</sup>A<sup>2</sup>DE** [IEEE TODS 2009]

- ▶ Stochastic
- ▶ Control-flow: Markov-chain
- ▶ Data models: clusters

# Deterministic Data-flow Anomaly Detection

The system calls generated by each process are examined

```
5866 8052110 unlink("/usr/local/var/proftpd/test.sock") = 0
5866 808ec46 capget(DONT_CARE, DONT_CARE) = 0
5866 80a0b7f timer_gettime (DONT_CARE, DONT_CARE) = 3
```






# Deterministic Data-flow Anomaly Detection

Different PCs means different process states


```
5866 8052110 unlink("/usr/local/var/proftpd/test.sock") = 0
5866 808ec46 capget(DONT_CARE, DONT_CARE) = 0
5866 80a0b7f timer_gettime (DONT_CARE, DONT_CARE) = 3
```



# Deterministic Data-flow Anomaly Detection

A system call changes the process' state...


```
5866 8052110 unlink("/usr/local/var/proftpd/test.sock") = 0
5866 808ec46 capget(DONT_CARE, DONT_CARE) = 0
5866 80a0b7f timer_gettime (DONT_CARE, DONT_CARE) = 3
```



# Deterministic Data-flow Anomaly Detection

...and so forth


```
5866 8052110 unlink("/usr/local/var/proftpd/test.sock") = 0
5866 808ec46 capget(DONT_CARE, DONT_CARE) = 0
5866 80a0b7f timer_gettime (DONT_CARE, DONT_CARE) = 3
```



# Deterministic Data-flow Anomaly Detection

This analysis is repeated until termination

```
5866 8052110 unlink("/usr/local/var/proftpd/test.sock") = 0
5866 808ec46 capget(DONT_CARE, DONT_CARE) = 0
5866 80a0b7f timer_gettime (DONT_CARE, DONT_CARE) = 3
```

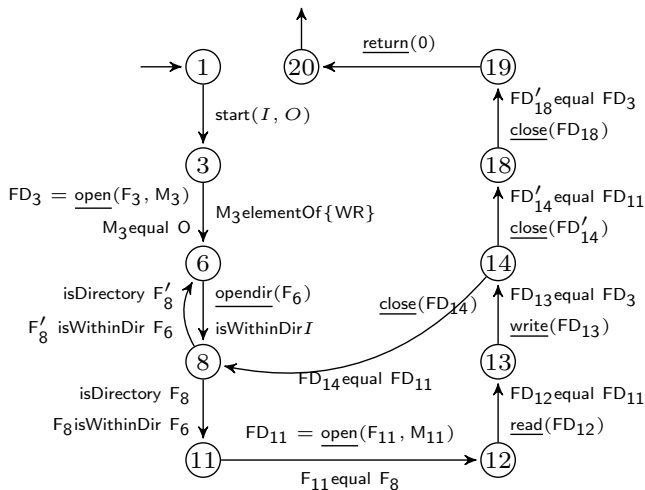


# Deterministic Data-flow Anomaly Detection

A network of unary/binary data-flow relations on top of the process' FSA

# Deterministic Data-flow Anomaly Detection

A network of unary/binary data-flow relations on top of the process' FSA



# Deterministic Data-flow Anomaly Detection

## Other types of relations

- ▶ **Unary**: capture properties of a single argument.
  - ▶ **equal**
  - ▶ **range**
  - ▶ **elementOf**
  - ▶ **isWithinDir**
  - ▶ **subsetOf**
  - ▶ **hasExtension**
- ▶ **Binary**: capture relations between two arguments.
  - ▶ **equal**
  - ▶ **hasSameDirAs**
  - ▶ **isWithinDir**
  - ▶ **hasSameBaseAs**
  - ▶ **contains**
  - ▶ **hasSameExtensionAs**

# Major Drawback: False Positives

Mostly due to the deterministic relations



# Major Drawback: False Positives

Mostly due to the deterministic relations

`open("/tmp/php1553", 0, 0x1b6) = 5`

unary **elementOf**({/tmp/php1553, /tmp/php9022})

`open("/tmp/php9022", 0, 0x1b6) = 5`

# Major Drawback: False Positives

Mostly due to the deterministic relations

`open("/tmp/php1553", 0, 0x1b6) = 5`

unary **elementOf**({/tmp/php1553, /tmp/php9022})

`open("/tmp/php9022", 0, 0x1b6) = 5`

What if `"/tmp/php1990"` is found?

# Major Drawback: False Positives

Mostly due to the deterministic relations

`open("/tmp/php1553", 0, 0x1b6) = 5`

unary **elementOf**({/tmp/php1553, /tmp/php9022})

`open("/tmp/php9022", 0, 0x1b6) = 5`

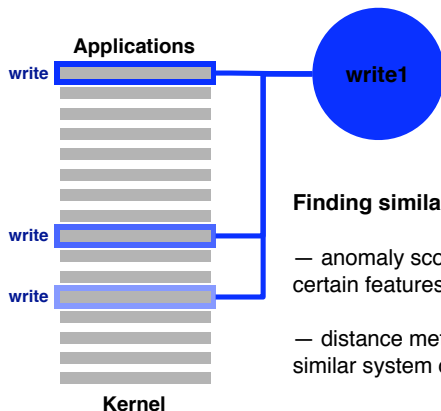
What if `"/tmp/php1990"` is found?  
These false positives occur pretty often.

# Stochastic Behavior Profiling of Processes

Clusters of similar system calls interconnected by Markov-chains

# Stochastic Behavior Profiling of Processes

Clusters of similar system calls interconnected by Markov-chains

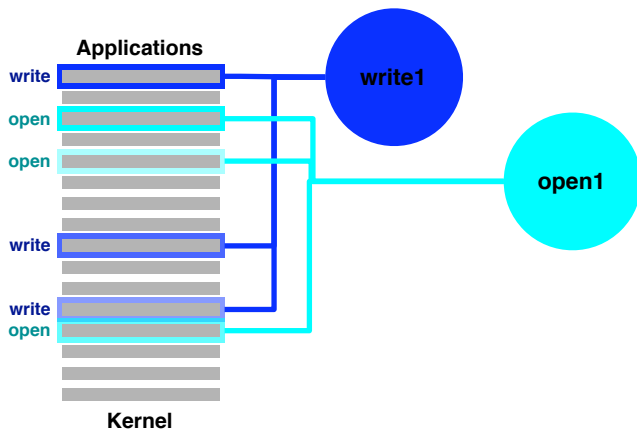


## Finding similar system calls

- anomaly scores [ACM TISSEC 2006] based on certain features of the arguments
- distance metrics [ACM TODS 2009] used to cluster similar system calls
- each application's process creates different clusters

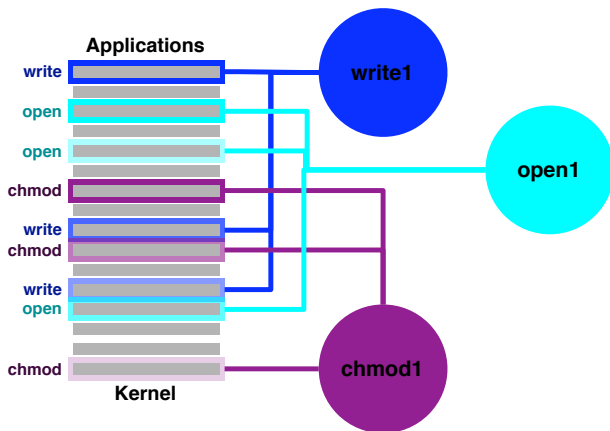
# Stochastic Behavior Profiling of Processes

Clusters of similar system calls interconnected by Markov-chains




# Stochastic Behavior Profiling of Processes

Clusters of similar system calls interconnected by Markov-chains



# Stochastic Behavior Profiling of Processes


Clusters of similar system calls interconnected by Markov-chains





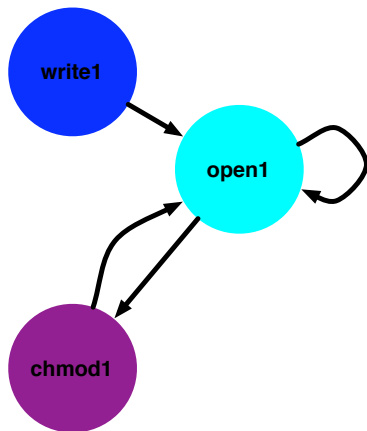
# Stochastic Behavior Profiling of Processes

Clusters of similar system calls interconnected by Markov-chains



# Stochastic Behavior Profiling of Processes

Clusters of similar system calls interconnected by Markov-chains




# Stochastic Behavior Profiling of Processes

Clusters of similar system calls interconnected by Markov-chains

## Markov-chains encode process behavior

- transitions between system calls can occur with different probabilities [ACM TODS 2009]
- a call is anomalous if either there is no matching state (i.e., cluster) or transition probability is violated



# Major Drawbacks: False Negatives

Mostly due to the stochastic nature of Markov-chains

# Major Drawbacks: False Negatives

Mostly due to the stochastic nature of Markov-chains

## Clustering

- ▶ clustering depends on configuration parameters
- ▶ different parameters  $\rightarrow$  different results


# Major Drawbacks: False Negatives

Mostly due to the stochastic nature of Markov-chains

## Clustering

- ▶ clustering depends on configuration parameters
- ▶ different parameters  $\rightarrow$  different results

## Markov-chains (example)




# Major Drawbacks: False Negatives

Mostly due to the stochastic nature of Markov-chains

## Clustering

- ▶ clustering depends on configuration parameters
- ▶ different parameters → different results

## Markov-chains (example)



$$\text{Threshold} = 0.5 * 1 * 0.5 = 0.25$$


# Major Drawbacks: False Negatives

Mostly due to the stochastic nature of Markov-chains

## Clustering

- ▶ clustering depends on configuration parameters
- ▶ different parameters → different results

## Markov-chains (example)



$$\text{Threshold} = 0.5 * 1 * 0.5 = 0.25$$




# Major Drawbacks: False Negatives

Mostly due to the stochastic nature of Markov-chains

## Clustering

- ▶ clustering depends on configuration parameters
- ▶ different parameters → different results

## Markov-chains (example)



$$\text{Threshold} = 0.5 * (1 * 0.5)^n * 0.5 \rightarrow 0$$


# Major Drawbacks: False Negatives

Mostly due to the stochastic nature of Markov-chains

## Clustering

- ▶ clustering depends on configuration parameters
- ▶ different parameters  $\rightarrow$  different results

## Markov-chains (example)



$$\text{Threshold} = 0.5 * (1 * 0.5)^n * 0.5 \rightarrow 0$$

No valid threshold can be found if cycles are not of fixed length.  
For instance, DoS attacks may not be detected.

# Pros and cons of the two approaches

# Pros and cons of the two approaches

	FSA-DF		S <sup>2</sup> A <sup>2</sup> DE	
FSA	<ul style="list-style-type: none"><li>● Perfectly models a software behavior</li><li>● No False Negatives</li><li>● Doesn't allow deviations</li></ul>	Control Flow	<ul style="list-style-type: none"><li>● Introduces a statistical approach</li><li>● False Negatives</li><li>● Few False Positives</li></ul>	MCM
Relations	<ul style="list-style-type: none"><li>● Deterministic approach</li><li>● No new input adaptation</li><li>● Prone to False Positives</li><li>● No False Negatives</li></ul>	Data Flow	<ul style="list-style-type: none"><li>● Stochastic approach</li><li>● Can adapt to new inputs</li><li>● Few False Positives</li><li>● False Negatives</li></ul>	Clusters

First contribution:  
combination of the two approaches

# Combining Complementary Approaches

Deterministic control-flow + stochastic data models

Hybrid IDS			
	FSA-DF	S <sup>2</sup> A <sup>2</sup> DE	
FSA	<ul style="list-style-type: none"><li>• Perfectly models a software behavior</li><li>• No False Negatives</li><li>• Doesn't allow deviations</li></ul>	Control Flow	
		Data Flow	<ul style="list-style-type: none"><li>• Stochastic approach</li><li>• Can adapt to new inputs</li><li>• Few False Positives</li><li>• False Negatives</li></ul>
			Models

# Combining Complementary Approaches

The learning algorithm is similar to that used in FSA-DF

- ▶  $\forall \text{couple} \langle \text{syscall}_{i-1}, \text{syscall}_i \rangle \in \{ \text{TrainingSet} \}$ 
  - ▶ make state
  - ▶ learn relations

# Combining Complementary Approaches

The learning algorithm is similar to that used in FSA-DF

- ▶  $\forall \text{couple} \langle \text{syscall}_{i-1}, \text{syscall}_i \rangle \in \{ \text{TrainingSet} \}$ 
  - ▶ make state
  - ▶ learn relations
    - ▶ equal
    - ▶ elementOf
    - ▶ subsetOf
    - ▶ range
    - ▶ isWithinDir
    - ▶ hasExtension
    - ▶ isWithinDir
    - ▶ contains
    - ▶ hasSameDirAs
    - ▶ hasSameBaseAs
    - ▶ hasSameExtensionAs



# Combining Complementary Approaches

The learning algorithm is similar to that used in FSA-DF

- ▶  $\forall \text{couple} \langle \text{syscall}_{i-1}, \text{syscall}_i \rangle \in \{ \text{TrainingSet} \}$ 
  - ▶ make state
  - ▶ learn relations
    - ▶ equal save model of similar strings
    - ▶ elementOf save model of similar strings
    - ▶ subsetOf
    - ▶ range
    - ▶ isWithinDir
    - ▶ hasExtension
    - ▶ isWithinDir
    - ▶ contains save model of similar strings
    - ▶ hasSameDirAs
    - ▶ hasSameBaseAs
    - ▶ hasSameExtensionAs

# Combining Complementary Approaches

The learning algorithm is similar to that used in FSA-DF

- ▶ learn string domains
- ▶  $\forall \text{couple} \langle \text{syscall}_{i-1}, \text{syscall}_i \rangle \in \{ \text{TrainingSet} \}$ 
  - ▶ make state
  - ▶ learn relations
    - ▶ equal save model of similar strings
    - ▶ elementOf save model of similar strings
    - ▶ subsetOf
    - ▶ range
    - ▶ isWithinDir
    - ▶ hasExtension
    - ▶ isWithinDir
    - ▶ contains save model of similar strings
    - ▶ hasSameDirAs
    - ▶ hasSameBaseAs
    - ▶ hasSameExtensionAs

# Paths And Filenames

How to find groups of good strings into `execve/open/read/...` args?

```
/var/log/http.0 ... /etc/ftp.conf ... /tmp/php1231  
... /var/run/nfsd.pid ... /etc/smb/samba.conf  
... /opt/local/lib/libncurses.a ... /usr/lib/libkmod.a  
... /tmp/uscreens/427.ttys000 ... /var/db/ntp.drift ...
```

# Paths And Filenames

How to find groups of good strings into `execve/open/read/...` args?

```
/var/log/http.0 ... /etc/ftp.conf ... /tmp/php1231  
... /var/run/nfsd.pid ... /etc/smb/samba.conf  
... /opt/local/lib/libncurses.a ... /usr/lib/libkmod.a  
... /tmp/uscreens/427.ttys000 ... /var/db/ntp.drift ...
```

## Self-Organizing Map

Type of artificial neural network, trained using unsupervised learning to produce a multi dimensional discretized representation of the input space of the training samples, called map.

### Idea

- ▶ SOM to capture classes of good strings.
- ▶ Model of good strings → nodes.
- ▶ Similar strings → neighbor nodes.

# Paths And Filenames

How to find groups of good strings into `execve/open/read/...` args?


```
/var/log/http.0 ... /etc/ftp.conf ... /tmp/php1231  
... /var/run/nfsd.pid ... /etc/smb/samba.conf  
... /opt/local/lib/libncurses.a ... /usr/lib/libkmod.a  
... /tmp/uscreens/427.ttys000 ... /var/db/ntp.drift ...
```


## Self-Organizing Map

Type of artificial neural network, trained using unsupervised learning to produce a multi dimensional discretized representation of the input space of the training samples, called map.

### Idea

- ▶ SOM to capture classes of good strings.
- ▶ Model of good strings → nodes.
- ▶ Similar strings → neighbor nodes.





OK, they look pretty nice. But why SOMs?

# Paths And Filenames

## Integration in Hybrid IDS - algorithm

- ▶ create SOM of all paths
  - ▶ SOM initialization with linux directory structure.
  - ▶ Extract all the paths from the syscalls
  - ▶ SOM training [Kohonen 2004] with a randomized subset of the paths.
- ▶  $\forall \text{couple} \langle \text{syscall}_{i-1}, \text{syscall}_i \rangle \in \{ \text{TrainingSet} \}$ 
  - ▶ make state
  - ▶ learn relations
    - ▶ if  $\text{syscall}_{i-1}$  contains a path argument
      - find BMU from the SOM
      - add BMU to the edge
    - ▶ subsetOf
    - ▶ range
    - ▶ isWithinDir
    - ▶ hasExtension
    - ▶ isWithinDir
    - ▶ hasSameDirAs
    - ▶ hasSameBaseAs
    - ▶ hasSameExtensionAs



## Second contribution: improved system call models

# Improved System Call Models


New models to reduce false detections

- ▶ **Goal 1: Resilience to spurious strings in the datasets.**

# Improved System Call Models

New models to reduce false detections


- ▶ **Goal 1: Resilience to spurious strings in the datasets.**
  - ▶ Long/short strings in the training data can bias interval based models.



# Improved System Call Models

New models to reduce false detections

- ▶ **Goal 1: Resilience to spurious strings in the datasets.**
  - ▶ Long/short strings in the training data can bias interval based models.




- ▶ **Goal 2: Detect simple DoS attacks.**


# Improved System Call Models

New models to reduce false detections

- ▶ **Goal 1: Resilience to spurious strings in the datasets.**
  - ▶ Long/short strings in the training data can bias interval based models.



- ▶ **Goal 2: Detect simple DoS attacks.**
  - ▶ i.e., process forced to execute the same code region until crash.



# Argument Length Using Gaussian Intervals

Yields to less false positives

# Argument Length Using Gaussian Intervals


Yields to less false positives

**Statistics:** to estimate the distribution of args length

- ▶  $|args| = X_{args} \sim \mathcal{N}(\mu, \sigma^2)$
- ▶ Sample Mean, Sample Variance.

**Model precision parameter:**

- ▶ Kurtosis  $\hat{\gamma}_X = \frac{\hat{\mu}_{X,4}}{\hat{\sigma}_X^4} - 3$
- ▶ If  $\gamma_{X_{args}} < 0$  the sample is spread on a big interval



**Anomaly threshold:** percentile  $T_{args}$  centered on the mean.

# Argument Length Using Gaussian Intervals

## Integration in Hybrid IDS - algorithm

- ▶ create SOM of all paths
- ▶  $\forall \text{couple} \langle \text{syscall}_{i-1}, \text{syscall}_i \rangle \in \{ \text{TrainingSet} \}$ 
  - ▶ make state
  - ▶ learn relations
    - ▶ save BMU
    - ▶ subsetOf
    - ▶ range save string length or num. value
    - ▶ isWithinDir
    - ▶ hasExtension
    - ▶ isWithinDir
    - ▶ hasSameDirAs
    - ▶ hasSameBaseAs
    - ▶ hasSameExtensionAs



# Mitigating DoS Using Edge Frequency Models


Yields to less false negatives

# Mitigating DoS Using Edge Frequency Models

Yields to less false negatives

## Given that:


- ▶ each FSA edge is traversed a variable number of times over multiple executions
- ▶ the traversal frequency has a range



**Idea:** estimate a validity interval to detect DoS attacks.

# Edge Traversal Frequency


## The Model




# Edge Traversal Frequency

## The Model

... multiple executions ...



- ▶  $f_{reqs} = X_{f_{reqs}} \sim \text{Beta}(\alpha, \beta)$
- ▶ Estimated  $\alpha$  and  $\beta$
- ▶ Interval from x-th percentile
- ▶ **Not** estimated if few values available



# Mitigating DoS Using Edge Frequency Models

## Integration in Hybrid IDS - algorithm

- ▶ create SOM of all paths
- ▶  $\forall \text{couple} \langle \text{syscall}_{i-1}, \text{syscall}_i \rangle \in \{ \text{TrainingSet} \}$ 
  - ▶ make state
  - ▶ learn relations
    - ▶ save BMU
    - ▶ subsetOf
    - ▶ save string length or num. value
    - ▶ isWithinDir
    - ▶ hasExtension
    - ▶ isWithinDir
    - ▶ hasSameDirAs
    - ▶ hasSameBaseAs
    - ▶ hasSameExtensionAs
  - ▶ save edge traverse count

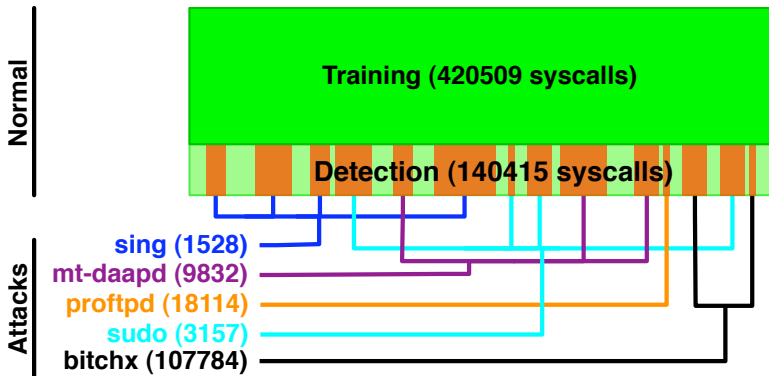
# How We Built The Evaluation Dataset

# How We Built The Evaluation Dataset

- ▶ “normal” execution of 5 tools (420509 syscalls)
- ▶ recent exploits from CVE (140415 syscalls)

# How We Built The Evaluation Dataset

- ▶ “normal” execution of 5 tools (420509 syscalls)
- ▶ recent exploits from CVE (140415 syscalls)





# Accuracy Evaluation

No false negatives (deterministic control-flow) + almost-zero false positives (stochastic data models)

	sing	mt-daapd	profdtpd	sudo	BitchX	mcweject	bsdtar	
<b>Traces</b>	22	18	21	22	15	12	2	
<b>Syscalls</b>	1528	9832	18114	3157	107784	75	102	
<b>S<sup>2</sup>A<sup>2</sup>DE</b>	10.0%	0%	0%	10.0%	0.0%	0.0%	8.7%	<b>S<sup>2</sup>A<sup>2</sup>DE</b>
<b>FSA-DS</b>	5.0%	16.7%	28%	15.0%	0.0%	0.0%	0.0%	<b>SOM-S<sup>2</sup>A<sup>2</sup>DE</b>
<b>Hybrid IDS</b>	0.0%	0%	0%	10.0%	0.0%			

**Table 2.** Comparison of the FPR of S<sup>2</sup>A<sup>2</sup>DE vs. FSA-DF vs. Hybrid IDS and S<sup>2</sup>A<sup>2</sup>DE vs. SOM-S<sup>2</sup>A<sup>2</sup>DE. Values include the number of traces used. Accurate description of the impact of each *individual* model is in Section 4.2 (first five columns) and 4.3 (last two columns).

# Accuracy Evaluation

No false negatives (deterministic control-flow) + almost-zero false positives (stochastic data models)

	sing	mt-daapd	proftpd	sudo	BitchX	mcweject	bsdtar		
<b>Traces</b>	22	18	21	22	15	12	2		
<b>Syscalls</b>	1528	9832	18114	3157	107784	75	102		
<b>S<sup>2</sup>A<sup>2</sup>DE</b>	10.0%	0%	0%	10.0%	0.0%	0.0%	8.7%	<b>S<sup>2</sup>A<sup>2</sup>DE</b>	
<b>FSA-DS</b>	5.0%	16.7%	28%	15.0%	0.0%	0.0%	0.0%	<b>SOM-S<sup>2</sup>A<sup>2</sup>DE</b>	
<b>Hybrid IDS</b>	0.0%	0%	0%	10.0%	0.0%				

**Table 2.** Comparison of the FPR of S<sup>2</sup>A<sup>2</sup>DE vs. FSA-DF vs. Hybrid IDS and S<sup>2</sup>A<sup>2</sup>DE vs. SOM-S<sup>2</sup>A<sup>2</sup>DE. Values include the number of traces used. Accurate description of the impact of each *individual* model is in Section 4.2 (first five columns) and 4.3 (last two columns).

- ▶ sing - write on arbitrary file (data-flow).
- ▶ mt-daapd - arbitrary code execution (data-flow + DoS).
- ▶ proftpd - arbitrary command execution (data-/control-flow).
- ▶ sudo - arbitrary command execution (control-flow).
- ▶ bitchx - arbitrary code execution (control-flow + DoS).

# Performance Evaluation

Not-so-negligible overhead, but mostly due to ptrace

	sing	sudo	BitchX	mcweject	bsdtar	Avg. speed
System calls	3470	15308	12319	97	705	
<b>S<sup>2</sup>A<sup>2</sup>DE</b>	0.4	0.8	1.9	0.1	0.1	8463
<b>FSA-DF</b>	1.3	1.5	1.2	-	-	7713
<b>Hybrid IDS</b>	29	5.8	27.7	-	-	1067
<b>SOM-S<sup>2</sup>A<sup>2</sup>DE</b>	-	-	-	8.8	19	25

**Table 3.** Detection performance measured in “seconds per system call”. The average speed is measured in system calls per second (last column).

# Conclusions and Future Works

Solve performance issues due to SOMs

- ▶ **deterministic** models accurately capture the **control**-flow
- ▶ **stochastic** models accurately capture **data**-flow features
- ▶ a **hybrid** approach lowers false detections
- ▶ performance issues:
  - ▶ the optimization of BMUs lookup is the first item on our TODO list
  - ▶ the use of a faster system call interceptor the second one ;)