

LAMBDA: Lightweight Assessment of Malware for emBeddeD Architectures

Research Objective

To propose a framework for runtime anomaly detection on embedded systems - The framework is capable of performing anomaly detection in a hierarchical manner (i.e. application level, operating system level and processor micro-architecture level) by harnessing the information available at various levels to detect malicious exploits.


Motivation




Measure the distance of a program under test from the characteristics of a given set of benign programs.

If the distance is less than a previously defined threshold value, the target program can be treated as a benign program otherwise the program is a malware.

HPC provides more sensitivity & Increased Protection



Run time Statistical T-test




$$t = \frac{\bar{X}_1 - \bar{X}_2}{\sigma \sqrt{\frac{1}{n_1} + \frac{1}{n_2}}}$$

The null hypothesis of two equal means is rejected when the test statistic |t| exceeds a threshold of 4.5, which ensures a confidence of 0.99999.

Detection Approach

Use (HPC, Indicator) for monitoring.

Control Flow




Dual core setup
- Watchdog Core: to monitor all the processes.
- Sanitized Core: to run non-malicious processes.

Data Collection

Total Target CPUs	13857	1878	156224	61535	8074
Feature	4874	322	335209	42281	1873
Feature	3661	325	323872	38559	3564
Feature	28424	218	130718	42039	273
Feature	3939	355	366466	5547	4785
Feature	4386	555	366466	6057	6057
Feature	4537	603	366466	6057	6057
Feature	7432	620	366466	6057	6057
Feature	5274	923	366466	32795	3817
Feature	5035	100	338325	52529	5253
Feature	5422	385	322298	39837	4104
Feature	27574	359	228230	49765	3403
Feature	4418	279	114252	36378	3274


Hardware Performance Counters
- More efficient to detect Kernel modifying rootkits.
- Easily accessible in most of the Linux based systems.

Normalized Weights



Determine weights during training
- Normalised weights help at runtime to determine distance of malware using statistical T-test.

Creating Bins



Emphasize on critical HPC-Indicator pair
- Performed to give more weightage on important performance counters and indicator programs.

Scoring at Runtime

Calculate the amount of maliciousness of a program under test
- Create bins for program under test at runtime.
- Multiplication of the trained weights with these bins produces score for the program under test.
- Score greater than a pre-defined threshold value signifies the malicious behaviour of the program.


Advantages for HPC observation:

- Difficult to manipulate HPC values by the malware.
- More sensitive when observed in conjunction with system calls.
- Results in better false positives and negatives

Advantages of the approach

- Monitoring only system calls doesn't provide any significant information but monitoring HPCs does. Significant changes can be observed in presence of malware.
- Enables semantic based malware detection.
- Supports multi-core environment.

T-test vs. Machine Learning based Detection



-- Box plot showing F1-Scores for various machine learning based techniques along with LAMBDA

-- It can be observed that variation in F1-Score is less for LAMBDA

-- Signifies LAMBDA is consistent over different train-test pairs

Advantage over Training and Detection Time

	Model Building Time (in milli seconds)	Detection Time (in milliseconds)
Multi Layer Perceptron	2036.91	10.24
Gaussian Naive Bayes	7.17	10.43
Random Forest	85.95	91.39
LAMBDA	43.72	15.38