

RAPPER: Ransomware Prevention via Performance Counters

Manaar Alam

Indian Institute of Technology Kharagpur
alam.manaar@iitkgp.ac.in

Debdeep Mukhopadhyay

Indian Institute of Technology Kharagpur
debdeep@cse.iitkgp.ernet.in

Sarani Bhattacharya

Indian Institute of Technology Kharagpur
sarani.bhattacharya@cse.iitkgp.ernet.in

Anupam Chattopadhyay

Nanyang Technological University Singapore
anupam@ntu.edu.sg

ABSTRACT

Ransomware can produce direct and controllable economic loss, which makes it one of the most prominent threats in cyber security. As per the latest statistics, more than half of malwares reported in Q1 of 2017 are ransomware and there is a potent threat of a novice cybercriminals accessing ransomware-as-a-service. The concept of public-key based data kidnapping and subsequent extortion was introduced in 1996. Since then, variants of ransomware emerged with different cryptosystems and larger key sizes though, the underlying techniques remained same. Though there are works in literature which proposes a generic framework to detect the crypto ransoms, we present a two step unsupervised detection tool which when suspects a process activity to be malicious, issues an alarm for further analysis to be carried in the second step and detects it with minimal traces. The two step detection framework RAPPER uses Artificial Neural Network and Fast Fourier Transformation to develop a highly accurate, fast and reliable solution to ransomware detection using minimal trace points.

KEYWORDS

Ransomware, Hardware Performance Counters, Time-Series, Fast Fourier Transformation, Autoencoder, Long-Short-Term-Memory

ACM Reference Format:

Manaar Alam, Sarani Bhattacharya, Debdeep Mukhopadhyay, and Anupam Chattopadhyay. 2018. RAPPER: Ransomware Prevention via Performance Counters. In *Proceedings of ACM Conference (Conference'17)*. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

If your organization has not been hit by ransomware yet, there are chances that it will soon be. The number of medium to large-scale business falling prey to ransom payment and extortion of their private databases have increased manifold. These malicious executables infect the victim machine and demands a ransom amount after encrypting the files and documents of the machine. In May 2017, WannaCry ransomware has affected approx. 200000 business across 150 countries. Identification, blocking of these ransoms at the earliest along with recovering the contents of the already encrypted files is already an open challenge.

Hardware Performance Counters (HPCs) were first introduced for checking the static and dynamic integrity of programs, for the purpose of detecting any malicious modifications to them as discussed in [8]. While in [2] performance counters are used to build a malware detector in hardware. Detecting malware which modifies the kernel control flow has been targeted in [10, 11]. The paper uses performance counters to monitor the system calls to detect

the vulnerability. However, detection of ransoms through the HPCs, to the best of our knowledge, has not been attempted so far. Though the underlying technique is similar [9], ransomware detection requires far more accuracy and faster response time to limit the damage.

A range of ransoms were studied in [5], which identified 15 different ransomware families. It is suggested that despite advancing encryption systems, the prominent ransoms leave a trait in the access of IO and file-systems. Accordingly, Kharraz *et al.* [4] proposed a technique of correlating high file system activity with the intrusion of ransomware, which, however, is susceptible to false positives and also can be defeated with a slow encryption process. Moreover, the technique requires modification in OS kernel, which may not be practical in many real scenarios. In a recent work, Kiraz *et al.* presented a technique, where large integer multiplication blocks are identified within an execution [6]. Since public-key cryptosystems rely on large integer multiplications, it can detect the threat at an early stage. Similar approaches, for detection of symmetric-key cryptographic primitives via data flow graph isomorphism [7] or by identifying characteristics of a cipher in a binary code [3], are also presented. In this paper, neither we target a specific family of ransoms nor the properties corresponding to a particular cipher implementation. Instead we develop a generic anomaly based approach based on the HPC statistics.

Motivation and Contribution

The primary contributions of this paper are listed below:

- The main objective of RAPPER is to learn the behavior of the system under observation with performance event statistics obtained from HPCs. Unlike other works in literature, which save the templates of malicious processes and matches it on its occurrence, here we allow our tool to learn the normal operating behavior of the system. The time-series data as observed from a selected cluster of HPC events is fed to an Artificial Neural Network to learn the specific characteristics of the data.
- Any deviation from this normal behavior as learned by the Autoencoder is considered as a suspect to RAPPER. Thus we observe that the performance statistics of the system in presence of ransomware are significantly dissimilar from the normal system behavior because of repeated encryption process.
- In the course of our study, we identified a benign benchmark application which can raise a false alarm. This is typically a benign process, but because of its high computational overhead the system behavior differ significantly from its normal behavior thus raising a false alarm.
- Thus in the final step, we transform the time series to frequency domain using Fast Fourier Transformation (FFT) and understand the repeatability of data with help of a second autoencoder.

RAPPER is a lightweight tool, which neither requires any root privilege, nor requires any hardware and kernel modification, thereby making it practical to use in almost every environment.

2 ANOMALY DETECTION BY ANALYSING THE SYSTEM BEHAVIOR

In this section, we first analyze the normal behavior of a system by monitoring some appropriately selected hardware performance counters in parallel. We then present a notion of anomalous activity in the system and demonstrate a detailed methodology for detecting those anomalies by using an Autoencoder.

2.1 Observing the system behavior using HPCs

The Hardware Performance Counters (HPCs) are a set of special purpose registers built into modern processors to dynamically observe the hardware related activities in a computer system. There are some recent works [10] which use these HPCs to detect malicious programs targeted for a particular system. The HPCs can be monitored dynamically with a user privilege using the well-known `perf` tool, available in Linux kernels 2.6.31+. One interesting property of the `perf` tool is that a user can observe the performance counters associated with a system with some time interval, thereby giving the benefit of observing the system behavior continuously in a succession of time. The command to monitor a particular HPC event in such way is as follows:

```
perf stat -e <event_name> -I <time_interval> <executable_name>
```

Since our objective is to detect the presence of ransomware, which mainly contains an encryption program, typically involving both symmetric and asymmetric key encryptions, we selected the hardware events which are more likely to change because of the encryptions. The hardware events selected for our study are instruction, cache-references, cache-misses, branches, and branch-misses. The events are self-explanatory by their names. Generally the symmetric encryption affects the cache based events while the asymmetric encryptions affect the instruction and branching events.

To represent the prototype of a typical system behavior, we designed a watchdog program, and collected the `perf stat` values with 10ms time interval for that executable. The effects of all the other processes including the ransomware running in the system will have an impact on the performance counters. We collected these values at the different point of time in the target system and created a dataset of regular observation. We articulate that any behavior which is not close to this dataset is an unusual activity, but may not be a malicious one.

We show the effect of a Ransomware Program (for example, a WannaCry) on the HPCs in Figure 1. The blue lines in Figure represent the effect of normal system programs on the watchdog executable for different HPCs, whereas the orange lines show the effect of WannaCry ransomware.

An important point to be observed from Figure 1(b) is that for a particular time interval the behavior of WannaCry does not change much from the typical system behavior, for example around time interval of 100 the effect of WannaCry on the hardware event branch misses is same as normal system behavior. So, instead of considering individual points for decision making, we select a window of

observations considering each of the five events collectively. Thus, we transform the problem into anomaly detection in multivariate time-series data.

2.2 Learning a Time-Series data using an Autoencoder

To present a generalized ransomware detection strategy we do not model the behavior of the ransoms as there can be potential new ransomware whose behavior is unknown and cannot be modeled. Instead, we model the normal system behavior, as we can get a majority of such instances. Another advantage of detecting anomalies by modeling normal behavior is that we do not need the necessity of labeled dataset as any activity with unusual behavior crossing a certain threshold can be detected as an anomaly. Thus, we propose an unsupervised approach to detect these anomalies. The HPC event counts observed over the watchdog application can be considered as the time-series data which is system dependent. An LSTM (Long-Short-Term-Memory) based autoencoder can efficiently implement the unsupervised anomaly detection for time-series, which we discuss below.

Autoencoder is an Artificial Neural Network used for efficient coding of the input space by unsupervised learning. The primary goal of an autoencoder is to induce a representation for a set of data by learning an approximate identity function, i.e., if the input data is \mathcal{X} , the goal of the autoencoder is to learn the function f , given by -

$$f : \mathcal{X} \rightarrow \mathcal{X}$$

An autoencoder always consists of two mapping, encoding and decoding, which are given as ϕ and ψ respectively.

$$\phi : \mathcal{X} \rightarrow \mathcal{F}, \quad \psi : \mathcal{F} \rightarrow \mathcal{X}$$

where \mathcal{F} is a vector referring to the decisive intermediate representation learned by the autoencoder, which is used to regenerate the original input data. The error incurred by the autoencoder to regenerate the input from vector \mathcal{F} is termed as Reconstruction Error, which is given as below:

$$\mathcal{L} = \|\mathcal{X} - (\psi \circ \phi)\mathcal{X}\|^2 \quad (1)$$

The learning goal of the autoencoder is to minimize these reconstruction cost for all the input samples, i.e., to find the mappings ϕ and ψ such that \mathcal{L} is minimum.

$$\arg \min_{\phi, \psi} \mathcal{L} = \arg \min_{\phi, \psi} \|\mathcal{X} - (\psi \circ \phi)\mathcal{X}\|^2 \quad (2)$$

In our case, the input \mathcal{X} is a multivariate time-series sequence, and the objective is to learn the structure of the sequence. LSTM networks, belonging to a class of Recurrent Neural Network Model, are typically used for modeling sequence data, which efficiently handles the dependencies within the sequence. Hence, we use the LSTM based autoencoder for our detection purpose. The anomaly detector model first takes a multivariate input sequence (\mathcal{X}), generates an intermediate feature vector (\mathcal{F}) related to the sequence, and then reconstructs the same sequence from the intermediate feature vector. The autoencoder is trained using all the input sequences by following the objective function mentioned in Equation (2).

The training dataset is constructed from the observed data for normal system behavior by taking a window of 100 trace points (i.e., a window trace points collected over 1 second, since each interval

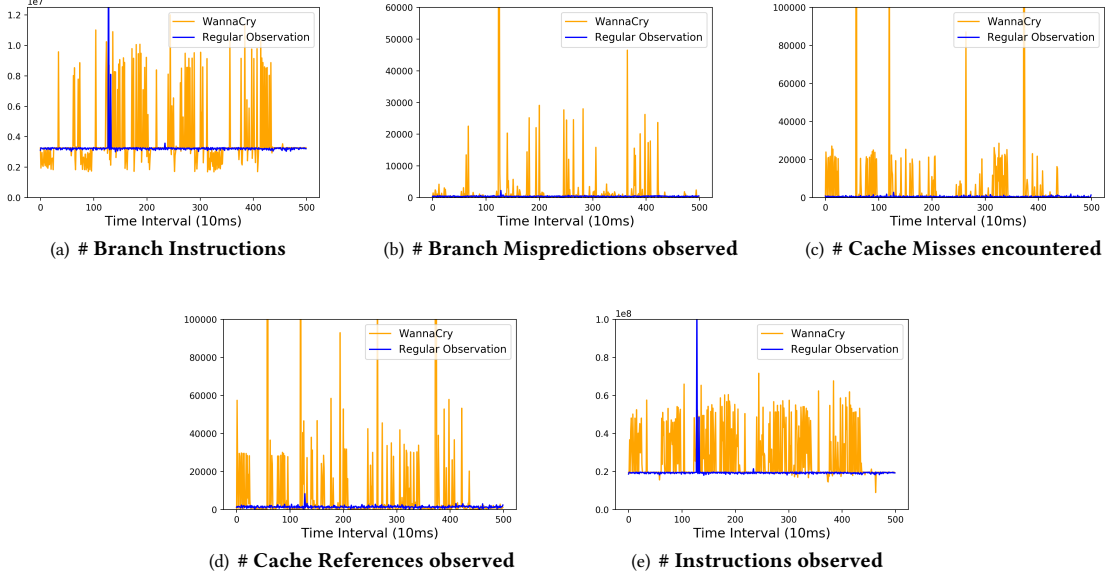


Figure 1: Variation of Performance Event Counters from HPCs in presence of Wannacry Ransomware

data is collected after 10ms). Without loss of generality, we have chosen 100 trace points for our experiments. We shift the window by one time-interval (i.e., 10ms) repeatedly to consider consecutive 100 sample point for learning. Once the learning of intermediate vector \mathcal{F} is completed, for an anomalous sequence, the autoencoder makes an attempt to reconstruct the original input sequence. Thus, the autoencoder maps it to the normal sequence, based on the intermediate vector \mathcal{F} . There is an inherent information loss in this process and hence will incur a substantial reconstruction error. Next, we quantify the amount of error to be incurred by a process to be termed as an anomaly.

2.3 Determining Threshold for Decision

To quantify the threshold for detecting anomalous activities, we calculate the reconstruction error distribution (\mathcal{R}) for all the training samples. According to the 3σ rule of thumb, all the values are considered within three standard deviations of the mean. Hence, we set the threshold for reconstruction error (\mathcal{R}_t) as below.

$$\mathcal{R}_t = \mu_{\mathcal{R}} + 3 * \sigma_{\mathcal{R}} \quad (3)$$

where $\mu_{\mathcal{R}}$ and $\sigma_{\mathcal{R}}$ are the mean and standard deviation of distribution \mathcal{R} . In our experimental setup \mathcal{R}_t came out to be 0.114.

2.3.1 Anomalous Behaviors of Ransomwares. In our study, we considered two ransomware programs - namely *WannaCry* and *Vipassana* to show the impact of selecting the threshold \mathcal{R}_t in detecting them as anomalies. Figure 2 shows the sequence of reconstruction errors for both the ransomwares. The first point on both the plot appears after observing the first window of 100 time-interval (equivalently 1 second after the start of execution of the ransomwares). The successive points come after each time-interval of 10ms since we are sliding the window by one time-interval for calculating the next reconstruction error. The blue line indicates the

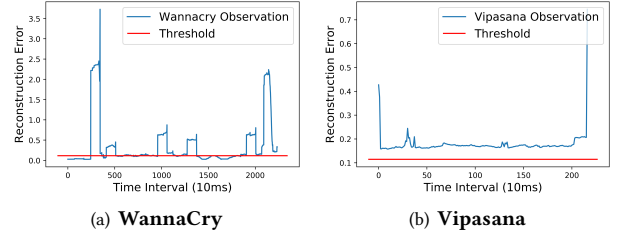


Figure 2: Sequence of Reconstruction Errors for Ransomware

reconstruction errors of each window whereas the red line signifies the threshold \mathcal{R}_t as calculated before.

We can observe from Figure 2(a), the execution of *WannaCry* starts behaving like a regular program (since the reconstruction error lies well below the threshold value), but the reconstruction error shoots over the threshold at 245th observation. Thus, the *WannaCry* is detected as anomaly $(1000 + 244 * 10) = 3440$ ms or 3.44 seconds after the start of execution. Whereas, from Figure 2(b), we can observe that the *Vipassana* is detected as an anomaly at the first window itself, i.e., 1 second after the start of execution. In both the cases there is an extra overhead of time due to the testing time of Autoencoder, which we discuss in Section 4.3.

3 HOW GOOD IS RECONSTRUCTION ERROR AS A DECIDER?

In the previous section, we suggest that a threshold as high as \mathcal{R}_t can be used to determine whether a particular application behavior deviates from the normal system behavior significantly. In this section, we explain why a single decision step is not enough to claim that the anomaly observed is from a malicious process.

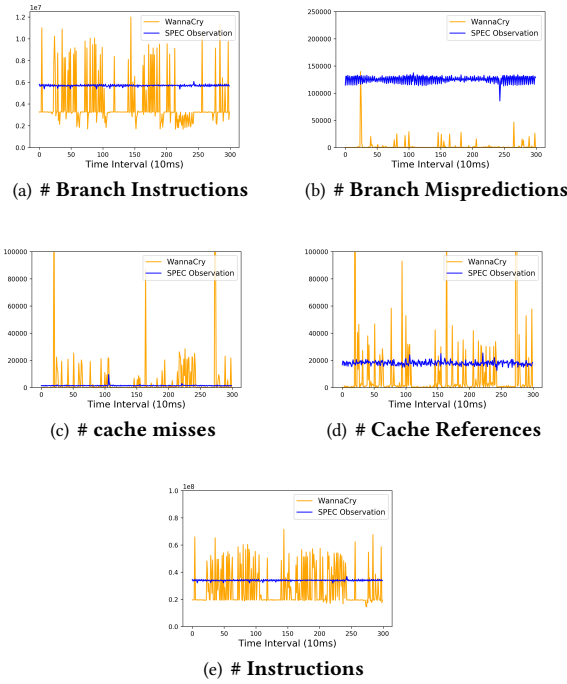


Figure 3: Comparison of the Effects on Performance Event Counters from HPCs in presence of Wannacry Ransomware and SPEC Benchmark Programs

3.1 Understanding the Ambiguity

In order to test the robustness of our detection scheme, we incorporate an analysis in presence of SPEC2006 server and multimedia benchmarks. We consider the Gshare predictor implementation as provided in (https://www.jilp.org/jwac-2/cbp3_framework_instructions.html) and observe the HPC sampling counts from a background process exactly like our previous setting. Figure 3 presents the variation of different hardware events in presence of both SPEC benchmarks and Wannacry ransomware. We can observe that the execution behavior for both the programs are significantly different from the normal observations. Thus, the sequences of data for the SPEC programs may also create considerable reconstruction errors. In Figure 4, it clearly shows that the reconstruction error for the sequences in presence of the SPEC benchmark programs is above the predetermined threshold at the first window itself. Though the error is very close to the threshold, this essentially raises an alarm to RAPPER that this benchmark program is a potential malicious program which deviates in an extent from the normal system behavior. But surely in this case, it is a false alarm, since the benchmark is composed of server and multimedia benchmarks and can be considered as the representative of the high computational processes which may deviate highly to the normally running processes in a system.

In the next subsection, we perform a transformation from the time domain to the frequency domain to differentiate actual malicious processes from false positives.

3.2 Introducing Fast Fourier Transformation

In the second phase of detection using RAPPER, we transform the traces from the time domain to the frequency domain using the Fast Fourier Transformation (FFT). FFT is the most efficient way to implement the Discrete Fourier Transformation. The primary reason to convert the analysis from the time domain to the frequency domain is to understand the repetitive pattern of the traces. The ransomware executable runs encryption repeatedly on multiple files thus it repeats a same set of operations of opening a file, encrypting and closing the file followed by deleting it for multiple files one after another. The transformation is illustrated in Figure 5, which typically indicates that the amplitude for each frequency bins are constantly higher for the ransomware in contrary to the SPEC benchmark.

We have applied FFT on the time domain values for different hardware events as mentioned in Section 2.1, to obtain the frequency domain values. Figure 5 presents the FFT plots for the normal system measurements in blue lines, along with the SPEC Observations in green lines and Wannacry Ransomware in orange lines for different hardware events. Figure 5 shows that for most of the hardware events (apart from the cache misses), the FFT plot behavior of the SPEC benchmark overlaps exactly with the FFT of the normal system behavior. Also, it is quite clear from the Figure 5(a), Figure 5(b), Figure 5(d), and Figure 5(e) that the amplitude of almost all the frequency bins are higher for Wannacry than the SPEC observation, which is eminent as the Wannacry program repeatedly encrypt multiple files.

The detection of these variations of amplitudes for different frequency bins can again be considered as a time-series data, and an LSTM based autoencoder, as discussed before, can be used to detect the anomaly. The amplitudes for SPEC benchmark programs are very close to that of regular observations for most of the hardware events. Thus, modeling the FFT data for regular sequences using an autoencoder will result in reconstruction errors close to the threshold (say \mathcal{R}_t') for SPEC benchmarks, and the error will be much higher in case of ransoms because of the repeated encryptions. We modeled an autoencoder as mentioned before and the calculated the threshold \mathcal{R}_t' to be 0.033. Fig 6 presents the sequence of reconstruction errors for both SPEC and Wannacry programs and we can verify that the reconstruction errors of the SPEC programs always lies below the threshold and thus discarded from false positives, whereas the reconstruction error of the Wannacry program always remains higher to the threshold.

4 ARCHITECTURE OF RAPPER

In this section, we present an overview of the architecture of proposed detection methodology - RAPPER. The basic diagram of the

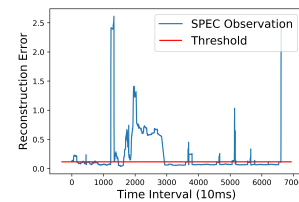


Figure 4: Sequence of Reconstruction Errors for SPEC Benchmark Programs

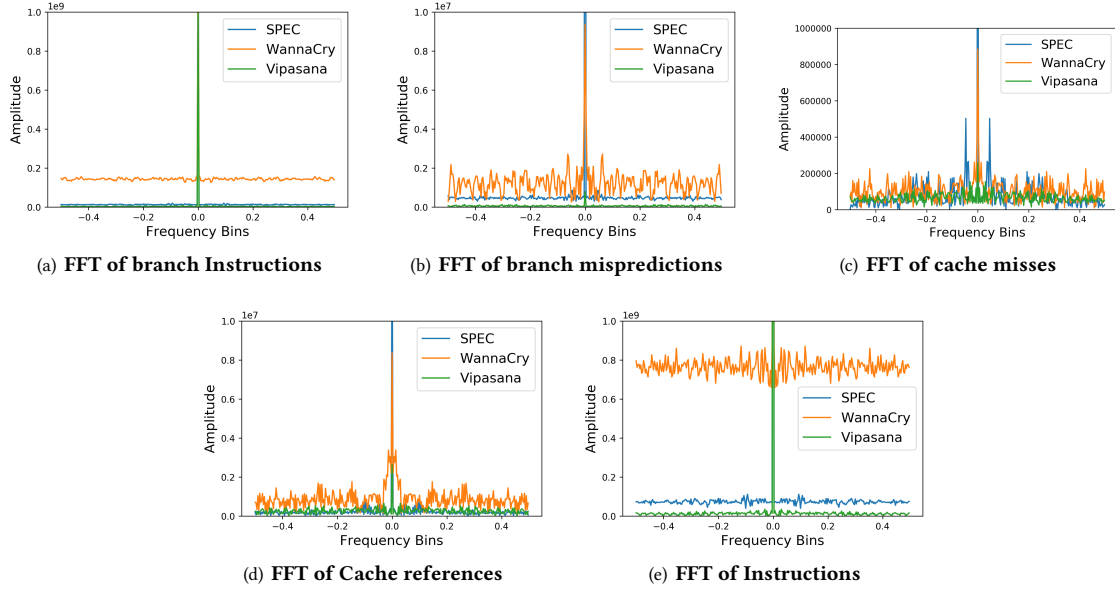


Figure 5: Variation of Amplitude in frequency domain of the performance counters from HPCs in the presence of SPEC observation and Wannacry Ransomware

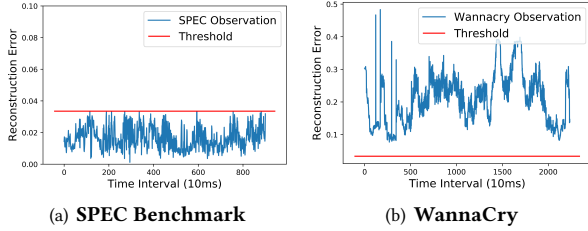


Figure 6: Sequence of Reconstruction Errors for SPEC Benchmark (a) and WannaCry (b)

system is shown in Figure 7. All the experimentation for this study have been performed in a sandbox environment, such that the ransomwares do not affect the actual filesystem.

The architecture contains four modules (*Watchdog Program*, *Autoencoder_1*, *FFT Converter*, and *Autoencoder_2*). The detection methodology works in two phases, namely *Offline Phase* and *Online Phase*. The functioning of each of the module in both the phases are described below:

4.1 Offline Phase

In the offline phase, the detection methodology is trained with the normal behavior of the sandbox environment, such that any unusual activity of a ransomware is properly detected in real-time scenario. The functioning of each of the modules in this phase are described below.

- (1) *Watchdog Program*: Monitors the HPCs of the Sandbox Environment continuously and forwards a window of data (calculated as described before) to the *Autoencoder_1* and the *FFT Converter* in parallel.
- (2) *Autoencoder_1*: Collects all the data forwarded by the watchdog program and an autoencoder is trained with the dataset as mentioned in Section 2.2.
- (3) *FFT Converter*: Converts the Fast Fourier Transformation of each window forwarded by the watchdog program and passes the results to the *Autoencoder_2*.

- (4) *Autoencoder_2*: Collects all the data passed by the *FFT Converter* and trains another autoencoder based on the *FFT* dataset.

4.2 Online Phase

In the online phase, the detection module is deployed in the sandbox system for real-time monitoring to detect ransomwares. The functioning of each modules in this phase are discussed below.

- (1) *Watchdog Program*: Monitors the sandbox system as performed in training phase, and forwards the data to the *Autoencoder_1* module. In this phase, watchdog program does not forward data to the *FFT converter*. This helps us to monitor the system with lower computational cost.
- (2) *Autoencoder_1*: Receives sequence window at each time interval from the watchdog program and calculates the reconstruction error of the sequence. If the error is higher than the predefined threshold \mathcal{R}_t , it sends a signal to the watchdog program to transmit the same window to the *FFT Converter*.
- (3) *FFT Converter*: Receives a sequence window from the watchdog module, converts the data into frequency domain, and forwards the transformed data to the *Autoencoder_2* as done in training phase, but with a condition imposed by the *Autoencoder_1* module.
- (4) *Autoencoder_2*: Calculates the reconstruction error of the received *FFT* data, and based on the predefined threshold, \mathcal{R}'_t , sends a warning whether the sequence belongs to a ransomware or not.

4.3 Evaluating the performance of RAPPER

We performed all the experiments in a sandbox system having specification Intel Core i3 M350 running Ubuntu 16.04 with 4.10.0-38-generic kernel. We used popular open source python based neural network library Keras [1] for the implementation of both the autoencoder. The architecture used to model the autoencoders are mentioned in Table 1. The Dropout Layer is added to regularize the neural network and prevent from overfitting (Dropout is a

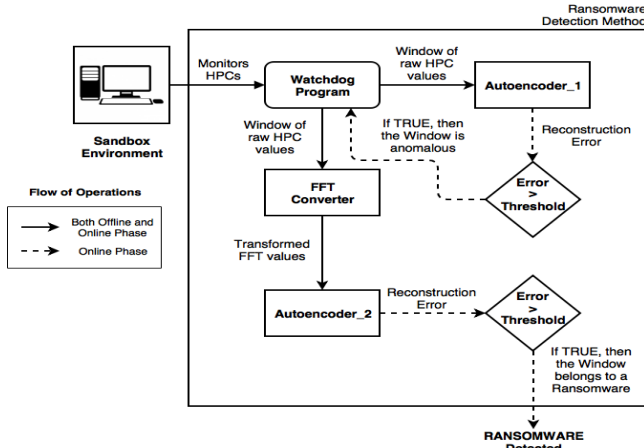


Figure 7: Detection Methodology of RAPPER

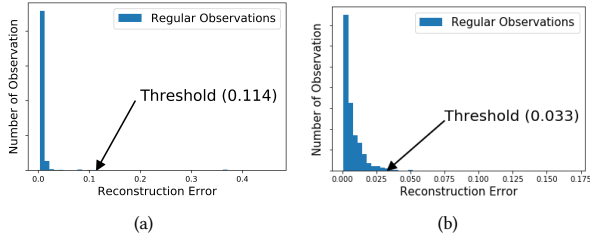


Figure 8: Distribution of reconstruction errors for (a) Autoencoder_1 (b) Autoencoder_2

technique where a set of randomly selected neurons are ignored during training).

The distribution of reconstruction errors for regular observations produced by both the autoencoders are shown in Figure 8. The thresholds thus calculated using Equation 3 are 0.114 and 0.033 respectively.

The FFT converter usually takes 0.0003 milliseconds to convert a sequence within a window into frequency domain. The model building times for Autoencoder_1 and Autoencoder_2 are on average 10 and 14 minutes respectively. Testing time to calculate whether a single window is an anomaly or not is 1.321 milliseconds for Autoencoder_1 and 1.699 milliseconds for Autoencoder_2 respectively. As shown in the Architecture of RAPPER in Figure 7, the testing of a regular observation only passes through the Autoencoder_1, thereby taking only 1.699 milliseconds, and an anomalous observation passes through all the three modules: Autoencoder_1, FFT Converter, and Autoencoder_2, thereby taking $1.321 + 0.0003 + 1.699 = 3.0203$ milliseconds to be detected. However, in both the cases, the detection time is less than the sampling interval, which is 10 milliseconds. Hence, the detection is performed seamlessly, without the need of any storage buffer, as a new window of data will be created after 10 milliseconds.

Without loss of generality, we measured the performance of RAPPER on most recent WannaCry ransomware. As shown in Section 2.3, the WannaCry is detected as an anomaly at the 245th window and instantly detected as ransomware at the same time because it's reconstruction error is always higher than the threshold of Autoencoder_2. Hence, the total time taken to detect WannaCry is equal to (Time taken to generate the first window) + 244

Table 1: Model Architecture for Autoencoders

Layer Number	Layer Type	Input Shape	Output Shape
Autoencoder_1			
1	LSTM	(None, 100, 5)	(None, 100, 32)
2	Dropout	(None, 100, 32)	(None, 100, 32)
3	LSTM	(None, 100, 32)	(None, 100, 5)
Autoencoder_2			
1	LSTM	(None, 100, 5)	(None, 100, 64)
2	Dropout	(None, 100, 64)	(None, 100, 64)
3	LSTM	(None, 100, 64)	(None, 100, 5)

* (time interval for each sample) + (Autoencoder_1 testing time) + (Time for single FFT Conversion) + (Autoencoder_2 testing time) = $1000 + 244 * 10 + 1.321 + 0.0003 + 1.699$ millisecond = 3443.0203 milliseconds. Thus, WannaCry is detected by RAPPER in approximately 3.443 seconds. As a sample run with RAPPER, out of 10000 files of approximately 21 bytes each, when the detection stops the execution, 68 files are encrypted. It may be noted that, the size of a typical file is much larger than 21 bytes, and hence, a lesser number of files will be encrypted.

5 FILE RECOVERY AND CONCLUSIONS

RAPPER is thus capable of detecting the presence of ransomwares fast, as we show for the case of WannaCry within a time of approximately 4 sec from its launch. Depending on the latency, the malware can encrypt few files (say n). We conclude with a suggested approach on data retrieval. A practical solution would be to take backups of the n -recently opened files. After the lapse of the time quantum required to encrypt these files, we delete the copies if no ransomware alarm is raised by RAPPER. This minimizes the storage necessary for the backup files. To further ensure that the backup files are not encrypted we perform locking operation, like in linux using `mlock`.

In this paper, thus we explored the effect of ransomware on normal system behaviors. We take the aid of the Artificial Neural Network to detect the presence of ransoms using a two-step detection framework. The entire detection procedure does not need any template of the malicious process beforehand. Instead it thrives on an anomaly detection procedure to detect the infectious ransoms in as less as 4 seconds with almost zero false positives, using a frequency analysis.

We also explored the opportunity of applying side channel techniques to recover the secret key used to encrypt the files from the perf statistics. Wlog. we found for ransoms like WannaCry, each file is encrypted using AES-128 CBC (Cipher Block Chaining) with a randomly generated distinct key. These keys are in turn encrypted using an infection specific RSA public key and stored in the memory. It would be indeed a challenging exercise to recover the AES key by targeting the AES CBC operation. However, we leave that as a future scope of work.

REFERENCES

- [1] 2015. Keras: The Python Deep Learning library. (March 2015). <https://keras.io/>
- [2] John Demme, Matthew Maycock, Jared Schmitz, Adrian Tang, Adam Waksman, Simha Sethumadhavan, and Salvatore Stolfo. 2013. On the feasibility of online malware detection with performance counters. In *ACM SIGARCH Computer Architecture News*, Vol. 41. ACM, 559–570.
- [3] Felix Gröbert, Carsten Willems, and Thorsten Holz. 2011. *Automated Identification of Cryptographic Primitives in Binary Programs*. Springer Berlin Heidelberg, Berlin, Heidelberg, 41–60. DOI: http://dx.doi.org/10.1007/978-3-642-23644-0_3

- [4] Amin Kharraz, Sajjad Arshad, Collin Mulliner, William Robertson, and Engin Kirda. 2016. UNVEIL: A Large-Scale, Automated Approach to Detecting Ransomware. In *25th USENIX Security Symposium (USENIX Security 16)*. USENIX Association, Austin, TX, 757–772. <https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/kharaz>
- [5] Amin Kharraz, William Robertson, Davide Balzarotti, Leyla Bilge, and Engin Kirda. 2015. Cutting the gordian knot: A look under the hood of ransomware attacks. In *DIMVA 2015, 12th Conference on Detection of Intrusions and Malware & Vulnerability Assessment, July 9–10, 2015, Milan, Italy*. Milan, ITALY. DOI: http://dx.doi.org/10.1007/978-3-319-20550-2_1
- [6] Mehmet Sabir Kiraz, Ziya Alper Genç, and Erdinç Öztürk. 2017. Detecting Large Integer Arithmetic for Defense Against Crypto Ransomware. *Cryptology ePrint Archive*, Report 2017/558. (2017). <http://eprint.iacr.org/2017/558>.
- [7] Pierre Lestringant, Frédéric Guihéry, and Pierre-Alain Fouque. 2015. Automated Identification of Cryptographic Primitives in Binary Code with Data Flow Graph Isomorphism. In *Proceedings of the 10th ACM Symposium on Information, Computer and Communications Security (ASIA CCS '15)*. ACM, New York, NY, USA, 203–214. DOI: <http://dx.doi.org/10.1145/2714576.2714639>
- [8] Corey Malone, Mohamed Zahran, and Ramesh Karri. 2011. Are hardware performance counters a cost effective way for integrity checking of programs. In *Proceedings of the sixth ACM workshop on Scalable trusted computing*. ACM, 71–76.
- [9] Adrian Tang, Simha Sethumadhavan, and Salvatore J. Stolfo. 2014. Unsupervised Anomaly-Based Malware Detection Using Hardware Features. In *Research in Attacks, Intrusions and Defenses - 17th International Symposium, RAID 2014, Gothenburg, Sweden, September 17–19, 2014. Proceedings (Lecture Notes in Computer Science)*, Angelos Stavrou, Herbert Bos, and Georgios Portokalidis (Eds.), Vol. 8688. Springer, 109–129. DOI: http://dx.doi.org/10.1007/978-3-319-11379-1_6
- [10] Xueyang Wang and Ramesh Karri. 2013. NumChecker: detecting kernel control-flow modifying rootkits by using hardware performance counters. In *DAC*. ACM, 79.
- [11] Xueyang Wang and Ramesh Karri. 2016. Reusing hardware performance counters to detect and identify kernel control-flow modifying rootkits. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 35, 3 (2016), 485–498.