

Fighting Ransomware with Guided Undo

Matthias Held and Marcel Waldvogel

University of Konstanz, Konstanz, Germany
{matthias.held, marcel.waldvogel}@uni-konstanz.de

Abstract

Ransomware attacks are rare, yet catastrophic. On closer inspection, they differ from other malware infections: Given appropriate preparation, they do not need to be caught on first sight, but can be undone later. However, current ransomware protection follows the beaten path of anti-malware copying their fallacies. We show how the move to personal cloud storage allows for a paradigm shift in ransomware protection: exceptional attack isolation, perfect elimination of false positive alerts, and simplified recovery.

In this paper, we analyze the necessary operations for ransomware, extend existing ransomware taxonomy, and verify them against real-world malware samples. We analyze the costs and benefits of moving ransomware detection to versioned personal cloud storage. Our content, meta data, and behavior analysis paired with a ‘guilt by association’ capability greatly improve the false positive rate, but the guided undo make this rate all but inconsequential. Even though the user now carries a new burden, it comes with clear responsibilities and benefits, while being freed from questionable duties, resulting in a win-win situation for user experience and detection quality.

1 Introduction

Ransomware attacks are most effective, when they strike at the worst possible moment, tempting the victims to pay ransom in the hope of getting relief, and thus financing organized crime. It is therefore imperative for users and society to reduce the impact of ransomware to a rare event, causing no more than a brief annoyance. Current ransomware detection is integrated with or modeled after other malware protection: The key to defense is software running on the victim’s machine, trying to identify malware when it is downloaded or first activated. Identifying the purpose of software from its bytestream is impossible [21], and once the software is running, anti-malware is not more powerful or bug-free than the operating system, whose duty includes protection from malicious software. This model cannot eliminate the problems of malware circumventing the detection software, thus leading to limitations which reduce the identification and recovery quality [20, 19].

Traditional ransomware detection analyzes program code and file operations, among other indicators, to detect potential malware activity. Either of them can result in false positives, triggering the user to react, often with panic. The anti-malware developer’s goal therefore is to minimize the number of false positives, resulting in false negatives. If malware goes undetected, the data, even on the backup disk, when connected, may be encrypted or destroyed; an essentially unsolvable dilemma.

The ongoing move to personal cloud storage, however, opens the possibility of a paradigm shift in ransomware detection. Having an up-to-date copy of all important data on an unrelated system where no ordinary user can install or run additional software provides for exceptional attack isolation, perfect elimination of false positive alerts, and simplified recovery. These benefits from attack isolation come at a cost: the weaker coupling to the events on the victim’s machine reduces information about ongoing file operations as part of the malware, such as the sequence of actual file operations, and details including file offsets.




Figure 1: Possible victim machines

The missing real-time information has some impact on *content-based* and *metadata-based* indicators, which may analyze short-term (on the order of seconds) event sequences; however, analyzing long-term *behavior* remains unaffected. We shift the classification towards this indicator, supported by the file versioning to observe file operations without data loss. This shift improves the separation between processes of benign software which work with high entropy files and ransomware attacks.

The independence gained by the weak coupling between victim machine and cloud storage more than offsets the analysis limitations, as the detection software and archived file versions may not be manipulated by the ransomware.

We propose the following data security model to protect against many incidents, including ransomware attacks (Figure 1): The user’s machine (left) does a traditional backup to a locally-attached storage (far left). It also performs a real-time synchronization to a personal cloud storage, such as provided by a personal machine of the Raspberry Pi class or a cloud provider. Important is that the cloud storage be versioned, so that previous, unaffected, states of the data may be restored. As long as at most one of the two systems (user machine, personal cloud) is infected, data can be recovered. In the unlikely event that the cloud storage is the one being infected, synchronization of destroyed data may happen back to the user machine. This is guaranteed not to affect the backup, as it is outside of the synchronization scope.

As a result, unlike most other security incidents, ransomware does not necessarily need to be detected on first sight. Relaxing the requirements by using a delayed detection and recovery approach assisted by the personal cloud storage, allows us to remove the need of having a 100% detection rate without any false positives, something infeasible with real-time detection.

This paper makes the following contributions:

1. We define the general problems fighting ransomware using generic indicators.
2. We formulate a taxonomy of requirements for operation classes as well as indicator categories and verify it with real-world malware samples.
3. We define a set of *content-based*, *metadata-based* and *behavior-based* indicators which are a set of already suggested indicators and indicators extended or improved by us.
4. Based on these indicators, we implement a ransomware detection in an app on top of *Nextcloud* cloud storage, which assists making the right selection when recovering from an attack.

2 Problems fighting ransomware

Ransomware attacks are rare events. However, is it necessary to analyze every file operation to detect them. Once we use an automatic recovery, we need to have a detection rate without any false positives since a – still so small – false positive rate would lead to many false recoveries due to many file operations performed on a computer system [15, 1].

Educating the user through an easy-to-use recovery, guiding the user with informations gathered by analysing and classifying file operation sequences, leads to perfect protection from false positives. Since ransomware attacks are so rare, it is easier for the user to take a spare moment to recover if an attack happens instead of taking many moments struggling with false positives.

This change can be compared with the usability step as in the 80s where the move from file-based manual versioning to the simple undo method was established. This reduced the effort for the user drastically. Furthermore, it opened up the file versioning for the group of unsophisticated users by simplifying the recovery. This point is important in companies, universities or organizations which offer a “bring your own device (BYOD)” integration, where our usability concept can reduce the effort of the IT-support by allowing the users to perform the recovery by themselves [5].

Deactivating the synchronization, and therefore disabling the ransomware detection, simplifies the recovery due to the double existence of the data. Only files which are modified by two or more users at the same time during a ransomware attack would be problematic, since it would lead to a loss of the latest changes performed on the victim’s host. However, the amount of changes would be minimal and can be easily restored since the attack would be detected on first sight and the user can remember the changes well.

The generic real-time ransomware detection approaches proposed in different papers offer a real-time detection together with an automatic recovery proclaiming a nearly perfect detection rate with very few false positives [2, 3, 10, 8]. The rarity of a ransomware attack and automatic recovery make these good detection rates problematic: Assuming we have a ransomware detection software, which monitors every of the 100,000 file operations a day on the file system with a false positive rate of 0.1%. This would result in 100 falsely triggered recoveries leading to 100 unnecessary user interactions reverting these false recoveries. Additionally, the efforts for reducing the false positives lead to a very concrete set of indicators allowing the ransomware authors to build the malware against these indicators avoiding the detection [1].

For blocking decisions, we therefore should include every possible generic property. It is unavoidable to integrate a file versioning mechanism to be able to consider more file operations. This can also be done in real-time and with automatic recovery [3] but is limited by the base rate fallacy.

This cannot be considered reliable due to ransomware circumventing local backup software and detection software by disabling them or overwriting the Master Boot Record to perform the encryption in a state of the pc where only the malware is running [20, 19].

3 Related work

Existing ransomware indicators were used in different attempts to detect malicious software, while having several limitations like false positives in form of benign software, which fulfilled the same indicators e.g. high entropy files created by compression tools [14] or bypassing the monitoring by avoiding the indicators thresholds or hiding the malicious process [3, 7]. Long-term observations of ransomware showed their proceeding and possible indicators to prevent zero-day attacks of ransomware [9]. This already constructed a wide-spread set of indicators for detecting ransomware. Possible improvements for the indicators were evaluated: Such as the analysis of the entropy of files to improve it by being able to distinguish between encryption and compression, which are both high entropy files, without leading to a satisfying result [4].

Additionally, the long-term observation *Symantec* published several reports with information and statistics of ransomware over the last several years showing the spreading, behavior, attack

methods and how to prevent attacks [16, 18, 17, 13].

Also other researchers published in-depth analyses of ransomware offering informations to many different families [12, 11].

Independent from such indicators, there are techniques like controlling the access to the command-and-control servers to prevent the execution of the ransomware [2] and the detection and escrowing of the encryption key to later decrypt the files without having to pay the money [10].

These indicators were already used to build techniques for detecting and preventing ransomware in real-time [7, 8, 10, 14]. In addition, there is a concept based on sequences and a backup strategy but it acts on the local system without the need of interaction of the user [3]. Limitations of this approach are mainly the bypassing of the monitor by tampering the Kernel or multiprocess malware.

Additionally, the base rate fallacy is a huge problem due to the rarity of ransomware attacks with just a few true positives and potentially many false positives due to plenty non-ransomware-specific file operations leading to a burden for the user [15, 1].

Our work is unlike the real-time approaches and can be described as *delayed detection and recovery*. It presents a method using a personal cloud storage by providing easy-to-use recovery from ransomware attacks including taking the users into control of the recovery. The need for such a concept was also indicated by a support page in the *Dropbox* help center [5]. They tackle this problem by requesting the user to recover every single file by hand or sending a list of encrypted files to the *Dropbox* support.

4 Background: Classes

Ransomware activity is not a static set of operations, which is performed in the same way in every ransomware family [9, 12, 11, 13]. The reasons for this are on one hand the race between security researchers and ransomware programmers to detect and hide the activities by changing the operations, and on the other hand the huge code base of existing ransomware to copy from. This leads to several implementations and many different families. Based on our analysis of ransomware operations, we also believe that some of the unusual, inefficient behavior we have seen is due to the high division of labor or just due to oversight or stupidity.

The process of constructing an efficient and working way to detect ransomware includes formulating classes for the different types of ransomware. These will then be used to define indicators.

The classes in this section are based on the behavior of existing ransomware families and possible – not yet existing – implementations. For a better overview, the ransomware classes are separated into the following subclasses: **File destruction**, **operation sequence**, **file type funnelling**, **file extension** and **file name**. In addition, the **operation sequence** classes are described for the local file system and the cloud storage.

4.1 File destruction

The main behavior is the encryption of user files, therefore the ransomware must read the data, write the data and remove the original data afterwards.

This behavior is implemented in various ways by different ransomware variants. The diverse implementations can be classified in two classes:




Figure 2: Flowchart of the behaviour of the file destruction classes A and B.

Class A Replaces the data of a file *in-place*, thus it first reads the data from the file then encrypts it, writes the encrypted data back to the file, and closes it. Afterwards, the file is optionally renamed.

Class B Creates a new file and moves the data to the new file by reading the content from the original file, encrypting it and writing the encrypted data to the new file, deleting the original file after closing the files.

Existing literature classifies ransomware in three classes, where **Class A** and **Class B** remain the same [14]. They additionally introduced a class, which is an extension of **Class A**, moves the files in an extra folder before encrypting the files *in-place* and then moving them back. We treat this class as a subclass of **Class A**, because the main indicator for this class is the *in-place* encryption and the movement of the files is an unnecessary extension, which does not change the fundamental characteristics.

4.2 Operation sequence

The file destruction can be performed in different operation sequences where every sequence is unlike the others due to the order of operations. The ordering can change with every system hence we divide the operation sequence classes into two different scenarios: The operations performed on a local file system and the operations performed during a synchronization process with a cloud storage.

The differentiation between those two scenarios have the simple reason, that the synchronization is performed by a client software, which reorders the file operations that happen during a synchronization interval leading to different operation orders.

Thus the classes for the local file system are defined as following:

1. Batch WRITE followed by batch DELETE/RENAME-with-overwrite.
2. Interleaved WRITE – DELETE/RENAME-with-overwrite.
3. In-place-overwrite

The classes for the file operations visible during the synchronization are defined like this:

1. Batch WRITE followed by batch DELETE/RENAME.
2. Batch DELETE/RENAME followed by batch WRITE.
3. Interleaved WRITE – DELETE/RENAME.
4. Chaos WRITE – DELETE/RENAME.

Class 2 represents a regroup of class 1 and class 4 depicts a regroup and arbitrary split of class 3. As mentioned in the previous section the operations can be preceded by first moving the files to a special directory.

4.3 File type funnelling

Other options include whether the file extensions of encrypted files are random strings, a ransomware-specific extension, or the file type is unchanged from the original file. Similar things may also happen to the file name [6].

File type funnelling describes the process of reading many files with known file extensions and writing files, which can be assigned to the following classes:

1. All file extensions are unknown and the same.
2. All file extensions are unknown and every extension is distinct.
3. All file extensions are unknown and mixed.
4. All file extensions are known, but files are corrupted.

We define "known" in this enumeration as: File extensions which are known well or file extensions with well-known file types. Additionally, class 3 and 4 are possible but so far non-existing implementations of file type funnelling.

4.4 Entropy funnelling

All ransomware families encrypt the data of the user files to extort the user. Hence the entropy of the files is increased – except if all the original files were compressed or also encrypted – thus all current ransomware families show the property of entropy funnelling, where the entropy level of all files is changed to become nearly the same.

In theory, there are options to reduce the entropy of the files written or have them stay nearly the same. Such malware has not been described nor seen in the wild.

5 Background: Indicators

The discussed indicators in this section capture the destructive properties of ransomware regarding a single file or a sequence of files and not modifications of the operation system. The described indicators were chosen to be general enough to work for all ransomware variants and try to ignore specific identifiers like specific strings or network traffic to the command-and-control server.

We divide these into three groups of *content-based*, *metadata-based* and *behavior-based* properties, where the *content-based* indicators capture the properties of the data, *metadata-based* indicators express the properties of the file and the *behavior-based* ones describe the properties of the sequence. They are described in more detail in [6].

Most of these indicators are well discussed in [7, 14, 3]. They also showed that the single indicator is not able to determine between benign software and ransomware, but several of those indicators together are able to do so.

The goal of the following chapter is to improve the significance of the Shannon Entropy and to append yet not discussed indicators, which are described in operation timing and operation quantities together with size quantities.

5.1 Content-based

5.1.1 Shannon entropy

The Shannon entropy is a measurement of information in a file. Files with a high entropy are compressed and encrypted files, where the information level is reduced to gain a specific




Figure 3: Boxplot of the standard deviation of files compressed with *Deflate*.

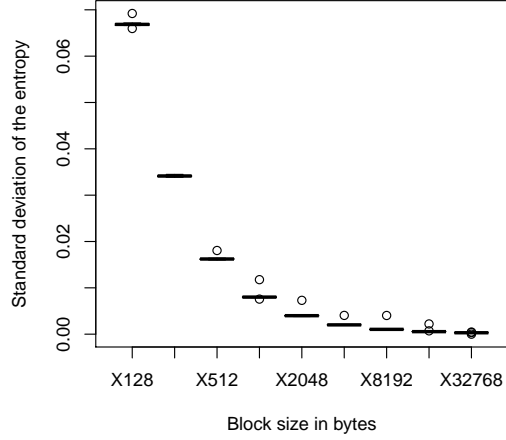


Figure 4: Boxplot of the standard deviation of files encrypted with *AES*.

property of compression or encryption. Thus a ransomware attack should change the entropy of files to a higher average due to the encryption of files.

The definition of the Shannon entropy for a byte array of 256 bytes is as follows:

$$e = \sum_{i=0}^{255} P_{B_i} \log_2 \frac{1}{P_{B_i}}$$

for $P_{B_i} = \frac{F_i}{totalbytes}$ and F_i , the number of instances of byte value i in the array. This sum expresses a value between 0 and 8, where 8 is the perfectly even distribution of the byte values in the array [14].

This is an indicator for malware files, as mentioned encrypted files have a nearly 8 distribution. Compression also tends to a nearly 8 distribution. Therefore, we analysed high entropy files with the goal to find a method to distinguish between encrypted and compressed files.

5.1.2 Improved Shannon entropy

To reduce the false positives of compressed files we use the standard deviation of the entropy of the data to distinguish between compressed and encrypted files. This method is based on the differences in the variance of high entropy blocks in the file data which are depicted in the box plots 4 and 3 of file encrypted with *AES* and files compressed with *Deflate*.

We use the property that the entropy of data blocks of encrypted files are very dense distributed, where in contrast the entropy of data blocks of compressed files is not very dense distributed.

This property is depicted by the standard deviation if we cut the data of the file in 2 kibibytes blocks and calculate the entropy of every block and the standard deviation of the entropy. The comparison of the standard deviation can be examined in figure 5. We choose a threshold of

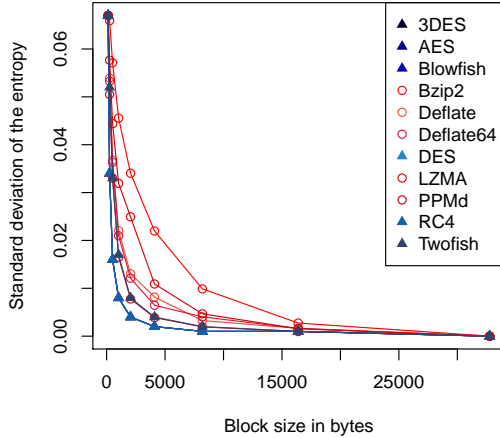


Figure 5: Comparison between the compressed files in red and encrypted files in blue.




Figure 6: Boxplot of the standard deviation of files compressed with *LZMA*.

0.04 for a data block size of 2 kibibytes for differentiating between encrypted and compressed files. The classification of this method does have a false positive rate of 25% but improves the indicator of the entropy of the data a lot due to the usage of a guilt by association presumption.

The false positive rate results from compression algorithms that have a very dense distribution with less outliers e.g. *LZMA* (see figure 6). In contradiction, all evaluated encryption algorithms generated a dense distribution hence we have less false negatives. Files classified as encrypted are more suspicious than other high entropy files thus we increase the suspicious classification of the file.

This means we have only a few false negatives and some false positives: If we classify a sequence we can use the ‘guilt by association’ conclusion to assure our suspected case.

6 Application implementation

The ransomware detection implementation is based on the personal cloud storage *Nextcloud*. The reason for this storage are the good file versioning and trash bin methods, which are integrated into this cloud storage and are used here to recover from a detected ransomware attack.

We defined all classification methods in section ??, hence the things left for the implementation is the integration into the existing *Nextcloud* server. Therefore we describe the monitoring of file operations in the first subsection, followed by the integration of the sequence analysis. The last subsection describes the recovery method which uses the integrated file versioning and trash bin of *Nextcloud*.

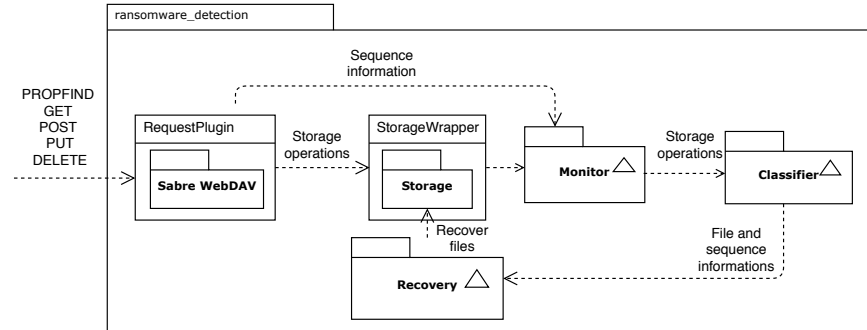


Figure 7: Architecture and integration of the ransomware detection application, which separates the concerns of monitoring, classifying and recovering files and file sequences.

6.1 Monitoring

To monitor all file operations, the data storage of *Nextcloud* is wrapped, piping all file operations through our analyser, gaining the needed operations by analysing the operations with indicators defined in section 5.

6.2 Sequence analysis

The sequence analysis puts file operations together to a sequence by looking for time intervals where nothing happens. In the context of *Nextcloud* this is equivalent with synchronization requests with no changes. Thus we are detecting all synchronization requests without file storage access and if there is a succession of six synchronization requests without changes, we start a new sequence of file operations.

6.3 Classification

The results of the analysis performed during the monitoring are used to classify file operations and sequences of them. These results are used in recovery interface to give the user a colour guide, which leads to the following depiction of file operation sequences (see figure 8, 9, 10 and 11).

6.4 Recovery

To recover from ransomware attacks the selected sequence of file operations will be reverted; that means new created files will be deleted and deleted or renamed files will be restored. This is done by using the integrated file versioning methods of *Nextcloud*.

7 Experiments

In this section the application is reviewed in different aspects to evaluate the functionality and effectiveness. Therefore the classification of the sequences and the effectiveness of the indicators are evaluated.

Name	Operation	Size	File class	File name class	Time
assignment01.pdf	W	88 KB	D	○	vor 5 Minuten
assignment01.pdf	A	88 KB	D	○	vor 5 Minuten
assignment02.pdf.plyblmUC	W	100 KB	A	▲	vor 5 Minuten
assignment02.pdf	W	100 KB	D	○	vor 5 Minuten
assignment02.pdf	A	100 KB	D	○	vor 5 Minuten
assignment01.pdf.plyblmUC	W	88 KB	A	▲	vor 5 Minuten
assignment05.pdf	W	82 KB	D	○	vor 4 Minuten
assignment05.pdf.plyblmUC	W	82 KB	A	▲	vor 4 Minuten
assignment05.pdf	W	479 KB	D	○	vor 4 Minuten
assignment05.pdf	A	479 KB	D	○	vor 4 Minuten
assignment05.pdf.plyblmUC	W	479 KB	A	▲	vor 4 Minuten
assignment03.pdf	W	99 KB	D	○	vor 4 Minuten
assignment03.pdf	A	99 KB	D	○	vor 4 Minuten
assignment03.pdf.v1523532259	A	479 KB	A	▲	vor 4 Minuten
assignment04.pdf	W	94 KB	D	○	vor 4 Minuten
assignment04.pdf	A	94 KB	D	○	vor 4 Minuten
assignment04.pdf.plyblmUC	W	99 KB	A	▲	vor 4 Minuten
assignment04.pdf.v1523530999	A	100 KB	A	▲	vor 4 Minuten
assignment03.pdf.plyblmUC	W	99 KB	A	▲	vor 4 Minuten
assignment05.pdf	A	82 KB	D	○	vor 4 Minuten

Figure 8: Colouring of a sequence of file operations created by the ransomware *GoldenEye*.

Name	Operation	Size	File class	File name class	Time
assignment02.pdf.axx	W	104 KB	A	▲	Gerade eben
assignment02.pdf	W	100 KB	D	○	Gerade eben
assignment03.pdf.axx	W	103 KB	A	▲	Gerade eben
assignment02.pdf	A	100 KB	D	○	Gerade eben
assignment01.pdf	W	88 KB	D	○	Gerade eben
assignment01.pdf	A	88 KB	D	○	Gerade eben
assignment01.pdf.axx	W	91 KB	A	▲	Gerade eben
assignment03.pdf	W	99 KB	D	○	Gerade eben
assignment04.pdf.axx	W	93 KB	A	▲	Gerade eben
assignment03.pdf	A	99 KB	D	○	Gerade eben
assignment03.pdf.v1523532259	A	479 KB	A	▲	Gerade eben
assignment05.pdf.axx	W	88 KB	A	▲	Gerade eben
assignment05.pdf	W	82 KB	D	○	Gerade eben
assignment06.pdf.axx	W	414 KB	A	▲	Gerade eben
assignment04.pdf	W	94 KB	D	○	Gerade eben
assignment04.pdf	A	94 KB	D	○	Gerade eben
assignment04.pdf.v1523530999	A	100 KB	A	▲	Gerade eben
assignment05.pdf	A	82 KB	D	○	Gerade eben
assignment06.pdf	W	479 KB	D	○	Gerade eben
assignment06.pdf	A	479 KB	D	○	Gerade eben

Figure 9: Colouring of a sequence of file operations created by the file encryption tool *AxCrypt*.

Name	Operation	Size	File class	File name class	Time
assignment04.pdf	W	94 KB	D	○	Gerade eben
assignment05.pdf	W	82 KB	D	○	Gerade eben
assignment05.pdf	A	82 KB	D	○	Gerade eben
assignment01.pdf	W	88 KB	D	○	Gerade eben
assignment01.pdf	A	88 KB	D	○	Gerade eben
assignment02.pdf	W	100 KB	D	○	Gerade eben
assignment02.pdf	A	100 KB	D	○	Gerade eben
assignment03.pdf	W	99 KB	D	○	Gerade eben
assignment03.pdf	A	99 KB	D	○	Gerade eben
assignment05.pdf	W	479 KB	D	○	Gerade eben
assignment06.pdf	A	479 KB	D	○	Gerade eben
assignment04.pdf	A	94 KB	D	○	Gerade eben
assignment03.pdf	W	99 KB	D	○	Gerade eben
assignment01.pdf	W	88 KB	D	○	Gerade eben
assignment04.pdf	W	94 KB	D	○	Gerade eben
assignment05.pdf	W	82 KB	D	○	Gerade eben
assignment05.pdf	W	479 KB	D	○	Gerade eben
assignment02.pdf	W	100 KB	D	○	Gerade eben

Figure 10: Colouring of a harmless sequence of file operations created by manual batch DELETE/RENAME followed by batch WRITE.

Name	Operation	Size	File class	File name class	Time
assignment02.pdf	W	100 KB	D	○	vor 3 Minuten
assignment05.pdf	W	82 KB	D	○	vor 3 Minuten
assignment04.pdf	W	94 KB	D	○	vor 3 Minuten
assignment01.pdf	W	88 KB	D	○	vor 3 Minuten
assignment03.pdf	W	99 KB	D	○	vor 3 Minuten
assignment06.pdf	W	479 KB	D	○	vor 3 Minuten
assignment01.pdf	W	88 KB	D	○	vor 2 Minuten
assignment02.pdf	W	100 KB	D	○	vor 2 Minuten
assignment06.pdf	W	479 KB	D	○	vor 2 Minuten
assignment05.pdf	W	82 KB	D	○	vor 2 Minuten
assignment03.pdf	W	99 KB	D	○	vor 2 Minuten
assignment04.pdf	W	94 KB	D	○	vor 2 Minuten

Figure 11: Colouring of a harmless sequence of file operations created by writing many files.

7.1 Classification

The sequence classification was tested with a dataset of multiple ransomware families which were submitted to *VirusTotal* in the last months. For every of those samples the classification reached a critical sequence suspicion score leading to a highly suspicious sequence.

AxCrypt also reached a higher suspicion level compared to the other benign processes. The reason for this is that this is file encryption software which has the same behavior as ransomware.

Family	File suspicion	Quantities	File type funnelling	Entropy funnelling	Sequence suspicion score
BTCWare	0.88	2.00	2.00	2.00	6.88
Cerber	0.75	2.00	2.00	2.00	6.75
Evasive	0.83	2.00	2.00	2.00	6.83
Evader	0.50	2.00	2.00	2.00	6.50
GlobeImposter	0.63	2.00	2.00	2.00	6.63
WannaCry	0.81	2.00	2.00	2.00	6.81
Mamba	0.79	2.00	2.00	2.00	6.79
GoldenEye	0.93	2.00	2.00	2.00	6.93
Median	0.63	2.00	2.00	2.00	6.63

Table 1: A list of ransomware families and their *behavior-based* indicator values with the sequence suspicion score.

Program	File suspicion	Quantities	File type funnelling	Entropy funnelling	Sequence suspicion score
Adobe Photoshop batch image resize	0.25	0.00	0.00	0.00	0.25
Application (FileZilla) update process	0.25	0.00	0.00	0.00	0.25
AESCrypt	0.72	0.00	2.00	2.00	4.72
AxCrypt	0.68	1.00	2.00	2.00	5.68
WinRAR compress	0.00	0.00	0.00	0.00	0.00
WinRAR decompress	0.38	0.00	0.00	0.00	0.38
Zip compress	0.00	0.00	0.00	0.00	0.00
Zip decompress	0.38	0.00	0.00	0.00	0.38
Git	0.25	0.00	0.00	0.00	0.25
Median	0.25	0.00	0.00	0.00	0.25

Table 2: A list of benign application and their *behavior-based* indicator values additionally with the sequence suspicion score.

7.2 Indicator effectiveness

Comparing the indicators and their score between the benign software and the ransomware we notice that the quantities and the file type funnelling clearly separate them. Additionally, the file suspicion level is lower by benign software but is only useful in the union of the others due to the classification of the file name, file extension and file class which are not free of false positives due to the simplicity of faking these properties.

In contrast the other two *behavior-based* indicators are less burdened with false positives, since it is much harder to hide those operations from the monitor, and capture multiple properties which are significant for ransomware.

8 Usability

To evaluate the usability of our approach we performed an online survey with 30 participants between the age of 20 and 46 years with 57,1% being male and 42,9% being female.

The goal was to evaluate the usability of our approach compared to the normal way of restoring files in a personal cloud storage and in addition the support of color guidance for detection ransomware attacks.

The survey was structured as following: Firstly, we collected personal information and the foreknowledge followed by the preference about single file recovery or sequence recovery. Thirdly, we asked the participants to decide, which of the given three sequences depicts a ransomware attack. This was done for three sets of three sequences each with and without color guidance. To assist the participants spotting the ransomware attacks, two behavior characteristics were given. Afterwards, the feedback and criticism of the color guidance and the whole approach was collected and we also asked – for participants who preferred single file recovery over sequence recovery – if they changed their mind about it.

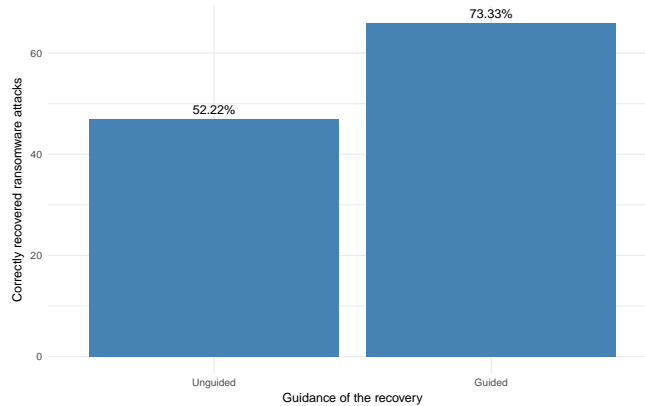


Figure 12: Correctly – by the survey participants – recovered ransomware attacks with and without guidance.

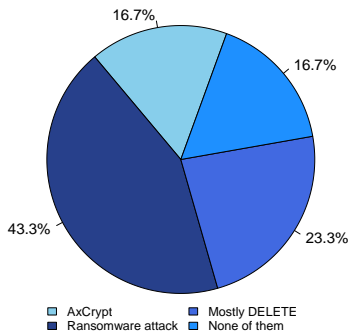


Figure 13: Unguided sequences classified as ransomware attack for sequence set 1.

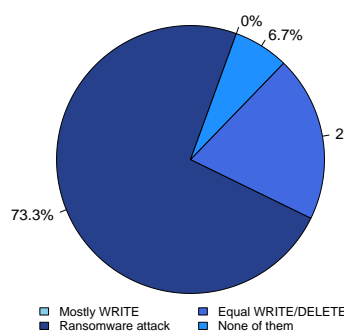


Figure 14: Unguided sequences classified as ransomware attack for sequence set 2.

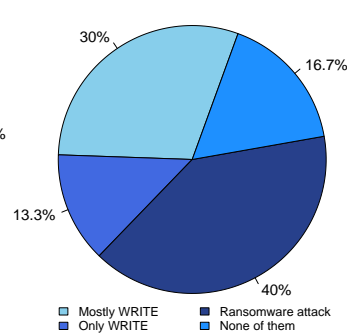


Figure 15: Unguided sequences classified as ransomware attack for sequence set 3.

50% did not know something about ransomware before taking the survey but two did not answer this question. Asked if the participants would prefer to recover all files file-by-file or as a sequence of files – 96,7% preferred the sequence recovery. Reasons for this were the easy usage and the quicker approach but it was also demanded to be able to select single files. The one person preferring single file recovery over sequence recovery changed their mind after learning about the sequence recovery.

As shown in figure 12 the correctly classified sequences of file operations increased by 21% with the help of the colour guidance. This is also confirmed by 73% of the participants who said that the colour guidance helped them spotting the ransomware attacks quicker and easier due to the clear colour choice and the direct comparison. One person was confused by the guidance and the approach as whole. Two participants were not able to spot the ransomware attack neither without guidance nor with the guidance and did not felt like the guidance made it easier.




Figure 16: Guided sequences classified as ransomware attack for sequence set 1.




Figure 17: Guided sequences classified as ransomware attack for sequence set 2.




Figure 18: Guided sequences classified as ransomware attack for sequence set 3.

Without the guidance the ransomware attack must fulfil one of the both given behavior characteristic very clear to be classified as ransomware attack by the participants. If this is not given or a sequence looks similar, the percentage of the correctly as ransomware classified sequence by the participants is only the half (see figures 13, 14 and 15).

In contrast the sequences with guidance were classified correctly in over 66% of all questions (see figures 16, 17 and 18).

The foreknowledge does not help the user when it comes to spotting the ransomware attacks. We assume the reason for this is that many reports about ransomware do explain the basic behavior but the most readers cannot transfer their knowledge into a concrete process on a computer.

Conclusively, the recovery interface with the file sequences and the colour coding are easily adopted and well understandable thus offer a very good usability. The color coding reduces the uncertainty of the user which seem to exist without it. The participants liked the quick recovery of file sequences making the extra amount of work – in contrast to automatic recovery approaches – not problematic.

More details of the survey are given in [6].

9 Limitations

Although we extract the false positives by handing over the decision to the educated user guided by the information gathered by our analysis and classification methods, we can separate the limitations in the non-technical user responsibilities and the technical ones.

The non-technical user responsibilities are to be creative, adaptive and sceptical to non-suspicious sequences which can be created with bypassing our technical approaches [15]. Furthermore, the possible false negatives deriving from short sequences which cannot be reliably classified by the *behavior-based* indicators.

Independent from the non-technical user responsibilities, there are some technical limitations, which can be used to evade our detection mechanisms and the according counter-measures.

9.1 Attacking the file operation monitor

Any interaction with the *Nextcloud* storage using the web interface or any client which uses the WebDAV protocol is monitored. Bypassing this monitor is only possible by having direct access to the file system of the server on which *Nextcloud* is installed. Assuming the user accesses the *Nextcloud* storage only with the web interface or any WebDAV client software, it is not possible to bypass the monitoring.

Additionally, the attack of the system of the server on which *Nextcloud* is installed would lead to an easy recovery since the files still exist on the client-side.

9.2 Attacking the file classification score

The file classification score consists of three parts: The file name classification, the file extension classification, including the file corruption, and the entropy of the file data.

For every part simulating a normal file takes effect and allows the attacker to keep the suspicion level below the threshold. The simulation of a normal file is simple for the file name and file extension but more extensive when it comes to the data entropy. This could be done by adding garbage data to it. This would lead to additional effort, for adding and removing the garbage data, and would also increase the file size. Hence we may gain additional detection patterns and encrypted files with the same reduced entropy or additional header informations about the entropy reduction if done randomly for every file.

All three bypassing methods are possible and would reduce the file suspicion level drastically but would be detected in the sequence classification score because the attacker has to encrypt many files on the system – optimally all files on the system – to launch an successful ransomware attack.

9.3 Attacking the sequence classification score

Attacking the sequence classification score is much more complex then attacking the file classification score since it is not based on *content-based* or *metadata-based* indicators but rather on *behavior-based* indicators. The behavior of ransomware is unique and is hardly changed.

We use four *behavior-based* detection techniques to classify a sequence of file operations: operation quantities and size quantities can be circumvented by delaying operations long enough that the monitor creates a new sequence or by writing additional files. However, delaying operations are contradictory to the needs of ransomware and increase the risk of being detected in the meantime. This risk is targeted by our approach by defining a long enough time interval between sequences. Although some ransomware families keep the file extension from the victims file, keeping the header informations of the file to bypass the corruption detection is special and would make the encryption and decryption process more complex. Writing additional files or increasing the file size by adding garbage data leads to additional effort because it must be added systematically that it can be removed before decrypting the file. The file type funnelling can be bypassed by using the same file extension and the same header informations as the victim file. To bypass the entropy funnelling the attacker must reduce the entropy of the file. The problems with this were already explained in the last subsection.

All these attacks are possible but are only effective as an union. Implementing only one or two would lead to a lower suspicion but the sequence could still be easily spotted as ransomware attack by the user.

10 Conclusions

In this paper, we defined general problems fighting ransomware with generic properties. This included the base rate fallacy and the problem of missing significant generic indicators regarding a single file for blocking them in real-time.

We also described ransomware classes based on the behavior of the file destruction, the operations sequence, the file extension of encrypted files, the file name of encrypted files and the file type funnelling. Following, we formulated ransomware indicators and categorised them into three classes: *Content-based*, *metadata-based* and *behavior-based* ones.

We observed that we are able to separate encrypted from compressed files on the basis of the standard deviation of the entropy once we split the file in data blocks, allowing us to increase the validity of the entropy indicator.

Justified through the base rate fallacy of automatic real-time detection and recovery of ransomware attacks, we propose a delayed detection and recovery based on a personal cloud storage with file versioning permitting us to remove false positives and increase the usability by taking the user into responsibility. This approach is supported by the guidance of the classification of the file operation sequence based on the ransomware indicators.

This method is supported by the idea that we have no need to recover from ransomware attacks in real-time as long as we do not lose data in the mean-time, we can increase the quality of the classification and the usability of the recovery for the user.

The usage of a strategy with file versioning opens up the use of *behavior-based* indicators to classify the whole sequence file operations according to a ransomware attack. This strategy reflects an improvement regarding the already proposed classification approaches.

Additionally, the implementation of the ransomware detection on the basis of a personal cloud storage removes the threat of the monitor being attacked and also takes care that there is always a backup of the files.

11 Acknowledgements

The authors would like to thank VirusTotal for providing malware samples.

References

- [1] BBC UK. A scanner to detect terrorists. [online], 2009. http://news.bbc.co.uk/2/hi/uk_news/magazine/8153539.stm (accessed March 28th, 2018).
- [2] Krzysztof Cabaj and Wojciech Mazurczyk. Using Software-Defined Networking for Ransomware Mitigation: The Case of CryptoWall. *Netw. Mag. of Global Internetwkg.*, 30(6):14–20, November 2016.
- [3] Andrea Continella, Alessandro Guagnelli, Giovanni Zingaro, Giulio De Pasquale, Alessandro Barengi, Stefano Zanero, and Federico Maggi. ShieldFS: A Self-healing, Ransomware-aware Filesystem. In *Proceedings of the 32Nd Annual Conference on Computer Security Applications*, ACSAC '16, pages 336–347, New York, NY, USA, 2016. ACM.
- [4] Craig from /dev/ttyS0. Differentiate Encryption From Compression Using Math. [online], 2013. <http://www.devttys0.com/2013/06/differentiate-encryption-from-compression-using-math/> (accessed March 28th, 2018).
- [5] Dropbox Inc. What to do if your files were corrupted or renamed by ransomware. [online], 2017. <https://www.dropbox.com/help/security/ransomware-recovery> (accessed March 28th, 2018).

- [6] Matthias Held. Detecting ransomware. Master's thesis, University of Konstanz, Konstanz, Germany, 2018.
- [7] Amin Kharaz, Sajjad Arshad, Collin Mulliner, William Robertson, and Engin Kirda. UNVEIL: A Large-Scale, Automated Approach to Detecting Ransomware. In *25th USENIX Security Symposium (USENIX Security 16)*, pages 757–772, Austin, TX, 2016. USENIX Association.
- [8] Amin Kharraz and Engin Kirda. Redemption: Real-Time Protection Against Ransomware at End-Hosts. In *RAID*, 2017.
- [9] Amin Kharraz, William Robertson, Davide Balzarotti, Leyla Bilge, and Engin Kirda. Cutting the Gordian Knot: A Look Under the Hood of Ransomware Attacks. In *Proceedings of the 12th International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment - Volume 9148*, DIMVA 2015, pages 3–24, New York, NY, USA, 2015. Springer-Verlag New York, Inc.
- [10] Eugene Kolodenker, William Koch, Gianluca Stringhini, and Manuel Egele. PayBreak: Defense Against Cryptographic Ransomware. In *Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security, ASIA CCS '17*, pages 599–611, New York, NY, USA, 2017. ACM.
- [11] Vadim Kotov and Mantej Singh Rajpal. Understanding Crypto-Ransomware: In-Depth Analysis of the Most Popular Malware Families. Technical report, Tech. rep. Bromium, 2014.
- [12] Segun I Popoola, Ujioghosa B Iyekekpolo, Samuel O Ojewande, Faith O Sweetwilliams, and AA Atayero. Ransomware: Current Trend, Challenges, and Research Directions. In *Proceedings of the World Congress on Engineering and Computer Science*, volume 1, pages 169–174, 2017.
- [13] Kevin Savage, Peter Coogan, and Hon Lau. The evolution of ransomware. *Symantec, Mountain View*, 2015.
- [14] N. Scaife, H. Carter, P. Traynor, and K. R. B. Butler. CryptoLock (and Drop It): Stopping Ransomware Attacks on User Data. In *2016 IEEE 36th International Conference on Distributed Computing Systems (ICDCS)*, pages 303–312, June 2016.
- [15] Bruce Schneier. *Beyond Fear: Thinking Sensibly About Security in an Uncertain World*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2003.
- [16] Symantec Corporation. Internet Security Threat Report: Ransomware and Businesses 2016. [online], 2016. https://www.symantec.com/content/en/us/enterprise/media/security_response/whitepapers/ISTR2016_Ransomware_and_Businesses.pdf (accessed March 28th, 2018).
- [17] Symantec Corporation. Internet Security Threat Report: Volume 22. [online], 2016. <https://www.symantec.com/content/dam/symantec/docs/reports/istr-22-2017-en.pdf> (accessed March 28th, 2018).
- [18] Symantec Corporation. Internet Security Threat Report: Ransomware 2017. [online], 2017. <https://www.symantec.com/content/dam/symantec/docs/security-center/white-papers/istr-ransomware-2017-en.pdf> (accessed March 28th, 2018).
- [19] Symantec Security Response. BadRabbit: New strain of ransomware hits Russia and Ukraine. [online], 2017. <https://www.symantec.com/connect/blogs/badrabbit-new-strain-ransomware-hits-russia-and-ukraine> (accessed March 28th, 2018).
- [20] Symantec Security Response. Petya ransomware outbreak: Heres what you need to know. [online], 2017. <https://www.symantec.com/blogs/threat-intelligence/petya-ransomware-wiper> (accessed March 28th, 2018).
- [21] Alan Turing. On computable numbers, with an application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society*, 42:230265, 1937.