Measurements for the Paranoid: The Effect of Encrypting Files in Cloud Storage

Stephen Geerlings
University of Twente
P.O. Box 217, 7500AE Enschede
The Netherlands
s.a.geerlings@student.utwente.nl

ABSTRACT

Cloud storage is increasing in popularity and has been for years. People upload personal data in large quantities to the cloud without concerning themselves with safeguarding of that data. This paper focusses on the encryption of data and the synchronization of that data. Different synchronization techniques are researched to gain an answer on how encryption will affect the techniques. Tests are performed with Dropbox, a service that uses many smart synchronization techniques, in a worst case manner. ECB and CBC modes of operation are compared to no encryption in light of the data exchange needed for synchronization. Conclusions are drawn based on the results of these performance measurements.

Keywords

Cloud Storage, Synchronization, Encryption, Performance measurements.

1. INTRODUCTION

The 'cloud' has been a buzzword for the last couple of years. It has been a major trend as a concept and users have adopted the idea of not knowing where their data is, for its benefits and ease of use. There are many different services that operate in the cloud to back-up files, grant acces to e-mail or offer virtual computing power. The cloud seems ideal for users but it has created many opportunities for agencies and hackers to exploit.

This research focusses on one aspect of the cloud, namely storage. Cloud storage is a technique for users to back-up, share or manage data from their local harddrives in the cloud. Most cloud storage providers (CSPs) offer services to automatically upload data from the user directory to the cloud. In this way, a user can safeguard their photo's or share a document with anyone they want. However, user data will be stored in the cloud and it will be stored somewhere on a server that is owned by the CSP. This user data could be harmless family photos but it could also be valuable information for a CSP or an government agency. The users should protect themselves against these possibilities.

Agencies of the United State have many rights to track or spy on internet users. The Patriot Act allows these agencies to use many measures to protect the United States. The documents of Edward Snowden have revealed that the NSA have advanced techniques to access user data on the internet. As many CSPs use servers that are based in the USA [14], these CSPs need to share user data with the US government if it is demanded. The privacy statements of Dropbox [5], Google [7] and Microsoft[10] all support this while these are the largest CSPs [14].

Techniques and applications exist to protect valuable data. The most used technique is encrypting the data. Encryption is a mathematical proces of transforming the data with a key. The key must be stored to ensure that the data can be tranformed back from the encrypted data. There are many proven usefull encryption techniques that can be used by average users. However, CSPs use advanced techniques to synchronize data with the cloud. These techniques, that are based on the structure and order of data, need to be able to keep track of the data and changes in the data. The synchronization techniques could be rendered useless when the data on which they operate is encrypted and unstructured.

The goal of the research is to find the effect that locally encrypting files has on the synchronization techniques. The scope of the research focusses on how users can safeguard the data that they store in the cloud and the possible disadvantages this would entail. A brief look into possible encryption programs, CSPs with local encryption and usage of cloud storage is given to shred a light on the background of the research. Experiments are performed to attain an idea whether local encryption is wasteful in data and time usage.

The paper starts an overview of the research in section 2 and with describing related work on the subject in section 3. An overview of encryption techniques and the CSPs is given in section 4. After that the experiment setup and and analysis method is described in section 5. Section 6 gives the results of the experiments and these are discussed in the section 7. The paper finishes with conclusions and future work in section 8 and 9.

2. RESEARCH

Research is performed investigate the previously described scenario. The research question is as follows.

- What is the effect of locally encrypting files on the performance of data synchronization of cloud storage?

To answer the main question the following subquestions were drafted.

20th Twente Student Conference on IT January 24th, 2014, Enschede, The Netherlands.

1. What are the different approaches for encryption in the cloud an end-user can take and how could they affect the CSP synchronization techniques?

2. What are the performance consequences in terms of bandwidth usage if users encrypt their cloud stored data?

The goal of the first question is to find the ways in which an average user could encrypt his cloud stored data. Furthermore a brief overview of encryption techniques is given. An evaluation of some of the most popular CSPs is done to clarify how they operate and what kind of smart synchronization techniques they employ. This is necessary to give a proper idea on how encryption could disrupt these smart encryption techniques.

The second question will yield an answer on the issue of how a cloud storage application behaves differently when files are encrypted. The development of a method for testing will be discussed as important design decisions will be made. The criteria that will be developed as part of the analysis method will be usefull for testing multiple CSPs and they will yield general results so CSPs could be compared. Therefore, the CSP applications will be tested as a black box. The results of these tests will be discussed and evaluated in light of current data usage of the cloud service. A conclusion about the benefits and disadvantages will be given based on data usage and encryption gains.

3. RELATED WORK

Based on our knowledge so far there is little research into the field of cloud storage for consumers. This research is split into multiple fields: Cloud Storage, synchronization, encryption and user file modifications. There is some research in cloud storage benchmarking which helped greatly. Synchronization and encryption are fields that are researched extensively. However, not much is known about how users use the cloud and how files in the cloud are modified.

The paper about the inner workings of Dropbox [14] gave an understanding of the system behavior of the Dropbox application. The paper reports that the files stored at Dropbox are actually on Amazon AWS (Amazon Web Services) servers that are placed in the United States.The research done on Dropbox revealed that it uses multiple servers to synchronize data with the cloud. This is confirmed by research done on synchronization delays and bottlenecks of Dropbox [20]. A following research by mostly the same authors tried to develop a way to benchmark personal cloud services [13]. The authors benchmarked the applications of Dropbox, Skydrive, Google Drive, Wuala and Amazon Cloud Drive. The setup to benchmark the services is similar to the one used in this research. The authors tested Wuala because it offers local encryption. The tests that were performed are more extensive en elaborate but they do not show the effect of encryption on the service of Wuala. The paper also contains an overview of synchronization techniques that are employed by the services that were benchmarked. The researchers showed the differences in approach of the CSPs and the effect that it had on the performance of the service.

Synchronization of data is a key aspect of this research. The benchmarking research sums up all the used techniques for efficient data synchronization [13] in cloud storage. Dropbox uses delta encoding amongst other techniques. This is a proces in which only the changes to a file are synchronized with the server [14]. Other techniques are chunking large files, bundling small files, compression and data deduplication. These techniques reduce the data traffic that is needed for synchronization. This is good as data traffic for Dropbox already accounts for a third of Youtube in some areas according to Drago. Dropbox is in the top ten of upstream data traffic globally in the first half of 2013 [9].

Encryption is deeply researched and algorithms and modes of operation are great in number. This research will focus on the laymans encryption: encryption that an average user could apply to their files. TrueCrypt is an encryption tool that is easy in usage [1] [18]. Users of TrueCrypt can encrypt their data using widely accepted algorithms such as AES (Advanced Encryption Standard). The authors of [19] describe the inner workings of these algorithms together with the modes of block operation. These different modes act differently on data and they are taken into consideration during experiments. Wuala uses Cryptree [16] to offer encryption to users.

To our knowlegde, no research is performed on how users change their files. There is little known about what files are kept in the cloud folder of a typical user. Cloud storage is used to store documents, photos and other data. These files are typically in the range of some kilobytes while photos are ranging around some megabytes. However, this is not the most important part. An understanding of typical changes made to a file is of interest. To benchmark a system designed for data deduplication, researchers in Korea have used a random patch that is placed in the file at a random location [17]. The researchers that developed the benchmarking system considered three cases: new data added/changed at the end, at the beginning or at a random position in the file [13]. Both approaches needed a great amount of tests to gain an appropiate certainty.

4. SYNCHRONIZATION AND ENCRYPTION

Synchronization techniques are an important part of this research. A background is needed to know how the synchronization techniques operate and how encryption could affect them. CSPs are analyzed on what techniques they use to synchronize files and if they offer encryption. Other techniques for file incryption are investigated to see how they could affect synchronization.

4.1 Overview of CSPs and their Synchronization Techniques

Many different CSPs exist and the offer of cloud storage is continuously increasing. Only a small number of CSPs will be discussed in this paper. These CSPs are: Dropbox [4], Skydrive[11], Google Drive [6], Wuala [12], Mozy [8] and Carbonite [3]. A choice is made for these as some have many users and others have used encryption in their service. The overview of CSPs and their properties can be found in table 1. The properties of interest are the synchronization techniques that are used and the encryption that is offered by the service.

Dropbox uses many of the smart synchronization techniques as these techniques decrease data exchange. This is beneficial as Dropbox only has servers in the United States and this decreases throughput [13]. Dropbox stores files on Amazon S3 servers which are encrypted. However employees have access to the keys and to data stored on the server. Dropbox employs many techniques but the most relevant is the Delta-Encoding. Dropbox is the only provider which uses this technique fully [13]. Skydrive and Google Drive have similar approaches except Google Drive uses compression. The rightmost three providers in table 1 are the CSPs that offer encryption. There has not been

research into the synchronization techniques they use but some things are noticable. Mozy and Carbonite do not offer sharing of files. The files are encrypted at client side with a key that is known only to the user. Wuala offers sharing as it employs Cryptree. This is a key management encryption paradigm which offers a tree like encryption for a folder structure. The key for a root folder can decrypt the child folders but a key for child folders cannot decrypt the root folder or other files in the folder [16].

4.1.1 Chunking

Chunking is the proces of dividing a large file in chunks. These chunks are then synchronized with the server. If synchronization stops at any moment the chunks that were already uploaded need not to be uploaded again. Similarly, if one chunk is changed then only that chunk needs to be synchronized as the other chunks are unchanged.

4.1.2 Bundling

Bundling small files into one batch to process will decrease the overhead per file. Costly operations like creating a SSL connection will be done for multiple files instead for every file. The research in [13] concludes that only Dropbox uses effective bundling. Wuala and Skydrive reuse connections but wait for acknowlegdement per file which causes unnecessary delays.

4.1.3 Compresson

Compression of the content of the file will deminisch the amount of data that needs to be exchanged if the data is compressable. Google Drive uses smart compression which will only compress a file if it is beneficial for the service.

4.1.4 Deduplication

Deduplication searches for copies of a file before uploading a file to the server. If a copy is already present in the cloud then the file does not need to be uploaded again. Deduplication can also be used in combination with chunking.

4.1.5 Delta-encoding

Instead of uploading files again after a change, delta-encoding will keep track of the changes and it will synchronize only the changed parts of a file with the cloud. This is not similar with chunking as the latter considers entire chunks and the first considers only changes. Dropbox is the only CSP that has implemented delta-encoding [14] [13].

4.1.6 Sharing and personal key(s)

Sharing of files can be done in multiple ways in cloud storage. Users can be part of a group which has access to a file or users could share a download-link to a file. When an CSP uses encryption, the process of sharing files becomes more complicated. Different users will need the key to decrypt the data that is shared. Keeping track of these keys is a responsibility for the users or for the CSP. However, it should be a user responsibility to ensure the privacy of users. Wuala uses an advanced key management system named Cryptree to encrypt files and folders locally and enable users to share files with others [12] [16].

4.1.7 Encryption transport and storage

Most CSPs use a layer of encryption to transport the data securely to the cloud. This is not necessary for Wuala as the service encrypts files before sending them. The CSPs use acces control to safeguard files from unauthorised users but not all of them encrypt the data stored on the servers.

4.2 Overview of encryption techniques

Encryption is a reversable process of changing files mathematically to obscure their contents. There are different methods of encrypting a file but most use a random key to change the data. The first that is discussed is BoxCryptor [2]. This is a specialized application for cloud storage to encrypt files locally. BoxCryptor services are designed for many different CSPs such as Dropbox, Skydrive, Google Drive and even Wuala. The keys for the encryption and decryption process are encrypted with the users public key and stored in the encrypted file. BoxCryptor uses 256-bit AES algorithm in CBC mode to encrypt files and RSA to manage the public and private keys. Using public and private keys enables BoxCryptor-users to share files. The public keys are stored on the BoxCryptor key server and users can download eachothers public keys to encrypt a file for another user. BoxCryptor also offers group keys and company keys to share files.

The report of Fahl et al [15] creates a system which can encrypt messages that will be posted on Facebook. They create a layer which is applied to the interface of a website to encrypt and decrypt data. The same service is developed for Dropbox and files are automatically encrypted and decrypted using AES, the mode of operation was not specified. Furthermore, no benchmarks were performed on the service.

The third method is TrueCrypt [1]. This is an example of a typical application that is able to create encrypted volumes on a users hard disk. The volumes can be decrypted and mounted as a drive with a password and a key. TrueCrypt manages the keys in a secure file on the hard disk. TrueCrypt volumes can be any size and can be stored anywhere in the file system. A user could make a TrueCrypt volume and store it in its cloud folder. The TrueCrypt volume will be synchronized with the cloud when it is unmounted.

Block encryption is common for file encryption and it is used in TrueCrypt and BoxCryptor. The algorithm chops up the file in blocks of the same size and processes one block at the time. Block ciphers work with a mode of operation. This mode of operation defines whether and how a previous block is used to encrypt the next block. A file modification in one block will then modify its following blocks. In a regular setting, this is known as error propagation. Some modes of operation, like ECB (Electronic Code Book), do not have this feature and a small change plaintext will yield an output of one changed block in the ciphertext. CBC is an example that has error propagation as it uses the cipher text of the previous block with the key to encrypt the next block.

Table 2. ECB and CBC modes of operation

	IV	Padding	Stream	Error Propagation
ECB	yes	yes	no	no
CBC	yes	yes	no	yes

4.3 Error propagation and Synchronization

A small change in the beginning of a file will change the cipher text almost completely. This behaviour is of significance in this research as some CSPs only upload changes to the cloud in stead of entire files. A small change in plaintext will cause large changes in ciphertext and these large changes will be synchronized. So while the user only makes a small change, this addition of data will cause a large amount of data traffic. Table 2 contains the properties of ECB and CBC. These two modes of operation are similar but ECB is weaker as identical plaintexts are mapped to identical ciphertexts which is vulnerable to exploits [19]. The research uses ECB and CBC for testing

Table 1. Summary of CSPs adapted from [13]

	Dropbox	Skydrive	Google Drive	Wuala	Mozy	Carbonite
Chunking	4 MB	var.	8MB	var.	-	-
Bundling	yes	no	no	no	-	-
Compression	always	no	smart	no	-	-
Deduplication	yes	no	no	yes	-	-
Delta-encoding	yes	versioning	versioning	versioning	versioning	versioning
Sharing	yes	yes	yes	yes	no	no
Personal key(s)	no	no	no	Cryptree	optional	optional
Encryption transport	SSL	SSL	SSL	no	SSL	SSL
Encryption storage	yes	no	no	256-bit AES	Blowfish/AES	128-bit Blowfish

the CSPs to investigate the difference these make on the performance of the CSPs.

5. METHODOLOGY

The development of a method to measure the performance of Dropbox is discussed. The designed method for measuring will be used for testing. The tests that are performed will be discusssed.

The CSP that is tested is Dropbox. As this service uses many strategies to avoid unnessecary data traffic it is be ideal to perform a worst case scenario test on. The tests compare the behavior of Dropbox when the files are not encrypted to the behavior when files are encrypted with ECB and with CBC. The tests without encryption show Dropbox behavior under normal circumstances.

5.1 Test setup

The test setup was hosted on a Windows 7 machine which runned Virtualbox. A clean Windows 7 virtual machine was runned on which only the client applications, Java and Wirehark where installed. The analysis of the pcap-files is done on a ubuntu-virtual machine. The specifications of the hardware and versions of the software can be found in table 3.

The setup consists of a script and the client application that is tested. The script connected to the CSP server to detect if a file is uploaded or revised at the server. The script uses the regular API of Dropbox whichl offers file search and versioning of files. The script first moves a new file to the local cloud folder. The client application will start synchronizing the file to the cloud. The script will start polling the server whether the file is completely uploaded and if so it will store a version number. The script will then overwrite the file in the local cloud folder with a file that has modifications and the client application will start synchronizing again. The script creates timestamps during each step so the flow data can be easily analyzed. Furthermore, Netstat, a regular Windows application, is used to keep track of all the open ports that are used by programs. This was necessary as both the script and the client application connected to Dropbox servers. When the script recieves a new version number from the server then the proces is repeated with a new file. An overview can be found in figure 1.

Wireshark is run during the tests to capture all the packets that are exchanged. The generated pcap-file is stored and analyzed on a ubuntu virtual machine with the tool Yaf. This creates flows from the packet data. The flows are split into small flows to distinquish between the upload of one file and the next. The next step is to combine the timestamps, the used ports and the flow analysis and generate a file which shows the IP-addresses and the exchange of data for each file. The results of the measurements are

Table 3. Versions of used software

Dropbox	2.0.22
Wireshark	1.10.3
Oracle VirtualBox	4.3.4 r91027
Yaf/Yafscii	2.4.0
Windows 7 (VM)	SP1, 32-bit, 2.0 GB RAM
ubuntu (VM)	12.04 LTS, 64-bit, 1.0 GB RAM
Windows 7 (host)	SP1, 64-bit, 8.0 GB RAM, i7-3517U

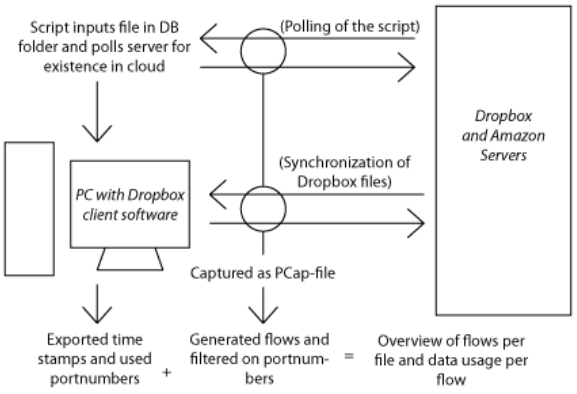

Figure 1. Experiment setup

averaged and then plotted in figures 2 to 5.

5.2 Test files

There is no research on the types, sizes, changes or movements of files in a cloud environment to our knowledge. This is important as it could have defined what size of testfiles should be used for testing. As there is no information on this subject a variety of filesizes is used for testing. A typical document amounts for a file size of 10 kB to 100 kB and a photo for 3 MB to 5 MB. As cloud storage is used for these files, the testfiles should be of similar size. It is necessary to keep in mind that CSPs could use chunking as a synchronization technique. A chunk could be synchronized entirely if a small part of the chunk is changed. Therefore the file sizes used for testing are somewhat larger than the chunksize of Dropbox. The sizes for the unencrypted test files can be found in table 4.

Since ECB and CBC differ in their property of error propagation these two modes of operation are used for testing. The unencrypted files will be encrypted with AES and padded accordingly with standard Java libraries. The changed unencrypted files are encrypted with the same key as the unchanged files. These encrypted files will be used in the experiments in the same way as the unencrypted files are. The ECB encrypted file sizes and changes can be found in table 5 and the CBC in 6. The file sizes are

Table 4. Files using no encryption

File	Filesize [B]	Change [B]	Change [%]
10 kB	10240	1038,33	10,14%
1024 kB	1048576	104450	9,96%
5120 kB	5242880	522256,33	9,96%
10240 kB	10485760	1044497	9,96%

Table 5. Files using ECB encryption

File	File size [B]	Change [B]	Change [%]
10 kB	10256	1068,33	10,42%
1024 kB	1048592	104476	9,96%
5120 kB	5242896	522342,33	9,96%
10240 kB	10485776	1044512,67	9,96%

Table 6. Files using CBC encryption

File	File size [B]	Change [B]	Change [%]
10 kB	10256	10234,67	99,79%
1024 kB	1048592	1044588,33	99,62%
5120 kB	5242896	5222319,67	99,61%
10240 kB	10485776	10444915	99,61%

slightly larger as the files where padded for encryption.

The changes that where made to the files are 10% of the file size. This was chosen as a suitable number as there is no knowlegde of file modifications in cloud environments. The placement of the change in the file also affects the encrypted result and the synchronization of the file modification [17]. The changes where made at the beginning of the file and the change itself were random bytes created with the Java Random function. This is a worst case scenario as the files that are encrypted with CBC as mode of operation will almost change completely while ECB encrypted files differ slightly compared to their unmodified files. In section 6 the effect of the chosen changes will be discussed.

6. RESULTS

The experiments were performed and the results are plotted in the figures. The results consist of upload and download traffic during the synchronization of a file. The traffic to Dropbox and Amazon is divided to give an idea of the signalling data and the actual file upload data exchange. The IP-addresses were identified using reverse nslookup and traffic to either Dropbox or Amazon was combined. The means were calculated with the results of three experiments. A logarithmic scale is used to plot the results clearly.

Figure 2 shows the outbound traffic while uploading a new file. The actual filesize is also plotted to give an idea that the measured traffic is proportional to the file size. The measured traffic is slightly larger than the file size, this is due to some overhead in the file upload. The inbound traffic during a new file upload is shown in figure 5. The amount of data send to and received from Dropbox does not grow proportionally to the file size. As there is no actual difference between uploading a new encrypted file and uploading a new unencrypted file all the upload and dowload traffic measured is roughly the same.

The behaviour of Dropbox after the modification of the file is shown in figures 4 and 5. The measured outbound traffic is lower for no encryption and ECB compared to the CBC mode. The actual filesize is plotted for reference. Note that the y-axis is again in logarithmic scale. The difference in the ECB and CBC mode is proportional to the amount of change shown in table 5 and 6. The upload

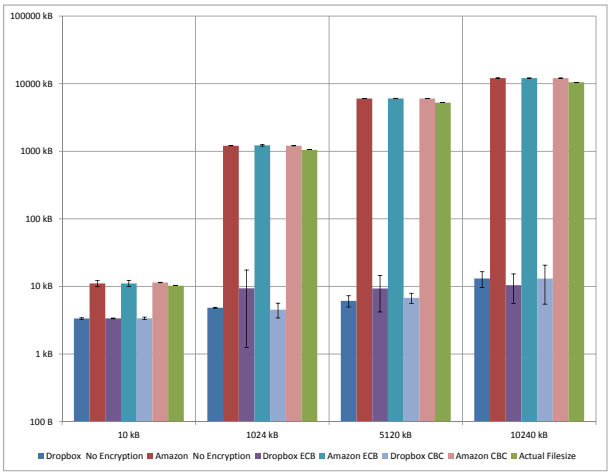

Figure 2. Send traffic with new file

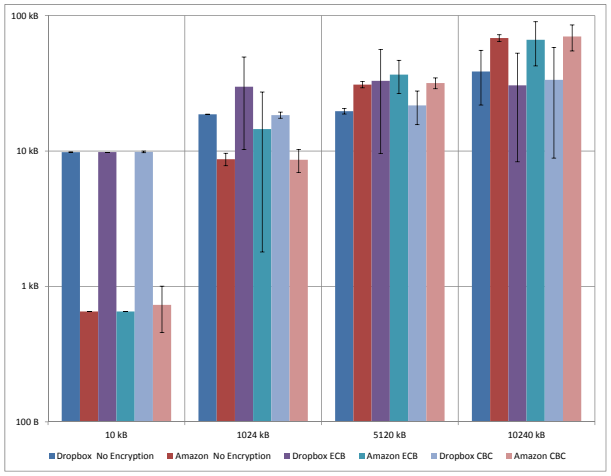

Figure 3. received traffic with new file

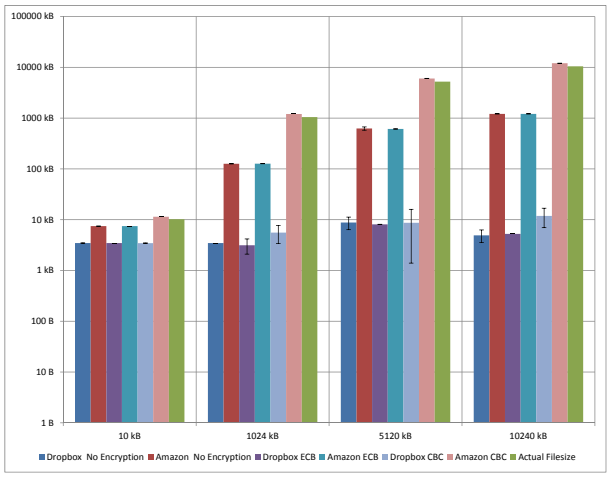

Figure 4. Send traffic with changed file

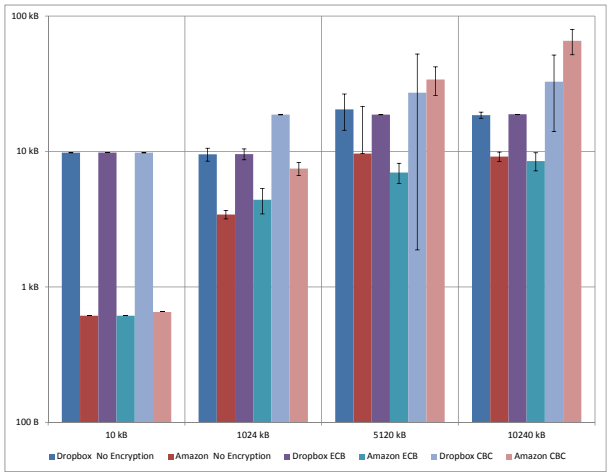

Figure 5. received traffic with changed file

of a change to the cloud using CBC uses almost the same amount of data of uploading a new file to the cloud. The files encrypted with CBC differ from the changed files for 99% as is shown in table 6

7. DISCUSSION

This research finds an answer to the question whether encryption of cloud stored data affects the performance of an cloud storage application. The performance is unaffected when uploading new files to the cloud. This is not suprising as the data that is uploaded is not present at the server and needs to be uploaded entirely whether it is encrypted or not. One could argue that the encryption proces takes time but this time is very small compared to the total amount of time it takes to upload a file to the cloud.

When the files are modified by a user, the CBC encryption proces changes the file entirely if this modification is at the start of the file. Errors propagate only through the rest of the blocks that need to be encrypted. If the file is changed in the middle then the encryption proces changes only the second half of the file. The ECB process only changes the blocks that are affected by the modification in the original file. Figure 5 clearly shows that CBC causes more data to be exchanged while there is only 10% of change. The client application performs almost the same with ECB encryption as it does without encryption.

The difference between ECB and no encryption and CBC encryption is noticable for the large files. But the test with the 10KB files show that Dropbox exchanges more then the 10% change. This could be due to overhead in the upload traffic.

With the fact that synchronization of ECB encrypted files operates as good as no encryption the benefit of ECB encryption is clear. While ECB is a method of security it has weaknesses such as a replay attack [19]. It does obscure data so that an attacker has to put effort into obtaining your private data.

8. CONCLUSIONS

This research aims to resolve the question whether the encryption of files interferes with synchronization techniques and whether this should be considered a loss or a gain. A method for test CSPs as a blackbox system was developed and Dropbox was tested with and without encryption. Encryption will interfere with the synchronization techniques. Encryption changes parts of a file and the cloud storage application sees this as modifications and uploads them to the cloud. Different modes of operation affect the data differently and error propagation is the key feature which defines how much data is changed. The synchronization techniques perform just as under normal circumstances except the data on which they operate is changed differently.

Tests showed that CBC will cause more data exchange as larger parts of the file are changed. ECB and no encryption operate equally good in the Dropbox case. The amount of data that is exchanged is proportional to the amount of change of a file. ECB could be usable in cloud storage as the application uses the same amount of data as it would without encryption while the file is still encrypted.

9. FUTURE WORK

The small files did not show the property that the data traffic is proportional to the change. This could have many reasons such as overhead or implementation of Dropbox. This should be further researched as many users have a great amount of small files compared to larger files. To give a clear estimation on the large scale effects an average case scenario should be researched with random amount of change on a random offset. It is also interesting on how the CSPs that offer encryption operate and how much data is exchanged during normal operation. This should bring to light some concessions that could be made by these providers.

10. AKNOWLEDGEMENTS

The research was conducted at the University of Twente at the chair for Design and Analysis of Communication Systems (DACS). I would like to thank A. Sperrotto and M. Karimzadeh who have guided me with the research for their feedback and ideas. It has been of invaluable help.

11. REFERENCES

[1] Truecrypt - free open-source on-the-fly disk encryption software for windows 7/vista/xp, mac os x and linux, Jan. 2013. http://www.truecrypt.org/.

[2] Boxcryptor | encryption for cloud storage, Jan. 2014. https://www.boxcryptor.com/.

[3] Carbonite support, Jan. 2014. http://www.carbonite.com/support.

[4] Dropbox frequently asked questions, Jan. 2014. https://www.dropbox.com/help.

[5] Dropbox privacy statement, Jan. 2014. https://www.dropbox.com/privacy.

[6] Google drive support, Jan. 2014. https://support.google.com/drive/?hl=nl#topic=14940.

[7] Google privacy statement, Jan. 2014. http://www.google.nl/intl/en/policies/privacy/.

[8] Mozy support, Jan. 2014. http://mozy.com/.

[9] Sandvine global internet phenomena report, Jan. 2014. https://www.sandvine.com/downloads/general/global-internet-phenomena/2013/sandvine-global-internet-phenomena-report-1h-2013.pdf.

[10] Skydrive support, Jan. 2014. http://www.microsoft.com/security/online-privacy/skydrive.aspx.

[11] Skydrive support, Jan. 2014. http://windows.microsoft.com/en-us/windows/files-folders-storage-help#files-folders-storage-help=windows-7&v2h=win8tab1&v3h=win7tab1.

[12] Wuala technology, Jan. 2014. http://www.wuala.com/nl/learn/technology.

[13] I. Drago, E. Bocchi, M. Mellia, H. Slatman, and A. Pras. Benchmarking personal cloud storage. In *Proceedings of the 2013 Conference on Internet Measurement Conference*, IMC '13, pages 205–212, New York, NY, USA, 2013. ACM.

[14] I. Drago, M. Mellia, M. Munafò, A. Sperotto, R. Sadre, and A. Pras. Inside dropbox: Understanding personal cloud storage services. pages 481–494, 2012.

[15] S. Fahl, M. Harbach, T. Muders, and M. Smith. Confidentiality as a service - usable security for the cloud. pages 153–162, 2012.

[16] D. Grolimund, L. Meisser, S. Schmid, and R. Wattenhofer. Cryptree: A folder tree structure for cryptographic file systems. In *SRDS*, pages 189–198. IEEE Computer Society, 2006.

[17] H. Jung, S. Park, J. Lee, and Y. Ko. Efficient data deduplication system considering file modification pattern. *International Journal of Security and its Applications*, 6(2):421–426, 2012.

[18] Q.-X. Miao. Research and analysis on encryption principle of truecrypt software system. pages 1409–1412, 2010.

[19] C. Paar and J. Pelzl. *Understanding Cryptography: A Textbook for Students and Practitioners*. Springer-Verlag New York Inc, 2010.

[20] H. Wang, R. Shea, F. Wang, and J. Liu. On the impact of virtualization on dropbox-like cloud file storage/synchronization services. 2012.

.