advances
in Technology

Full Paper

Robust Efficiency Evaluation of NextCloud and GoogleCloud

Nicholas Singh, Kevin Bui and Akalanka B. Mailewa*

Department of Computer Science & Information Technology, St. Cloud State University, St. Cloud, Minnesota 56301, USA

E-mail correspondence: amailewa@stcloudstate.edu (A. B. Mailewa)

Received: 18 December 2021; Revised: 31 January 2022; Accepted:15 February 2022; Published: 4 March 2022

Abstract

Cloud storage services such as GoogleCloud and NextCloud have become increasingly popular among Internet users and businesses. Despite the many encrypted file cloud systems being implemented worldwide today for different purposes, we are still faced with the problem of their usage, security, and performance. Although some cloud storage solutions are very efficient in communication across different clients, others are better in file encryption, such as images, videos, and text files. Therefore, it is evident that the efficiency of these algorithms varies based on the purpose and type of encryption and compression. This paper focuses on the comparative analysis of NextCloud with composed end-to-end solutions that use both an unencrypted cloud storage and an encrypted solution. In this paper, we measured the network use, file output size, and computation time of given workloads for two different services to thoroughly evaluate the efficiency of NextCloud and GoogleCloud. Our findings concluded that there is similar network usage and synchronization time. However, GoogleCloud had more CPU utilization than NextCloud. On the other hand, NextCloud had a longer delay when uploading files to their cloud service. Our experimental results show that the evaluation model is considered robust if its output and forecasts are consistently accurate, even if one or more of the input variables or assumptions are drastically changed due to unforeseen circumstances.

Keywords: Cloud efficiency, data encryption, GoogleCloud, NextCloud, security and privacy

Introduction

As tools for personal storage, file synchronization and data sharing become more in demand. The market providers are quickly evolving, with well-established providers like NextCloud [1] having quickly gained popularity is now competing with the likes of GoogleCloud [2] and Microsoft [3], which results in them now offering more and more integrated solutions into Windows, Android, and Linux Operating Systems. These services provide users with ubiquitous, reliable data storage that can be automatically synchronized across multiple devices that can be shared and accessed among groups of users anywhere in the world at any time. Understanding the typical usage of these services is of primary importance to improve their end-user experience by identifying bottlenecks, security vulnerabilities, and privacy concerns [4] [5].

Therefore, it is not surprising that the research community has put increased effort into understanding how personal and private cloud storage works. In our preliminary research, we identified several related works that are based on active measurements that will aid us in developing a test suite. In our paper, we performed a set of active experiments that measure the network use, file output size, and computation time of given workloads for two different services, which we will use to thoroughly evaluate the efficiency of NextCloud and GoogleCloud. We hypothesize that the tech giant GoogleCloud will have more robust efficiency than NextCloud in both CPU use and network usage.

https://doi.org/10.31357/ait.v1i2.5392

Cloud Storage and Features

There are many cloud service providers, and these service providers usually give out free storage space to a certain number of gigabytes, after which monthly fee subscriptions begin [6]. The cloud storage service providers supply drags and drop accessing options and synchronize folders and files between desktop and mobile devices and the cloud drive. They also allow all account users to collaborate on documents. The two providers will be conducting our research on are GoogleCloud and NextCloud.

GoogleCloud

The security of encryption in GoogleCloud is vital to them. Hence, they take body encryption seriously every day to protect the user's data, whether travelling over the internet, moving between their data centers, or storing it on their servers [7]. The use of encryption at GoogleCloud is usually combined with integrity protection. For example, someone with access to the ciphertext can neither understand it nor modify it without knowing the key. So, what does GoogleCloud consider customer data? Customer data is referred to the contents in GoogleCloud directly or indirectly. This data includes the customer's contents and metadata [8-9]. The customer content is data that GoogleCloud customers generate or provide to GoogleCloud. Data stored in cloud storage, disk snapshots used by Compute Engine, and Identity and Access Management policies. The customer metadata makes up the rest of the customer's data. That refers to all data that cannot be classified as customer content like auto-generated project numbers, timestamps, and IP-Addresses [10-11].

Encryption

GoogleCloud encrypts all the customer's content using one or more encryption mechanisms. All the data is encrypted into multiple layers to ensure protection and select the optimal approach based on application requirements [12]. As shown in Figure 1, all the data here is stored at rest, and the layers include Application, Platform, Infrastructure, and Hardware.

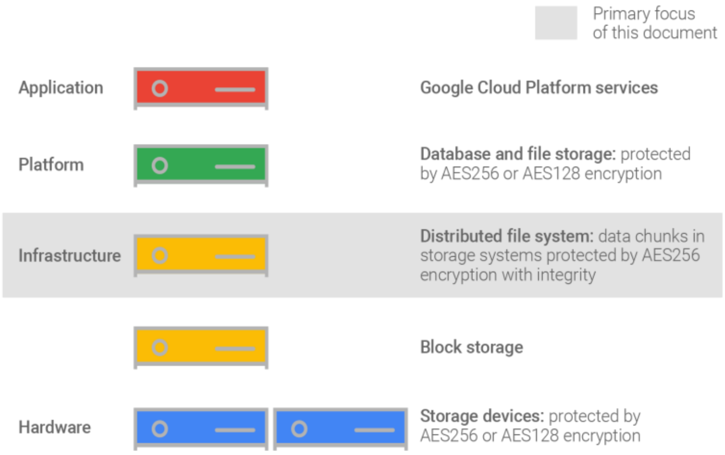
Primary focus of this document
Application
Google Cloud Platform services
Platform
Database and file storage: protected by AES256 or AES128 encryption
Infrastructure
Distributed file system: data chunks in storage systems protected by AES256 encryption with integrity
Block storage
Hardware
Storage devices: protected by AES256 or AES128 encryption

Figure 1. GoogleCloud Infrastructure [12]

During transit, GoogleCloud employs several security measures to help ensure authentication, integrity, and encryption. First, the authentication verifies the data source, whether human or electronically. The integrity makes sure the data that the user sends arrives at the destination without being changed [13]. And lastly, the encryption makes the user's data private during transport. The encryption has three states: rest, transit, and use [14-15]. Encryption during rest protects the user's data from a system compromise while being stored. The encryption during transit protects the user's data from being intercepted between the cloud or between two services. This protection is achieved by encrypting the data before transmission, authenticating the endpoints, and decrypting and verifying the data on arrival [14-16].

User-side vs Server-side

User-side (Client-side) encryption [17] has many downfalls when users do it on their own versus server-side encryption. As the user-side encryption, the user must create and manage their encryption keys. They also must use their tools to encrypt the data before sending it to cloud storage. The data you encrypted yourself arrives at the cloud storage in an encrypted state. However, the cloud storage does not know the user's keys to encrypt the data. When cloud storage receives the user's data, it is encrypted a second time. Cloud storage then removes the server-side layer of encryption; the user must decrypt it themselves [18]. If the user did send it through the cloud storage without encrypting it beforehand, the user would not need to decrypt it later in the process.

Storage

As shown in Figure 2, the data is broken up into subfile chunks for storage. Each chunk can be several GB in size and encrypted at the storage level (Block Storage in the infrastructure layer). Every chunk has a different key, even if the chunk belongs to the same layer. If a particular chunk is updated, it will encrypt a new key instead of reusing the key before. Since each chunk has a unique key, the Access control list (ACLs) ensures that the GoogleCloud services can only decrypt each chunk under authorized roles [19].

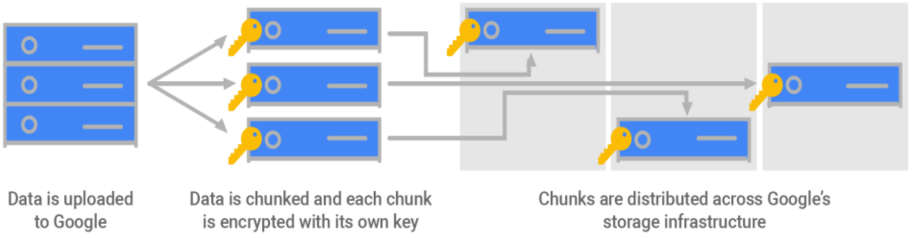
Data is uploaded
to Google
Data is chunked and each chunk
is encrypted with its own key
Chunks are distributed across Google's
storage infrastructure

Figure 2. GoogleCloud Data Storage [19]

The encryption that GoogleCloud uses is Advanced Encryption Standards (AES). In each storage, GoogleCloud uses both AES 256 and AES 128 for encryption. In AES 256, the storage is stored in hard disk drives (HDDs) and solid-state drives (SSDs). And as for AES 128, there is a small number of legacy HDDs for it (legacy HDDs support older software or data) [20] [21].

Supported Files, Clients, and Features

GoogleCloud allows text, images, PDFs, and Word documents. Certain files are not allowed, and that is Binary files. Binary files are unsupported files (during a storage scan) and images that cannot be scanned using optical character recognition (OCR). In the storage scan, if the file is not recognizable, it will scan it as a binary file and attempt to convert it to UTF_8 [22]. The supported clients of GoogleCloud are Linux, Windows, Mac, and mobile devices such as iOS and Android. To use GoogleCloud, there is a free feature and a paid feature. The free feature has limited uses, while the paid features have many more uses [23]. There are different types of Linux operating systems. A few can be CentOS, Debian, RHEL, SUSE, and Ubuntu. The version and VMware also matter when dealing with GoogleCloud storage. As for Windows, if the user has 2008 R2 and older versions, it only allows Offline migration. Otherwise, 2012 to the current version allows storing [24] [25]. For Mac, it is different in a way. Instead of using GoogleCloud storage, it uses GoogleCloud Drive as cloud storage. It requires the user to download GoogleCloud Drive to the Mac operating system and use the backup and sync from GoogleCloud. GoogleCloud Drive is integrated with other GoogleCloud services such as cloud storage. The only downfall of GoogleCloud Drive is that it has limited storage. GoogleCloud Drive gives the user a free 15 GB storage [26]. GoogleCloud is available on GoogleCloud Play and the App Store for mobile devices. The key features in this are incident management, alerts, error reporting, cloud storage, and a few more others [27].

Nextcloud

NextCloud protects the user's data with built-in controls from granular permissions to stronger user authentication. Every file is encrypted using AES 256 encryption. NextCloud also avoids fines and meets the most demanding global compliance [28][29]. NextCloud monitors how work happens inside and outside a user's company, with insights and complete audit trails. NextCloud machine also learns how to defend against threats. It uses 2-factor authentication to prevent data leaks. Users can manage their encryption keys using NextCloud's KeySafe, and further reduce risk with NextCloud Shield's classification-based policies and intelligent threat detection [30][31][32].

Supported Files, Clients, Features

NextCloud allows all kinds of files. As shown in figure 3; text, presentation, design documents, videos, and photos can be stored in NextCloud. NextCloud's cloud allows multiple people to collaborate without the risk of version-control issues [33].

Figure 3. NextCloud File Transformation [30]

NextCloud supports Mac, Windows, and mobile devices currently. Unfortunately, Linux is not supported right now, or possibly ever. NextCloud has a free feature but has many downfalls, such as limited uploads per file and limited storage. Mac and Windows can download NextCloud Drive to use NextCloud's cloud storage [34]. All of the files on the desktop can be transferred directly to the cloud and vice versa. Alternatively, if the user does not want to use the downloadable NextCloud drive, users can use the NextCloud app in the browser. As for mobile devices, users can download the NextCloud app to store their files from there. All three clients are compatible with third-party applications to get work done, such as docs, word, excel, and PowerPoint [35].

In addition to the aforementioned GoogleCloud and NextCloud, there are many other reliable and efficient cloud storage such as Dropbox, Amazon Cloud (AWS), and Microsoft OneDrive is available for end-users to use those providers' services and storage directly. However, we need a strong evaluation of these services' security and performance to make a solid conclusion so that the end-users decide which option provides the best services for a specific requirement.

Related Work

As the popularity of cloud storage services has continued to grow, so too have the number of research papers relating to evaluating these services:

- Hu et al. [36] performed the first measurement study on cloud storage services, focusing on Dropbox, Mozy, CrashPlan, and Carbonite. Their goal was to decide the relative download and upload performance of the cloud storage services. They concluded that Dropbox performs best while Mozy performs worst.
- Drago et al. [37] performed extensive research on the architecture of the Dropbox service and conducted experiments based on IP level traces of Dropbox network traffic
- Drago et al. [38] then went on to further compare the system capabilities of Dropbox, GoogleCloud Drive, SkyDrive, Wuala, and Amazon Cloud Drive. They concluded by outlining each service's limitations and advantages.
- Li et al. [39] created a tool called "CloudCmp" to thoroughly compare and evaluate the performances of Amazon AWS, Microsoft Azure, GoogleCloud AppEngine, and Rackspace cloud servers. They concluded that the performance of cloud storage could vary significantly across providers. More Specifically, Amazon S3 was found to be suitable for handling large data objects rather than small data objects
- Jackson et al. [40] studied and revealed that the scalability of Dropbox is limited by their use of Amazon's EC2 hosting service and proposed novel solutions for overcoming these bottlenecks.

Methodology

We rely on active measurements collected by performing experiments using a benchmarking tool used to measure and compare cloud providers' services. PerfKit Benchmarker - Streamlines running benchmark tests on supported cloud providers with unified, simple commands. PerfKit Benchmarker measures the end-to-end time to provision resources and generates reports on standard peak performance metrics, such as latency, throughput, time-to-complete, and IOPS.

Additionally, Cacti is an open-source, web-based network monitoring and bandwidth monitoring tool that polls services at predetermined intervals and graphs the resulting data. We used it to graph and export time-series data of metrics such as CPU load and network bandwidth use. For each TCP flow observed in the network, Cacti exported more than 100 metrics. However, what relevant to our experiment were:

(i) The total number of bytes exchanged with servers.
(ii) The timestamp of the first and the last packet with payload
(iii) The Fully Qualified Domain Name (FQDN) the client resolved via DNS queries before transmitting packets

Since we are focusing on GoogleCloud and NextCloud cloud storage, we needed to isolate traffic to allow these applications to have a controlled environment. Next, we retrieved a list of FQDNs used by cloud servers. Then used this list to filter the records exported by Cacti. For instance, upload.drive.GoogleCloud.com isolates the traffic sent from users to GoogleCloud Drive servers, while the *.storage.nexcloud.com domain is used when connecting to NextCloud.

For our experiment on the low-level behavior of the GoogleCloud and NextCloud application, we reveal complex interactions between CPU time and network traffic to the cloud. In this section, we dive deeper into this relationship by performing carefully controlled microbenchmarks of cloud storage applications. Our goal is to quantify the relationship between the size of file updates and frequency with the amount of traffic generated by GoogleCloud and NextCloud.

Our benchmarks are conducted on two test systems located in the United States in 2021.

- The first is a laptop with an 11th Generation Intel® Core™ i5-1135G7 processor, 2.4 gigahertz, 8 gigabytes of RAM, and a 5400RPM, 512 GB hard drive disk (HDD).

- The second is a desktop with an Intel® Core™ i5-11400 processor, 2.6 gigahertz, 8 GB of RAM, and a 7200 RPM, 1 TB HDD.

Results and Discussion

We conducted tests on different machines with different hard drive rotational speeds because this affects the time it takes for cloud storage software to index files.

- Both machines run Windows 10
- Both machines run Windows GoogleCloud Drive application version 53.0
- Both machines run Windows NextCloud application version 22.2.3
- Both machines are connected to a 100 Mbps Internet connection, which gives both GoogleCloud and NextCloud sufficient resources for synchronizing files to the cloud.

Table 1. Test Systems

Type	Processor	Processor Speed (GHz)	RAM (GB)	HDD RPM	Storage Size (GB)
Laptop	11th Generation Intel® Core™ i5-1135G7	2.4	8	5400RPM	512
Desktop	Intel® Core™ i5-11400	2.6	8	7200 RPM	1000

File Creation and Traffic Analysis

First, we examine the amount of network traffic generated by GoogleCloud and NextCloud then new files are created in the Sync folder.

Table 2. Amount of traffic sent to the index server on the 5400 RPM machine

New File Size	Index Server Traffic	GoogleCloud Traffic	NextCloud Traffic	α	Sync Delay (s)
1 B	29.8 KB	6.5 KB	6.8 KB	38200	4.0
1 KB	31.3 KB	6.8 KB	8.0 KB	40.1	4.0
10 KB	31.8 KB	13.9 KB	14.0 KB	4.63	4.1
100 KB	32.3 KB	118.7 KB	130.8 KB	1.528	4.8
1 MB	35.3 KB	1.2 MB	1.5 MB	1.22	9.2
10 MB	35.1 KB	11.5 MB	12.2 MB	1.149	54.7
100 MB	38.5 KB	112.6 MB	128.7 mb	1.1266	496.3

Table 2 shows the amount of traffic sent to the index server, GoogleCloud, and NextCloud when files of different sizes are placed in the Sync folder on the 5400 RPM machine. We used ZIP for our experiment since ZIP files are a compressed file format. This prevents the GoogleCloud and NextCloud applications from further compressing the data when it uploads to the cloud. The α column in Table 2 shows the ratio of the GoogleCloud and NextCloud traffic to the size of the newly created file. An α close to 1 is ideal since that indicates that NextCloud and GoogleCloud have very little overhead beyond the size of the user's file. However, for small files, α is large because the fixed size of the index server meta-data skews the actual size of the file. On the other hand, for larger files, α can be seen as more reasonable because overhead increases with the file size.

Furthermore, several interesting findings can be deduced about GoogleCloud traffic. First, regardless of the size of the created file, the size of the meta-data sent to the index server remains almost the same. Additionally, the amount of data sent to GoogleCloud closely tracks the size of the created file. This result makes sense since the actual file data (plus some checksumming and HTTP overhead)

Lastly, the final column of Table 2 finds the average time taken to complete the cloud synchronization of GoogleCloud and NextCloud. These experiments report that all cloud synchronizations take at least 4 seconds on average regardless of file size. This minimum time interval is dictated by GoogleCloud and NextCloud's cloud infrastructure and is not due to the hard drive speed, RTT, or Internet connection speed. Moreover, we found that the synchronization time grows proportionately larger for the larger files. The delay is dominated by its time to upload the file to NextCloud.

File update timings

Further research from our experiment reports the timing of file updates can affect both NextCloud and GoogleCloud network use. We conducted an additional experiment where the time interval between 1-byte file appends varied from 100 ms to 10 seconds. The goal of this analysis is to determine the relationship between update timing and network traffic.

The amount of network traffic generated by GoogleCloud and NextCloud during each experiment on the 5400 and 7200 RPM machines. Our findings revealed a clear trend: faster file updates result in less network traffic. This is due to GoogleCloud and NextCloud being able to batch updates that occur very quickly. This batching minimizes the total number of meta-data updates sent to the index server and allows multiple appended bytes in the file to be aggregated into a single binary diff. However, GoogleCloud was determined to perform less batching as the time interval between appends grows.

CPU usage Evaluation

We begin by evaluating the CPU usage characteristics of GoogleCloud and NextCloud storage applications by themselves. As previously mentioned, we will be using our Desktop test setup. For our experiment on this platform, we conducted a benchmark test where 3k random bytes are appended to an initially empty file in GoogleCloud and NextCloud's Sync folder every 300 ms for 1000 seconds. Thus, the final size of the file is 20 MB. During this process, we record the CPU use for both processes.

Our results reveal that the percentage of CPU resources used by GoogleCloud’s and NextCloud’s applications throughout the benchmark is around 57% and 30%. Both applications are multi-threaded. Therefore, it uses all resources of the desktop CPUs. There are two main findings.

- The NextCloud application showed two major jumps in CPU usage that occur around 500 seconds (5 MB file size) and 900 seconds (8 MB). These increases are a result of NextCloud segmenting files into 4 MB chunks
- The average CPU use of GoogleCloud’s application was determined to be is 57% during the experiment, which is relatively high. It also saw periods where it went up to 100%.

Conclusion

This paper evaluated GoogleCloud and NextCloud's using active measurements collected in a controlled environment. Although GoogleCloud dominates the market, we can show that the times are changing due to the increasing usage of cloud storage in both competitors and mobile terminals. Furthermore, our cloud storage usage and performance studies revealed new insights. More specifically, we saw that CPU usage across providers is distinct, with high numbers recorded in GoogleCloud drive application client and low numbers on NextCloud when creating workload and performing benchmark tests. Furthermore, we determined that performance bottlenecks were highlighted due to the integration of faster hard drive speeds. Finally, numerous, small updates to files occur at intervals on the order of several seconds. Under these conditions, cloud storage applications cannot batch updates together, causing the amount of sync traffic to be several orders of magnitude larger than the actual size of the file.

Future Works

While measurements in this paper are tied to evaluating File Creation, Traffic Analysis, and CPU use, we believe they supply valuable information into overall trends and are of interest to understand and track the evolution of personal cloud storage systems and applications. For example, it could be interesting to consider the underlying issue that causes cloud storage applications to create massive amounts of traffic to the cloud: many times, more data than the actual content of the user's files [41].

Conflicts of Interest

The authors declare no conflicts of interest.

References

[1] S. Kariyattin, S. Marru, and M. Pierce, "Evaluating NextCloud as a File Storage for Apache Airavata," in Proceedings of the Practice and Experience on Advanced Research Computing - PEARC '18, 2018.

[2] P. Garraghan, P. Townend, and J. Xu, "An analysis of the server characteristics and resource utilization in Google cloud," in 2013 IEEE International Conference on Cloud Engineering (IC2E), 2013.

[3] M. Filer et al., "Elastic optical networking in the Microsoft cloud [invited]," J. Opt. Commun. Netw., vol. 8, no. 7, p. A45, 2016.

[4] A. Singh and K. Chatterjee, "Cloud security issues and challenges: A survey," J. Netw. Comput. Appl., vol. 79, pp. 88–115, 2017.

[5] A. Mailewa, J. Herath, and S. Herath, "A survey of effective and efficient software testing." The Midwest Instruction and Computing Symposium (MICS), Grand Forks, ND. 2015.

[6] N. Ghosh, S. K. Ghosh, and S. K. Das, "SelCSP: A framework to facilitate selection of cloud service providers," IEEE trans. cloud comput., vol. 3, no. 1, pp. 66–79, 2015.

[7] K. R. Sajay, S. S. Babu, and Y. Vijayalakshmi, "Enhancing the security of cloud data using hybrid encryption algorithm," Journal of Ambient Intelligence and Humanized Computing, Jul. 2019, doi: 10.1007/s12652-019-01403-1.

[8] T. Deriyenko, O. Hartkopp, and D. C. Mattfeld, "Supporting product optimization by customer data analysis," in Operations Research Proceedings, Cham: Springer International Publishing, 2017, pp. 491–496.

[9] R. R. Shetty, A. M. Dissanayaka, S. Mengel, L. Gittner, R. Vadapalli, and H. Khan, "Secure NoSQL based medical data processing and retrieval: The exposome project," in Companion Proceedings of the10th International Conference on Utility and Cloud Computing, 2017.

[10] J. M. Palma-Ruiz and R. Gómez-Martínez, "Google Trends metadata as a revenue indicator for digital marketing activities in Spanish businesses," in Handbook of Research on Digital Marketing Innovations in Social Entrepreneurship and Solidarity Economics, IGI Global, 2019, pp. 281–292.

[11] A. Mailewa Dissanayaka, R. R. Shetty, S. Kothari, S. Mengel, L. Gittner, and R. Vadapalli, "A review of MongoDB and singularity container security in regards to HIPAA regulations," in Companion Proceedings of the10th International Conference on Utility and Cloud Computing, 2017.

[12] T. Halabi and M. Bellaiche, "A broker-based framework for standardization and management of Cloud Security-SLAs," Computers & Security, vol. 75, pp. 59–71, Jun. 2018, doi: 10.1016/j.cose.2018.01.019.

[13] A. Albugmi, M. O. Alassafi, R. Walters, and G. Wills, "Data security in cloud computing," in 2016 Fifth International Conference on Future Communication Technologies (FGCT), 2016.

[14] A. M. Dissanayaka, S. Mengel, L. Gittner, H Khan, "Dynamic & portable vulnerability assessment testbed with Linux containers to ensure the security of MongoDB in Singularity LXCs." In Companion Conference of the Supercomputing-2018 (SC18). 2018.

[15] A. M. Dissanayaka, S. Mengel, L. Gittner, and H. Khan, "Security assurance of MongoDB in singularity LXCs: an elastic and convenient testbed using Linux containers to explore vulnerabilities," Cluster Comput., 2020.

[16] D. S. Raghuwanshi and M. R. Rajagopalan, "MS2: Practical data privacy and security framework for data at rest in cloud," in 2014 World Congress on Computer Applications and Information Systems (WCCAIS), 2014.

[17] A. Tachikawa and A. Kanaoka, "Private cloud storage: Client-side encryption and usable secure utility functions," in HCI for Cybersecurity, Privacy and Trust, Cham: Springer International Publishing, 2020, pp. 652–670.

[18] L. Siwik, AGH-UST University of Science and Technology, Krakow, 30-059, Poland, and L. Mozgowoj, "Server-side encrypting and digital signature platform with biometric authorization," Int. j. comput. netw. inf. secur., vol. 7, no. 4, pp. 1–13, 2015.

[19] R. Mendes, T. Oliveira, V. Cogo, N. Neves, and A. Bessani, "Charon: A secure cloud-of-clouds system for storing and sharing big data," IEEE trans. cloud comput., vol. 9, no. 4, pp. 1349–1361, 2021.

[20] A. M. Dissanayaka, S. Mengel, L. Gittner, and H. Khan, "Vulnerability prioritization, root cause analysis, and mitigation of secure data analytic framework implemented with MongoDB on singularity Linux containers," in Proceedings of the 2020 the 4th International Conference on Compute and Data Analysis, 2020.

[21] D. M. A. B. Mailewa, T. D. B. Weerasinghe, S. P. J. Perera, and C. A. Munasinghe, "Types and Modes Combined Algorithm for Data Encryption and Decryption." proceedings, 13th Peradeniya University Research Sessions, Peradeniya, Sri Lanka (2008): 181-182.

[22] D. Vaithiyanathan and M. Muniraj, "Cloud based Text extraction using Google Cloud Vison for Visually Impaired applications," in 2019 11th International Conference on Advanced Computing (ICoAC), 2019.

[23] R. Aljamal, A. El-Mousa, and F. Jubair, “A user perspective overview of the top infrastructure as a service and high performance computing cloud service providers,” in 2019 IEEE Jordan International Joint Conference on Electrical Engineering and Information Technology (JEEIT), 2019.

[24] S. P. T. Krishnan and J. L. U. Gonzalez, Building your next big thing with Google cloud platform: A guide for developers and enterprise architects. Apress, 2015.

[25] A. Mailewa, and J. Herath. “Operating systems learning environment with VMware.” The Midwest Instruction and Computing Symposium (MICS), Verona, WI. 2014.

[26] B. Sengupta, A. Dixit, and S. Ruj, “Secure Cloud Storage with Data Dynamics Using Secure Network Coding Techniques,” IEEE Transactions on Cloud Computing, pp. 1–1, 2020, doi: 10.1109/tcc.2020.3000342.

[27] I. Malavolta, S. Ruberto, T. Soru and V. Terragni, "Hybrid Mobile Apps in the Google Play Store: An Exploratory Investigation," 2015 2nd ACM International Conference on Mobile Software Engineering and Systems, 2015, pp. 56-59, doi: 10.1109/MobileSoft.2015.15.

[28] R. Hristev and M. Veselinova, “ICT for Cyber Security in Business,” IOP Conference Series: Materials Science and Engineering, vol. 1099, no. 1, p. 012035, Mar. 2021, doi: 10.1088/1757-899x/1099/1/012035.

[29] S. Thapa, and A. Mailewa. “The Role of Intrusion Detection/Prevention Systems in Modern Computer Networks: A Review.” The Midwest Instruction and Computing Symposium (MICS), vol. 53, pp. 1-14. 2020.

[30] A. Fadaeddini, B. Majidi, and M. Eshghi, “Secure decentralized peer-to-peer training of deep neural networks based on distributed ledger technology,” The Journal of Supercomputing, vol. 76, no. 12, pp. 10354–10368, Mar. 2020, doi: 10.1007/s11227-020-03251-9.

[31] M. Akintaro, T. Pare, and A. Mailewa. “Darknet and black market activities against the cybersecurity: a survey.” The Midwest Instruction and Computing Symposium (MICS), North Dakota State University, Fargo, ND. 2019.

[32] H. Mazi, F. N. Arsene, and A. Mailewa. “Build Data Backup with Nextcloud Based Infrastucture as A Service (IAAS) Concept on Budi Darma University.” The Midwest Instruction and Computing Symposium (MICS), Milwaukee School of Engineering and Northwestern Mutual, Milwaukee, Wisconsin. 2020.

[33] S. Kariyattin, S. Marru, and M. Pierce, “Evaluating NextCloud as a File Storage for Apache Airavata,” in Proceedings of the Practice and Experience on Advanced Research Computing - PEARC ’18, 2018.

[34] J.-J. Yang et al., “Emerging information technologies for enhanced healthcare,” Comput. Ind., vol. 69, pp. 3–11, 2015.

[35] X. Cheng, X. Zhou, C. Jiang, and J. Wan, “Towards computation offloading in edge computing: A survey,” in High-Performance Computing Applications in Numerical Simulation and Edge Computing, Singapore: Springer Singapore, 2019, pp. 3–15.

[36] W. Hu, T. Yang, and J. N. Matthews, “The good, the bad and the ugly of consumer cloud storage,” Oper. Syst. Rev., vol. 44, no. 3, pp. 110–115, 2010.

[37] I. Drago, M. Mellia, M. Maurizio, A. Munafo, R. Sperotto, and A. Sadre, “Inside dropbox: understanding personal cloud storage services,” in Proceedings of the 2012 Internet Measurement Conference, 2012, pp. 481–494.

[38] 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, 2013.

[39] A. Li, X. Yang, S. Kandula, and M. Zhang, “CloudCmp: Comparing public cloud providers,” in Proceedings of the 10th annual conference on Internet measurement - IMC ’10, 2010.

[40] K. R. Jackson et al., “Performance analysis of high performance computing applications on the Amazon web services cloud,” in 2010 IEEE Second International Conference on Cloud Computing Technology and Science, 2010.

[41] A. Mailewa, S. Mengel, L. Gittner, and H. Khan, “Mechanisms and techniques to enhance the security of big data analytic framework with mongodb and Linux containers,” SSRN Electron. J., 2021.