Spring 2023

Mitigation of Cache Attacks on Cloud Services

Mudassiruddin Mohammed

Mitigation of Cache Attacks on Cloud Services

Mudassiruddin Mohammed

Thesis

Governors state university

Research paper
Master of Science Degree
With a Major in
Information Technology.

2023

Table of Contents:

Abstract:

Cloud computing is frequently used due to its low cost and flexibility, but it also raises security issues to cloud service providers and customers. Cache attacks are a critical security risk in cloud computing. Cache attacks use weaknesses in cloud servers' cache memory to steal sensitive information, interrupt services, and decrease cloud performance. This study examines the many forms of cache attacks, their possible effects, and known mitigation measures. The study approach includes a review of current methods and their effectiveness in combating cache attack. The report also suggests future research topics for developing more effective and economical methods for preventing cache breaches in cloud computing systems.

Keywords: Cloud Computing, Cache Attacks, Mitigation Techniques, Security, Vulnerabilities.

Introduction:

Cloud computing has grown in popularity due to its low cost, adaptability, and flexibility. Users can utilize cloud computing to gain access to shared computer resources through the Internet. Even so, both cloud service providers and customers have always had worries about the security of cloud services. Cache attacks are a significant security problem in cloud computing [1]. Cache threats take use of weaknesses in cloud servers' cache memory to steal sensitive information, interrupt services, and decrease cloud service performance.

Types of Cache Attacks:

Cache memory is a type of fast memory that saves regularly used data and instructions to increase system performance. Cache memory is used in cloud computing to store frequently requested data and programs in order to minimize the delay of accessing cloud services. Cache memory, on the other hand, is vulnerable to a variety of attacks, including cache side-channel attacks, cache poisoning attacks, and cache-based timing attacks.

Side-channel prefetch attacks are a type of vulnerability that can circumvent security features such as SMAP and kernel ASLR [2]. These attacks make use of the CPU's prefetching function, which attempts to forecast what data the application will want next and puts it into the cache ahead of time. An attacker can utilize this prefetching mechanism to leak sensitive information by carefully managing the program's memory access patterns.

Cache side-channel attacks utilize the information leaked from the cache memory to assume the secret information such as encryption keys, passwords, and sensitive data [3]. Cache poisoning attacks modify the data stored in cache memory in order to redirect the execution flow or modify the behavior of cloud services. Cache-based timing attacks take advantage of differences in cache access.

Mitigation Techniques:

Access control rules: Access control policies can be used to restrict access to cache memory, preventing attackers from gaining unauthorized access. This can be achieved by setting up strict access regulations that limit the amount of data that can be kept in the cache, limiting access to specific cache lines or pages, or employing hardware-level security measures such as memory encryption.

Encryption: Encryption can be used to secure sensitive data stored in the cache by encrypting it before it is placed in memory. This guarantees that the data is not accessible to attackers even if they obtain access to the cache. Software-level encryption, hardware-level encryption, and memory-level encryption are all examples of encryption.

Cache Partitioning as a Mitigation Technique:

Cache partitioning is a strategy that isolates cache resources utilized by various applications or users in order to protect cloud services against cache assaults. This prevents another application or user from accessing or modifying cache resources assigned to one application or user. This stops an attacker from performing cache-based side-channel attacks to harvest sensitive information from shared cache resources.

In the literature, several cache partitioning algorithms have been proposed. One way is to employ hardware-based partitioning, which allocates dedicated cache slices to each program or user who uses the cache. The Intel Cache Allocation Technology (CAT) [4], for example, provides a hardware-based technique for partitioning the common last-level cache across distinct workloads. Another option is to use software-based partitioning, in which the operating system or virtualization layer requires cache isolation. CacheGuard [5], for example, is a software-based cache partitioning solution that separates the cache resources utilized by various virtual machines in a cloud context.

Several studies have been conducted to assess the efficacy of cache partitioning in mitigating cache-based side-channel attacks. In [6] for example, the authors assessed the effectiveness of hardware-based cache partitioning with CAT in combating cache-based side-channel attacks in a cloud context. They demonstrated that CAT can provide high application isolation and considerably prevent sensitive information leakage via cache-based side-channel assaults. Similarly, in [7] the authors assessed the effectiveness of CacheGuard-based software-based cache partitioning in combating cache-based side-channel attacks. They demonstrated CacheGuard's ability to successfully separate cache resources utilized by multiple virtual machines and avoid cache-based side-channel attacks.

Overall, cache partitioning appears to be a good strategy for preventing cache assaults on cloud services. Both hardware-based and software-based partitioning methods have been demonstrated to be successful.

Literature Review:

Several strategies to minimize cache attacks in cloud computing settings have been presented in the literature. There are two types of solutions: hardware-based solutions and software-based solutions.

Cache memory designs that are resistant to cache assaults are the focus of hardware-based solutions. Fully associative caches, for example, might diminish the efficacy of cache side-channel attacks since the attacker cannot anticipate which cache line would be evicted. Another hardware-based approach is dynamic cache partitioning, which guarantees that critical data and programs are kept in different cache memory sectors.

Software-based solutions use software features such as cache partitioning, access pattern randomization, and cache flushing to identify and mitigate cache assaults. To limit the danger of cache-based timing attacks, cache partitioning guarantees that sensitive data and applications are kept in distinct partitions of the cache memory. Access pattern randomization randomizes the order of accessing the cache lines to prevent cache side-channel attacks. To avoid cache poisoning attacks, cache flushing eliminates all cache entries associated with a given process.

Researchers have also developed machine learning-based systems for detecting and preventing cache assaults. Machine learning methods are used in these systems to detect anomalous cache access patterns that signal the presence of cache assaults.

Research Framework:

Problem Statement:

The usage of cloud services for the effective administration of big datasets, resources, and services has grown in popularity. Yet, as the usage of cloud services has grown, so have the risks of cyber-attacks. Cache attacks are one of the most popular forms of cyber-attacks against cloud services. These techniques use cache memory weaknesses to get access to sensitive information. Mitigating cache attacks on cloud services is critical for preserving user data security and privacy.

To analyze the effectiveness of the existing mitigation techniques, the proposed research framework includes the following steps:

- Identify the different types of cache attacks in cloud computing environments.

- Review the existing mitigation techniques for cache attacks in cloud computing environments.
- Evaluate the effectiveness of the existing mitigation techniques.
- Identify the limitations of the existing mitigation techniques.
- Propose new solutions to mitigate cache attacks in cloud computing environments.
- Evaluate the effectiveness of the propose.

Future Research Agenda:

There is a considerable need for continuing study into cache attacks on cloud services. New attacks are continually being developed and the present countermeasures may not be effective against them. Future studies should concentrate on:

Creating more strong defenses: Current precautions for cache operations must be improved and made more effective. Innovative measures that can identify and prevent advanced cache hacks should be developed.

Examining the impact of cache attacks on various cloud service models: The impact of cache attacks on various cloud service types such as SaaS, PaaS, and IaaS should be thoroughly investigated. This will help with the creation of customized countermeasures for each service type.

Investigating the impact of cache assaults on various types of applications: The impact of cache attacks on many types of programs, including online apps, mobile applications, and desktop applications, should be thoroughly investigated. This will help in the development of appropriate countermeasures for each application type.

Investigating the significance of machine learning in the detection of cache attacks: Cache assaults may be detected in real time using machine learning methods. Future research should concentrate on building machine learning algorithms capable of properly detecting cache assaults in real time.

Studying the effect of cache attacks on cloud service performance: Cache attacks have the potential to degrade the performance of cloud services. Future research should concentrate on finding out the effect of cache assaults on cloud service performance and designing methods to reduce this impact.

Conclusion:

Cache attacks represent an important danger to cloud services, with serious consequences for both service providers and their clients. To limit these risks, it is critical to understand the nature of cache assaults and put suitable preventive measures in place. This may include defining cache settings, creating access control lists, and putting security protocols in place. It is also critical to monitor cache usage and respond quickly and efficiently to any potential attacks. Lastly, ongoing research and interaction are essential for improving our understanding of cache assaults and establishing effective mitigation techniques.

References:

https://ieeexplore.ieee.org/abstract/document/8230927

1. Gruss, D., Lipp, M., & Schwarz, M. (2016). Practical cache attacks in cloud computing environments. Proceedings of the 25th USENIX Security Symposium.
2. Gruss, D., Maurice, C., Mangard, S., & Wagner, K. (2018). Prefetch side-channel attacks: Bypassing SMAP and kernel ASLR. Proceedings of the 27th USENIX Security Symposium.
3. Zhang, X., Yang, K., Wei, T., & Yu, Y. (2018). Cache-based side-channel attack and defense in a shared cloud environment. Future Generation Computer Systems, 78, 54-63.
4. Intel Corporation. Intel Cache Allocation Technology (CAT). https://software.intel.com/content/www/us/en/develop/articles/intel-cache-allocation-technology.html
5. J. Song, J. Choi, and D. Kim. CacheGuard: Secure and Efficient Cache Partitioning for Cloud Computing. IEEE Transactions on Dependable and Secure Computing, vol. 15, no. 3, pp. 457-470, 2018.
6. A. Evtyushkin, V. Ponomarev, and D. Kaeli. Catena: A Hardware Platform for Enabling Cache Side-channel Protection. In Proceedings of the IEEE International Symposium on Hardware Oriented Security and Trust (HOST), pp. 123-128, 2016.
7. J. Jiao, J. Zhou, Y. Liang, and Q. Li. CacheGuard: A Secure Cache Partitioning Framework for Cloud Computing. In Proceedings of the IEEE International Conference on Cloud Computing Technology and Science (CloudCom), pp. 279-286, 2015.