International Journal On Engineering Technology and Sciences – IJETS™
ISSN(P): 2349-3968, ISSN (O): 2349-3976
Volume III, Issue IV, April- 2016

DATA PROTECTION IN CLOUD FEDERATIONS

Shruthi P[1], Indulekha K[2]
Dept. of Computer Science
Malabar Institute of Technology,
Anjarakandy.
spanand.nambiar@gmail.com[1], indusreesan@gmail.com[2]

ABSTRACT

Cloud federation is an upcoming concept that addresses the challenges faced in cloud computing. It allows one or more cloud providers to share or rent computing resources to other cloud providers which otherwise maybe idle or underutilized. This is achieved by forming cloud federations that consists of multiple cloud providers interconnected to each other to server a common purpose. Although there are proven cloud federation formation mechanisms, data protection or security within cloud federations needs to be addressed. The nature of privacy and security challenges in clouds requires that cloud providers design data protection mechanisms that work together with their resource management systems. Here, I consider the privacy requirements when outsourcing data and computation within a federation of clouds and propose methods in achieving data protection among cloud providers.

Keywords- *Data protection; federation formation; cloud computing;*

I. INTRODUCTION

Cloud computing can significantly reduce the cost and complexity of owning and operating computers and networks. If an organization uses a cloud provider, it need not spend money on information technology infrastructure, or buy hardware or software licenses. Cloud services can often be customized and flexible to use, and providers can offer advanced services that an individual company might not have the money or expertise to develop. The variability of users' demands increases when it comes to their requests for data-intensive applications. The amount of computing resources that data intensive applications require can dramatically increase, and cloud providers' available resources may not be sufficient enough to cope with such demands. This emerging service management problem in cloud computing necessitates that cloud providers reshape their business structures and seek to improve their dynamic resource scaling capabilities. Federated clouds offer a practical platform for addressing this service management problem. A cloud provider can dynamically scale- up its resource capabilities by forming a cloud federation with other cloud providers. On the other hand, other cloud providers that have unused capacities can make profit by participating in a federation. Users' requests can be satisfied by federating resources belonging to several cloud providers. Forming cloud federations helps achieve greater scalability and performance. If a cloud provider does not have enough resources to provide all the requested resources to the customer, it will reject the requests which leads not only to profit loses, but also to reputation losses. However, forming cloud federations presents a host of new challenges resulting from the current lack of efficient cloud federation formation mechanisms. One of the major challenges is that a federation formation mechanism needs to address the data protection concerns that arise from outsourcing computations. A key aspect of cloud federations, however, is that their infrastructure is shared among cloud providers and it is off the premises of a single cloud provider. Therefore, there exists a significant threat to data privacy and security associated with remote storage

and processing of data.CloudSecurity is a leading concern in the cloud environment. Fear of data breaches or outages continue to keep executives up at night.

According to Gartner Inc., the market for cloud-security services is expected to reach nearly $4 billion in revenue in 2016, up from $2.1 billion last year. As more businesses move to the cloud, it's essential that companies work with partners that understand best practices of cloud security and provide transparency when it comes to their solutions.

Cloud federation inherits this problem from existing cloud computing setups with add on issues while sharing resources among different cloud providers within a cloud federation. Given the importance of data privacy, customers' fear of sensitive data leakage should be the primary concern when providing cloud services. One of the benefits of clouds is the possibility of dynamically enhancing their resource capabilities by forming federations with other cloud providers in order to rapidly scale-up when required.

In this paper, I consider the privacy requirements when outsourcing data and computation within a federation of clouds. We design a data protection framework for cloud federations that minimizes the cost of outsourcing. The benefits of employing our framework on a cloud federation are threefold. First, it helps align the cloud services with the users' concerns regarding their data protection. Second, it reduces the cost of upgrading the system software to support future data protection needs. Third, it avoids future costs and penalties resulting from leakage of sensitive data.

The federation creates a pool of virtualized resources which are offered to users as different types of VM instances. When the demand exceeds the capacity of available resources, the cloud provider can scale up by forming a cloud federation with other cloud providers, and flexibly mapping and moving VMs using VM migration technologies to the other cloud providers. In this study, we focus on data protection when the computation is outsourced to other cloud providers within the federation via such VM migrations, and propose a framework to minimize the total cost of outsourcing while considering two restrictions, as follows.

First, there are several limitations for a cloud provider in assigning VMs to other cloud providers. Such limitations could be due to trust and reliability issues, or due to the geographical location of the other cloud provider. The location of a cloud provider can arise privacy and security issues since transferring a VM over to specific regions (e.g., crossing national boundaries), raises legal concerns. Therefore, we consider such restrictions, called trust restrictions, during the VM assignment and migration. The trust restrictions specify the VMs that should not be assigned to specific cloud providers. Second, if several VMs are co-located on the same cloud provider they can reveal sensitive information. However, the user or the cloud customer requires through agreements that such information be accessible only to the main cloud provider and not to others. The cloud provider is accountable for protecting such information, and revealing it is against the agreement. Therefore, we consider such restrictions on co-locating VMs, called disclosure restrictions. These disclosure restrictions, specify which VMs can never be co-located on the same cloud provider when outsourcing the cloud services. This represents another level of data protection (after the encryption level) which subsequently reduces the need to encrypt all data.

2. SYSTEM MODEL

I first describe the system model considering that a

cloud provider A wants to outsource its workload consisting of a pool of N VMs to other cloud providers. I consider a federation of cloud providers F = {P0, P1, P2, PM} that are available to provide services. Each cloud provider Pj 2 F has restricted computing capacity, denoted by Rj, available to provide to other cloud providers. Each provider Pj incurs cost when providing resources. For a cloud provider Pj , we denote by cj , the cost associated with each VM instance executed on Pj , and by mj , the cost associated with VM migration from P0 to Pj .

Cloud provider P0 does not have enough resources to fulfill the requested VMs, and needs to outsource some of the requested VMs to other cloud providers within the federation in order to execute the jobs and more importantly, to minimize its cost while satisfying the data protection restrictions imposed by the users. Therefore, the outsourcing decisions have to be made considering both data protection restrictions and cost minimization. As we mentioned in the introduction section, there are two data protection restrictions that have to be considered when outsourcing VMs within a federation

of cloud providers: (i) the trust restrictions, specifying that some VMs cannot be outsourced to specific cloud providers; and, (ii) the disclosure restrictions, specifying that some VMs cannot be outsourced to the same cloud provider.

In this section, I introduce my proposed algorithm that solves the DPCF problem. I first describe my proposed partitioning algorithm, called VMPA, which uses the conflict graph of disclosure restrictions to determine a Partitioning of the VMs. I describe my proposed algorithm, called DPCFA that solves the DPCF problem.

DPCFA algorithm uses the VM partitioning algorithm (VMPA) in order to minimize the cost of federation formation while satisfying the trust and disclosure restrictions

2.1 VM Partitioning Algorithm

In this section, I propose the VM Partitioning Algorithm (VMPA). VMPA partitions the VMs in a way that conflicting VMs are not assigned to the same cloud provider. VMPA uses a conflict graph G′(V ′,E′) which is a subgraph of the conflict graph G, G′ _ G, as an input. The algorithm builds a max-heap V to order the VMs in V ′ based on their number of conflicting VMs, denoted by fi. The max-heap has two main functions associated with it: (i) enqueue (), that inserts a VM along with its priority into the heap; and (ii) extractMax(), that extracts the VMwith the highest priority. A VM with the highest priority, i.e., with the most conflicts is always at the top of the heap. The algorithm also creates a subset S0 of VMs that do not have any conflicts (i.e.,fi = 0). Considering S0 is critical for the algorithm in order to minimize cost. We will discuss this in more details in the next subsection. The algorithm extracts the VM with the highest priority (i.e., with the highest number of conflicts),and assigns it to S1, where K = 1 tracks the number of current subsets without considering S0. The assignment of the rest of the VMs in V based on their priorities is as follows: For each of the current subsets, it checks if there exists any conflicting VMs. The algorithm assigns the VMto the first subset that does not have any conflicting VMs.If there is no subset without conflict, it exists, and creates new subset for that VM. The result of the partitioning of the VMs is S0 along with other subsets Si. Then, the algorithm sorts the partitioned VMs based on the number of VMs in each subset such that S1 is the largest subset. Finally, the algorithm returns the sets Si, which represent the partitioning of the set of VMs. The VMs that are part

of a setSi do not conflict with each other, and can be assigned to thesame cloud provider.

2.1.1 VM Partitioning Algorithm

1: Input: G′(V ′,E′)
2: Create an empty max heap V
3: for all i2 V ′do
4: fi the number of connected VMs to VMi
5: if fi>1 then
6: V.enqueue(i, fi)
7: else
8: S0 S0 [{i}]
9: (i, fi) = V.extractMax()
10: S1 = {i}
11: K = 1
12: while V is not empty do
13: (i, fi) V.extractMax()
14: flag FALSE
15: for all k = 1, . . . ,Kdo
16: for all j,<i, j >2 E′, j 2 Skdo
17: flagTRUE
18: break
19: if !flagthen
20: SkSk[{i}
21: break
22: if flag then
23: K = K + 1
24: SK = {i}
25: Sort S1, . . . , SK
based on descending orderof their size
26: Output: S0, S1, . . . , SK

3. CONCLUSION

The benefits from using cloud services should not comeat the cost of compromising the privacy and security ofusers' data. Data protection in terms of legal complianceand user trust are major issues in clouds, and they shouldbe a top priority when designing cloud systems. On theother hand, the ever-growing demand for cloud servicesalong with the demand dynamics require cloud providers top scale-up when needed. A practical platform to cope withsuch demand is the cloud federation. In this paper, I proposed a data protection framework for cloud federationsthat minimizes the cost of outsourcing the computationunder data protection constraints. I proposed to representthe data protection restrictions as conflict graphs, and modelthe data protection problem as an integer program. In theabsence of computationally tractable optimal algorithms forsolving this problem, I designed an algorithmin co-operating a novel VM placement strategy in order tofind close to optimal solutions.

4. REFERENCES

[1] MahyarMovahedNejad, Lena Mashayekhy and Daniel Grosu, Cloudfederations in the sky: Formation game and mechanism, Cloud Computing,IEEE Transactions on 3 (2015), no. 1, 14–27.

[2] S. Pearson, "Taking account of privacy when designing cloudComputing services," in Proc. of the IEEE ICSE Workshop onSoftware Engineering Challenges of Cloud Computing, 2009,pp. 44–52.

[3] J. Zheng, T. Ng, K. Sripanidkulchai, and Z. Liu, "Pacer: Aprogress management system for livevirtualmachinemigrationin cloud computing," IEEE Transactions on Network andService Management, vol. 10, no. 4, pp. 369–382, 2013.

[4] Antonio Celesti, Francesco Tusa, Massimo Villari, and Antonio Puliafito,How to enhance cloud architectures to enable cross-federation, Cloud Computing(CLOUD), 2010 IEEE 3rd International Conference on, IEEE, 2010,pp. 337–345.

[5] L. Mashayekhy and D. Grosu, "A coalitional game-basedmechanism for forming cloud federations," in Proc. of the5th IEEE Intl. Conf. on Utility and Cloud Computing, 2012,pp. 223–227.

[6] Hongxing Li, Chuan Wu, Zongpeng Li, and Francis Lau, Profitmaximizingvirtual machine trading in a federation of selfish clouds, INFOCOM,2013 Proceedings IEEE, IEEE, 2013, pp. 25–29.