IJARCCE

ISSN (Online) 2278-1021
ISSN (Print) 2319 5940

International Journal of Advanced Research in Computer and Communication Engineering
Vol. 5, Issue 4, April 2016.

Reduction of Migration Cost in Cloud Applications

Korimani Radha[1], Vasudev S Shahapur[2]

M.Tech Student, Computer Science Department, Alva's Institute of Engineering & Technology, Shobhavana campus, Mijar, Moodbidiri, DK[1]

Professor, Computer Science Department, Alva's Institute of Engineering & Technology, Shobhavana campus, Mijar, Moodbidiri, DK[2]

Abstract: Enabled by virtualization technologies, various multi-tier applications are hosted by virtual machines (VMs) in cloud data centers. Live migration of multi-tier applications across geographically distributed data centers is important for load management, power saving, routine server maintenance and quality-of-service. Different from a single-VM migration, VMs in a multi-tier application are closely correlated, which results in a correlated VM migrations problem. In this paper, we explore performance of multi-tier applications during the virtual machine migration in cloud as well as issues and solutions during VMs migration.

Index Terms: Cloud, live migration, multi-tier application, virtual machine.

INTRODUCTION

INTERNET applications have been prosperous in the era of cloud computing, which are usually hosted in virtual machines in geographically distributed data centers. Live migration of Internet applications across data centers is important for different scenarios including load management, power saving, routine server maintenance and quality- of-service [1], [2]. Additionally, Internet applications tend to have dynamically varying workloads that contain long-term variations such as time-of-day effects in different regions. It is desirable to move the interactive/ web application to the data center that has better network performance to users for lower response time [2]. Also, workloads can be migrated across different data centers to exploit time-varying electricity pricing. Applications are hosted by virtual machines. Virtual machine is nothing but a program which runs on operating system.

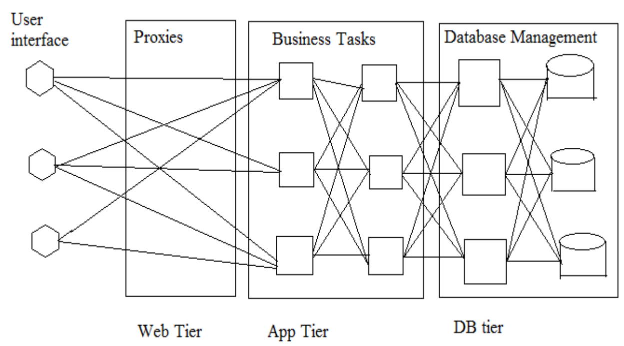

Fig 1.Three-tier Architecture of web application

Typical Internet applications employ a multi-tier architecture, with each tier providing certain functionality. Specific to multi-tier applications, we need to transfer several tightly-coupled VMs in multi-tiers, instead of a single VM. Previous studies have demonstrated the potential performance penalty of multi-tier applications during migration.

A typical multi-tier web application consists of three tiers: presentation layer (web tier), business logic layer (App tier) and data access layer (DB tier) which is shown in fig 1. Different layers usually run on different VMs and have different memory access patterns.

CORRELATED VM MIGRATION PROBLEM

Live migration of VMs has been an effective approach to manage workloads in a non-disruptive manner. As shown in Fig. 2, VM live migration i.e. memory pre-copying is conducted in several iterative rounds. The VM's physical memory is first transferred from Datacenter A to Datacenter B, while the source VM continues running in Datacenter A. VMs are correlated because only when all VMs of the multi-tiers are migrated to another data center, they can completely and efficiently serve requests in that data center. We call this problem correlated VM migrations. Correlated VM migrations can cause significant performance penalty to multi-tier applications. Consider the following scenario: if the middle tier is first migrated, then the other two tiers must redirect the communication and data access traffic to another data center and wait for the processing results to be sent back. Moreover, because the multi-tier application and migration processes share the same link for data transferring, given the data-intensive nature of multi-tier applications and limited network bandwidth between two data centers, network bandwidth contention may cause significant performance degradation for both applications and VM migrations.

IJARCCE

ISSN (Online) 2278-1021
ISSN (Print) 2319 5940

International Journal of Advanced Research in Computer and Communication Engineering
Vol. 5, Issue 4, April 2016.

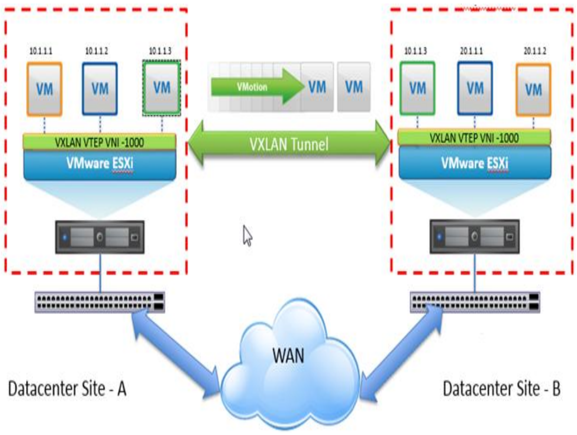

Fig 2. VM live migration

There are a number of factors affecting the migration cost in terms of migration downtime, migration completion time and total network traffic. The major factors include the size of VM memory, memory dirtying rate, network transmission rate and configuration of migration algorithm (e.g., conditions for starting the stop-and-copy phase). Among these factors, the size of VM memory and the memory dirtying rate are mostly determined by the VM and workloads.

PERFORMANCE PENALTY OF APPLICATION

As illustrated in Fig. 3, a multi-tier web application is migrated from DC 1 to DC 2. Each tier is running on multiple VMs, and thus the VMs across different tiers are correlated with data dependency. Network bandwidth is a critical resource across distributed data centers. It is usually much smaller than the network bandwidth within a data center.

Previous studies (e.g., [1],[2],[3]) assumed that the network bandwidth between two data center was 465 Mbps. Without loss of generality, we assume that the peak network bandwidth between DC 1 and DC 2 reserved for the migration processes is B. As discussed in Introduction, the application traffic and migration traffic share the same links between two data centers. The bandwidth contention between them may result in significant performance degradation in both application and VM migration.

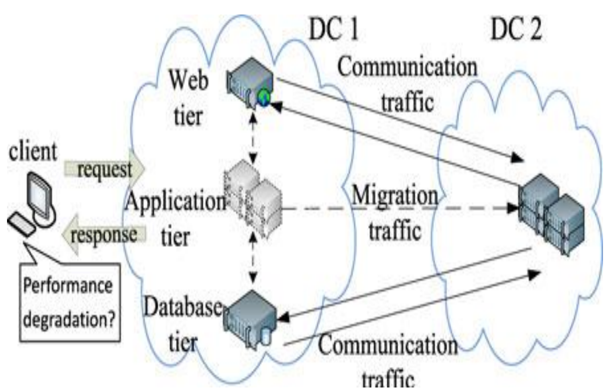

Fig 3. Performance penalty due to live migration of a multi-tier web application

PROPOSED MODEL

To overcome above problems, we should avoid splitting multiple tiers between VMs so that the performance penalty of VM migrations on the multi-tier application is minimized. Meanwhile, we should diminish the VM migration cost in terms of migration completion time, network traffic and migration downtime. Here synchronization protocol is developed to orchestrate all VMs to proceed the stop-and-copy phase at the same time.
As shown in Fig. 4, each VM migration may reach its stop-and-copy phase at different points of time (called pseudo-synchropoint). The pseudo-synchropoint depends on the termination conditions of pre-copying algorithm. In our synchronization protocol, we postpone the stop-and-copy phase until all VMs reach the stop-and-copy phase (called synchropoint). However, all VMs are still running during the synchronization, and the dirtied memory pages still need to be transmitted to the destinations. We call this phase "wait-and-copy". The bandwidth consumed in this phase is determined by the memory dirtying rates of the VMs. Algorithm 1 shows the pseudo code of the synchronization protocol for correlated VM migrations. The synchronization protocol relies on an arbitrator for control. The arbitrator implements a message-passing mechanism for controlling the VM migrations. When a VM reaches the pseudo-synchropoint, it should immediately send a message "reach_pseudo-synchropoint" to the arbitrator, and then proceed the "wait-and-copy" phase until it receives the "start_stop-and-copy" message from the arbitrator. The arbitrator uses a variable p to record the number of VMs that have reached the pseudo-synchropoint. Once all VMs have reached the synchropoint, the arbitrator broadcasts a message "start_stop-and-copy" to all VMs. To handle the potential migration failures, we adopt a simple approach for fault tolerance. We view the coordinated migration processes as a transaction in a batch. In case of failures, all correlated VM should resume at their original host, aborting the migration. More advanced fail-tolerant VM migration techniques will be studied in our future work.

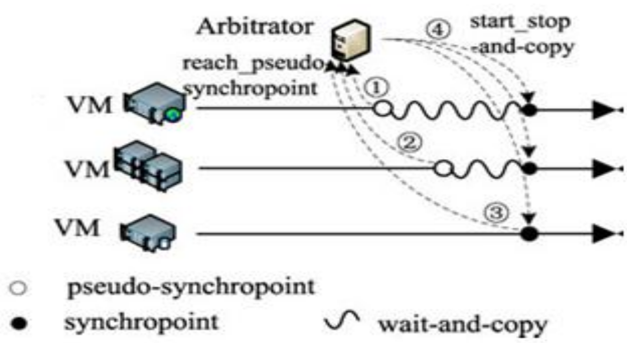

Fig. 4 Synchronization of VMs

RELATED WORK

A number of studies had investigated the live migration performance of multi-tier applications for both intra- and interdatacenter scenarios. Voorsluys et al. [3] evaluated

DOI 10.17148/IJARCCE.2016.54202

the performance degradation of 2-tier web 2.0 applications running within VMs. The similar evaluations were conducted in [2]. Their experimental results all showed significant performance degradation in terms of request response time, even when the multi-tier applications are migrated in the same data center. While those studies had given preliminary results on the performance problem of correlated VM migrations in a multi-tier application, they had not formulated or solved the problem.

Recently, Zheng et al. proposed Pacer [2], a progress management system for VM migration in the cloud. Pacer controls VM migration completion time based on analytic models of progress prediction and online adaptation.

Rajawasim [1], His paper reviews state-of-the art live and non-live VM migration schemes. Through an extensive literature review, a detailed thematic taxonomy is proposed for the categorization of VM migration schemes. Critical aspects and related features of current VM migration schemes are inspected through detailed qualitative investigation. He extracted significant parameters from existing literature to discuss the commonalities and variances among VM migration schemes. Finally, open research issues and challenges with VM migration that require further consideration to develop optimal VM migration schemes in Cloud Data Centers are briefly addressed. E.

Gustafsson[4], present a technique to reduce the total time required to migrate a running VM from one host to another while keeping the downtime to a minimum. Based on the observation that modern operating systems use the better part of the physical memory to cache data from secondary storage, his technique tracks the VM's I/O operations to the network-attached storage device and maintains an updated mapping of memory pages that currently reside in identical form on the storage device. A.

Mashtizadeh[5], describe the evolution of live storage migration in VMware ESX through three separate architectures, and explore the performance, complexity and functionality trade-offs of each. These works all focused on migrating a single VM in LANs or over WANs. None has considered the problem of correlated VM migration in multi-tiered applications across distributed data centers.

Deshpande et al[6], investigated live gang migration of VMs in LAN environments. They proposed page and sub-page level memory de-duplication among co-located VMs and compression strategies to optimize memory migration of multiple VMs. As the network bandwidth is sufficiently high in LANs, their work did not consider the correlated VM migration problem.

CONCLUSION

The paper discussed the notion of Cloud Data Centers, VM migration process. It analyzed current VM migration schemes based on a thematic taxonomy and highlighted the commonalities and variances among VM migration schemes based on the selected performance parameters. It has also discussed the issues and challenges in exist during VM migration and solution to the challenges.

REFERENCES

[1] RajaWasim Ahmad , Abdullah Gani, Siti Hafizah Ab, Hamid , Muhammad Shiraz , Feng Xia, Sajjad A and Madani, "Virtual machine migration in cloud data centers: a review, taxonomy, and open research issues SURVEY PAPERS," in proc. Springer Science+Business Media , jun. 2015, pp. 2473–2515.
[2] J. Zheng, T. S. Eugene Ng, K. Sripanidkulchai, and Z. Liu, "Pacer: A progress management system for live virtual machine migration in cloud computing," IEEE Trans. Netw. Serv. Manage., vol. 10, no. 4, pp. 369–382, Dec. 2013.
[3] W. Voorsluys, J. Broberg, S. Venugopal, and R. Buyya, "Cost of virtual machine live migration in clouds: A performance evaluation," in Proc. Int. Conf. Cloud Comput., Dec. 2009, pp. 254–265.
[4] C. Jo, E. Gustafsson, J. Son, and B. Egger, "Efficient live migration of virtual machines using shared storage," in Proc. ACM Conf. Virtual Execution Environ., Mar. 2013, pp. 41–50.
[5] A. Mashtizadeh, E. Celebi, T. Garfinkel, and M. Cai, "The design and evolution of live storage migration in VMware ESX," in Proc. USENIX Ann. Tech. Conf., Jun. 2011, pp. 1–14.
[6] U. Deshpande, X. Wang, and K. Gopalan, "Live gang migration of virtual machines," in Proc. Int. Symp. High Perform. Distrib. Comput., Jun. 2011, pp. 135–146.