ISSN: 2278 – 1323
International Journal of Advanced Research in Computer Engineering & Technology (IJARCET)
Volume 2, Issue 1, January 2013

Survey on Live VM Migration Techniques

Anju Mohan, Shine S

***Abstract*— Virtual Machines(VMs) refers to the software implementation of a computer that runs its own OS and applications as if it was a physical machine(PM). Live migration of VMs allows a server administrator to move a running virtual machine or application among different physical machines without disconnecting the client or application. Total migration time and downtime are two key performance metrics that the clients of a VM service care about the most, because they are concerned about service degradation and the duration that the service is completely unavailable. When a VM is migrating, it is important that this transfer occurs in a manner that balances the requirements of minimizing both the downtime and the total migration time.**

***Index Terms*— Virtual machine, live migration, post-copy, pre-copy, stop and copy**

I. Introduction

Cloud computing can be considered as a new computing paradigm with many exciting features like greater flexibility and lower expenses for installation and maintenance. Cloud computing is defined by NIST[1] as a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources that can be rapidly provisioned and released with minimal management effort or service provider interaction. Because of this, cloud computing has been receiving a good deal of attention lately.

Compared with general network service, cloud computing is easy to extend, and has a simple management style. Cloud is not only a collection the computer resources, but also provides a management mechanism and can provide services for millions of users simultaneously. The concept of virtualization is now applied to cloud data centres. When the storage and computing capacity of the server cluster are surplus, we need not purchase servers, all we need to do is to add a virtual machine running on the server. If the cluster is large enough, the request of adding server will have marginal effect, and then we can save the money that should be used for purchasing new servers.

VMs refer to one instance of an operating system along with one or more applications running in an isolated partition within the computer. There will be multiple virtual machines running on top of a single physical machine. When one physical host gets overloaded, it may be required to dynamically transfer certain amount of its load to another machine with minimal interruption to the users. This process of moving a virtual machine from one physical host to another is termed as migration. In the past, to move a VM between two physical hosts, it was necessary to shut down the VM, allocate the needed resources to the new physical host, move the VM files and start the VM in the new host. Live migration makes possible for VMs to be migrated without considerable downtime.

The transfer of a VM actually refers to the transfer of its state. This includes its memory, internal state of the devices and that of the virtual CPU. Among these, the most time-consuming one is the memory transfer. Two parameters to be considered while performing the live VM-migration are downtime and migration time. Migration time refers to the total amount of time required to transfer a virtual machine at source node to destination node without affecting its availability. Down time refers to the time during which the service of the VM is not available.

II. Live VM Migration

Virtualization technology is the one that hides the details of physical hardware and provides virtualized resources for high-level applications. An essential characteristic of a virtual machine is that the software running inside it is limited to the resources and abstractions provided by the VM. The software layer that provides the virtualization is called a Virtual Machine Monitor (VMM) or hypervisor. It virtualizes all of the resources of a physical machine, thereby defining and supporting the execution of multiple virtual machines. Virtualization can provide significant benefits in cloud computing by enabling virtual machine migration to balance load across the data centres.

A. Migration Goals

Migration is mainly done for dynamic resource management[2]. Its main goals are as follows:

Load balancing: This reduces the inequality of resource usage levels across all the PMs in the cluster. This prevents some machines from getting overloaded in the presence of lightly loaded machines with sufficient spare capacity. Live migration can be used to balance the system. The overall system load can be balanced by migrating VMs from overloaded PMs to under-loaded PMs.

Server Consolidation: In order to reduce server sprawl in data centres, server consolidation algorithms are required. These algorithms are VM packing heuristics which try to pack as many VMs as possible on a PM so that resource usage is improved and unused or under-utilized machines can be turned off. Consolidation will result in reduced power

Manuscript received Jan , 2013.
***Anju Mohan**, Department of Computer Science and Engineering, College of Engineering, Trivandrum.*
***Shine S,** Department of Computer Science and Engineering, College of Engineering, Trivandrum.*

consumption and thus reducing overall operational costs for data centre administrators. Live migration of VMs could achieve this. Based on the load conditions, under-utilized machines having resource usage below a threshold and overloaded machines having resource usage above a certain threshold are identified, and migrations are triggered to tightly pack VMs to increase overall resource usage on all PMs and free up resources/PMs if possible.

Hotspot & Coldspot Mitigation: The detection of hotspots and coldspots are always based on thresholds which are set by the data center owner or based on the Service Level Agreements specified by the clients. Usually, a higher resource usage value close to maximum is set as the upper threshold and a very low resource usage value is set as the lower threshold. PMs having resource usage values beyond the upper threshold are said to have formed hotspots, and whose usage values below the lower threshold are said to have formed coldspots. The former implies over-utilization and the latter implies under-utilization, applicable across any resource dimension. These conditions are inherently taken care of in the above mentioned consolidation and load balancing algorithms.

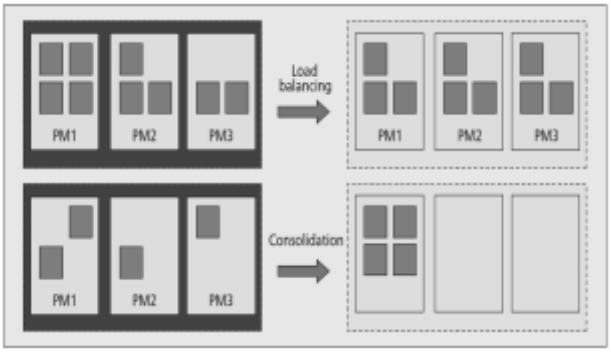

Fig. 1 Load balancing and consolidation scenarios

Fig. 1 shows two conditions: load balancing and consolidation of VMs based on migration. In the first case, either the goal is to distribute load evenly across PMs, or a VM needs more resources and hence is migrated to another PM. With consolidation, machines are migrated to fewer PMs to reduce server sprawl.

B. Live VM Migration Techniques

Initially, pure stop-and-copy approach[3] was used for VM migration. This involves halting the original VM, copying all pages to the destination, and then starting the new VM. This has advantages in terms of simplicity but the service downtime is proportional to the amount of physical memory allocated to the VM. This can lead to an unacceptable outage if the VM is running a live service.

Another option is post-copy approach[4],[5] in which a short stop-and-copy phase transfers essential kernel data structures to the destination. The destination VM is then started, and other pages are transferred across the network on demand. This results in a much shorter downtime, but produces a much longer total migration time; and in practice, performance after migration is likely to be unacceptably degraded until a considerable set of pages have been faulted across. Until this time the VM will fault on a high proportion of its memory accesses, each of which initiates a synchronous transfer across the network.

The pre-copy[6]-[8] migration iteratively copies the memory pages from the source machine to the destination host, without ever stopping the execution of the VM being migrated. The iterative nature of the algorithm is due to the dirty pages, i.e. memory pages that have been modified in the source host since last page transfer must be sent again to the destination host. If the rate of updation of pages is very high, migration time will rise to a very high value. But the advantage of this approach is that all updations are available at the destination host. It can be activated any time. Every VM will have some (hopefully small) set of pages that it updates very frequently and which are therefore poor candidates for pre-copy migration. The pre-copy approach is shown in fig 2.

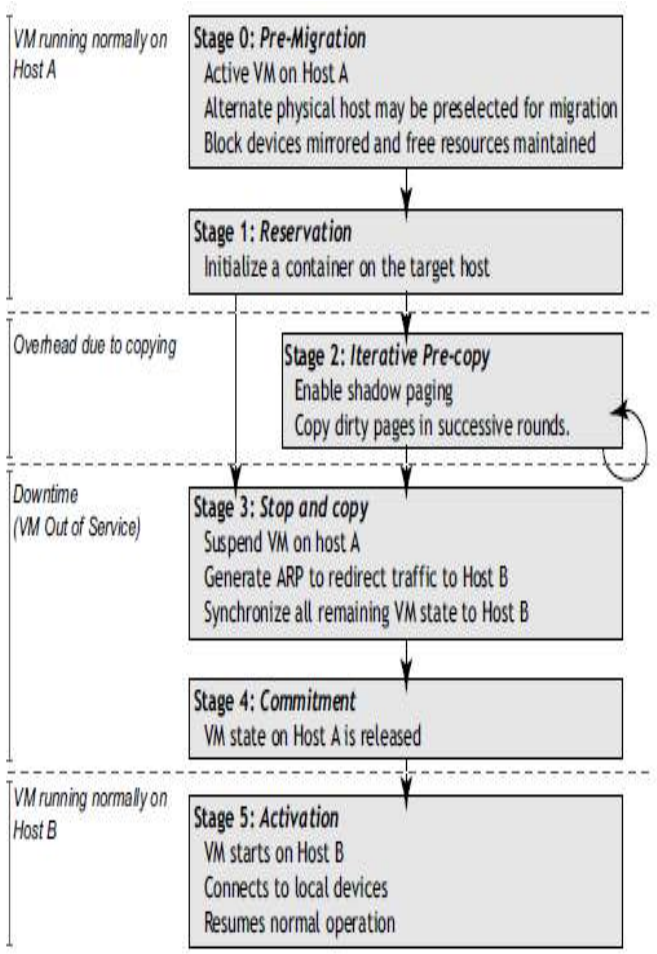

Fig. 2 Pre-copy algorithm

In another method, which is a modification of pre-copy approach[9], a framework is constructed that includes pre-processing phase in traditional pre-copy based live migration so that the amount of transferred data is reduced. For preprocessing, they proposed the working set prediction algorithm. Applying the proposed algorithm which combined LRU(Least Recently Used) cache with splay tree algorithm, the system can reduce the amount transferred memory page.

In an approach to optimize live virtual machine migration[10] based on pre-copy algorithm, memory compression is used. Virtual machine migration performance

is greatly improved by cutting down the amount of transferred data. A compression algorithm has been designed for live migration. In the source node, data being transferred in each round are first compressed by the algorithm. When arriving on the target, compressed data are then decompressed. Compression time is a performance bottleneck of additional overhead introduced by compression operations. If the compression overhead outweighs the advantage of memory compression, live VM migration would not get any benefit from it.

III. CONCLUSION

With the increase in the popularity of cloud computing systems, virtual machine migrations across data centers and diverse resource pools will be greatly beneficial to data center administrators. Live virtual machine migration is an indispensable tool for dynamic resource management in modern day data centers. In this paper, we present a survey of the current research efforts on live virtual machine migration. But this process takes considerable amount of migration time and downtime. A new model for live migration of VMs should be designed with reduced overheads of downtime and migration time.

REFERENCES

[1] Cloud Computing for Dummies, Wiley Publishing, Inc.
[2] Mayank Mishra, Anwesha Das, Purushottam Kulkarni, and Anirudha Sahoo, "Dynamic resource management using virtual machine migrations", IEEE Communications Magazine, vol. 50, pp. 34-40, September 2012.
[3] C. P. Sapuntzakis, R. Chandra, B. Pfaff, J. Chow, M. S. Lam, and M.Rosenblum, "Optimizing the migration of virtual computers", In Proc. of the 5th Symposium on Operating Systems Design and Implementation (OSDI-02), December 2002.
[4] Christopher Clark, Keir Fraser, Steven Hand, Jacob Gorm Hansen, Eric Jul, Christian Limpach, Ian Pratt, and Andrew Warfield. "Live migration of virtual machines", NSDI'05: Proceedings of the 2nd conference on Symposium on Networked Systems Design & Implementation, pages 273-286, 2005.
[5] Diego Perez-Botero, "A Brief Tutorial on Live Virtual Machine Migration From a Security Perspective", University of Princeton, USA
[6] B. Hu, Z. Lei, Y. Lei, D. Xu, and J. Li, "A time-series based precopy approach for live migration of virtual machines", in 2011 IEEE 17th International Conference on Parallel and Distributed Systems (ICPADS), pp. 947-952, Dec. 2011.
[7] Chen Xianqin, Gao Xiaopeng, Wan Han, Wang Sumei, Long Xiang,"Application-Transparent Live Migration for Virtual Machine on Network Security Enhanced Hypervisor", China communications, 2011.
[8] Sherif Akoush, Ripduman Sohan, Andrew Rice, Andrew W.Moore, Andy Hopper, "Predicting the Performance of Virtual Machine Migration", The 18th Annual meeting of the IEEE International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication systems (MASCOTS'10), Miami, USA Aug 2010.
[9] Ei Phyu Zaw and Ni Lar Thein, "Improved Live VM Migration using LRU and Splay Tree Algorithm", International Journal of Computer Science and Telecommunications [Volume 3, Issue 3, March 2012].
[10] Hai Jin, Li Deng, Song Wu, Xuanhua Shi, Xiaodong Pan, Live Virtual Machine Migration with Adaptive Memory Compression, 2009 IEEE.