Survey on VM Migration and Consolidation in Cloud Computing

Deepak Mewada[1], Prof. Bhupchandra Kumar[2]

[1]*M. Tech. Scholar,* [2]*Research Guide*

Department of CSE M.Tech, IES Institute of Technology & Science Management, Bhopal

Abstract - Cloud computing provides processing of huge data stored on centralized data storage using shared computing resources on the internet. It works on client-server model. The data centers consume huge amount of power. Even the idle server consumes 70% of the energy. Moreover, Virtual Machine migration is not properly balanced because, moving of virtual machine from one physical machine to another machine doesn't consider already existing load on the new machine. So, resource is not utilized properly and may lead to resource congestion and resource contention. To tackle these drawbacks in this paper a survey is done on intelligent VM Migration Schemes. It manages load across the machines such that resources are properly utilized. By properly utilizing resources, the energy consumption reduces. In idle state, the servers consume energy which is very less than the existing systems.

Keywords: Virtual Machine, Virtual Machine Migration, Load Balancing, Pre-copy Live Migration, Post-copy Live Migration.

I. INTRODUCTION

Cloud computing is used when the task is needed to be computed in distributed manner on the resources available on a large node of computers. It includes on demand computation and storage of Platforms and software applications over the internet by the various users [1] as an emerging distributed computing technology, Cloud computing has attracted a rapidly increasing number of service providers, users, and business opportunities. A cloud is a cluster of virtualized machines to provide scalability which are dynamic in nature and act as computing resources and provides service level agreement between clients and service providers [2] and on-demand resource provision. It provides an approach to deliver the infrastructure, platforms and software as services accessible to consumers in a pay-as you-go manner [3].such usual commercial service providers include Amazon, Google, and Microsoft There are lots of virtual machines interconnected with each other to form cluster on which lots of physical nodes are running to compute huge datasets on a cloud computing environment. To offer virtualization services which seem to be from provider, virtual machine migration is used. VMM is responsible for the movement of VM from one hardware environment to another. It is a part of hardware virtualization system. A virtual machine is an operating system or application environment. It is installed on software which resembles dedicated hardware. The same thing is experienced by end user when compared with virtual machine to dedicated software. VMs are joined to make a cluster using logical and physical hardware pieces to build hardware virtualization. Logical pieces are also called as programmed pieces. In a virtualization setup, a central supervisor tool is responsible to allocates resources like CPU, Memory, network I/O etc. to virtual machines. It was impractical to move virtual pieces between physical and other pieces in the earlier days using interface of VMM. It was done by system

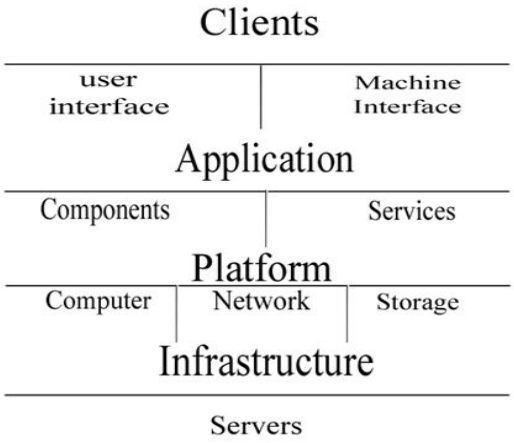

Figure 1: Cloud Architecture

administration itself. This motivates a new kind of migration technique which can be used to move the VMs without shutting down the client system. This new technique is called as live VMM. To reduce the power consumption, Vm migration technology makes vm on smaller number of server and make idle server into power saving mode. Virtualization enables workloads to be dynamically allocated resource in the form of virtual machines. These VMs can be relocated to a different set of physical resource while retaining their sate through migration.

1.1 Types of cloud computing

1.1.1 Web-based: Instead of developing complete applications, web-based cloud services let us develop conclusive web service functionality.

1.1.2 Software based: Software as a service (Saas) is a service which uses internet browser to provide applications to multiple users.

1.1.3 Platform based: An extension of Saas, which executes its programs only on the platform provided by cloud providers.

1.1.4 Managed services: Instead of end user, the application is utilized by a cloud provider.

1.1.5 Utility based: These services are used by the organizations to store the data virtually on cloud on demand.

1.1.6 Service commerce: It is a combination of SaaS and managed services which is used to provide services to the end user.

1.2 Types of migrations

This part represents three usual range of virtual machine migration: non-live migration, pre-copy live migration, and post-copy live migration.

1.2.1 Non-live migration: The execution of VM which runs on supply machine is kept on standby at the start. Without movement between running vms or applications and physical machines can be done by this process. Memory the storage and network property of the virtual machine square measure transferred from the first guest machine of the position.

1.2.2 Pre-copy live migration: Pre-copy live migration is implemented in production VMMs, having attracted researchers looking to reduce migration-related downtime [4], [5], [6]. Migrations of this sort proceed in three stages:

1) Pre-copy memory contents while continuing execution on the original physical machine, retransmitting any changed pages until this phase ends

2) Suspension of the VM on the original physical machine and transfer of the final pages of memory

3) In pre-copy phase, the VM executes on Physical machine. The VMM start to move pages to new machine. The iteration goes on and copies it to receiver machine. Memory usage of the running VM is monitored and any time that a page of memory that has been transferred is overwritten this page is marked as “dirty” and flagged to be retransmitted. If additional pages of memory are allocated by the still-active virtual machine then these are also flagged to be transmitted. During the movement, the pages will be sending or transmitted many time.

Eventually, according to a set of VMM-specific conditions, the execution of the VM on the source physical machine will be suspended and the second phase of the live migration begun. In Xen, this phase will be initiated when any of the following conditions hold [7]: few pages of memory were modified during the last pre-copy iteration, much pre-copy iteration has been carried out, or a large amount of network transfer has occurred. KVM uses different criteria to determine when to enter the second phase, doing so when it estimates that this will not result in downtime above a certain threshold. Such an approach may lead to KVM migrations lasting indefinitely and may still have adverse effects on the performance of these applications [8]. KVM also throttles the rates at which VMs may be migrated, with the default throttle defined as 32 MB/s, limiting the speed at which migration occurs and reducing the likelihood that memory-intensive applications can be successfully migrated.

In 2nd stage VM shutting down will lead to transfer of pages which have not executed to receiver machine. In 3rd stage, VM will continue running on receiver machine and remaining will be on source machine. Downtimes of maximum 1-2 seconds and 3.5 second (in worst case)[6] for WAN migrations [5].

1.2.3. Post-copy live migration [9], [10]:

Before sending files to receiver to execute it on receiver machine, instead of copying data, it let the running of VM in waiting on receiver system. The virtual machine uses memory content which then transferred to receiver machine from source. This process makes the virtual machine migration faster because the memory needs to be copied only once. This faster execution will be effective for green computing environment because it can allow more rapid adaptation to change everything. However, the performance of the migrated virtual machine may be more adversely affect during the transition due to waiting for the arrival of needed memory pages. Post-copy migration also increases the risk of VM failure during the migration, as if either node fails before the transfer of memory contents is completed the VM fails, but if the destination machine crashes for either non-live or pre-copy live migrations, the VM can be resumed instead on the original physical machine.

II. RELATED WORK ON MIGRATION PERFORMANCE

Because of frequently use of VMs, it is been discovered by the users or researchers that what are the constraints of migration and how this migration will save energy by enabling it. The bandwidth is affected by efficiency and performance of VMM. When to finish the phase of pre-copy state was also termed as another factor in migration.

The re-transmission of pages was done once the pages are overwritten in the memory.

In addition to examining the speed and data transfer of virtual machine migrations, ways in which to use migration strategies to improve performance or reduce energy consumption have also be studied. The typical focus of these works has been total migration time and impact upon the workload's quality of service rather than explicitly accounting for energy consumption added by migration. In 2008 Verma *et al* [11] presented pMapper, a tool to conserve power by dynamically consolidating HPC workloads on servers using live migration. They noted that HPC applications might take a minute to migrate and saw a 20–25% decline in performance during this period. Due to these factors and the large number of migrations required for workload consolidation they concluded that migration costs needed to be taken into account but looked at this in terms of time requirements and impacts upon throughput rather than incremental power consumption. Others [12] have found greater adverse effects on HPC workloads, leading sometimes to application failures. Ye *et al* [13] also looked at optimizing energy consumption by performing virtual machine migrations, focusing on the metrics of downtime, total migration time, total data transmitted and the impact of this on the quality of service experienced by the workload rather than the incremental power consumption of the migration process. Srinivasan *et al* [14] extended the SPECvc sc2010 benchmark to multiple machines by adding in automated virtual machine migrations, but, although their paper mentioned that "the benchmark should provide metrics to show power consumption relative to its performance", the focus of this work was primarily on throughput. Le *et al* [15] studied how to optimize migrations for high performance computing jobs, focusing primarily on exploiting variability in electrical pricing in different regions at on and off peak times as well as interacting with cooling systems to ensure that data centres can safely deal with a large influx of workload. Their migration cost analysis appears largely focused on the estimated transfer time required to move a virtual machine rather than any energy costs associated with doing so. Researchers have Evaluated a variety of strategies for migrating virtual machines between data centres to reduce energy consumption but assumed that machines could be migrated instantaneously. Researchers have introduce a benchmark providing reproducible results for analysing migrations. Their results show overhead for migrations ranging from 0.36%–11.31% of system resources, that migrations using the KVM virtual machine manager have higher overhead than those done using Xen. However they do not detail why KVM and Xen might have different overheads nor examine how overhead impacts power consumption. Determining how much power individual virtual machines are responsible for consuming is difficult, yet reasonable estimates can be made . Understanding the power impact of migration processes and other functions in the virtual machine manager thus is also feasible. We differ from the above in looking at the incremental power consumption of migration which have generally been ignored even in those studies looking at migration as a strategy to reduce power consumption.

III. CONCLUSIONS

The main issues with cloud computing are: (1) cloud servers consume huge amount of energy even in their idle state; (2) Migration of VM for load balancing also imparts energy consumption. In this paper a survey is done on intelligent VM migration approaches. And different VM migration types which include live VM migration and non live VM migration. Techniques on resource utilization are also discussed. This effective resource utilization also leads to less power consumption.

REFERENCES

[1] I. Foster, Z. Yong, I. Raicu, and S. Lu, "Cloud computing and grid computing 360-degree compared," Proc. computing environments workshop, pp. 1-10, 2008.

[2] R. Buyya, C. S. Yeo, and S. Venugopa, "Market-oriented cloud computing: vision, hype, and reality for delivering it services as computing utilities," Proc. 10th IEEE Int. Conf. on High Performance Computing and Communications (HPCC-08), pp. 5-13, 2008.

[3] W. Wang, Y. Zhang, B. Lin, X. Wu, K. Miao, "Secured and Reliable VM Migration in Personal Cloud", In Proceeding of the 2nd International Conference on Computer Engineering and Technology (ICCET), Apr. 2010, vol.1, pp:705-709.

[4] A. Gupta, L. V. Kal´e, D. Milojicic, P. Faraboschi, and S. M. Balle, "Hpcaware vm placement in infrastructure clouds," in Cloud Engineering (IC2E), 2013 IEEE International Conference on. IEEE, 2013, pp. 11–20.

[5] A. Murtazaev and S. Oh, "Sercon : Server consolidation algorithm using live migration of virtual machines for green computing," Iete Technical Review, vol. 28, no. 3, pp. 212–231, 2011. [Online]. Available: http://www.scopus.com/inward/record.url?eid=2-s2.0-9958712219&partnerID=40&md5=c44ad1129701c05756f92cb667af7374

[6] Girish Metkar Sanjay Agrawal Dr. Shailendra Singh "A Live Migration of Virtual Machine Based on the Dynamic Threshold at Cloud Data Centres" Volume 3, Issue 10, October 2013 ISSN: 2277 128X.

[7] R.S.C. Ho, C.L. Wang, and FC Lau. Lightweight process migration and memory prefetching in openmosix. In Parallel

and Distributed Processing, 2008. IPDPS 2008. IEEE International Symposium on, pages 1–12. IEEE, 2008.

[8] Anton Beloglazov, RajkumarBuyya, Young Choon Lee, and Albert Y. Zomaya. 2011. A taxonomy and survey of energy-efficient data centers and cloud computing systems. Advances in Computers 82 (2011), 47–111.

[9] R. Jansen and P. R. Brenner, "Energy efficient virtual machine allocation in the cloud," in 2011 International Green Computing Conference and Workshops (IGCC), Jul. 2011, pp. 1 –8.

[10] M. R. Hines and K. Gopalan, "Post-copy based live virtual machine migration using adaptive pre-paging and dynamic self-ballooning," in Proceedings of the 2009 ACM SIGPLAN/SIGOPS international conference on Virtual execution environments, New York, NY, USA, 2009, pp. 51–60.

[11] A. Verma, P. Ahuja, and A. Neogi, "pMapper: power and migration cost aware application placement in virtualized systems," in Proceedings of the 9th ACM/IFIP/USENIX International Conference on Middleware, New York, NY, USA, 2008, pp. 243–264.

[12] K. Srinivasan, S. Yuuw, and T. J. Adelmeyer, "Dynamic VM migration: assessing its risks & rewards using a benchmark," SIGSOFT Softw. Eng. Notes, vol. 36, no. 5, pp. 42–42, Sep. 2011.

[13] D. Huang, D. Ye, Q. He, J. Chen, and K. Ye, "Virt-LM: a benchmark for live migration of virtual machine," in Proceeding of the second joint WOSP/SIPEW international conference on Performance engineering. New York, NY, USA: ACM, 2011, pp. 307–316.

[14] V. Krishna Reddy, B. Thirumala Rao, Dr. L.S.S. Reddy, P. Sai Kiran, "Research Issues in Cloud Computing" in Global Journal of Computer Science and Technology 2011 Volume 11 Issue 11 Version 1.0 Online ISSN: 0975-4172 & Print ISSN: 0975-4350.

[15] B. Krishnan, H. Amur, A. Gavrilovska, and K. Schwan, "VM power metering: feasibility and challenges," SIGMETRICS Perform. Eval. Rev., vol. 38, pp. 56–60, Jan. 2011.