Energy Efficient Security Preserving VM Live Migration In Data Centers For Cloud Computing.

Korir Sammy, Ren Shengbing, Cheruiyot Wilson

School of Information Science and Engineering, Central south University
Changsha, Hunan, PR. China

Abstract

Virtualization is an innovation that has widely been utilized in modern data centers for cloud computing to realize energy-efficient operations of servers. Virtual machine (VM) migration brings multiple benefits such as resource distribution and energy aware consolidation. Server consolidation achieves energy efficiency by enabling multiple instances of operating systems to run simultaneously on a single machine. With virtualization, it is possible to consolidate severs through VM live migration. However, migration of virtual machines brings extra energy consumption and serious security concerns that derail full adoption of this technology.

In this paper, we propose a secure energy-aware provisioning of cloud computing resources on consolidated and virtualized platforms. Energy efficiency is achieved through just-right dynamic Round-Robin provisioning mechanism and the ability to power down sub-systems of a host system that are not required by VMs mapped to it. We further propose solutions to security challenges faced during VM live migration.

We validate our approach by conducting a set of rigorous performance evaluation study using CloudSim toolkit. The experimental results show that our approach achieves reduced energy consumption in data centers while not compromising on security.

Keywords: *Virtualization, VM-security, Energy-Efficiency, Data Center, Cloud Computing.*

1. Introduction

Dynamic consolidation of virtual machines (VMs) is a promising technology that can be used to reduce energy consumption in data centers. The number of power-on server nodes is kept to a minimum at any time, so that the excessive power used for running idle server nodes can be eliminated. The locations of VMs are continuously optimized in response to resource requirements of VMs. When there are many idle VMs, a management system consolidates them onto fewer server nodes, and temporarily shuts down the rest of the server nodes. When these idle VMs become active, the system restarts the powered off servers, and relocates VMs onto them.

Recently, with the rapid development of virtualization technology, majority of data centers have adopted this technology to design new generation data center architecture [1, 24].The benefits offered by this technology include, improved resource utilization, reduced operation costs and easier server management. Server consolidation and live migration of virtual machine is also used to achieve load balancing and energy saving. Server consolidation allows underutilized physical servers to be turned off after VMs running on it have been migrated to other unsaturated physical servers. Although virtual machine technology can improve the energy efficiency in data centers, the overheads caused by virtualization, its efficiency and security strategies used in consolidation and migration need to be investigated further.

The main contribution of this paper is to provide a secure way of reducing energy consumption in data centers through sever consolidation using VM live migration, without compromising on VM security. We propose strategy for deploying virtual machines to servers and migrating virtual machines among servers clusters based on sever workload utilization. The objective is to minimize the number of physical machines used to run all virtual machines. This goal is very important because the number of physical machines used strongly affects the overall power consumption. We propose strategies to protect VMs during migration. We further investigate various power-aware VM provisioning schemes like DVFS and DNS.

The remainder of this paper is structured as follows. In section 2 we discuss our methodology followed by the proposed workload placement approaches in Section 3. Section 4 describes security aware migration strategies. In Section 5, we perform a comprehensive evaluation and analysis our proposed solution. Section 6 presents the simulation results and related work is discussed in section 7.Conclusions and future work are reviewed in Section 7.

2. Methodology

Dynamic Round-Robin algorithm [22] is proposed as a virtual machine deployment method for power saving purpose. The objective of using this method is to reduce the number of physical servers used to run VMs. To guarantee VM security during migration, we prose the use of a security module that is incorporated into Virtual machine monitor (VMM).

2.1 Dynamic Round-Robin

We propose the use of Dynamic Round-Robin as an extension to the Round-Robin method. The method uses two rules to help consolidate virtual machines for maximum resource usage. The first rule is, if a virtual machine has finished and there exist other virtual machines hosting on the same physical machine, this physical machine will accept no more new virtual machine. We refer to such physical machines to be in "retirement" state, meaning that when the rest of the virtual machines finishes, this physical machine can be shutdown.

The second rule of Dynamic Round-Robin is that if a physical machine is in the "retirement" state for a sufficiently long period of time, instead of waiting for the hosting virtual machines to finish its jobs on its own, the physical machine will be forced to migrate the rest of active virtual machines to other physical machines. The machine is shut down after VM migration has been completed. This is achieved by invoking the workload controller functionality. The waiting time threshold is denoted as "retirement threshold". A physical machine which is in the retirement state but could not finish all virtual machines after the retire threshold has been exceeded, will be forced to migrate its virtual machines and shutdown.

Our Dynamic Round-Robin method uses two basic rules in order to consolidate virtual machines deployed by the original Round-Robin method. The first rule is to avoid adding extra virtual machines to a retiring physical machine so it can be shut down. The second rule is to speed up the consolidation process and enable dynamic round-robin method to shutdown physical machines, so that it can reduce the number of physical machine used to run all virtual machines.

3. Management Services

The following section describes the workload placement controller and the reactive workload migration controller [4]. We refer to active VMs on a given sever as workloads. We employ the use of two controllers to manage workload placements in the sever cluster.

3.1 Workload Placement Controller

The controller has two components;

- A component that simulates the assignment of several application workloads on a single server. It traverses the per-workload time varying traces of historical demand to determine the peak of the aggregate demand for the combined workloads. If for each capacity attribute, e. g., CPU and memory, the peak demand is less than the capacity of the attribute for the server then the workloads fit on the server.
- An optimizing search component examines many alternative placements of workloads on servers and reports the best solution found. The optimizing search is based on genetic algorithm [2, 5].

The workload placement controller is based on the Capman tool that is described further in [6, 7]. It supports both consolidation and load leveling exercises. Load leveling balances workloads across a set of resources to reduce the likelihood of service level agreement violations. Capman supports the controlled overbooking of capacity that computes a required capacity for workloads on a server that may be less than the peak of aggregate demand. It is capable of supporting a different quality of service for each workload [8]. Without loss of generality, this paper considers the highest quality of service, which corresponds to a required capacity for workloads on a server that is the peak of their aggregate demand.

We exploit Capman's multi-objective optimization functionality in this paper [23]. Instead of simply finding the smallest number of servers needed to support a set of workloads, Capman evaluates solutions according to a second simultaneous objective. The second objective aims to minimize the number of changes to workload placement. When invoking Capman, an additional parameter specifies a target *t* as a bound for the number of workloads that it is desirable to migrate. Limiting the number of migrations limits the migration overhead and reduces the risk of incurring a migration failure. If it is possible to find a solution with fewer than *t* migrations, then Capman reports the workload placement that needs the smallest number of servers and has t or fewer migrations. If more changes are needed to find a solution, then Capman reports a solution that has the smallest number of changes to find a feasible solution. A data center operator could choose a value *t* based on experience

regarding the overhead that migrations place upon network infrastructure and servers.

3.2 Workload Migration Controller

Migration controller is a fuzzy-logic based feedback control loop. An advisor module of the controller continuously monitors the servers' resource utilization and triggers a fuzzy-logic based controller whenever resource utilization values are too low or too high. When the advisor detects a lightly utilized, i. e., under-load situation, or overload situation the fuzzy controller module identifies appropriate actions to remedy the situation. For this purpose, it is initialized with information on the current load situation of all affected servers and workloads and determines an appropriate action. For example, as a first step, if a server is overloaded it determines a workload on the server that should be migrated and as a second step it searches for a new server to receive the workload. Furthermore, these rules initiate the shutdown and startup of nodes.

The implementation of the workload migration controller uses the following rules:

- A server is defined as overloaded if its CPU or memory consumption exceeds a given threshold. In an overload situation, first, a fuzzy controller determines a workload to migrate away and then it chooses an appropriate target server. The target server is the least loaded server that has sufficient resources to host the workload. If such a server does not exist, we start up a new server and migrate the workload to the new one.
- An under-load situation occurs whenever the CPU and memory usage averaged over all servers in the server pool drops below a specified threshold. While an overload condition is naturally defined with respect to a particular server, the under-load situation is different. It is defined with respect to the average utilization of the overall system involving all the nodes. In this way, we try to avoid system thrashing: e.g., a new server generally starts with a small load and should not be considered immediately for consolidation. In an under-load situation, first, the fuzzy controller chooses the least loaded server and tries to shut it down. For every workload on this server, the fuzzy controller determines a target server. If a target cannot be found for a workload then the shutdown process is stopped. In contrast to overload situations, this controller does not ignite additional servers.

4. Secure Live Migration

Live migration of VMs imposes critical security issues. Simply using cryptographic means and hashing to protect the sensitive data and meta-data is not enough. Critical issues lie in the VM live migration, in which the migrating VM is still running while migration is in process. Time-of-Check to time-of-use (TOCTTOU) [9, 10] attack and replay attack can be launched against the VM if the protection for migration is not well planned.
In security preserving VM live migration, the challenges lie in the following three aspects:

- Preserving the privacy and integrity of protected contents.
- Packing the maintenance metadata in the Virtual Machine Monitor (VMM), solving the namespace conflict, and re-establishing the protection base on the target platform.
- Eliminating the security vulnerabilities imposed by live migration.

This section presents a secure migration strategy that provides live migration capability to VM without exposing it to attacks. Our goal is to preserve integrity and privacy protection during and after VM live migration.

4.1 Workload Migration Controller

Our design goal is to provide VM migration capability while preserving strict protection during and after the migration without significant performance degradation and migration downtimes. We propose a design which is mainly based on Xen system. Three modules are added to VMM .Figure 1 gives an overview of our design approach.

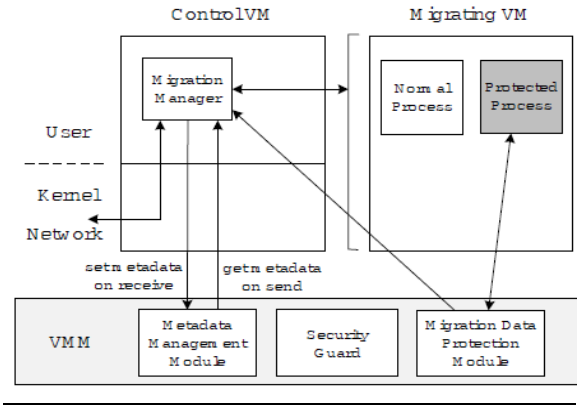

Figure 1: Migration Model.

From figure 1, the migration data protection module is responsible for intercepting and protecting contents that

belong to the protected processes inside the migrating VM.
Metadata management module is responsible for serializing the metadata for transmission and re-constructing the metadata in the migration target machine. The security module protects against vulnerabilities discussed in the next section.

4.2 Security Module

The protection approach presented in previous subsection works well when the VM is suspended before it is migrated. In the context of live migration [11], the protection is not sufficient. Here, we describe the security problems existing in live migration, including time-of-check to time-of-use (TOCTTOU) attack, resumption ordering problem, and replay attacks.

4.2.1 TOCTTOU

There is TOCTTOU security vulnerability in the map-and send mechanism when live migrating the VM. The key difference is the VM is running when the map-and-send is in process. When the migration manager issues a hyper-call to map a batch of memory pages of the migrating VM, the check is performed to see if the pages are protected. Then the VMM encrypts the protected pages. However, if a page is not protected when the migration manager performs the checking and mapping, and then the page is assigned to a secure process in the migrating VM, that page turns into a sensitive page but is still mapped in the space of the migrating tools.

We fix this security hole by revoking and redirecting the memory mapping to a zero page when the page is assigned as a secure page. The moment a page is allocated as a secure page and the Security module will check the page table of migration manager and force any existing mappings to the secure page to be unmapped. To prevent the page fault, the Security module will temporarily map a zero page to the page table entry. This forced redirection is performed no matter the mapped content has been sent, is being sent, or has not been sent yet. It is possible that part or the entire page may be sent in zero value. This is safe because the page is assigned to another process, which indicates the page content is obsolete. And if the protected process writes the page later, it will be dirtied and sent again. This is safe because the transmitted content is in encryption form and will not be decrypted by the VMM at the target machine because the page is no longer a secure page.

4.2.2 VM Resumption Ordering

Adopting a wrong resuming order may also incur serious security problems. For example, the Xen migration manager will pin all page table pages when a VM is received and about to resume. Page pinning is a process that Xen validates the mappings in the page tables and guarantees the VM has the privilege to map the memory pages. Meanwhile, the metadata that indicate the security level of pages are also set in the resuming phase. If the page tables are pinned before the correct security level of pages are set, it is possible that an unauthorized page table can bypass the security check and illegally map the secure pages. Although we can implement a correct hyper-call invocation sequence in the migration manager, it is possible that the migration manager is compromised, and hence breaks the sequence.

This security hole is fixed by enforcing the resuming order in the VMM. Specifically, the VMM ensures the pinning operation cannot be done before it has fully received and restored the metadata of page security information.

4.2.3 Replay Attack

Pages are sent in a round based fashion in the pre-copy phase. In each round, the migration manager will get the encrypted image of the protected page and send it over network if the page is dirtied in previous round. If a protected page is dirtied multiple times, it is possible for a replay attack that a malicious migration manager sends the old content instead of the new version. To prevent this attack, we hash the content of each protected page in the stop-and-copy phase. This guarantees the hashing values are finalized. The VMM on the target node will confirm that all the secure pages match the final hashing before resuming the VM.

5. Evaluation and Analysis

We have evaluated the proposed algorithms through simulations using the CloudSim toolkit [13, 14] with an extension enabling secure power-aware simulations. We have chosen CloudSim toolkit as a simulation framework, as it is built for simulation of Cloud computing environments. We have extended the framework in order to enable our proposed energy aware algorithm simulations as the core framework does not provide this capability. In addition, we have incorporated a security module to ensure VM security during migration.

The simulated data center consists of 100 heterogeneous physical nodes. Each node is modeled to have one CPU core with performance equivalent to 1000, 2000 or 3000 MIPS, 8 Gb of RAM and 1 TB of storage. Users submit requests for provisioning of 290 heterogeneous VMs that fill the full capacity of the data center. We simulated a Non Power Aware policy (NPA) and Dynamic Voltage Frequency Scaling (DVFS) that adjusts the voltage and frequency of CPU according to current utilization. We simulated a Single Threshold policy (ST) and two-threshold policy aimed at Minimization of Migrations (MM). Besides that, the policies have been evaluated with different values of the thresholds.

6. Simulation Results

Table 1: Simulation Results

Policy	Energy	SLA	Migrations	AVG SLA
NPA	9.14KWh	-	-	-
DVFS	4.42KWh	-	-	-
ST	2.04KWh	5.41%	35,225	81%
MM	1.46KWh	9.05%	34,230	88%
DRR	1.30KWh	8.01%	334,110	84%

The simulation results are presented in Table I. Our results show that dynamic reallocation of VMs using extended Round Robin algorithm saves more energy compared to static allocation policies. DRR policy allows to achieve the best energy savings: by 82%, 67% and 24% less energy consumption relatively to NPA, DVFS,MM and ST policies respectively with thresholds 30-70% and ensuring percentage of SLA violations of 1.1%; and by 87%, 74% and 43% with thresholds 50-90% and 6.7% of SLA violations. MM policy leads to more than 10 times fewer VM migrations than ST. The results show the flexibility of the algorithm, as the thresholds can be adjusted according to SLA requirements. Strict SLA (1.11%) allows achievement of the energy consumption of 1.48 KWh. However, if SLA requirements are relaxed (6.68%), the energy consumption is further reduced to 1.14 KWh.

7. RELATED WORK

Cloud power efficiency is a hot topic that is receiving an increasing attention due to both environmental [15] and economic issues [16]. Many works focused on the analytical modeling of the VMs allocation problem, and proposed different solutions to increase data center power efficiency. However, differently from our work, they tended to leave out real-world implementation and evaluation of proposed solutions. In the following, with an order of increasing similarity with our work, we present some important proposals dealing with Cloud power efficiency.

In [17], authors address the problem of Cloud resource provisioning for real-time services. They exploit SLAs to drive resource allocation, and strive to increase power efficiency as tradeoff between task completion times and powered on servers. However, they assume that VM workloads are known a-priori, and hence they do not deal with VMs migration and dynamic server consolidation.

In [18], authors present four different migration techniques to consolidate VMs. The experimental results, obtained in the CloudSim simulator [19], show that all the proposed policies can greatly reduce the Cloud power consumption. Unfortunately, proposed algorithms do neither mention nor consider service providers SLAs.

In [20], given a particular workload, authors present a mathematical model to find the exact number of physical servers to ensure power efficiency. In addition, the proposed solution considers also dynamic CPU frequency scaling, and supplies the optimal value for this parameter. However, authors assume that Cloud jobs can be partitioned among different physical servers. Even if this assumption could be viable in particular scenarios, it does not fit well the general VM allocation problem.

In [21], authors present a two phase solution to the problem of SLAs enforcement and VMs allocation. The first phase finds the optimal number of VMs required to meet service providers' SLA, and exploits high level directives to increase applications utility while reducing power consumption. Then, the second phase places VMs with the main goal of minimizing the number of powered on servers. The experimental results support the technical soundness of the proposal; however, authors do not present any result on the total power consumption of the test bed during the tests.

Finally, [22] presents Mistral, a novel solution that optimizes both VMs performance and power consumption, while considering transient costs associated with run-time reconfigurations. Authors adopt the A*-search technique to find complex reconfiguration actions, and consider indicators on the stability of the next data center configuration in the decision process. Presented experimental results, obtained in a real Cloud test bed, make the proposed solution extremely solid; however, to the best of our knowledge, Mistral does not consider power capping and dynamic CPU frequency scaling.

While significant amount of work has been done provide energy efficient VM allocation, VM security has not been extensively been researched on with respect to energy consumption. Our research strives to fill this gap that has been overlooked by researchers. Hence, our work is complementary to others, and offers some useful insights on server consolidation and security aspects in real world deployments.

8. Conclusion and Future Work

In this paper, we have proposed a secure energy-aware provisioning of cloud computing resources in virtualized platforms. Our simulation results convinced us that VMs migration using Dynamic Round robin algorithm for server consolidation is an extremely feasible solution to reduce energy consumption in a data center without compromising on security. Our strategy also has lower SLA violations compared with existing strategies like ST. We are convinced that proposed security mitigation strategies during migration guards against TOCTTOU, VM Resumption Ordering and Replay Attacks.

Future work includes further optimization of our approach and analysis and measuring of VM migration cost in a cluster. Other security threats facing VM migration shall be investigated further. We also plan to incorporate Quality of Services (QoS) monitoring capability at VM level.

References

[1] H. Liu, H. Jin, X. Liao, L. Hu, and C. Yu, "Live migration of virtual machine based on full system trace and replay, "in Proceedings of the 18th ACM international symposium on High performance distributed computing, 2009, pp. 101–110.

[2] Cheruiyot Wilson, Guan-Zheng Tan, Joseph Cosmas Mushi, Felix Musau, " Genetic Algorithm-Enhanced Retrieval process for Multimedia Data", International journal of Advancements in Computing Technology, Volume 3, number 3, 2011.

[3] R.Buyya,A.eloglazov,and J.Abawajy "Energy-Efficient Management of Data Center Resources for Cloud Computing: A Vision, Architectural Elements, and Open Challenges" In proceedings of of 2010 international conference on parellel and distributed Processing Techniques and Applicatins (PDPTA 2010),Las Vegas,USA,July 2010.

[4] D. Gmach, J. Rolia, L. Cherkasova, G. Belrose, T. Turicchi, A. Kemper, An Integrated Approach to Resource Pool Management: Policies, Efficiency and Quality Metrics, in: Proc. of the 38th IEEE/IFIP Int. Conf. on Dependable Systems and Networks (DSN), Anchorage, Alaska, USA, 2008.

[5] J. H. Holland, Adaptation in Natural and Artificial Systems, University of Michigan Press, Ann Arbor, 1975.

[6] D. Gmach, J. Rolia, L. Cherkasova, G. Belrose, T. Turicchi, A. Kemper, An Integrated Approach to Resource Pool Management: Policies, Efficiency and Quality Metrics, in: Proc. of the 38th IEEE/IFIP Int. Conf. on Dependable Systems and Networks (DSN), Anchorage, Alaska, USA, 2008.

[7] J. Rolia, L. Cherkasova, M. Arlitt, A. Andrzejak, A Capacity Management Service for Resource Pools, in: Proc. of the 5th Int. Workshop on Software and Performance (WOSP), Palma, Illes Balears, Spain, 2005, pp. 229–237.

[8] L. Cherkasova, J. Rolia, R-Opus: A Composite Framework for Application performability and QoS in Shared Resource Pools, in: Proc. of the Int. Conf. on Dependable Systems and Networks (DSN), Philadelphia, USA, 2006.

[9] W. McPhee. Operating system integrity in OS/VS2. IBM Journal of Research and Development, 13(3):230, 1974.

[10] M. Bishop and M. Dilger. Checking for Race Conditions in File Accesses. Computing Systems, 2(2):131–152, 1996.

[11] C. Clark, K. Fraser, S. Hand, J. G. Hansen, E. Jul, C. Limpach, I. Pratt, and A. Warfield. Live migration of virtual machines. In Proc. of NSDI, pages 273–286, 2005.

[12] Buyya R, Ranjan R, Calheiros RN. Modeling and simulation of scalable Cloud computing environments and the CloudSim toolkit: Challenges and opportunities. In Proceedings of the 7th High Performance Computing and Simulation (HPCS 2009). Leipzig, Germany, June 2009.

[13] Calheiros RN, Ranjan R, Beloglazov A, De Rose CAF, and Buyya R. CloudSim: A Toolkit for Modeling and Simulation of Cloud Computing Environments and Evaluation of Resource Provisioning Algorithms.

[14] "Gartner Estimates ICT Industry Accounts for 2 Percent of Global CO2 Emissions",http://www.gartner.com/it/page.jsp?id=503867, Available the 4 Mar. 2011

[15] "Gartner Says Energy-Related Costs Account for Approximately 12 Percent of Overall Data Center Expenditures", http://www.gartner.com/it/page.jsp?id=1442113, Available the 4 Mar. 2011.

[16] K. H. Kim, A. Beloglazov, and R. Buyya, "Power-aware Provisioning of Cloud Resources for Real-time Services", In Proc. Of the 7th International Workshop on Middleware for Grids, Clouds and e-Science (MGC 2009)

[17] A. Beloglazov, R. Buyya, "Energy Efficient Resource Management in Virtualized Cloud Data Centers", In Proc. of the 10th IEEE/ACM International Conference on Cluster, Cloud and Grid Computing, IEEE Press, 2010, pp.826-831.

[18] R. Buyya, R. Ranjan, and R. N. Calheiros, "Modeling and simulation of scalable cloud computing environments and the CloudSim toolkit: Challenges and opportunities", in Proceedings of the 7th High Performance Computing and Simulation Conference (HPCS'09), IEEE Press, NY, USA, 2009.

[19] H. S. Abdelsalam, K. Maly, R. Mukkamala, M. Zubair, D. Kaminsky, "Analysis of Energy Efficiency in Clouds", In Proc. of the Computation World: Future Computing, Service Computation, Cognitive, Adaptive, Content, Patterns, 2009, pp.416-421.

[20] H. N. Van, F. D. Tran, J.-M. Menaud, "Performance and Power Management for Cloud Infrastructures", In Proc. of the IEEE 3rd International Conference on Cloud Computing (CLOUD'10), IEEE Press, Jul. 2010, pp.329-336.

[21] G. Jung, M.A. Hiltunen, K.R. Joshi, R. D. Schlichting, C. Pu, "Mistral: Dynamically Managing Power, Performance, and Adaptation Cost in Cloud Infrastructures", In Proc. of the IEEE 30th International Conference on Distributed Computing Systems (ICDCS'10), IEEE Press, Jun. 2010, pp.62-73.

[22] Ching-Chi,Pangffeng,Jan-jan Wu, "Energy-Aware Virtual machine Dynamic provision and scheduling for cloud computing." In Proc. of the 2011 IEEE 4th International Conference on cloud computing.

[23] Dun-wei Gong; Na-na Qin; Xiao-yan Sun," Multi-objective optimization with uncertainty: Probabilistic and fuzzy approaches" Bio-Inspired Computing: Theories and Applications (BIC-TA), 2010

[24] Matthias Schmidt, Niels Fallenbeck, Matthew Smith, Bernd Freisleben,"Efficient Distribution of Virtual Machines for Cloud Computing" 18th Euromicro Conference on Parallel, Distributed and Network-based Processing, 2010