International Journal of Current Research and Modern Education (IJCRME)
ISSN (Online): 2455 - 5428
(www.rdmodernresearch.com) Volume I, Issue I, 2016

HYBRID SWIPER: SYNCHRONIZATION OF IP AND MAC ADDRESS TO OVERCOME THIRD PARTY ATTACKS IN CLOUD

M. Punitha Valli* & R. Selvakumar**

* PG Scholar, Department of Master of Computer Applications, Dhanalakshmi Srinivasan Engineering College, Perambalur, Tamilnadu

** Assistant Professor, Department of Master of Computer Applications, Dhanalakshmi Srinivasan Engineering College, Perambalur, Tamilnadu

Abstract:

The emerging paradigm of cloud computing, e.g., Amazon Elastic Compute Cloud (EC2), promises a highly flexible yet robust environment for large-scale applications. Ideally, while multiple virtual machines (VM) share the same physical resources (e.g., CPUs, caches, DRAM, and I/O devices), each application should be allocated to an independently managed VM and isolated from one another. Unfortunately, the absence of physical isolation inevitably opens doors to a number of security threats. In this paper, we demonstrate in EC2 a new type of security vulnerability caused by competition between virtual I/O workloads - i.e., by leveraging the competition for shared resources, an adversary could intentionally slow down the execution of a targeted application in a VM that shares the same hardware. In particular, we focus on I/O resources such as hard-drive throughput and/or network bandwidth - which are critical for data-intensive applications. We design and implement Swiper, a framework which uses a carefully designed workload to incur significant delays on the targeted application and VM with minimum cost (i.e., resource consumption). We conduct a comprehensive set of experiments in EC2, which clearly demonstrates that Swiper is capable of significantly slowing down various server applications while consuming a small amount of resources.

Introduction: (Project Description)

Cloud computing is a computing paradigm, where a large pool of systems are connected in private or public networks, to provide dynamically scalable infrastructure for application, data and file storage. With the advent of this technology, the cost of computation, application hosting, content storage and delivery is reduced significantly. Cloud computing is a practical approach to experience direct cost benefits and it has the potential to transform a data center from a capital-intensive set up to a variable priced environment. The idea of cloud computing is based on a very fundamental principal of „reusability of IT capabilities'. The difference that cloud computing brings compared to traditional concepts of “grid computing”, “distributed computing”, “utility computing”, or “autonomic computing” is to broaden horizons across organizational boundaries. Forrester defines cloud computing as: “A pool of abstracted, highly scalable, and managed compute infrastructure capable of hosting end customer applications and billed by consumption.” Cloud Computing is a technology that uses the internet and central remote servers to maintain data and applications. Cloud computing allows consumers and businesses to use applications without installation and access their personal files at any computer with internet access. This technology allows for much more efficient computing by centralizing data storage, processing and bandwidth. A simple example of cloud computing is Yahoo email, Gmail, or Hotmail etc.

Cloud Computing Services:

Cloud Providers offer services that can be grouped into three categories.

- ✓ Software as a Service (SaaS): In this model, a complete application is offered to the customer, as a service on demand. A single instance of the service runs on the cloud & multiple end users are serviced. On the customers‟ side, there is no need for upfront investment in servers or software licenses, while for the provider, the costs are lowered, since only a single application needs to be hosted & maintained. Today SaaS is offered by companies such as Google, Salesforce, Microsoft, Zoho, etc.

- ✓ Platform as a Service (Paas): Here, a layer of software, or development environment is encapsulated & offered as a service, upon which other higher levels of service can be built. The customer has the freedom to build his own applications, which run on the provider's infrastructure. To meet manageability and scalability requirements of the applications, PaaS providers offer a predefined combination of OS and application servers, such as LAMP platform (Linux, Apache, MySql and PHP), restricted J2EE, Ruby etc. Googles App Engine, Force.com, etc are some of the popular PaaS examples.
- ✓ Infrastructure as a Service (Iaas): IaaS provides basic storage and computing capabilities as standardized services over the network. Servers, storage systems, networking equipment, data centre space etc. are pooled and made available to handle workloads. The customer would typically deploy his own software on the infrastructure. Some common examples are Amazon, GoGrid, 3 Tera, etc.

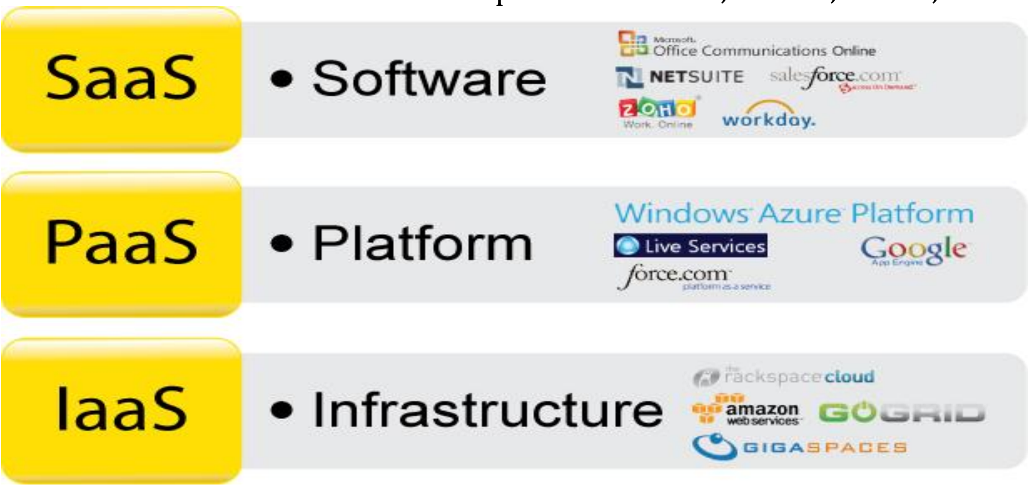

Figure: Types of Cloud Services

Cloud Computing Models:

Enterprises can choose to deploy applications on Public, Private or Hybrid clouds. Cloud Integrators can play a vital part in determining the right cloud path for each organization.

Public Cloud: Public clouds are owned and operated by third parties; they deliver superior economies of scale to customers, as the infrastructure costs are spread among a mix of users, giving each individual client an attractive low-cost, "Pay-as-you-go" model. All customers share the same infrastructure pool with limited configuration, security protections, and availability variances. These are managed and supported by the cloud provider. One of the advantages of a Public cloud is that they may be larger than an enterprises cloud, thus providing the ability to scale seamlessly, on demand.

Private Cloud: Private clouds are built exclusively for a single enterprise. They aim to address concerns on data security and offer greater control, which is typically lacking in a public cloud. There are two variations to a private cloud:

On-premise Private Cloud: On-premise private clouds, also known as internal clouds are hosted within one's own data center. This model provides a more standardized process and protection, but is limited in aspects of size and scalability. IT departments would also need to incur the capital and operational costs for the physical resources. This is best suited for applications which require complete control and configurability of the infrastructure and security.

Externally hosted Private Cloud: This type of private cloud is hosted externally with a cloud provider, where the provider facilitates an exclusive cloud environment with full guarantee of privacy. This is best suited for enterprises that don't prefer a public cloud due to sharing of physical resources.

Hybrid Cloud: Hybrid Clouds combine both public and private cloud models. With a Hybrid Cloud, service providers can utilize 3rd party Cloud Providers in a full or partial manner thus increasing the flexibility of computing. The Hybrid cloud environment is capable of providing on-demand, externally provisioned scale. The ability to augment a private cloud with the resources of a public cloud can be used to manage any unexpected surges in workload. Enterprises would need to align their applications, so as to exploit the architecture models that Cloud Computing offers. Some of the typical benefits are listed below:

Reduced Cost: There are a number of reasons to attribute Cloud technology with lower costs. The billing model is pay as per usage; the infrastructure is not purchased thus lowering maintenance. Initial expense and recurring expenses are much lower than traditional computing.

Increased Storage: With the massive Infrastructure that is offered by Cloud providers today, storage & maintenance of large volumes of data is a reality. Sudden workload spikes are also managed effectively & efficiently, since the cloud can scale dynamically.

Flexibility: This is an extremely important characteristic. With enterprises having to adapt, even more rapidly, to changing business conditions, speed to deliver is critical. Cloud computing stresses on getting applications to market very quickly, by using the most appropriate building blocks necessary for deployment.

Cloud Computing Challenges:

Despite its growing influence, concerns regarding cloud computing still remain. In our opinion, the benefits outweigh the drawbacks and the model is worth exploring. Some common challenges are:

Data Protection: Data Security is a crucial element that warrants scrutiny. Enterprises are reluctant to buy an assurance of business data security from vendors. They fear losing data to competition and the data confidentiality of consumers. In many instances, the actual storage location is not disclosed, adding onto the security concerns of enterprises. In the existing models, firewalls across data centers (owned by enterprises) protect this sensitive information. In the cloud model, Service providers are responsible for maintaining data security and enterprises would have to rely on them.

Data Recovery and Availability: All business applications have Service level agreements that are stringently followed. Operational teams play a key role in management of service level agreements and runtime governance of applications. In production environments, operational teams support

- ✓ Appropriate clustering and Fail over
- ✓ Data Replication
- ✓ System monitoring (Transactions monitoring, logs monitoring and others)
- ✓ Maintenance (Runtime Governance)
- ✓ Disaster recovery
- ✓ Capacity and performance management

If, any of the above mentioned services is under-served by a cloud provider, the damage & impact could be severe.

Management Capabilities: Despite there being multiple cloud providers, the management of platform and infrastructure is still in its infancy. Features like „Auto-scaling‟ for example, is a crucial requirement for many enterprises. There is huge potential to improve on the scalability and load balancing features provided today.

Regulatory and Compliance Restrictions: In some of the European countries, Government regulations do not allow customer's personal information and other sensitive information to be physically located outside the state or country. In order to meet such requirements, cloud providers need to setup a data center or a storage site exclusively within the country to comply with regulations. Having such an infrastructure may not always be feasible and is a big challenge for cloud providers. With cloud computing, the action moves to the interface — that is, to the interface between service suppliers and multiple groups of service consumers. Cloud services will demand expertise in distributed services, procurement, risk assessment and service negotiation — areas that many enterprises are only modestly equipped to handle.

Existing System:

Existing system implements Markov decision process (MDP) formalization of reinforcement learning, a single adaptive agent interacts with an environment defined by a probabilistic transition function. In this solipsistic view, secondary agents can only be part of the environment and are therefore fixed in their behavior. The MM policy did slightly better. In the limit, this should not be the case since an agent trained by the mini-max-Q algorithm should be not sensitive to the enemy against which it was trained and always behave so as to maximize its score in the worst case and explores the empathy between realm scheduling in a virtual machine monitor (VMM) and I/O performance. Habitually, VMM schedulers have inattentive on fairly distribution the processor resources along with domains while leave-taking the scheduling of I/O property as a less important anxiety. The disadvantages are managing the scarce bandwidth could be viewed as a global optimization problem servers from all applications must be placed with great care to ensure the sum of their traffic does not saturate any of the network links. The followed by a lengthy period of highly busty load and low utilization. Cloud computing platforms enable users to rent computing and storage resources on demand to run their networked applications and employ virtualization to multiplex virtual servers belonging to different customers on a shared set of servers. In this paper, empirically evaluate the efficacy of cloud platforms for running latency-sensitive multimedia applications. Having the disadvantages of Techniques such as server-side and client side buffering can help mask network fluctuations.

Disadvantages:

- ✓ Difficult to analyze third party attacks in virtual machine allocation
- ✓ Only maintain single VM vulnerabilities
- ✓ Limit the I/O and CPU performance

Proposed System:

A cloud computing system offers to its users the illusion of "infinite" computing and storage capacities on an on-demand basis. New variety of security vulnerability caused by competition between virtual I/O workloads - i.e., by investing the competition for shared resources associate degree individual may designedly curtail the execution of a targeted application during a VM that shares an equivalent hardware. Specially, we have a tendency to specialize in I/O resources like hard-drive turnout and/or network information measure - that area unit essential for data-intensive applications. Implement an SWIPER framework on I/O resources such as hard-drive throughput and/or network bandwidth - which are critical for data-intensive applications. We design and implement SWIPER, a framework which uses a carefully designed workload to incur significant delays on the targeted application and VM with minimum cost (i.e., resource consumption).While there are more number of users uses the application except an active tab other tab loading bandwidth to be stopped. The speeds of the ideal tab are reallocated to the new requested user. Then the user can uses the same application with the same speed.

Advantages:

- ✓ Multiple VM scheduling is performed effectively.
- ✓ Network and bandwidth authentication to overcome the third party attacks in cloud system.
- ✓ Resources are consumed to decrease the server performance.

Literature Survey:

Resource-Freeing Attacks: Improve Your Cloud Performance (at Your Neighbor's Expense)

Techniques:

- – Contention measurements

Demerits:

- – Loose overall efficiency because of the load caused by the extraneous
- – There is no performance isolation in cloud environments

Resource-Freeing Attacks: Improve Your Cloud Performance (at Your Neighbor's Expense)

Techniques:

- – Contention measurements

Demerits:

- – Loose overall efficiency because of the load caused by the extraneous
- – There is no performance isolation in cloud environments

Resource-Freeing Attacks: Improve Your Cloud Performance (at Your Neighbor's Expense)

Techniques:

- – Contention measurements

Demerits:

- – Loose overall efficiency because of the load caused by the extraneous
- – There is no performance isolation in cloud environments

The Impact of Virtualization on Network Performance of Amazon EC2 Data Center

Techniques:

- – Spatial experiment approach

Demerits:

- – Unstable to provide TCP/UDP throughput
- – Difficult to analyze characteristics of virtualized data centers

Eliminating the Hypervisor Attack Surface for a More Secure Cloud

Techniques:

- – No Hype framework

Demerits:

- – Major modifications to the guest OS to perform all system discovery during boot up

Implementation and Results:

Risk Aware Overbooking Analysis:

Overbooking:

The overbooking strategies for cloud infrastructure providers lead to increase their profit and competitiveness. The objective of overbooking is to improve the expected resource allocation. Instead of allocating each resource once, high performance can be achieved by allocating them several times. Overbooking offers a solution for the system utilization problem, by allowing the resource provider to accept more reservations than the capacity. Hence, it can be effectively used to minimize the loss of revenue. However, the challenging issues in using overbooking are determining the appropriate number of excess reservations, minimizing total compensation cost, addressing legal and regulatory issues, and dealing with market acceptance, especially the ill-will or negative effects from users who have been denied access. The Risk aware overbooking approach can be used to improve system utilization. Within fuzzy assessment putting a job in a gap is acceptable if the first job in the queue is not delayed. This preserves starvation and leads to an increased utilization of the underling system. However, this assessment has to be used with caution in systems guaranteeing QoS aspects, since jobs in the queue might be delayed.

$$\mathrm{OBF} = \frac{\mathrm{UsageRequested} - \mathrm{Realusage}}{\min(\mathrm{UseageRequested}, \mathrm{Realcapacity})}$$

Overbooking process analyzed based on following formula such as

$$\mathrm{risk}_i = \begin{cases} 0 & \text{if } \mathrm{Req}_i < \mathrm{Unreq}_i \\ \mathrm{Req}_i - \mathrm{Unreq}_i & \text{if } \mathrm{Unreq}_i < \mathrm{Req}_i < \mathrm{Free}_i \\ 1 & \text{if } \mathrm{Req}_i > \mathrm{Free}_i \end{cases}$$

Risk Admission Control Mechanisms:

```
Let Req be the incoming service request
Let get Fuzzy Risk Assessment the function that obtains the associated risk of accepting app
Let Threshold (CPU, mem, IO) the current risk thresholds of the data center
App_Risk [CPU, mem, IO] = get Fuzzy Risk Assessment (req)
Accept Req
Else
Reject Req
End if
```

Algorithm Description:

Resource overbooking is an admission control technique to increase utilization in cloud environments. However, due to uncertainty about future application workloads, overbooking may result in overload situations and deteriorated performance. We mitigate this using brownout, a feedback approach to application performance steering, that ensures graceful degradation during load spikes and thus avoids overload. Additionally, brownout management information is included into the overbooking system, enabling the development of improved reactive methods to overload situations. One way of addressing those problems and increasing resource utilization is resource overbooking. In essence, the provider allocates more capacity than the real capacity of the data center. In other words, a new VM is admitted although the sum of requested cores or memory exceeds the number of cores or total memory in the data center. However, such an approach may lead to resource overload and performance degradation. Therefore, besides carefully choosing how to place VMs on physical machines, a new resource management challenge appears: estimating the appropriate level of overbooking that can be achieved without impacting the performance of the cloud services. Admission control techniques are therefore needed to handle this tradeoff between increasing resource utilization and risking performance degradation. Combining statistical multiplexing of resource demands, server consolidation and economy of scales, cloud providers are able to offer users resources at competitive prices. Users often exaggerate the sizes of the Virtual Machines (VMs) they lease, either because the provider forces them to use predefined sizes, common practice, or to compensate for uncertainty. Hence, a provider could practice overbooking: An autonomic admission controller selects whether to accept a new user application or not, based on predicted resource utilization, which is likely smaller than the requested amount of resources. Overbooking is beneficial both to the provider, who can gain a competitive advantage and increase profits, and the user, who may observe lower prices. Although combining overbooking and brownout may seem straight-forward, the two approaches should not be used without thorough evaluation. Indeed, the two autonomic feedback loops, belonging to the brownout application and the overbooking provider, may take conflicting decisions, which may degrade performance. By contrast, if both approaches are effectively combined, the overbooking system may take advantage of the application performance knowledge from brownout, and use both reactive and proactive methods to avoid overload situations.

This algorithm first evaluates the risk associated to the new incoming request by calling the fuzzy risk assessment module. Once the associated risk is known, the admission control obtains the current (new) risk thresholds for the whole data center. Finally, it is checked, for each capacity dimension, if the risk of accepting the new incoming request is below the currently acceptable level and if so, the request is accepted. The process to calculate the service acceptance risk and the data center risk thresholds. The risk assessment module provides the Admission Control with the information needed to take the final decision of accepting or rejecting the service request, as a new request is only admitted if the final risk is bellow a pre-defined level (risk threshold). Calculating the risk of admitting a new service includes many uncertainties. Furthermore, choosing an acceptable risk threshold has an impact on data center utilization and performance. High thresholds result in higher utilization but the expense of exposing the system to performance degradation, whilst using lower values leads to lower but safer resource utilization. This method of choosing the representative risk thresholds for the data center balances utilization in all capacity dimensions. If capacity is imbalanced, e.g., CPU utilization is greater than memory; the admission control can act on this fact and admit applications that request more capacity of the type that is further from the target utilization level.

Results:

Efficient resource management in the virtualized data center is always a practical concern and has attracted significant attention. In particularly, economic allocation mechanism is desired to maximize the revenue for commercial cloud providers. This paper uses overbooking from Revenue Management to avoid resource over-provision according to its runtime demand. We propose an economic model to control the overbooking policy while provide users probability based performance guarantee using risk estimation. To cooperate with overbooking policy, we optimize the VM placement with traffic-aware strategy to satisfy application's QoS requirement. We design fuzzy assessment and algorithm to achieve traffic localization in order to reduce network bandwidth consumption, especially the network bottleneck bandwidth, thus to accept more requests and increase the revenue in the future. The simulation results show that our approach can greatly improve the request acceptance rate and increase the revenue by up to 87% while with acceptable resource confliction.

Conclusion:

Cloud computing allow customers to scale up and down their resource usage based on desires. In this project, we present a system that uses virtualization technology to allocate data center resources dynamically based on application demands and support green computing by optimizing the number of servers in use. And design SPRNT system to predict the loads that can capture the future resource usages of applications accurately without looking inside the VMs. The algorithm can capture the rising trend of resource usage patterns and help reduce the placement churn significantly. We have implemented the resource management concept in cloud computing in which we have reached the goal of achieving the overload avoidance and green computing concept successfully. And avoid performance attack at the time resource sharing. Performance degradation directly increases the cost of per workload completed in cloud-computing systems.

Future Work:

In future, we present a novel I/O workload based performance attack which uses a carefully designed workload to incur significant delay on a targeted application running in a separate VM but on the same physical system. Such a performance attack poses an especially serious threat to data intensive applications which require a large number of I/O requests. Performance degradation directly increases the cost of per workload completed in cloud-computing systems. In future we can extend our approach to provide improved security level based MAC address based authentication and implement various performance metrics to overcome various attacks in VM scheduling.

References:

1. V. Varadarajan et al., “Resource-freeing attacks: improve your cloud performance (at your neighbor’s expense),” in CCS, 2012, pp. 281–292.
2. S. K. Barker et al., “Empirical evaluation of latency-sensitive application performance in the cloud,” in Proceedings of the first annual ACM SIGMM conference on Multimedia systems. ACM, 2010, pp. 35–46.
3. K. Ye et al., “Virtual machine based energy-efficient data center architecture for cloud computing: a performance perspective,” in Proceedings of the 2010 IEEE/ACM Int’l Conference on Green Computing and Communications & Int’l Conference on Cyber, Physical and Social Computing. IEEE Computer Society, 2010, pp. 171–178.
4. J. Szefer et al., “Eliminating the hypervisor attack surface for a more secure cloud,” in Proceedings of the 18th ACM conference on Computer and communications security. ACM, 2011, pp. 401–412.
5. G. Wang et al., “The impact of virtualization on network performance of amazon ec2 data center,” in INFOCOM. IEEE, 2010, pp. 1–9.