



## A Survey over Virtualization and Hypervisor

**Kirti Raj Bhatele**

Department of Computer Science & Engineering,  
University Institute of Technology( UIT),  
RGPV, Bhopal (M.P.), [INDIA]  
Email : Kirtirajbhatele8@gmail.com

**Dr. Mukta Bhatele**

Department of Computer Science & Engineering,  
Jai Narain College of Technology  
RGPV, Bhopal (M.P.), [INDIA]  
Email : mukta\_bhatele@rediffmail.com

**Arundhati Arjaria**

CSE Department,  
University Institute of Technology( UIT),  
RGPV, Bhopal (M.P.), [INDIA]  
Email : Arundhati\_arjaria@yahoo.com

**Akhilesh Pahade**

CSE Department,  
University Institute of Technology( UIT),  
RGPV, Bhopal (M.P.), [INDIA]  
Email : akhi005@gmail.com

**Abstract:** This survey paper simply contains the recent developments in the field of Cloud computing. Cloud computing is an adoption and diffusion which are threatened by unresolved security issues that affect both the cloud provider and the cloud user. In particular, carried out a survey over Virtualization and Hypervisor Protection System (VHPS), aimed at guaranteeing increased security to cloud resources. Virtualization and Hypervisor Protection System (VHPS) can be deployed on several cloud solutions and can effectively monitor the integrity of guest and infrastructure components while at the same time being fully transparent to virtual machines and to cloud users. We simply present the survey over VHPS hypervisor security architecture and examine in detail, its mandatory access control architecture. While existing hypervisor security approaches aimed at high assurance have proven useful for high-security environments Virtual Machines (VM).

**Keywords:** Cloud Computing, Virtualization, Hypervisor, DDOS attacks.

### 1. INTRODUCTION

As general-purpose workstation and server-class computer systems increase in

available processing power and decrease in cost, it becomes cost-effective to aggregate the functionality of multiple standalone systems onto a single hardware platform. This minimizes costs for system management and maintenance and maximizes system utilization. Virtualization technology, which enables single system hardware to support multiple operating systems, is quickly becoming a commodity. This technology creates multiple virtual machines (VM) out of one real machine and carefully multiplexes multiple virtual resources onto a single real resource. The broad availability and use of virtualization technology is driven by improved hardware support, such as fully virtualizable CPUs and IO-MMU controlling direct memory access to devices, which enables very efficient implementation of virtual machines. Suddenly, multiple operating systems can be efficiently co-located inside virtual machines on a single general-purpose hardware platform. In addition to its availability, the potential impact of virtualization on workload consolidation and load balancing is getting the attention of key industry players. Microsoft has announced that their next generation security architecture NGSCB will be based on virtualized environments and Intel hopes to run home entertainment in virtualized environments, while large companies selling servers have

very successfully used virtualization for server consolidation, service provisioning and workload-balancing for decades. Although co-locating operating systems and their workloads on the same hardware platform offers great opportunities, it also requires us to carefully consider possible undesirable interactions between those systems sharing resources. Therefore, VMM environments by default do not allow to share real resources directly. Real system resources are virtualized by the hypervisor layer (e.g., memory, CPU) and can be accessed by VMs exclusively through their virtualized counterpart (e.g., virtual memory, virtual CPU). This hypervisor layer is strongly protected against the operating systems running in VMs on top of it and enforces isolation of these virtual resources. Peripherals, such as disk or network adapters, are exclusively assigned to a single VM. If necessary, such a VM can in turn virtualize its real resources to share it with other VMs (e.g., virtual disk server, VLAN). We will carefully examine under which conditions such VMs are allowed to share peripherals with other VMs without violating the isolation properties between VMs. Consequently, virtual machines that do not share virtual resources are considered isolated from each other.

There are currently at least two challenging security problems when broadly deploying virtualization technology. (1) The sharing of virtual resources among co-operating virtual machines is defined statically and resulting isolation properties of VMs are a side-effect of administration rather than of well-defined security management. However, today's environments depend more on sharing of resources and interconnection of workloads than ever before and this trend promises to increase. Consequently, there is need for an architecture that efficiently defines and enforces access control between related groups (coalitions) of virtual machines.(2) The isolation of virtual resources, while sufficient for commercial environments, is insufficient for high security environments where leaking even of very small amounts of data is unacceptable. Such leaks are introduced by

covert channels, which are based on observing system behavior (timing of events or storage patterns) rather than by explicit data sharing.

The first problem concerns the (explicit) sharing of virtual resources between VMs. On one hand, the current framework for controlling sharing is extremely static, offering only limited VM-isolation guarantees. Such guarantees are often a side-effect of a particular system configuration instead of a consciously architected and designed policy that can be reasoned about. On the other hand, co-operating workloads running in different virtual machines offer a unique opportunity to implement access control in the generic virtualization layer very efficiently. By enforcing access control in the self-protecting virtualization infrastructure, related access controls are protected against misbehavior of operating systems and workloads. The coarse-granular resources and VMs enable simple security policies that control their interactions. The second problem concerns

Covert channels. While controlling the explicit information flows between VMs is efficient, preventing implicit information flows comes at the cost of increased complexity, rewriting of hypervisor code, and decreased performance. These disadvantages of eliminating covert channels outweigh the interests of most customers. We believe, that the existing isolation of virtual resources is commercial-grade, meaning that controlling explicit data flows from one to another virtual machine and minimizing covert storage channels by careful resource management is sufficient in commercial environments. Our position is not to eliminate covert channels but (i) to minimize them through careful resource management, and (ii) to enable users through configuration options to mitigate remaining covert channels where necessary. To mitigate remaining covert channels, we introduce security rules guaranteeing that certain workloads never run on the same real platform; protection against covert channels between these workloads thus approximates the protection by air-gaps as they exist between

non-virtualized environments. The main focus in this project report is on the controlled sharing of resources, which is of broad interest in commercial environments. The sharing of virtual resources is currently not controlled by any formal policy. This makes it extremely difficult to measure the effectiveness of isolation between VMs and current approaches do not scale when considering the management of groups of systems and workload-balancing through VM migration.

We explore in this survey paper, the design and implementation of VHPS, a security architecture for virtualization environments, which leverages this virtualization layer to control the sharing of resources among VMs according to formal security policies. The major goals are (i) non-intrusiveness with regard to existing code, (ii) near-zero overhead on the performance-critical path, (iii) scalability regarding the management of many machines (simple policies) and the migration of VMs between them (machine independent policies). We implemented the core hypervisor security architecture (VHPS) into the Xen hypervisor [4] where it controls all inter-VM communication according to formal security policies. Our modifications to the Xen hypervisor are small and add about 2000 lines of code. The secure hypervisor architecture is designed to achieve medium assurance (Common Criteria) for hypervisor implementations. Our hypervisor security enhancement achieves near-zero overhead on the performance-critical path. While this project report describes VHPS for the Xen hypervisor, the presented architecture proves flexible; it was originally implemented for the rHype research hypervisor and is being implemented into the PHYP hypervisor.

## 2. LITERATURE SURVEY

Apart from security, there are reliability-related issues in virtualization [30] that can affect performance of cloud. For example, the provider may combine too many Virtual Machines onto a physical server. This can result in performance problems caused by impact factors such as limited CPU cycles or I/

O bottlenecks. These problems can occur in a traditional physical server, but they are more likely to occur in a virtualized server because of the connection of a single physical server to multiple Virtual Machines such that they all compete for critical resources. Thereby, management tasks such as performance management and capacity planning management are more critical in a virtualized environment than in a similar physical environment. This means that IT organizations must be able to continuously monitor the utilization of both physical servers and Virtual Machines in real time. This capability allows IT organizations to avoid both over- and underutilization of server resources such as CPU and memory and to allocate and reallocate resources based on changing business requirements. This capability also enables IT organizations to implement policy-based remediation that helps the organization to ensure that service levels are being met. Another challenge in Virtualization is that cloud organizations must now manage Virtual Machine sprawl. With Virtual Machine sprawl, the number of Virtual Machines running in a virtualized environment increases because of the creation of new Virtual Machines that are not necessary for business. Worries about Virtual Machine sprawl include the overuse of infrastructure. To prevent Virtual Machine sprawl, Virtual Machine managers should analyze the need for all new Virtual Machines carefully and ensure that unnecessary Virtual Machines migrate to other physical servers. In addition, an unnecessary virtual machine will be able to move from one physical server to another with high availability and energy efficiency. However, consider that it can be challenging to ensure that the migrated Virtual Machine keeps the same security, QoS configurations, and needed privacy policies. It must be ensured that the destination maintains all the required configurations of migrated Virtual Machines.

As mentioned before, there are at least two levels of virtualization, Virtual Machines and the hypervisor. Virtualization is not as new a technology as cloud, but it contains several

security issues that have now migrated to cloud technology. Also, there are other vulnerabilities and security issues which are unique to cloud environment or may have a more critical role in cloud.[11]

In a virtualization environment, there are several Virtual Machines that may have independent security zones which are not accessible from other virtual machines that have their own zones. A hypervisor has its own security zone, and it is the controlling agent for everything within the virtualization host. Hypervisor can touch and affect all acts of the virtual machines running within the virtualization host. There are multiple security zones, but these security zones exist within the same physical infrastructure that, in a more traditional sense, only exists within a single security zone. This can cause a security issue when an attacker takes control over the hypervisor. Then the attacker has full control over all data within the hypervisor's territory. Another major virtualization security concern is "escaping the Virtual Machine" [11] or the ability to reach the hypervisor from within the Virtual Machine level. This will be even more of a concern as more APIs are created for virtualization platforms As more APIs are created, so are controls to disable the functionality within a Virtual Machine that can reduce performance and availability.

#### ***Benefits and weakness of hypervisor-based systems:***

The hypervisor, apart from its ability to manage resources, has the potential to secure the infrastructure of cloud. Hypervisor-based virtualization technology is the best choice of implementing methods to achieve a secure cloud environment. The reasons for choosing this technology:

Hypervisor controls the hardware, and it is only way to access it. This capability allows hypervisor-based virtualization to have a secure infrastructure. Hypervisor can act as a firewall and will be able to prevent malicious users to from compromising the hardware infrastructure.

Hypervisor is implemented below the guest OS in the cloud computing hierarchy, which means that if an attack passes the security systems in the guest OS, the hypervisor can detect it.

The hypervisor is used as a layer of abstraction to isolate the virtual environment from the hardware underneath.

The hypervisor-level of virtualization controls all the access between the guests' OSs and the shared hardware underside. Therefore, hypervisor is able to simplify the transaction-monitoring

***Security management in hypervisor-based virtualization:*** As mentioned before, hypervisor is management tools and the main goal of creating this zone is building a trust zone around hardware and the VMs. Other available Virtual Machines are under the probation of the hypervisor, and they can rely on it, as users are trusting that administrators will do what they can to do provide security. There are three major levels in security management of hypervisor as mentioned below:

***Authentication:*** users must authenticate their account properly, using the appropriate, standard, and available mechanisms.

***Authorization:*** users must secure authorization and must have permission to do everything they try to do.

***Networking:*** the network must be designed using mechanisms that ensure secure connections with the management application, which is most likely located in a different security zone than the typical user. Authentication and Authorization are some of the most interesting auditing aspects of management because there are so many methods available to manage a virtual host auditing purpose. The general belief is that networking is the most important issue in the transaction between users and the hypervisor, but there is much more to virtualization security than just networking. But it is just as

important to understand the APIs and basic concepts of available hypervisor and virtual machines and how those management tools work. If security manager can address Authentication, Authorization, and Virtual Hardware and hypervisor security as well as networking security, cloud clients well on the way to a comprehensive security policy [6]. If a cloud provider at the virtualization level depends only on network security to perform these tasks, then the implemented virtual environment will be at risk. It is a waste of money if a cloud provider spends too much on creating a robust, secure network and neglects communication among virtual machines and the hypervisor.

**Traditional Intrusion Detection Techniques in VMs:** The IDSs can use in hypervisor level, because all the communication between the VMs and the hardware is under the control of hypervisor. If there is an IDS in the hypervisor, it can detect attacks better than the same IDS, running on the guest OS. The guest OS cannot monitor events in cloud, only events within its VM. However, it is possible for the guest OS to monitor VM events if the cloud provider performs this feature or if the cloud is IaaS. Using IDSs, the HIDS has more performance than the NIDS. However, there are direct attacks against the IDS, and if the attack succeeds, the whole cloud is at risk, because the attacker can access all the information that NIDS has gathered, which can include a lot of important and useful data about the cloud users. In addition, in the cloud environment, all the cloud users may prefer to use encryption methods to prevent access to their data. This causes NIDSs to become less effectiveness, because it can't probe information within cloud, due to the encryption. In addition, NIDS generally runs outside of the hypervisor in the individual VM, and the NIDS won't be able to access privileged data that is accessible only by the hypervisor in cloud technology. In traditional networks, this is achievable by NIDS, however. In addition, if the attacker is in the same cloud as his victim is, the NIDS is unable to detect him. It seems NIDS may be best solution for cloud environment but using

NIDS has serious problems that one of the main problems when using NIDS for monitoring is the encrypted data.

### 3. RELATED WORK

As general-purpose workstation- and server-class computer systems grow in available power and capability, it becomes more attractive to aggregate the functionality of multiple standalone systems onto a single hardware platform. For example, a small business that originally used three computer systems—perhaps to take customer orders using a web server front-end, a database server in the middle, and a file server back-end—can reduce the required physical space, configuration complexity, management complexity, and overall hardware cost by running all three applications on a single system. Taking this one step further, several small businesses could achieve an even lower-cost solution by contracting out the management of their respective business computing applications to a centralized server managed by a nonpartisan third party. This idea of virtualization of standalone computer systems on a single system has been around for decades, often being employed in “big iron” mainframe systems whose hardware was explicitly designed with virtualized operation in mind. However, until recently it has not been feasible to build systems out of commodity PC hardware that meet the security guarantees required by mutually distrusted parties— i.e., that the data and execution environment of one party's applications are securely isolated from those of a second party's applications. For example, such systems were often vulnerable to Direct Memory Access (DMA) attacks where one party's application could break isolation by issuing DMA instructions to effect a copy into or out of the memory used by the second party's applications. Such systems were vulnerable no matter what software mechanisms were used for isolation—whether the property was enforced by the operating system, or by a virtual machine monitor (VMM) controlling multiple virtual machines

(VMs). Emerging technology, such as the I/O-MMU, eliminates these previous limitations on isolation for commodity systems and makes it feasible to ensure a VMM can control all memory accesses, especially those between mutually distrusted parties. This development, combined with the inability to make definitive statements about resource sharing among heterogeneous and potentially mutually distrusted operating systems running as guests in VMs, motivates us to claim that VMMs will not only need to provide isolation, but also they will need to provide a basis for control of information flows and sharing of resources among VMs which was formerly expected of operating systems.

### ***Minimizing the hypervisor***

Work on minimizing hypervisors aims to reduce the amount of code within the hypervisor, which should translate to fewer bugs and vulnerabilities. One example is SecVisor, a hypervisor which supports a single guest VM and protects that VM from rootkits. Another example is TrustVisor which is a special-purpose hypervisor for protecting code and data integrity of selected portions of the application. Previous minimal hypervisors are not practical for deployment in the hosted cloud computing model where multiple VMs from multiple customers run on the same server. With VHPS we show how to remove attack vectors (in effect also reducing the hypervisor load) while still being able to support the hosted cloud computing model.

### ***Hardening the hypervisor***

Much of hypervisor-related work has centered around hardening of the hypervisor. Especially interesting is HyperSafe which aims to protect a hypervisor against control-flow hijacking attacks. They use a non-bypassable memory lockdown technique (only a special routine in the hypervisor can write to memory) coupled with a restricted pointer indexing technique (all function calls in the hypervisor are transformed to jumps from a special table). While making it more difficult to subvert the hypervisor, these additions add about a 5%

performance overhead and any bugs in the hypervisor could still be exploited through one of the attack vectors. Recently, HyperSentry used the SMM (system management mode) to bypass the hypervisor for integrity measurement purposes. Unfortunately, the integrity measurements only reveal traces of an attack after it has already happened and are limited to protecting against attacks which persistently modify the hypervisor executable. While the authors report being able to invoke the measurement every 8 seconds in HyperSentry, this still leaves a window for attackers. Furthermore, their approach results in a 2.4% overhead if HyperSentry protections are invoked every 8 seconds. In contrast, VHPS prevents the attacks from happening in the first place, and does this with about a 1% performance improvement.[7]

## **4. PROBLEM DESCRIPTION**

The problem we address in this project report is the design of a VMM reference monitor that enforces comprehensive, mandatory access control (MAC) policies on inter-VM operations. A reference monitor is defined to ensure mediation of all security sensitive operations, which enables a policy to authorize all such operations. A MAC policy is defined by system administrators to ensure that system (i.e., VMM) security goals are achieved regardless of system user (i.e., VM) actions. This contrasts with a discretionary access control (DAC) policy which enables users (and their programs) to grant rights to the objects that they own.

We apply the reference monitor to control all references of shared virtual resources by VMs and to allow coalitions of workloads to communicate or share efficiently within a coalition while efficiently confining workloads of different coalitions.

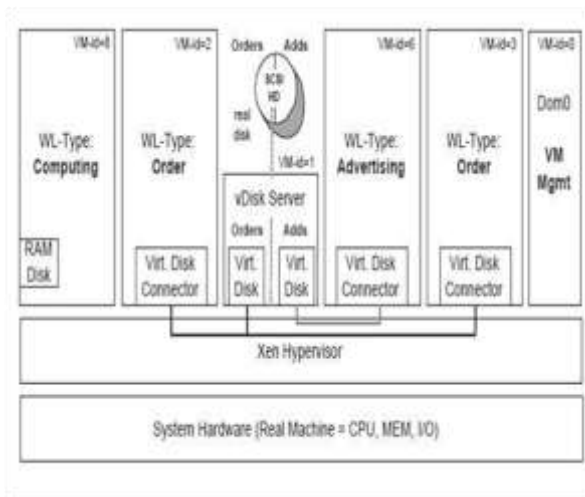


Figure 1: VM Coalitions and payloads in Xen

Above figure shows an example of VM coalitions. Domain 0 has started 5 user domains (VMs), which are distinguished inside the hypervisor by their domain id (VM-id in Figure 2). Domains 2 and 3 are running order workloads. Domain 6 is running an advertising workload, and domain 8 is running an unrelated generic computing workload (provisioned CPU time). Finally, domain 3 runs the virtual block device driver that offers two isolated virtual disks `vdiskorder` and `vdiskadds` to the Order coalition and the Advertising domain. In this example, we want to enable most efficient communication and sharing among VMs of the Order coalition but contain communication of VMs inside this coalition. For example, no VM with Order workload is allowed to communicate or share information with any VM running Computing or Advertising workloads and vice versa. While the hypervisor controls the ability of the VMs to connect to the device domain, device domain is trusted to keep data of different virtual disks securely isolated inside its VM and (on the real SCSI disk) and to assign them correctly to the coalitions. This is a reasonable requirement since device domains are not application specific and can run minimized run-time environments (e.g., micro-kernel).

### Threats and Attacks in Virtualization

**Threats:** In the hypervisor, all users see their systems as self-contained computers isolated from other users, even though every user is

served by the same machine. In this context, a Virtual Machine is an operating system that is managed by an underlying control program. [30]

**Virtual machine level attacks:** Potential vulnerabilities are the hypervisor or Virtual machine technology used by cloud vendors are a potential problem in multi-tenant architecture. These technologies involve "virtual Machines" remote versions of traditional on-site computer systems, including the hardware and operating system. The number of these virtual Machines can be expanded or contracted on the fly to meet demand, creating tremendous efficiencies.

**Cloud provider vulnerabilities:** These could be platform-level, such as an SQL-injection or cross-site scripting vulnerability that exist in cloud service layer which cause insecure environment.

**Expanded network attack surface:** The cloud user must protect the infrastructure used to connect and interact with the cloud, a task complicated by the cloud being outside the firewall in many cases.

**Authentication and Authorization:** The enterprise authentication and authorization framework does not naturally extend into the cloud. Enterprises have to merge cloud security policies with their own security metrics and policies.

**Lock-in:** It seems to be a lot of angst about lock-in in cloud computing. The cloud provider can encrypt user data in particular format and if user decides to migrate to another vendor or something like.

**Data control in cloud:** For midsize businesses used to having complete visibility and control over their entire IT portfolio, moving even some components into the Cloud can create operational "blind spots", with little advance warning of degraded or interrupted service.

**Communication in virtualization level:**

Virtual machines have to communicate and also share data with each other. If these communications didn't meet significant security parameters then they have potential of becoming attacks target.

**Attacks**

Nowadays, there are several attacks in the IT world. Basically, as the cloud can give service to legal users it can also service to users that have malicious purposes. A hacker can use a cloud to host a malicious application for achieve his object which may be a DDoS attacks against cloud itself or arranging another user in the cloud. For example an attacker knew that his victim is using cloud vendor with name X, now attacker by using similar cloud provider can sketch an attack against his victim (s). This situation is similar to this scenario that both attacker and victim are in same network but with this difference that they use virtual machines instead of physical network.

**DDoS attacks**

Distributed Denial of Service (DDoS) attacks typically focus high quantity of IP packets at specific network entry elements; usually any form of hardware that operates on a Blacklist pattern is quickly overrun. In cloud computing where infrastructure is shared by large number of VM clients, DDoS attacks make have the potential of having much greater impact than against single tenanted architectures. If cloud has not sufficient resource to provide services to its VMs then maybe cause undesirable DDoS attacks. Solution for this event is a traditional solution that is increase number of such critical resources. But serious problem is when a malicious user deliberately done a DDoS attacks using bot-nets.

It may be more accurate to say that DDoS protection is part of the Network Virtualization layer rather than Server Virtualization. For example, cloud systems use virtual machines can be overcome by ARP

spoofing at the network layer and it is really about how to layer security across multivendor networks, firewalls and load balances.

**Client to client attacks**

One malicious virtual machine could infect all Virtual Machines that exist in physical server. An attack on one client VM can escape to other VM's that hosted in the same physical, this is the biggest security risk in a virtualized environment. When malicious user puts the focus on virtual machines become easy to access, the attacker has to spend time attacking one virtual machine, which can lead to infecting other VMs, and thereby escaping the hypervisor and accessing the environment level that officially it can't accessible from VM level. Hence, the major security risk in virtualization environments is "client to client attacks".

**5. CONCLUSION**

This survey paper simply presents the recent developments and security related issues in the field of Virtualization and Hypervisor, so that protocols can be proposed to resolve them and ensure the quotient of security in the cloud computing environment.

**REFERENCES :**

- [1] L. Litty, "Hypervisor-based Intrusion Detection," M.S. thesis, Dept. Computer Science, University of Toronto, 2005.
- [2] G. Rowel, "Virtualization: The next generation of application delivery challenges," 2009.
- [3] G. Texiwill, "Is Network Security the Major Component of Virtualization Security?", 2009.
- [4] D. E. Y. Sarna, "Implementing and Developing Cloud Computing Applications: Taylor and Francis Group, LLC, 2011.
- [5] T. Ristenpart and e. al, "Hey, you, get



- off of my cloud: exploring information leakage in third-party compute clouds," presented at the 16th ACM conference on Computer and communications security, Chicago, IL, November 9-13, 2009.
- [6] "Securing Virtualization in Real-World Environments," White paper, 2009.
- [7] F. Sabahi, "Intrusion Detection Techniques performance in Cloud Environments " in Proc. Conf. on Computer Design and Engineering, Kuala Lumpur, Malaysia, 2011, pp. 398-402.
- [8] R. Chow, P. Golle, M. Jakobsson, E. Shi, J. Staddon, R. Masouka, and J. Molina, "Controlling Data in the Cloud: Outsourcing Computation without Outsourcing Control," presented at the ACM Cloud Computing Security Workshop, Chicago, Illinois, USA., 2009.
- [9] P. Sefton, "Privacy and data control in the era of cloud computing."
- [10] D. Rowe, "The Impact of Cloud on Mid-size Businesses," 2011.
- [11] C. Almond, "A Practical Guide to Cloud Computing Security," 2009. [12] F. Sabahi, "Security of Virtualization Level in Cloud Computing," in Proc. 4th Intl. Conf. on Computer Science and Information Technology, Chengdu, China, 2011, pp. 197-201.
- [12] P. R. Gallagher, A Guide to Understanding Data Remanence in Automated Information Systems: The Rainbow Books, ch.3 & ch.4, 1991.
- [13] Software Engineering Institute reports, N. Mead, E. Hough, and T. Sehny, "Security quality requirements engineering (SQUARE) methodolgy," Carnegie Mellon Software Engineering Institute, 2005.
- [14] K. K. Fletcher, "Cloud Security requirements analysis and security policy development using a high-order object-oriented modeling," M.S. thesis, Dept. Computer Science, Missouri Univ. of Science and Technology, Rolla, MS, 2010.
- [15] E. Keller, J. Szefer, J. Rexford, and R. B. Lee. NoHype: Virtualized cloud infrastructure without the virtualization. In International Symposium on Computer Architecture (ISCA), June 2010.
- [16] J. Kelsey, B. Schneier, D. Wagner, and C. Hall. Side channel cryptanalysis of product ciphers. In J.-J. Quisquater, Y. Deswarte, C. Meadows, and D. Gollmann, editors, Computer Security: ESORICS 98, volume 1485 of Lecture Notes in Computer Science, pages 97–110. 1998.
- [17] G. Klein, K. Elphinstone, G. Heiser, J. Andronick, D. Cock, P. Derrin, D. Elkaduwe, K. Engelhardt, R. Kolanski, M. Norrish, T. Sewell, H. Tuch, and S. Winwood. seL4: Formal verification of an OS kernel. In Symposium on Operating Systems Principles (SOSP), pages 207–220, October 2009.
- [18] T. Kooburat and M. Swift. The best of both worlds with on-demand virtualization. In Workshop on Hot Topics in Operating Systems (HotOS), May 2011.
- [19] M. A. Kozuch, M. Kaminsky, and M. P. Ryan. Migration without virtualization. In Workshop on Hot Topics in Operating Systems (HotOS), May 2009.
- [20] R. B. Lee, P. C. S. Kwan, J. P. McGregor, J. Dwoskin, and Z. Wang. Architecture for protecting critical secrets in microprocessors. In International Symposium on Computer Architecture (ISCA), June 2005.

- [21] I. Leslie, D. McAuley, R. Black, T. Roscoe, P. Barham, D. Evers, R. Fairbairns, and E. Hyden. The design and implementation of an operating system to support distributed multimedia applications. *IEEE Journal on Selected Areas in Communication*, 14(7), Sept. 1996.
  
- [22] C. Li, A. Raghunathan, and N. K. Jha. Secure virtual machine execution under an untrusted management OS. In *Proceedings of the Conference on Cloud Computing (CLOUD)*, July 2010.
  
- [23] D. Lie, C. Thekkath, M. Mitchell, P. Lincoln, D. Boneh, J. Mitchell, and M. Horowitz. Architectural support for copy and tamper resistant software. In *Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, November 2000.