

# Deep-diving into an easily-overlooked threat: Inter-VM attacks

Su Zhang

Kansas State University  
Manhattan, KS  
zhangs84@ksu.edu

**Abstract.** ‘Cloud computing’ has been embraced more than ever before due to a number of obvious advantages (*e.g.* elasticity, space-saving, hardware consolidation and cost-saving) over traditional computing infrastructure. Virtualization is the core part of cloud computing as it provides a multi-tenancy model which is the foundation of cloud computing. Typically there are two types of threats inherent with virtualization environments. One widely noticed threat is vm-escape which means a user on a virtual machine can ‘escape’ from it and take control over the whole hypervisor. The other threat known as inter-VM attacks have often been overlooked by the research community. In this paper we want to propose a framework which can systematically detect/analyze a number of inter-VM attacks. We want to build a system capturing such threats with a low false positive rate while providing valuable recommendations to system administrators.

**Keywords:** inter-VM attacks, Cloud Computing, Virtualization, Hypervisor Security

## 1 Introduction

Generally speaking, cloud services can be categorized into three models: Software as a service (SaaS), Platform as a service (PaaS) and Infrastructure as a service (IaaS). A cloud service becomes more vulnerable as more freedom is given to the users since any one of them may be an attacker. Furthermore, it is more error-prone for normal users to manage a whole virtual machine (IaaS) than just a single piece of software (SaaS). Therefore, security threats under IaaS model deserves more attention. There are a number of security issues that need to be considered under the IaaS model such as multi-jurisdictional issue, redundancy checking, privacy-preserving calculations, and hardening virtualization environment. However, it is the last one that requires the most attention because virtualization technology is the enabler and core component of cloud computing.

There are typically two types of threats brought by virtualization. One is called VM-escape, and the other is inter-VM attack. VM-escape means that attackers launch attacks from a virtual machine to compromise the hypervisor it is running on. This could allow attackers to monitor or compromise other co-residing virtual machines. VM-escape attacks are catastrophic threats to the cloud platform as the hypervisor is a single point of failure. A number of works [20, 19] have introduced a various of VM-escape vulnerabilities, even though no such attack in real world has been reported by any cloud service providers (CSPs). All VM-escape attacks are very sophisticated and ad-hoc while hard to detect. Consequently, hypervisor hardening has become a hot topic in research and several works [42, 32, 37, 33] have proposed techniques to realize this. inter-VM attacks refer to attacks launched from one virtual machine to another co-residing virtual machine directly, typically by bypassing the virtual machine monitor (VMM). This type of attack can be easier to launch and harder to detect. In this paper we focus on inter-VM attacks for the following reasons. Firstly, a certain number of inter-VM communication mechanisms [36, 17, 18] needs to bypass the VMM in order to achieve high performance. Since system administrators mostly monitor the VMM the VMM bypass attack may not be detected. Secondly, compromising a hypervisor would be more difficult than penetrating into individual virtual machines due to the fact that the hypervisor is typically well monitored and better hardened by people both in industry and in academia. Moreover, the VMM typically has a smaller attack surface than an ordinary operating system (*e.g.* Xen has only 200 KSLOC while Linux kernel has over 1,000 KSLOC). Thirdly, researcher have focused on hypervisor-security rather than inter-VM threats. Lastly, inter-VM threats are potentially more damaging than previously thought as once a virtual machine within a cluster (a group of machines communicating with each other through memory-sharing channel) is compromised, the rest of the machine may be more vulnerable. The rest of this report is structured as follows. Section 2 provides background about Xen – an open source hypervisor which will probably be used at the experimental section. Section 3 talks about related works including several state-of-arts inter-VM communication mechanisms, existing hardening techniques and several threat models. Section 4 asks research questions and illustrates a speculative framework to effectively detect and response to inter-VM attacks. Section 5 states the directions of future work.

## 2 Related works

A number of inter-VM communication mechanisms [56, 21, 44, 17, 34, 36, 43, 18] have been proposed. They mostly strive for a better performance (*e.g.* high throughput, low latency, *etc.*). They are mostly based on shared-memory channel (bypassing virtual network interface (VIF)) because there is a need for high efficiency inter-VM communication (services requested by customers may be provided by highly collaborative virtual machines like web service needs independent web server and database server both of which need to communicate with each other efficiently). Packets go through the VIF between different virtual machines on the same physical machine is one factor of magnitude slower than shared-memory-based approaches [56]. A number of summaries [14, 35] are in existence but they also stress more on performance than security. However, we want to focus on security evaluation over this new type of inter-VM communication mechanism (shared-memory-based channel). Performance penalty will be taken into consideration as well. Besides, there are tools detecting inter-VM threats, but they merely rely on low-level data (*e.g.* resource usage) to identify the threat other than utilize anything related to the new type of mechanisms.

### 2.1 State-of-the-art inter-VM communication mechanisms

**Xen Default Virtual Network Interface** The default virtual network interface (VIF) is a full emulation of a network stack as a normal physical machine. VM communicates with each other as if they are on different physical machines. Packets go through the whole network stack of one virtual machine and then cross another one in a reversed order. Even though this mechanism can provide a highest isolation, it performs poorly when bridging co-residing virtual machines. The reason why the virtual network interface has such performance limitation has been partly stated in [27]. The performance penalty couldn't be affordable if tasks count on intensive inter-VM communications. The inter-VM security level under this mechanism is the same as networked environment as all of the hosts are treated as fully-isolated physical machines.

**XenStore [12]** XenStore can be regarded as a database storing configuration information of each virtual machine (domain) running on the same hypervisor. It also can be used as an inter-domain communication channel, but the channel is narrow (the maximum size of each message is 1024 bytes). Message here typically refers to a variable-value pair. Therefore XenStore may mostly be used as a place for data collection (for the configuration information (resource (memory, CPU cycles) usage) of each virtual machine). VMs still need to count on other channels (memory-share based channel, IO buffer or VIF) exchanging large amount of data.

**XenLoop [36]** Mechanism: XenLoop relies upon a self-contained Linux module (netfilter) in each virtual machine. It can intercept and inspect each outgoing packets. If the destination host is a co-residing vm, then the packet will be delivered through a shared memory channel (two FIFO pipes (one for each direction allocated from receiver's memory space) and one event channel as a notifier to let each authenticated (through grant table references) end-point realize the data/space is ready to read/write in the pipe) bypass the VIF, otherwise the it will go through the default VIF.

Performance: This mechanism is user-transparent because even if the host is migrated to another physical machine, it is still able to use VIF to reach the other VMs on its previous physical machine. The two mechanisms are switchable which makes live-migration easier. However, the throughput and latency may not as good as other fully-mapped approaches (*e.g.* [4]) because it is limited by the size of the shared buffer.

Security Concern: A hypervisor-bypassing communication reduces the attack surface since the VMM is a single point of failure. However, it increases the chance of inter-VM attack (by letting the VM to manage its own grant tables). Along with the fact that shared-memory based communication doesn't check the content of exchanging memory. Therefore a client-side attack could be easily triggered since the shared-memory channel is a proper place for malicious code injection. Furthermore, DDoS attacks could be potentially launched. If attackers owning a number of VMs, then He can establish as many as number of services from the VM he has. The guests can only tear down the inactive communication (by checking the existence of other communication parties' tokens storing in XenStore) channels. However, there is no defined rules for user to identify if the other party connecting to its machine is a malicious or benign user. Current standard is far from enough from the security's perspective.). The FIFO (shared ring buffer) can be filled up easily by co-residing attackers.

**XenSocket [56]** Mechanism: XenSocket is another memory-sharing-based communication mechanism like XenLoop. Similar to XenLoop, it doesn't include per packet page-flipping in order to have a better performance (reduce the number of hypercalls which could generate lots of performance overhead). However, it only has one circular shared buffer rather than two (unlike XenLoop). So it is a unidirectional (rather than bidirectional) communication mechanism. It adapted a Unix socket-like channel to establish the connection. The server (receiver) first create a socket, allocate a share memory (including a descriptor page and a circular shared buffer), obtain the grant table reference to the descriptor page, create an event channel for event notifications between the two endpoints. At the client (sender) side, it connect to the server side through the function `connect()`, the function could also help it obtain the address of the physical pages of the shared circular buffer (the address was placed in shared address memory when `bind()` is called at the server's side, so that `bind()` could return the grant table reference address to the server). Meanwhile, the sender also hooks up with the event channel created by server and the very beginning for further events notification (*e.g.* arrival of new data or more space available due to the deletion of newly consumed data). In order to avoid DoS attack, resources (memory) must be shared by a lower-trustworthy party and mapped by a higher-trustworthy party. Therefore the server side must be lower than client side in terms of trustworthy level (However, this wouldn't guarantee a secure communication as the shared memory from the lower trustworthy party main contain malicious code which could put the higher trustworthy host in trouble once the connection established). Also in order to keep the synchronization feature between the two parties, the connection must be first terminated from the client side as the server side owned all of the resources (descriptor page, event channel and shared circular buffer). So if server side terminate first, parts of data in the air may have been lost, making the synchronization impossible. There is a functionality drawback of XenSocket since it only allows lower trusted host as server (receiver) while higher trusted host as client (sender). It is not convenient if two parties want to have a bidirectional communication. XenLoop resolved this issue gracefully.

**Fido [4]** Fido is especially designed for highly-collaborative VMs over the same hypervisor. It assumes that each virtual machine on the hypervisor is benign. Therefore all of the communication Virtual Machine's memory are readable to each other. They utilize the advantage of 64 bit addressing strategy to resolve the throughput limitation from circular shared producer and customer buffer. Because a host typically only uses a minor part of the spaces mappable by all of the 64-bit addresses. The initial memory mapping information is exchanged through bootstrap step via I/O ring (The circular producer and consumer buffer). Besides memory mapping, they have another two components included in their framework: connection module and signaling module. The connection module is managed through `xenStore` (a centralized key-value configuration database in `dom0`). `XenStore` makes the connection dynamically pluggable since communicating virtual machines synchronize the connection information through a connection token storing on `XenStore`. Once there is a change on the token, the communication VMs will be notified. The connection management is similar to XenLoop. The signaling module used here is the same as event channel in Xen, which is used as a notifier of an arrival of an event. The security concern of this work is about the assumption that each of the virtual machine on the hypervisor is trustworthy. This assumption may be true under certain circumstance. However, one physical server may hold several clusters of VMs. VMs within the same cluster are trustworthy to each other. But if a VM compromised by attackers outside of its cluster, then the rest of VMs in the same cluster will be completely exposed to the attackers. The data transfer instead of copying everything, they just share the pointers and size of each chunk of shared data (in source VM) to the destination VM. The author introduces a global virtual address by mapping all communication virtual machines into 64-bit addresses, so each virtual machine will have a global page table including all of the co-residing virtual machine's address.

**Inter-OS Communication on Highly Parallel Multi-Core Architectures [44]** The authors break the multi-trusted level rules among all of the virtual machines, instead, they treat all of the virtual machine equally by extending the functionality of grant table. Grant table is used only for memory shared based communication between privileged domain (`Dom0`) and normal user domain (`DomU`) for the data used by device driver. They allow arbitrary virtual machine within the same physical host to share resources with each other based on the grant table. However, this new settings would result in potential DoS attack if there is any user in a virtual machine is malicious. Even though this somehow increase the performance, it could only work for internal enterprise network while all of the users are trustworthy.

**A High-efficient Inter-Domain Data Transferring System for Virtual Machines [21]** The author combines several advantages of previous work and overcome a number of difficulties of them. It adapted bidirectional

communication mechanism like XenLoop (but unlike XenSocket). It is also a VIF-bypass mechanism. While this feature somehow increased the performance (if all of the traffics are belong to the virtual machine within the same physical server), it compromise the migration compatibility. Also it seems that dom0 is a via-point for all of the traffics which would potentially result in DoS attacks if there are malicious virtual host continuously establish connections to dom0. All of the rest components (grant table, event channel and shared buffer) utilized are the same as other popular communication mechanisms.

**Virtual machine aware communication libraries for high performance computing [17]** This work is based on a high performance computing library developed by the authors previously. Under high performance computing environment, processes need to cooperate while maintain a certain level of isolation. Therefore it would be good if the communication between processes on two different virtual machines could be comparable to the IPC within the same machine. The communication mechanism is a combination of message passing and shared-memory based inter-VM communication. The process will invoke a function within IVMC lib which will allocate memory pages as shared buffer pages. Meanwhile, it will utilize a socket-liked communication channel to notify the other end-point about the address of the shared memory, grant table reference and event channel port number. The rest of the inter-VM communication will be the same as the others. Their work could also avoid data lose during vm's migrating by copying the unread data into a cushion buffer before the terminate the connection. Another advantage of this work is VMM-bypass which dramatically reduced the performance overhead.

**Inter-domain Socket Communications Supporting High Performance and Full Binary Compatibility on Xen [18]** The advantages of this work including binary compatibility, bypassing VIF, avoiding page flipping, offering a direct, accelerated communication path between domains. Bypassing VIF could minimize the latency while save the CPU cycles. Avoid page-flipping could reduce the number of hypercalls which could take much time. Also, page table and TLB need to be flushed as the mapping between virtual and physical addresses changed. A virtual device – XWAY has been created for inter-domain data exchange. It manages XWAY channels, transfer data, and inform upper layers of any changes in the status of queues. XWAY channel is created while both endpoints agreed on tokens of shared memory, port number of event channel through another control channel across two arbitrary domains. The XWAY packets are TCP format compilable (with a XWAY specific extension appended to each TCP packet).

**VMware VMCI [34]** VMware adapted hardware-assisted virtual machine (HVM) mechanism other than the para-virtualization used by Xen. A device called VMCI needs to be installed. The application needs to be re-compiled against with existing communication libraries. Another disadvantage of it is it prohibits network connection therefore live-migration may be affected if a virtual machine is migrated to another physical host (As it may want to change the network mode to standard TCP/IP communication mechanism).

## 2.2 Threat modeling

There are a number of possible inter-VM attacks. Since not all of the co-residents are trustworthy to each other, we need to dig into the details of inter-VM attacks in order to know how could these attacks be detected or prevented.

**Denial of Service** Denial of service attack (DoS) is a catastrophic threat from the server's perspective. As we can see from the previous analysis, both privilege host (Dom0) and normal guest hosts (DomUs) are under the threat of such type of attacks due to the weak authentication with current communication mechanisms. Sometimes in order to pursue a higher degree of performance, system designer would rather sacrifice certain level of security. This maybe okay if the whole system is isolated and used by trustworthy users (*e.g.* for internal use of a corporation, or by a small community with none of them has conflict of interest.) Nowadays, virtualization is the core technology of cloud computing, and the tenants on the same physical server may have conflict of interest, or maybe there are malicious users as there are no strict authentication before the tenants are allowed to use the service. Therefore strong isolation, good performance and effective tools (can detect/analysis aforementioned threats) are of vital importance.

**Side-Channel Threat [47, 30]** There are several works related to side-channel attacks. Since more than one tenants could share a same set of hardware, therefore certain sensitive information (like the pattern of accessing data) may be leaked to malicious users. For example, encryption keys could be derived based on side (timing)-channels.

**Threat to Dormant VM** And even if VMs are dormant, attackers may still be able to access them. Because the data center is always on, an offline VM is not equivalent to an off-powered computer at home. Also pre-build images need to be carefully scanned in order to avoid some legacy vulnerable applications or trapdoors (*e.g.* AWS pre-build images store builder's SSH keys in it which means all of the hosts using such type of images are accessible from their publisher.).

### 2.3 Existing Protecting Works

There are existing hardening or detecting tools to identify virtualization-related threats. However, they are mostly focus on the hypervisor security. While the hypervisor is a single point of failure, it could be easily detected by system admin once it got compromised. Therefore even though inter-VM attack is not as bad as vm-escape, it is hard to be discovered from the system admin's position. As an attack may have already known which vm he want to attack, therefore he only need to target the potential virtual machine rather than the whole hypervisor in order to reduce the chances of being discovered and increase the chance of compromising the target successfully.

**Virtual Machine Monitor-Based Lightweight Intrusion [2]** The author build an intrusion detection system over the VMM. Based on the resource consumption information, they define several rules as the trigger of the alerts based on historical data. They trained the data by using data-mining algorithms and continuous applying new data with the model. Alerts will be raised once outlier is caught.

**Building a MAC-based security architecture for the Xen open-source hypervisor [31]** This framework is built upon VMM and incorporated with a security policy defined by the system administrator/user. It is a role-based access control for privileged actions, it will first trap each privileged action and check it against with security policy see if it is legal, illegal actions will trigger alerts. System admins can be ensured that each action conforms to the security policy as long as there is no alert raised. This can somehow guarantee each virtual machine can only access the resource it ought to.

**HomeAlone: Co-Residency Detection in the Cloud via Side-Channel Analysis [57]** Side-channel is regarded as a catastrophic threat as aforementioned. The author proposed a way of utilizing side-channel as a detector to identify illegal co-residency based on the timing channel (the response time of accessing L2 cache). They set up a threshold (response time in seconds). If the response time is longer than the threshold, then it'll be regarded that the cache (so do the other hardware on the same machine) has been used by other illegal co-resident, otherwise they conclude that the user is using the hardware resource exclusively. They evaluated their tool installed on different VMM environments with all true positive rate about 0.8.

**A Security-Aware Scheduler for Virtual Machines on IaaS Clouds [1]** This work tries to solve the conflict of interests among different VMs at a high-level. Users are allowed to provide a list of names including all of their conflict of interest parties. The scheduler (system administrator) will schedule all of these virtual machines based on their priority. The VM with higher priority could be scheduled right away while the VMs with lower priority and VMs with conflict of interests with the current running VMs will be scheduled later or migrate to other physical computers. This work stresses the importance of resolving potential cross-vm attacks at high level.

## 3 Framework

We here propose a framework to resolve the security issues raised under the memory sharing based inter-VM communication environment.

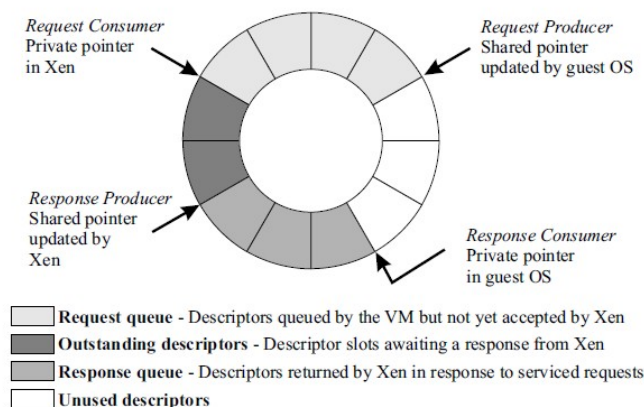
### 3.1 Objective

As mentioned earlier, there are a number of reasons why inter-VM threats deserve more attention. The first reason is because this type of attack is harder to be detected than vm-escape threat and system admin or defenders pay most of their attention over the hypervisor. Based on the related works, we can find out that VIF-bypassing memory-sharing based inter-VM communication will become mainstream mechanism as highly interactive VMs require better performances than current settings. However, this type of mechanism needs to be inspected carefully because co-residing VMs may not be trustworthy to each other.

### 3.2 Generalization over existing inter-VM communication mechanisms

We have seen that most state-of-the-art inter-VM communications include three essential components: shared memory buffer(s), grant table and event channel.

**shared memory** Shared memory is the actual communication channel between VMs. It is typically like a producer and consumer circular buffer shared between communication VMs. Generally it is a bidirectional asynchronous communication. VM can be both writer and reader. Once it finished writing, it will let the other party to know that the data is ready. If it finished reading, it will delete the data it just read on the shared buffer and let the other party know that more space is available for new input. Every time before the writer starts to write, it will check if the buffer has enough space available for the next input. If not enough available space, it will then wait for the reader to ‘consume’ the data on the buffer until the available space is enough. Therefore the size of the shared memory channel needs to be carefully designed in order to reduce latency. Figure 1 is a classical I/O shared buffer design cited from [3].



**Fig. 1.** The structure of asynchronous I/O rings

**Event Channel** Event channel is like a signal channel to let communication parties know the arrival of a new event. The writer tell the reader for the data it just wrote. On the other hand, once a reader finished reading it will delete the data and tell the writer that new space is ready for the next input through event channel. Therefore, event channel must be well protected, otherwise it will mess up the whole communication. Delivering wrong information may cause the share-memory channel out-of-sync.

**Grant table** Grant table provides two types of grants between different VMs. One is page-flipping and the other is page-sharing. Since per-packet page-flipping will have too much performance overhead (due to the high frequency of hypercalls). Therefore the new communications dropped page-flipping but preserve page-sharing grant. Xen memory sharing mechanism is at a page granularity. Shared pages are identified by an integer, known as a grant reference. The hypervisor keeps the grants information and pass the grant reference to communication

VMs and signal it via event channel. The hypervisor will be the authority to authenticate the communication parties. Under certain circumstance [36], the system may delegate communication parties to manage the grant table by themselves for performance reasons.

### 3.3 Speculative Framework

To systematically build a framework to detect or defend against such threats, we need to understand the whole system completely. Also we need to know what data needs to be collected for analysis and what correlations need to be built based on the possible threat models as we discussed earlier. We draw a draft framework ((see figure 2))to give a high level impression to the audience while make it extensible to more scenarios (threat models) discovered in the future.

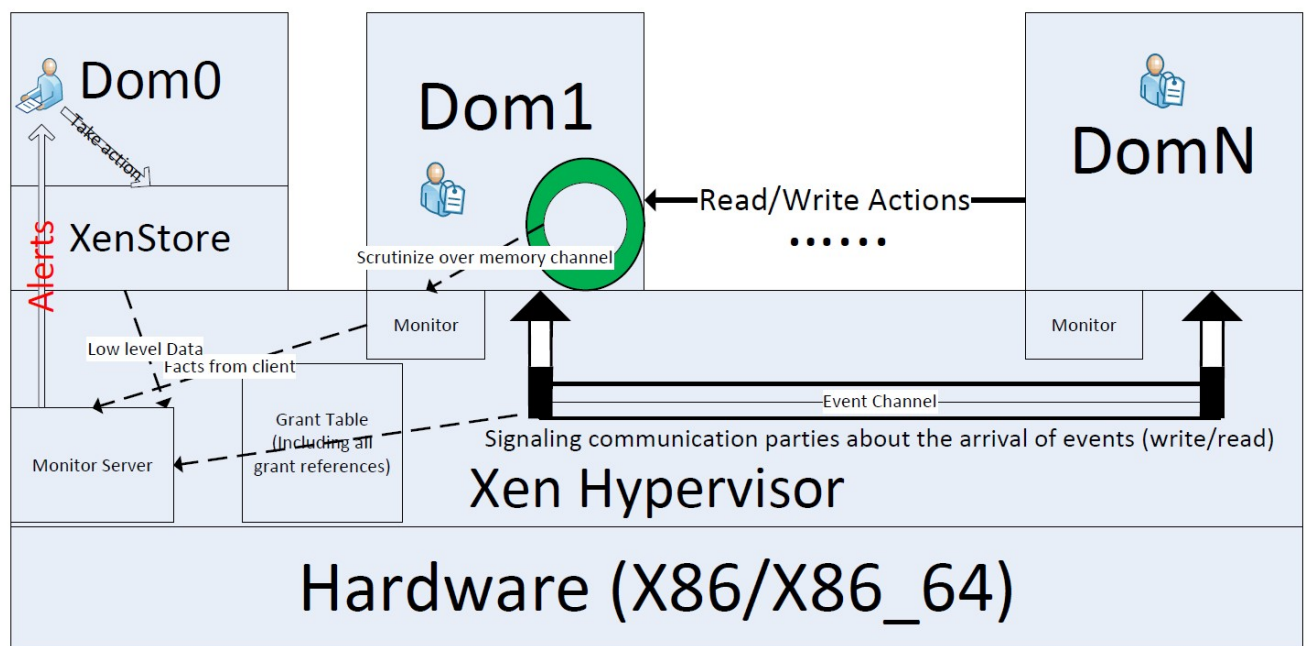


Fig. 2. Speculative Framework

**Data Collection** Data used for correlations could be obtained from various places like event channel, XenStore, grant table and shared memory channel. The event channel needs to be well monitored because once it couldn't function correctly, the shared memory channel may be out-of-sync, then the connection will be massed up completely. The information stored in grant table must be archived as well because these record can be used for future audit. This type of information is analogous to network log on apache server. The shared-memory channel needs to be monitored as it is like packets transmitted over network. For this new mechanism, there is no IPS or firewall residing between the communication VMs. However, the transmitted content must be scrutinized in order to secure the whole communication. XenStore can be used as a place for low level data collection (e.g. resource usage and configuration parameters of VMs).

**Alerts Correlation** Based on the data collected and threat models discussed, correlations can be built to raise alerts when anomalous situation detected. We are going to use a MulVAL [52, 53, 15, 16] like logical-based engine to analyze over the data and aim at raising alerts with low false positive rate. Furthermore, more scenarios (correlations) can be added easily to our engine therefore as more threat models discovered, we can make our tool more completed.

## 4 Conclusion and Future Work

This paper enumerates several inter-VM attack scenarios and proposes an approach of mitigating such attacks. It consists of the monitoring flow including how to collect data for alert correlation. In the future, we are going to implement this framework in conjunction with specific contexts. For example, advanced defenses [62, 61, 13, 6, 5, 39], risk assessment approaches [50, 49, 51, 7, 10], software dependency [55], SaaS [28, 45, 9, 29, 59], PaaS [58, 38], IaaS [54, 48, 41, 11, 8, 23, 22, 40, 46] and hardware [24–26, 60]

## References

1. Zaina Afoulki, Zaina Afoulki, and Jonathan Rouzaud-Cornabas. A security-aware scheduler for virtual machines on iaas clouds. Technical Report RR-2011-08, LIFO, ENSI de Bourges, August 2011.
2. Fatemeh Azmandian, Micha Moffie, Malak Alshawabkeh, Jennifer Dy, Javed Aslam, and David Kaeli. Virtual machine monitor-based lightweight intrusion detection. *ACM SIGOPS Operating Systems Review*, 45, July 2011.
3. Paul Barham, Boris Dragovic, Keir Fraser, Steven Hand, Tim Harris, Alex Ho, Rolf Neugebauer, Ian Pratt, and Andrew Warfield. Xen and the art of virtualization. In *the nineteenth ACM symposium on Operating systems principles*, New York, NY, December 2003.
4. A. Burtsev, K. Srinivasan, and P. Radhakrishnan. Fido: Fast inter-virtual-machine communication for enterprise applications. In *2009 USENIX*, San Diego, CA, June 2009.
5. Q. Chen, S. Abdelwahed, and A. Erradi. An autonomic detection and protection system for denial of service attack. IASTED/ACTA Press 2012, 2012.
6. Qian Chen, S. Abdelwahed, and A. Erradi. A model-based validated autonomic approach to self-protect computing systems. *Internet of Things Journal, IEEE*, 1(5):446–460, Oct 2014.
7. Qian Chen and Sherif Abdelwahed. Towards realizing self-protecting scada systems. In *Proceedings of the 9th Annual Cyber and Information Security Research Conference, CISR '14*, pages 105–108, New York, NY, USA, 2014. ACM.
8. Qian Chen, Sherif Abdelwahed, and Abdelkarim Erradi. A model-based approach to self-protection in computing system. In *Proceedings of the 2013 ACM Cloud and Autonomic Computing Conference, CAC '13*, pages 16:1–16:10, New York, NY, USA, 2013. ACM.
9. Qian Chen, Sherif Abdelwahed, and Weston Monceaux. Towards automatic security management: A model-based approach. In *Proceedings of the Eighth Annual Cyber Security and Information Intelligence Research Workshop, CSIIRW '13*, pages 56:1–56:4, New York, NY, USA, 2013. ACM.
10. Qian Chen, Robert K. Abercrombie, and Frederick T. Sheldon. Risk assessment for industrial control systems quantifying availability using mean failure cost (mfc). *Journal of Artificial Intelligence and Soft Computing Research*, 5(3):205–220, September 2015.
11. Qian Chen, R. Mehrotra, A. Dubey, S. Abdelwahed, and K. Rowland. On state of the art in virtual machine security. In *Southeastcon, 2012 Proceedings of IEEE*, pages 1–6, March 2012.
12. David Chisnall. *The Definitive Guide to the Xen Hypervisor*. Prentice Hall, 2007.
13. Scott A DeLoach, Xinming Ou, Rui Zhuang, and Su Zhang. Model-driven, moving-target defense for enterprise network security. In *Models@ run. time*, pages 137–161. Springer International Publishing, 2014.
14. C. Gebhardt and A. Tomlinson. Challenges for inter virtual machine communication. Technical Report Technical Report RHUL-MA-2010-12, Department of Mathematics Royal Holloway University of London, August 2010.
15. John Homer, Su Zhang, Xinming Ou, David Schmidt, Yanhui Du, S Raj Rajagopalan, and Anoop Singhal. Aggregating vulnerability metrics in enterprise networks using attack graphs. *Journal of Computer Security*, 21(4):561–597, 2013.
16. Heqing Huang, Su Zhang, Xinming Ou, Atul Prakash, and Kareem Sakallah. Distilling critical attack graph surface iteratively through minimum-cost sat solving. In *Proceedings of the 27th Annual Computer Security Applications Conference*, pages 31–40. ACM, 2011.
17. W. Huang, M. J. Koop, Q. Gao, and D. K. Panda. Virtual machine aware communication libraries for high performance computing. In *ACM/IEEE conference on Supercomputing*, New York, NY, November 2007.
18. K. Kim, C. Kim, S.-I. Jung, H.-S. Shin, , and J.-S. Kim. Inter-domain socket communications supporting high performance and full binary compatibility on xen. In *the fourth ACM SIGPLAN/SIGOPS international conference on Virtual execution environments*, New York, NY, March 2008.
19. KostyaKortchinsky. Cloudburst: A vmware guest to host escape, August 2009.
20. Invisible Things Lab. Vm-escape: Xen0wning trilogy, August 2008.
21. D. Li, H. Jin, Y. Shao, , and X. Liao. A high-efficient inter-domain data transferring system for virtual machines. In *the 3rd International Conference on Ubiquitous Information Management and Communication*, New York, NY, January 2009.
22. Tonglin Li, Xiaobing Zhou, Kevin Brandstatter, and Ioan Raicu. Distributed key-value store on hpc and cloud systems. In *2nd Greater Chicago Area System Research Workshop (GCASR)*. Citeseer, 2013.
23. Tonglin Li, Xiaobing Zhou, Kevin Brandstatter, Dongfang Zhao, Ke Wang, Adithya Rajendran, Zhao Zhang, and Ioan Raicu. Zht: A light-weight reliable persistent dynamic scalable zero-hop distributed hash table. In *Parallel & Distributed Processing (IPDPS), 2013 IEEE 27th International Symposium on*, pages 775–787. IEEE, 2013.



24. Chen Liu, J. Rajendran, Chengmo Yang, and R. Karri. Shielding heterogeneous mpsocs from untrustworthy 3pips through security-driven task scheduling. *Emerging Topics in Computing, IEEE Transactions on*, 2(4):461–472, Dec 2014.
25. Chen Liu and Chengmo Yang. Exploiting heterogeneity in mpsocs to prevent potential trojan propagation across malicious ips. In *Proceedings of the 24th Edition of the Great Lakes Symposium on VLSI, GLSVLSI '14*, pages 335–340, 2014.
26. Chen Liu and Chengmo Yang. Secure and durable (sedura): An integrated encryption and wear-leveling framework for pcm-based main memory. In *Proceedings of the 16th ACM SIGPLAN/SIGBED Conference on Languages, Compilers and Tools for Embedded Systems 2015 CD-ROM, LCTES'15*, pages 12:1–12:10, 2015.
27. A. Menon, J. R. Santos, Y. Turner, G. J. Janakiraman, and W. Zwaenepoel. Diagnosing performance overheads in the xen virtual machine environment. In *First International Conference on Virtual Execution Environments*, Chicago, IL, June 2005.
28. Liang Peng, Yimin Yang, Xiaojun Qi, and Haohong Wang. Highly accurate video object identification utilizing hint information. In *Computing, Networking and Communications (ICNC), 2014 International Conference on*, pages 317–321, Feb 2014.
29. Hao Qian and Daniel Andresen. Jade: Reducing energy consumption of android app. *the International Journal of Networked and Distributed Computing (IJNDC)*, Atlantis press, 2015.
30. Thomas Ristenpart, Eran Tromer, Hovav Shacham, and Stefan Savage. Hey, you, get off of my cloud! Exploring information leakage in third-party compute clouds. In Somesh Jha and Angelos Keromytis, editors, *Proceedings of CCS 2009*, pages 199–212. ACM Press, November 2009.
31. R. Sailer, T. Jaeger, E. Valdez, R. Caceres, R. Perez, S. Berger, J. Griffin, and L. van Doorn. Building a mac-based security architecture for the xen open-source hypervisor. In *Annual Computer Security Applications Conference (ACSAC)*, Washington, DC, December 2005.
32. Arvind Seshadri, Mark Luk, and Ning Qu Adrian Perrig. Secvisor: a tiny hypervisor to provide lifetime kernel code integrity for commodity oses. In *twenty-first ACM SIGOPS symposium on Operating systems principles (SOSP)*, New York, NY, December 2007.
33. J. Szefer, E. Keller, R. B. Lee, and J. Rexford. Eliminating the hypervisor attack surface for a more secure cloud. In *the 18th ACM conference on Computer and communications security (CCS)*, Chicago, IL, October 2011.
34. VMware. Vmci sockets programming. [www.vmware.com/pdf/ws65\\_s2\\_vmci\\_sockets.pdf](http://www.vmware.com/pdf/ws65_s2_vmci_sockets.pdf), 2008.
35. J. Wang. Survey of state-of-the-art in inter-vm communication mechanisms. Technical Report Research Proficiency Report, Binghamton University, September 2009.
36. J. Wang, K.-L. Wright, and K. Gopalan. Xenloop: a transparent high performance inter-vm network loopback. In *the 17th international symposium on High performance distributed computing*, New York, NY, June 2008.
37. Jiang Wang, Angelos Stavrou, and Anup K. Ghosh. Hypercheck: A hardware-assisted integrity monitor. In *the 13th International Symposium on Recent Advances in Intrusion Detection*, Ottawa, Canada, September 2010.
38. Ke Wang, Xiaobing Zhou, Hao Chen, Michael Lang, and Ioan Raicu. Next generation job management systems for extreme-scale ensemble computing. In *Proceedings of the 23rd international symposium on High-performance parallel and distributed computing*, pages 111–114. ACM, 2014.
39. Ke Wang, Xiaobing Zhou, Kan Qiao, Michael Lang, Benjamin McClelland, and Ioan Raicu. Towards scalable distributed workload manager with monitoring-based weakly consistent resource stealing. In *Proceedings of the 24rd international symposium on High-performance parallel and distributed computing*, pages 219–222. ACM, 2015.
40. Ke Wang, Xraobing Zhou, Tonglin Li, Dongfang Zhao, Michael Lang, and Ioan Raicu. Optimizing load balancing and data-locality with data-aware scheduling. In *Big Data (Big Data), 2014 IEEE International Conference on*, pages 119–128. IEEE, 2014.
41. Xiaoguang Wang, Yue Chen, Zhi Wang, Yong Qi, and Yajin Zhou. Secpod: a framework for virtualization-based security systems. In *2015 USENIX Annual Technical Conference (USENIX ATC 15)*, pages 347–360, 2015.
42. Zhi Wang and Xuxian Jiang. A lightweight approach to provide lifetime hypervisor control-flow integrity. In *the 2010 IEEE Symposium on Security and Privacy (Oakland)*, Oakland, CA, May 2010.
43. M. Wegiel and C. Krintz. Xmem: type-safe, transparent, shared memory for cross-runtime communication and coordination. In *the 2008 ACM SIGPLAN conference on Programming language design and implementation*, New York, NY, June 2008.
44. Lamia Youseff, Dmitrii Zagorodnov, and Rich Wolski. Inter-os communication on highly parallel multi-core architectures. Technical Report TR 2008-17, UCSB, October 2008.
45. Jie Yu, Sandra Skaff, Liang Peng, and Francisco Imai. Leveraging knowledge-based inference for material classification. In *Proceedings of the 23rd Annual ACM Conference on Multimedia Conference*, pages 1243–1246. ACM, 2015.
46. Jiaan Zeng, Guangchen Ruan, Alexander Crowell, Atul Prakash, and Beth Plale. Cloud computing data capsules for non-consumptiveuse of texts. In *Proceedings of the 5th ACM Workshop on Scientific Cloud Computing*, ScienceCloud '14, pages 9–16, New York, NY, USA, 2014. ACM.
47. Kehuan Zhang, Zhou Li, Rui Wang, XiaoFeng Wang, and Shuo Chen. Sidebuster: Automated detection and quantification of side-channel leaks in web application development. In *Proceedings of CCS 2010*. ACM Press, October 2010.
48. Su Zhang. Deep-diving into an easily-overlooked threat: Inter-vm attacks. *Whitepaper, provided by Kansas State University, TechRepublic/US2012*, available online: <http://www.techrepublic.com/resourcelibrary/whitepapers/deep-diving-into-an-easilyoverlooked-threat-inter-vm-attacks>, 2013.

49. Su Zhang. *QUANTITATIVE RISK ASSESSMENT UNDER MULTI-CONTEXT ENVIRONMENTS*. PhD thesis, Kansas State University, 2014.
50. Su Zhang, Doina Caragea, and Xinming Ou. An empirical study on using the national vulnerability database to predict software vulnerabilities. In *Database and Expert Systems Applications*, pages 217–231. Springer, 2011.
51. Su Zhang, Xinming Ou, and Doina Caragea. Predicting cyber risks through national vulnerability database. *Information Security Journal: A Global Perspective*, 0(0):1–13, 0.
52. Su Zhang, Xinming Ou, and John Homer. Effective network vulnerability assessment through model abstraction. In *Detection of Intrusions and Malware, and Vulnerability Assessment*, pages 17–34. Springer, 2011.
53. Su Zhang, Xinming Ou, Anoop Singhal, and John Homer. An empirical study of a vulnerability metric aggregation method. In *The 2011 International Conference on Security and Management (SAM'11), special track on Mission Assurance and Critical Infrastructure Protection (STMACIP'11)*, 2011.
54. Su Zhang, Xinwen Zhang, and Xinming Ou. After we knew it: empirical study and modeling of cost-effectiveness of exploiting prevalent known vulnerabilities across iaas cloud. In *Proceedings of the 9th ACM symposium on Information, computer and communications security*, pages 317–328. ACM, 2014.
55. Su Zhang, Xinwen Zhang, Xinming Ou, Liqun Chen, Nigel Edwards, and Jing Jin. Assessing attack surface with component-based package dependency. In *Network and System Security*, pages 405–417. Springer International Publishing, 2015.
56. X. Zhang, S. McIntosh, P. Rohatgi, , and J. L. Griffin. Xensocket: A high-throughput interdomain transport for virtual machines. In *the ACM IFIP USENIX 2007 International Conference on Middleware.*, New York, NY, November 2007. Springer-Verlag.
57. Y. Zhang, A. Juels, A. Oprea, and M. Reiter. Homealone: Co-residency detection in the cloud via side-channel analysis. In *the 2011 IEEE Symposium on Security and Privacy (Oakland)*, Oakland, CA, May 2011.
58. Dongfang Zhao, Zhao Zhang, Xiaobing Zhou, Tonglin Li, Ke Wang, Dries Kimpe, Philip Carns, Robert Ross, and Ioan Raicu. Fusionfs: Toward supporting data-intensive scientific applications on extreme-scale high-performance computing systems. In *Big Data (Big Data), 2014 IEEE International Conference on*, pages 61–70. IEEE, 2014.
59. Cong Zheng, Shixiong Zhu, Shuaifu Dai, Guofei Gu, Xiaorui Gong, Xinhui Han, and Wei Zou. Smartdroid: An automatic system for revealing ui-based trigger conditions in android applications. In *Proceedings of the Second ACM Workshop on Security and Privacy in Smartphones and Mobile Devices, SPSM '12*, pages 93–104, New York, NY, USA, 2012. ACM.
60. Yajin Zhou, Xiaoguang Wang, Yue Chen, and Zhi Wang. Armlock: Hardware-based fault isolation for arm. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, pages 558–569. ACM, 2014.
61. Rui Zhuang, Su Zhang, Alex Bardas, Scott A DeLoach, Xinming Ou, and Anoop Singhal. Investigating the application of moving target defenses to network security. In *Resilient Control Systems (ISRCs), 2013 6th International Symposium on*, pages 162–169. IEEE, 2013.
62. Rui Zhuang, Su Zhang, Scott A DeLoach, Xinming Ou, and Anoop Singhal. Simulation-based approaches to studying effectiveness of moving-target network defense. In *National Symposium on Moving Target Research*, 2012.