

Application performance and resource consumption in virtualized environments might be quite different compared to their execution on bare metal hardware because of additional virtualization overheads, which are typically caused by I/O processing and the application interactions with the underlying virtual machine monitor (VMM). Different papers describe various VMM implementations and analyze virtualization overhead when executing specially selected microbenchmarks or macrobenchmarks (e.g., [6, 8, 12, 22]). The reported virtualization overhead greatly depends on the server hardware used in such experiments. This extensive body of previous benchmarks characterizes performance in virtualized environments from a very different angle compared to the goals and functionality of *HT-vmbench*.

Nested Virtualization Technique: During the last decade software virtualization solutions for x86 systems were broadly adopted, forcing both Intel and AMD to add virtualization extensions to their x86 platforms [5, 17]. There was a stream of efforts to incorporate nested virtualization in Xen hypervisor [4, 10]. Nested virtualization has many potential uses: e.g., platforms with hypervisors embedded in firmware might need to support other hypervisors as guest virtual machines. In the Cloud, IaaS providers might offer a user the ability to run the user-controlled hypervisor as a VM. In such a way, the user can manage his own virtual environment with the choice of his favorite hypervisor. This might significantly simplify many management tasks, such as the live migration of their virtual machines as a single entity, e.g., for disaster recovery or load balancing. It could also be used for testing, demonstrating, benchmarking and debugging hypervisors and virtualization setups.

Nested virtualization enables new approaches to security in virtualized environments, such as honeypots capable of running hypervisor-level rootkits [14], hypervisor-level rootkit protection [13, 15], hypervisor-level intrusion detection [9, 11] for both hypervisors and operating systems. Nested virtualization is a foundation of the AERIE reference architecture [16]: it supports a set of components and services in a managed platform to reduce the level of trust required for IaaS providers. It helps to increase control and isolation and improve the system security and data protection.

In our work, we applied nested virtualization for creating a large scale virtualized environment using a limited number of physical servers to perform scalability assessment of security management solution (HTCC). This approach is of interest to many startups, small companies, and research organizations, which might not have access to a production size virtual environment needed for their scalability studies and performance experiments.

7 CONCLUSIONS

Engineering teams face many challenges when they implement new management and security solutions in large-scale virtual environment. They need to assess performance and scalability of such management solutions, analyze their performance overheads, and perform solution's capacity planning and resource sizing. In this paper, we introduce a novel approach for accomplishing these performance goals. We offer an extensible benchmark, called *HT-vmbench*, which allows users to mimic *session-based* activities of system administrators in virtualized environments. To perform scalability studies with *HT-vmbench*, the users need access to large-scale testbeds (that mimic the production virtual environments). To solve this challenge, we describe and promote an approach, based on a nested virtualization technique, which enables us to create

a large scale virtualized environment (with 30,000 VMs) using a limited number of physical servers (4 servers in our experiments).

We believe that more interesting opportunities are available for constructing specialized virtual environment using this approach. Combined with an extensible nature of *HT-vmbench*, the proposed framework offers a powerful solution for performance assessment of different management and security solutions in large-scale virtual environments.

8 ACKNOWLEDGMENT

The research presented in the paper has been partially supported by NSF grant CCF-1649087.

REFERENCES

- [1] HyTrust Inc. Cloud Control Virtual and Private Cloud Security. <https://www.hytrust.com/products/cloudcontrol/>.
- [2] HyTrust Inc. HyTrust Cloud Control: Security, Compliance and Control for Virtual Infrastructure. <https://www.hytrust.com/uploads/HyTrust-CloudControl.pdf>.
- [3] HyTrust Inc. Protecting Sensitive Data and achieving compliance in a multi-cloud world. https://www.hytrust.com/uploads/Compliance-in-a-Multi-Cloud-World_WP.pdf.
- [4] Nested Virtualization on Xen. https://wiki.xenproject.org/wiki/Nested_Virtualization_in_Xen.
- [5] AMD. Secure Virtual Machine Architecture Reference Manual. <https://www.mimuw.edu.pl/~vincent/lecture6/sources/amd-pacifica-specification.pdf>.
- [6] P. Barham, B. Dragovic, K. Fraser, S. Hand, T. Harris, A. Ho, R. Neugebauer, I. Pratt, and A. Warfield. Xen and the Art of Virtualization. In *Proceedings of the Nineteenth ACM Symposium on Operating Systems Principles, SOSP '03*, 2003.
- [7] M. Ben-Yehuda, M. D. Day, Z. Dubitzky, M. Factor, N. Har'El, A. Gordon, A. Liguori, O. Wasserman, and B.-A. Yassour. The Turtles Project: Design and Implementation of Nested Virtualization. In *Proceedings of the 9th USENIX Conference on Operating Systems Design and Implementation, OSDI'10*, 2010.
- [8] L. Cherkasova and R. Gardner. Measuring CPU Overhead for I/O Processing in the Xen Virtual Machine Monitor. In *Proceedings of the Annual Conference on USENIX Annual Technical Conference, ATEC '05*, 2005.
- [9] T. Garfinkel and M. Rosenblum. A Virtual Machine Introspection Based Architecture for Intrusion Detection. In *Proceedings of Network and Distributed Systems Security Symposium*, 2003.
- [10] Q. He. Nested Virtualization on Xen. In *Proceedings of Xen Summit Asia, 2009*.
- [11] J.-C. HUANG, M. MONCHIERO, and Y. TURNER. Ally: OS-Transparent Packet Inspection Using Sequestered Cores. In *Proceedings of Seventh ACM/IEEE Symposium on Architectures for Networking and Communications Systems (ANCS)*, 2011.
- [12] S. T. King, G. W. Dunlap, and P. M. Chen. Operating System Support for Virtual Machines. In *Proceedings of the Annual Conference on USENIX Annual Technical Conference, ATEC '03*, 2003.
- [13] R. Riley, X. Jiang, and D. Xu. Guest-Transparent Prevention of Kernel Rootkits with VMM-Based Memory Shadowing. In *Proceedings of the 11th International Symposium on Recent Advances in Intrusion Detection, RAID '08*, 2008.
- [14] J. Rutkowska. Subverting Vista Kernel for Fun and Profit. In *Proceedings of SyScan'06 and Black Hat Briefings*, Aug, 2006.
- [15] A. Seshadri, M. Luk, N. Qu, and A. Perrig. SecVisor: A Tiny Hypervisor to Provide Lifetime Kernel Code Integrity for Commodity OSES. In *Proceedings of Twenty-first ACM SIGOPS Symposium on Operating Systems Principles, SOSP '07*, 2007.
- [16] M. Shtern, B. Simmons, M. Smit, and M. Litoiu. An Architecture for Overlaying Private Clouds on Public Providers. In *Proceedings of 8th International Conference on Network and Service Management (CNSM)*, 2012.
- [17] L. Smith, A. Kagi, F. C. Martins, G. Neiger, F. Leung, D. Rodgers, A. Santoni, S. Bennett, R. Uhlig, and A. Anderson. Intel virtualization technology. *Computer*, 38, 2005.
- [18] VMware. Running Nested VMs | VMware Communities. <https://communities.vmware.com/docs/DOC-8970>.
- [19] VMware. VMware vSphere 6.5 Nested Virtualization. <http://jermismit.com/vmware-vsphere-6-5-nested-virtualization-create-and-install-esxi-6-5/>.
- [20] VMware. VMware vSphere Web Services SDK Documentation. <https://www.vmware.com/support/developer/vc-sdk/>.
- [21] VMware. VMmark Virtualization Benchmark. <http://www.vmware.com/products/vmmark.html>.
- [22] T. Wood, L. Cherkasova, K. Ozonat, and P. Shenoy. Profiling and Modeling Resource Usage of Virtualized Applications. In *Proceedings of the 9th ACM/IFIP/USENIX International Conference on Middleware, Middleware '08*, 2008.
- [23] L. Yang, L. Cherkasova, R. Badgajar, J. Blancaflor, R. Konde, J. Mills, and E. Smirni. Evaluating Scalability and Performance of a Security Management Solution in Large Virtualized Environments.