

Experimental Systems Research on Virtual Machine Migration and Ecosystem of Hypervisors

Kartik Gopalan and Ping Yang

State University of New York at Binghamton

{kartik,pyang}@binghamton.edu, <http://osnet.cs.binghamton.edu>

Our research group at Binghamton University investigates virtualization technologies for cloud infrastructures. We conduct experimental systems research on a range of problems related to new functionality, better performance, and higher security/privacy in the cloud. Below we describe some of our current projects and how they can benefit from the NSF cloud testbed platforms.

Scalable and Efficient Live Migration of Virtual Machines

State-of-the-art approaches for live migration of VMs are considered too heavyweight for use in production clouds. Our work aims to remove the obstacles to wider adoption of live VM migration. We are developing new techniques for scalable, lightweight, and agile migration of VMs in data center environments to achieve higher consolidation, energy savings, and faster response to scale-out workloads. Our recent work [1][2] on mass migration of multiple virtual machines within data centers — also called *Gang Migration* — reduces the network overhead of simultaneously migrating a large number of VMs. Our group was the first to propose [3] the notion of *Post-copy Live Migration* of VMs to reduce their total migration time and network overhead. We have also developed a new technique called *Scatter-Gather live migration* [4] for faster de-provisioning of servers using intermediate staging nodes. Finally, we are investigating ways use virtualized cluster-wide memory for greater flexibility and agility in live VM migration.

An Ecosystem/Marketplace of Hypervisors-level Services in the Cloud

Currently, virtual machines are restricted to run on a single hypervisor. This limits the VMs to use only the “baked-in” features exposed by a single hypervisor. Conversely, a single hypervisor platform may be unable (or its developers unwilling) to support an ever-growing list of hypervisor-level features proposed by the research community. Our group (in collaboration with IBM Research Labs) has proposed and prototyped the systems support for *Multi-Hypervisor Virtual Machines* [5], i.e. VMs that can simultaneously run on multiple co-located hypervisors. Our prototype, called *Span*, can allow cloud providers to enable a “marketplace of hypervisors” on their cloud platform wherein third parties could leverage nested virtualization to support a diverse array hypervisor-level features on a minimal base hypervisor. *Span* VMs can simultaneously run on multiple hypervisors and benefit from a variety of features such as VM introspection, live guest patching, network monitoring, or high availability. We are also investigating new hypervisor-level services to improve security and privacy in virtualized clouds. These deal with services such as privacy-preserving VM checkpointing [6][7] to limit the lifetime of confidential data in the cloud and support for better isolation and access control among co-located VM having different trust levels.

Experimental Needs from NSF Cloud Testbed

The Chameleon and CloudLab testbeds will enable us to investigate our solutions at larger scales and in settings more representative of real-world infrastructures. Availability of specialized platforms with hardware support for virtualization will allow us to address newer problems that we cannot investigate with our current in-house testbed besides being able to compare our solutions on heterogeneous hardware platforms.

The facility to fully customize the software stack to our needs is critical in that it will allow us to experiment with novel hypervisor-level and kernel-level solutions that existing public testbeds do not allow. Access to cloud usage and performance data from the testbed will be valuable in evaluating our solutions against realistic workload traces.

The scale and effectiveness of our new VM migration approaches described above can be better demonstrated using large number nodes having customizable software stack. One particular feature that would be valuable for our VM migration experiments would be the ability to customize the cluster topology, possibly using low-overhead SDNs. This would allow us to demonstrate the effectiveness of our approaches over a variety of datacenter topologies and workloads.

Our Span VM technology can potentially complement the NSF Cloud testbed to support, in a modular fashion, a variety of features needed by users of the testbed such VM monitoring, benchmarking, checkpointing, and mirroring, without having to build these features statically into the base hypervisor. Span VMs also provide the ability to insert and remove hypervisors from underneath a running VM (so-called *ephemeral* hypervisors), thus minimizing any runtime overhead of these hypervisor-level services. The ability to create custom clouds using customized software stack will allow us to demonstrate the feasibility and usability of new cloud architectures, such as an eco-system of hypervisors that support Span VMs.

References:

- [1] Umesh Deshpande, Brandon Schlinker, Eitan Adler, and Kartik Gopalan, Gang Migration of Virtual Machines using Cluster-wide Deduplication, In 13th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid), Delft, The Netherlands, May 2013.
- [2] Umesh Deshpande, Xiaoshuang Wang, and Kartik Gopalan, Live Gang Migration of Virtual Machines. Proceedings of the 20th international symposium on High Performance Parallel and Distributed computing (HPDC), San Jose, CA, June 2011.
- [3] Michael Hines, Umesh Deshpande, and Kartik Gopalan, Post-Copy Live Migration of Virtual Machines, In SIGOPS Operating Systems Review, Volume 43, Number 3, pages 14--26, 2009.
- [4] Umesh Deshpande, Yang You, Danny Chan, Nilton Bila, Kartik Gopalan, Fast Server Deprovisioning through Scatter-Gather Live Migration of Virtual Machines, In IEEE Cloud 2014 (Research Track), Anchorage, Alaska.
- [5] Yaohui Hu, Siddhesh Phadkhe, Kartik Gopalan, Michael Hines, Multi-hypervisor Nested Virtual Machines, Poster in USENIX Annual Technical Conference 2014, Philadelphia, PA, USA.
- [6] Yaohui Hu, Tianlin Li, Ping Yang, and Kartik Gopalan, An Application-Level Approach for Privacy-preserving Virtual Machine Checkpointing, In the 6th IEEE International Conference on Cloud Computing, Research Track, 2013.
- [7] Mikhail Gofman, Ruiqi Luo, Ping Yang, and Kartik Gopalan, SPARC: A Security and Privacy Aware Virtual Machine Checkpointing Mechanism, Proc. of the 10th annual ACM Workshop on Privacy in the Electronic Society (WPES), In conjunction with the ACM Conference on Computer and Communications Security (CCS), full paper, Chicago, IL, 2011.