

VIRTUALIZING FILE TRANSFER AGENTS FOR INCREASED THROUGHPUT ON A SINGLE HOST

Amanda Bonnie*, Zach Fuerst†, Thomas Stitt^Δ
Mentors: Brad Settlemyer*, Christopher Mitchell*, Michael Mason*
*Los Alamos National Lab, †Dakota State University, ^ΔPennsylvania State University

OVERVIEW

Single Lustre File Transfer Agent (FTA) performance is known to underutilize the bandwidth of Infiniband (IB) cards. The utility and viability of multiple FTA Virtual Machines (VMs) for improved network throughput on a single host was investigated. It was proposed that having multiple VMs on a single host will help achieve better usage of the IB card. Single Root - I/O Virtualization (SR-IOV) was configured so that the IB card could be divided up among the VMs.

RESULTS

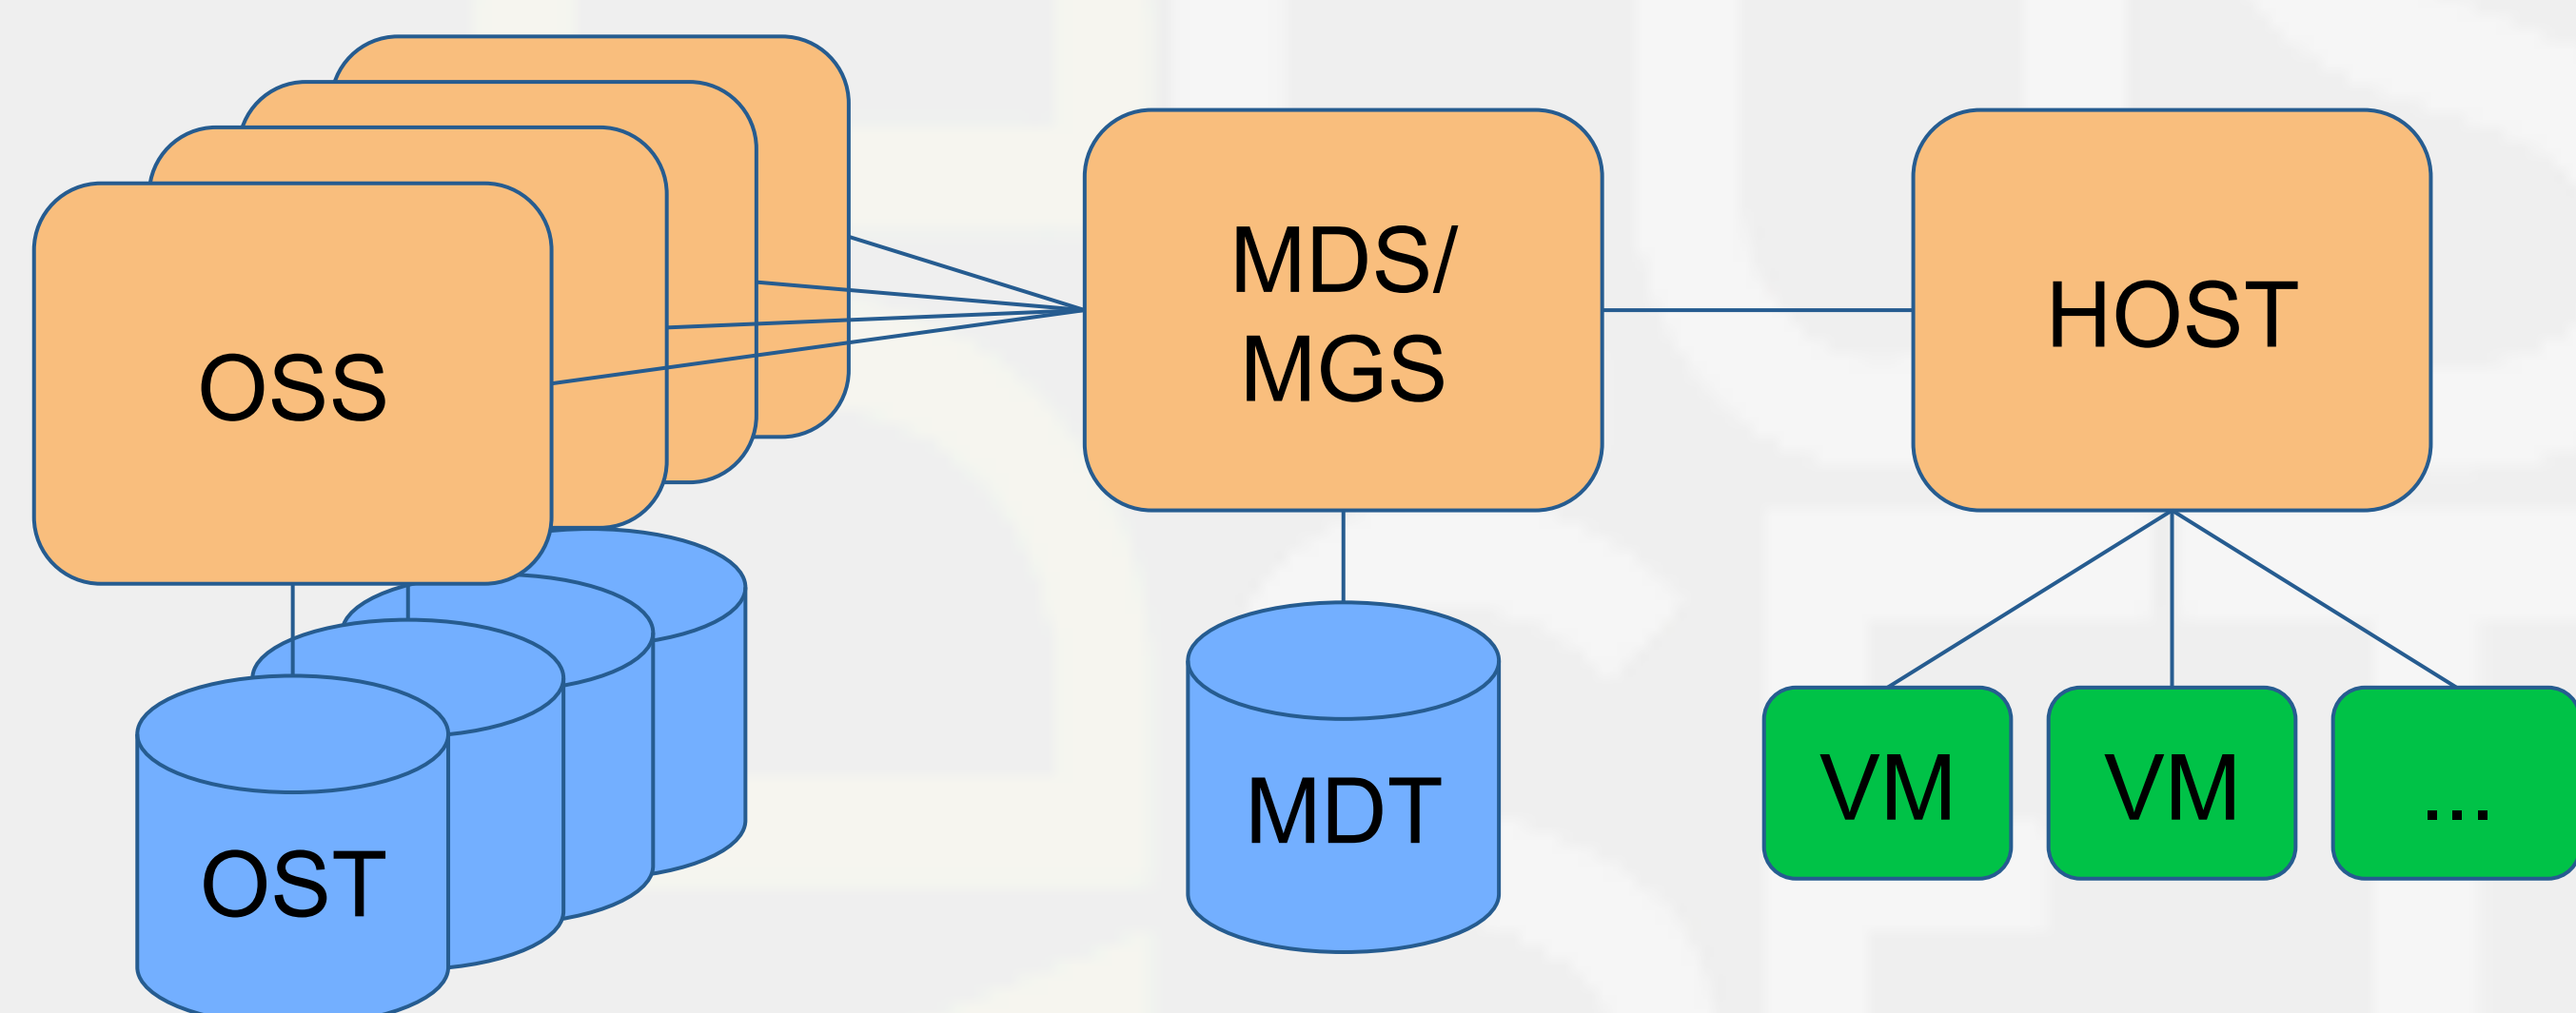
TEST SETUP

- IOR as the benchmark
 - 44.7 GiB file per client
 - 10K, 100K, 1M byte transfer sizes
 - Only N to N file reads/writes
- KVM as the hypervisor
- Centos 6.6 VMs
 - 1 core, 4GB Memory, 20GB Disk

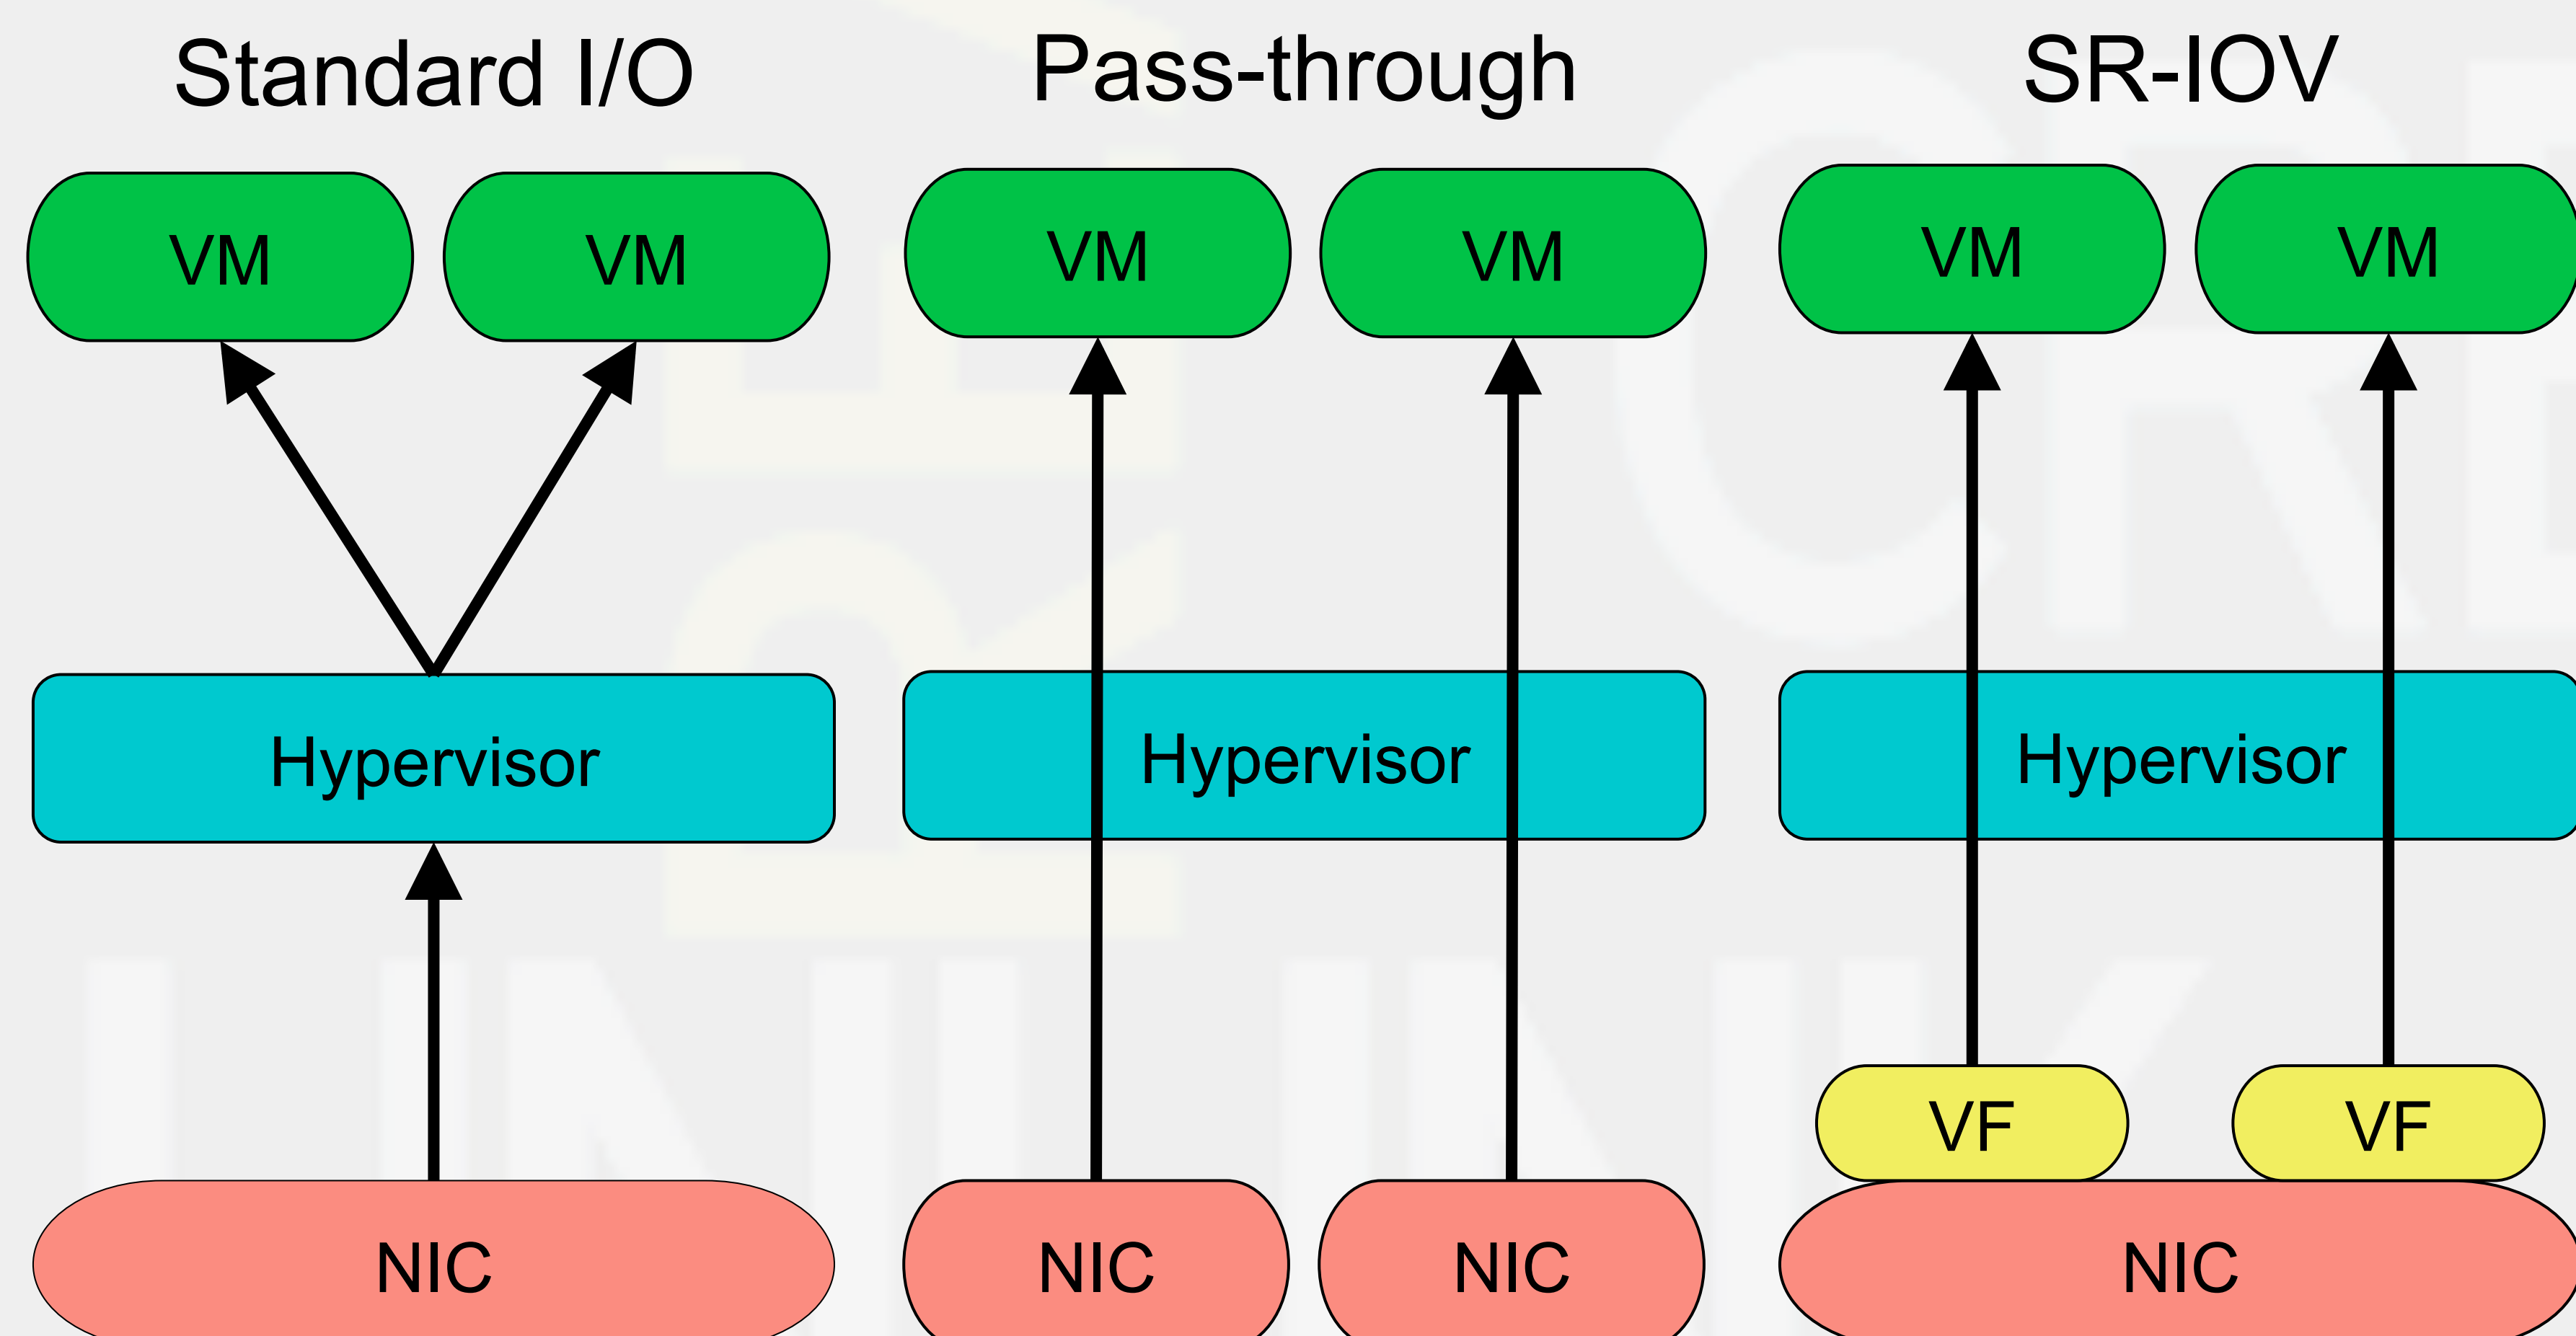
HARDWARE & SR-IOV

- Mellanox ConnectX-3 IB Card
- SR-IOV allows one to share 1 PCIe interface to many VMs
- Create 1 Virtual Function (VF) per VM
 - 1-11 VMs (only 12 cores on Host)
- VF count is a boot argument Reboot to change

NETWORK TOPOLOGY



VM DEVICE ACCESS COMPARISON

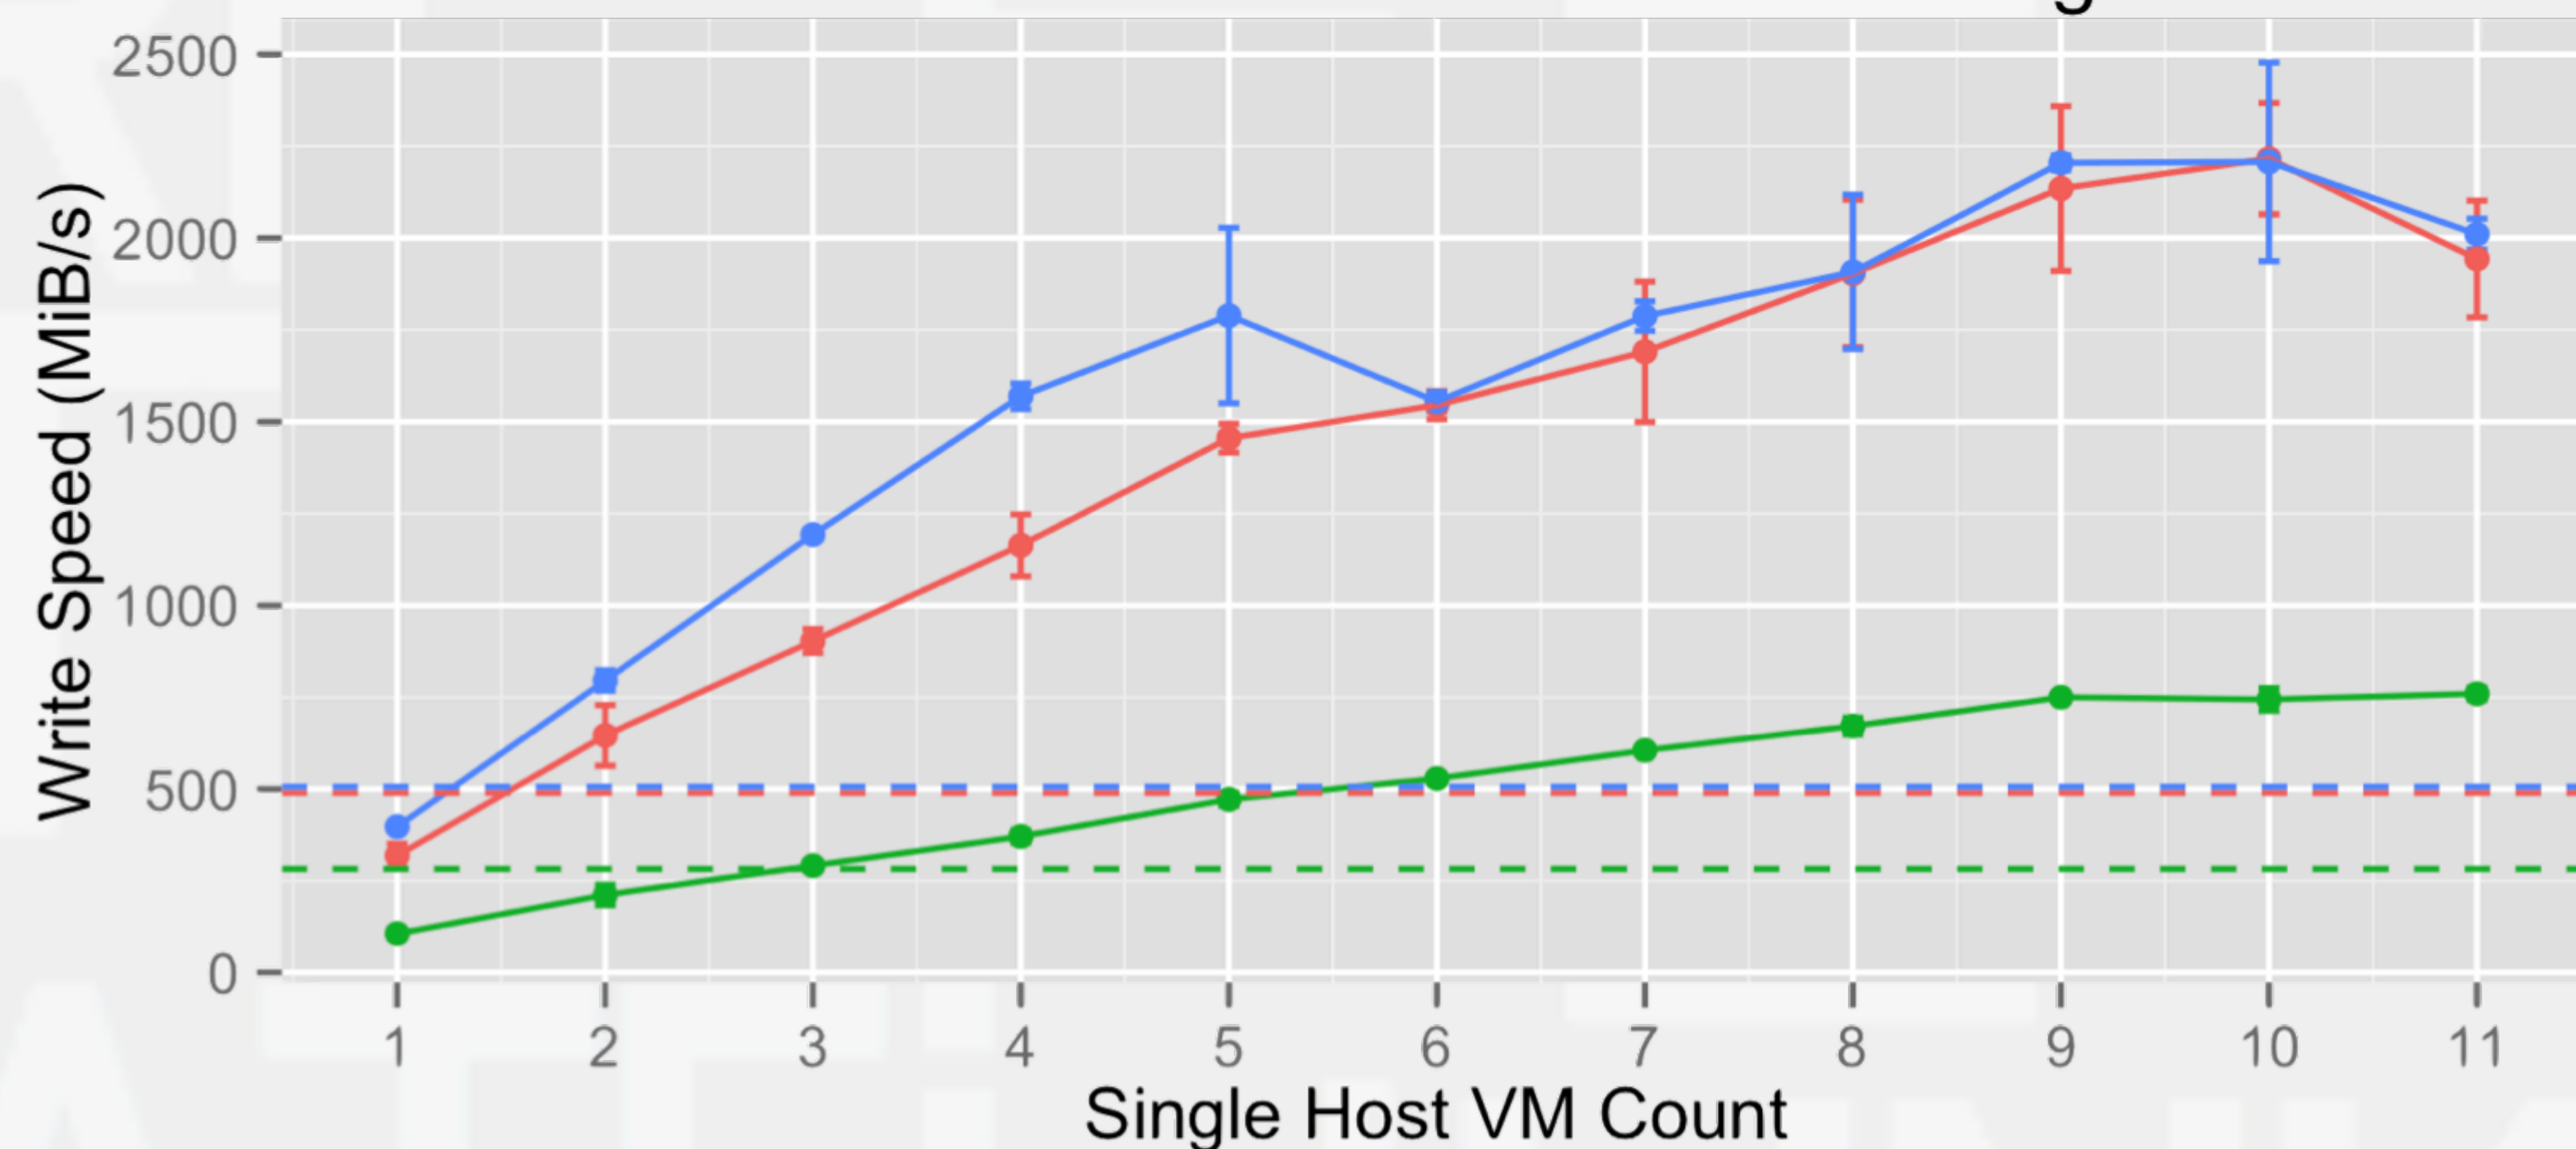


Transfer Size (Bytes) 100K 10K 1M (dashed lines are native installs)

IOR Read Results for 1-11 VMs on a Single Host



IOR Write Results for 1-11 VMs on a Single Host



We observed maximum performance increases of about 4.5x over a native install. It should be noted that transfer size is important for maximizing throughput, where in our case larger transfer sizes lead to better performance. Our results lead us to conclude that this method of provisioning file transfer agents should be further explored for HPC production use.

FUTURE WORK

- Greater number of VMs
- Local subnet per host
- Other hypervisors
- Containers instead of VMs