

SuperCloud Storage Measurement and Optimization

CS5413 - High Performance Systems and Networking

Sneha Prasad - sh824@cornell.edu

Lu Yang - lu77@cornell.edu

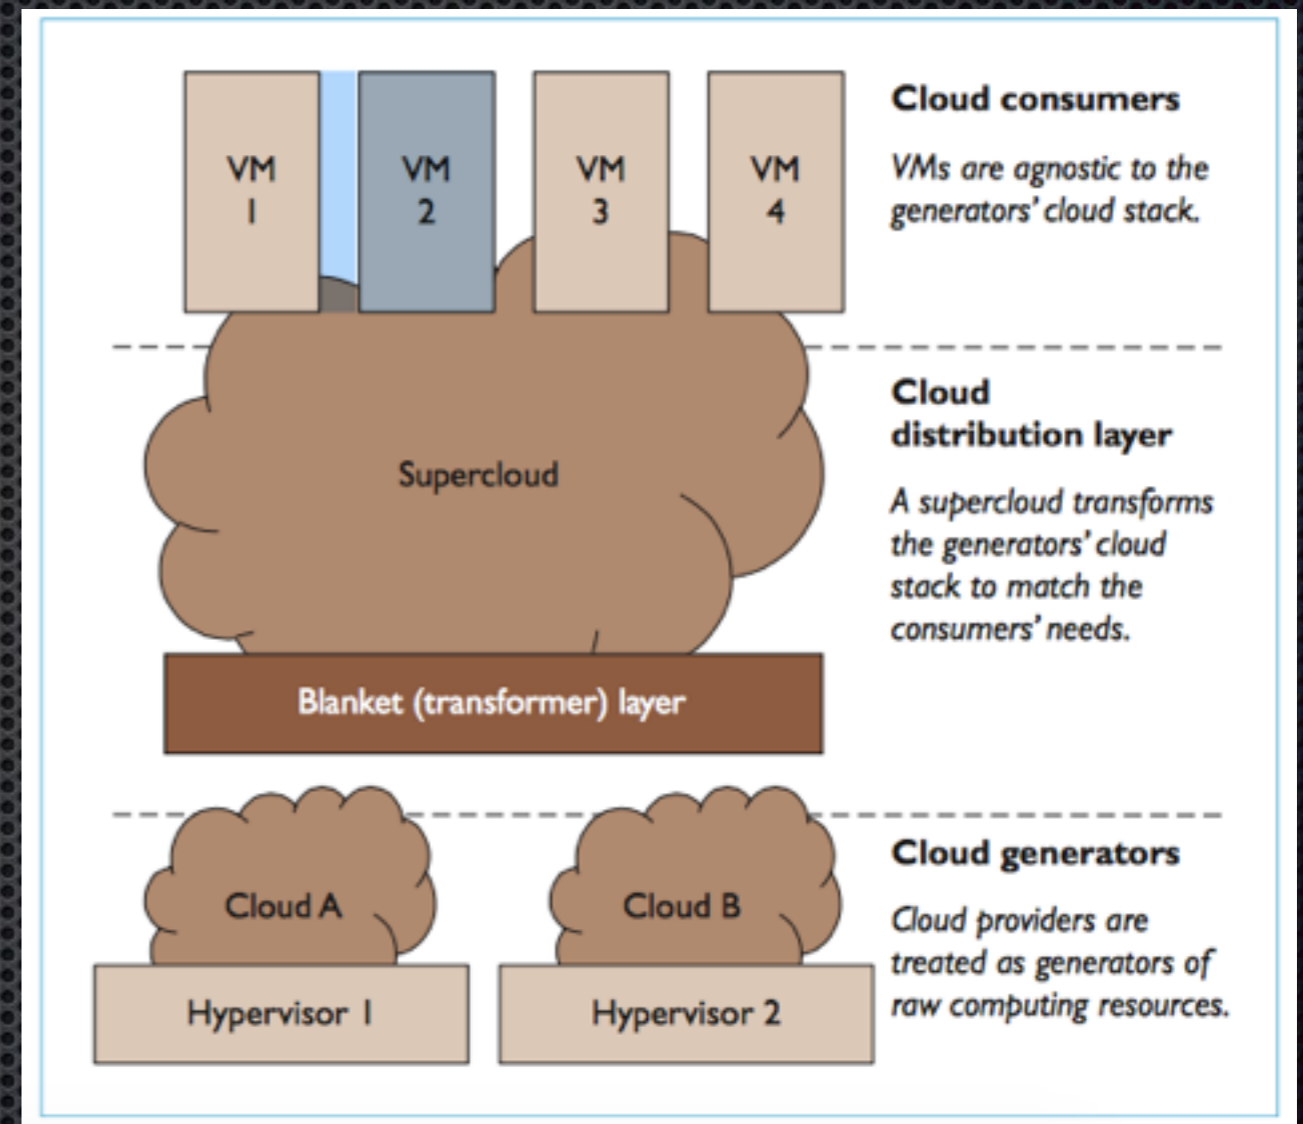
Motivation

- Super Cloud is a super cool idea!
- Server virtualization decouples vm from physical hardware
- More VM more I/O
- VM's are moved around sequential to Random I/O
- Multiple VM's single Hypervisor



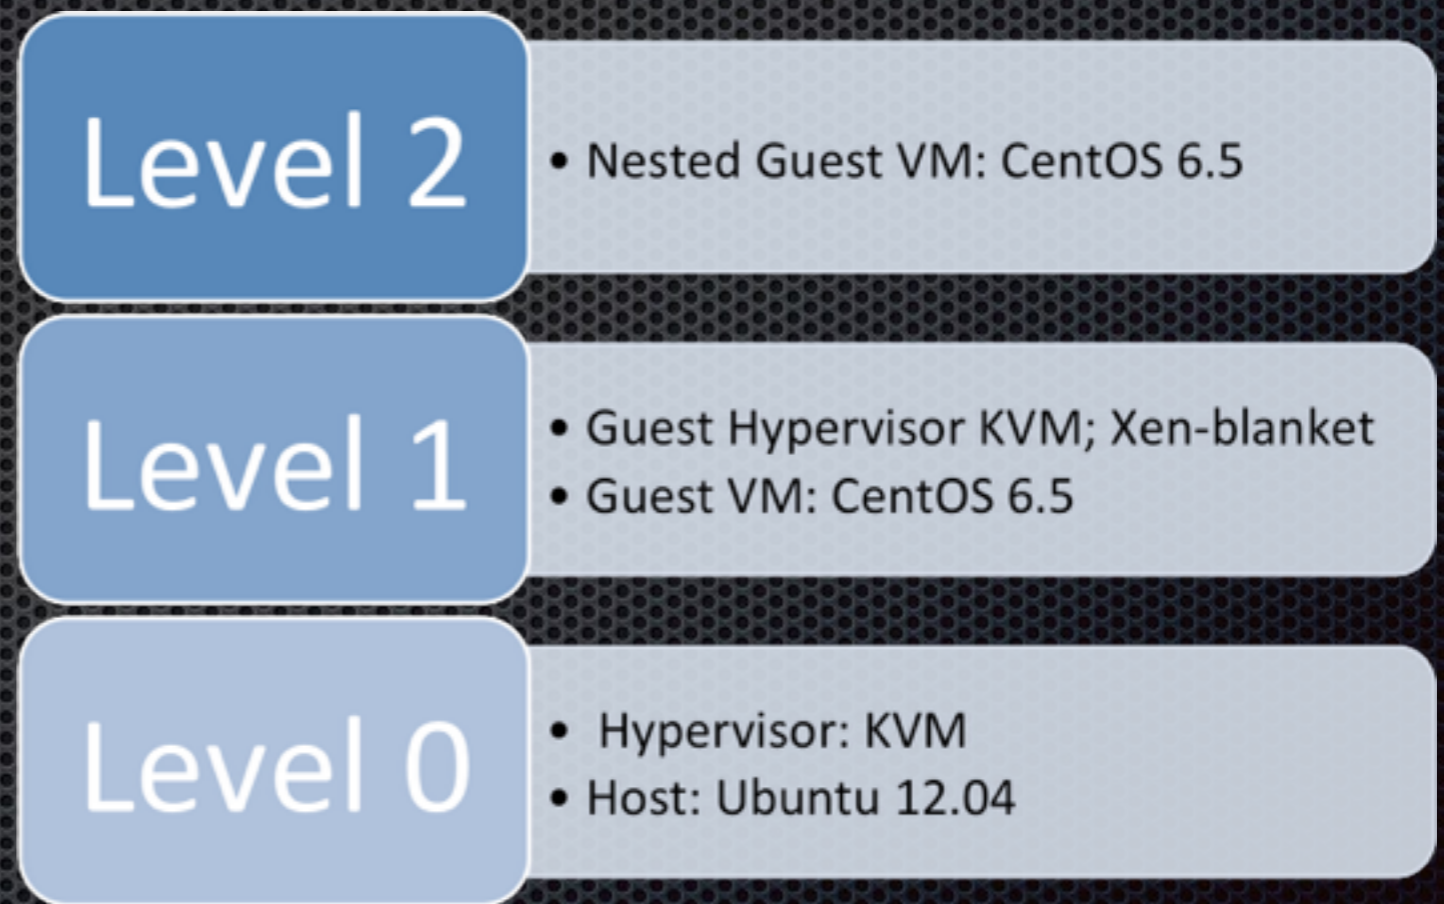
Supercloud

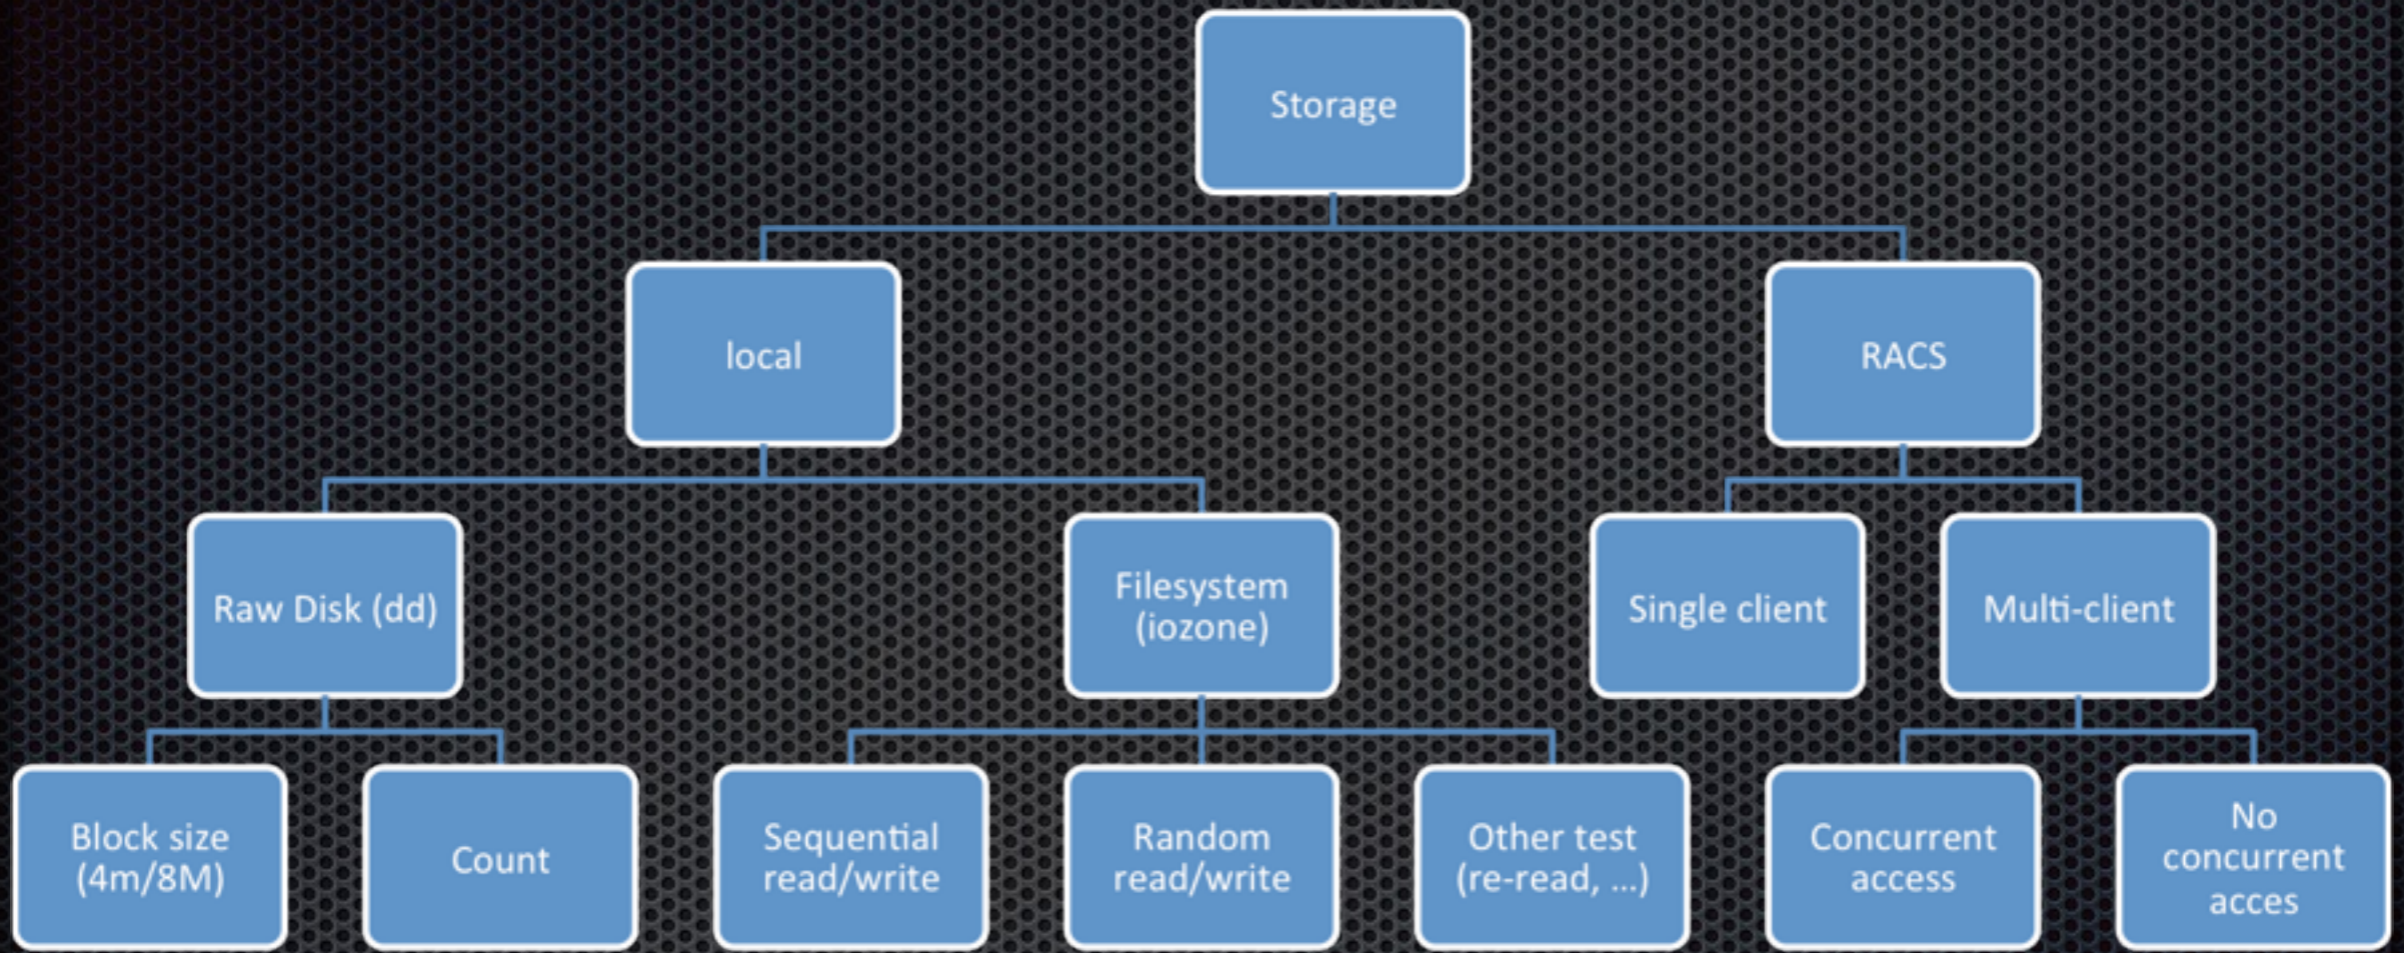
Nested Virtualization
May impact
performance
Need for performance
study



Test Setup


- Nested Guest/ Nested VM - CentOS 6.5
- Guest- Xen Blanket
- Guest Hypervisor - KVM
- Guest - CentOS 6.5
- Hypervisor - KVM
- Host - Ubuntu 12.04 - Baremetal server



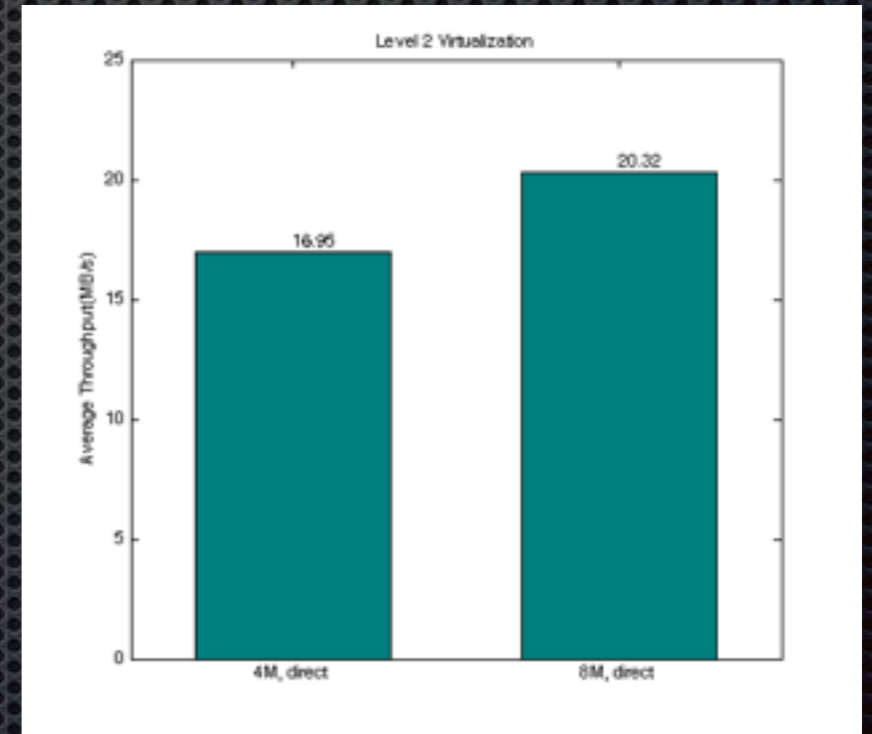

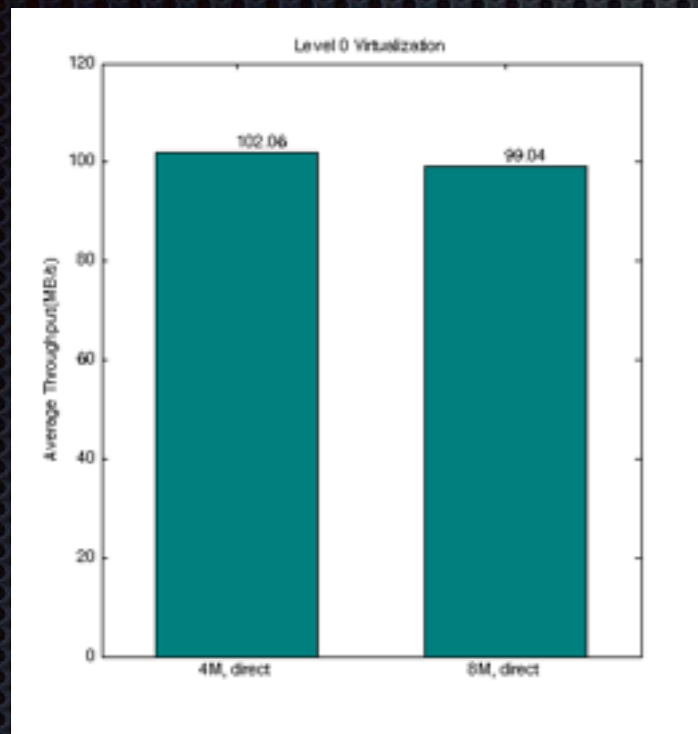


Methodology

DD command




DD Command

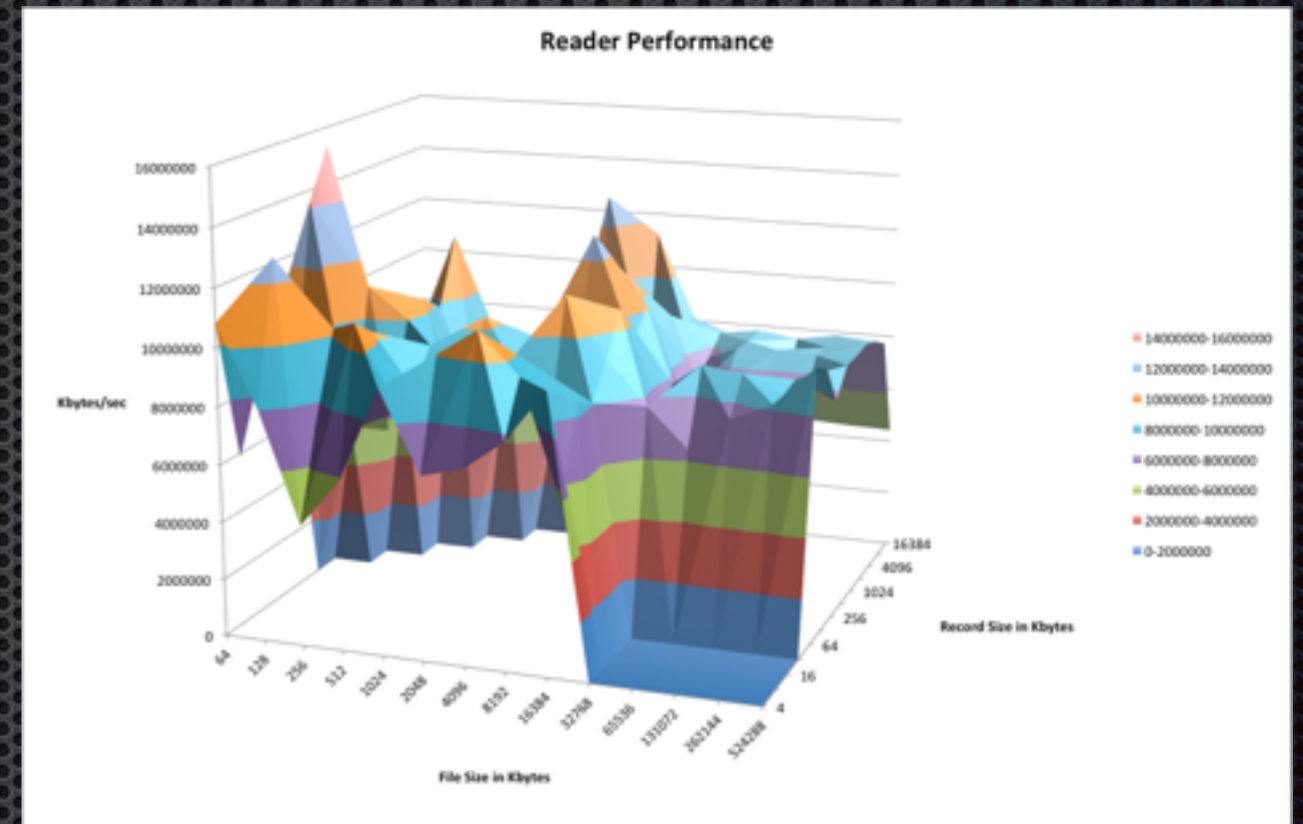
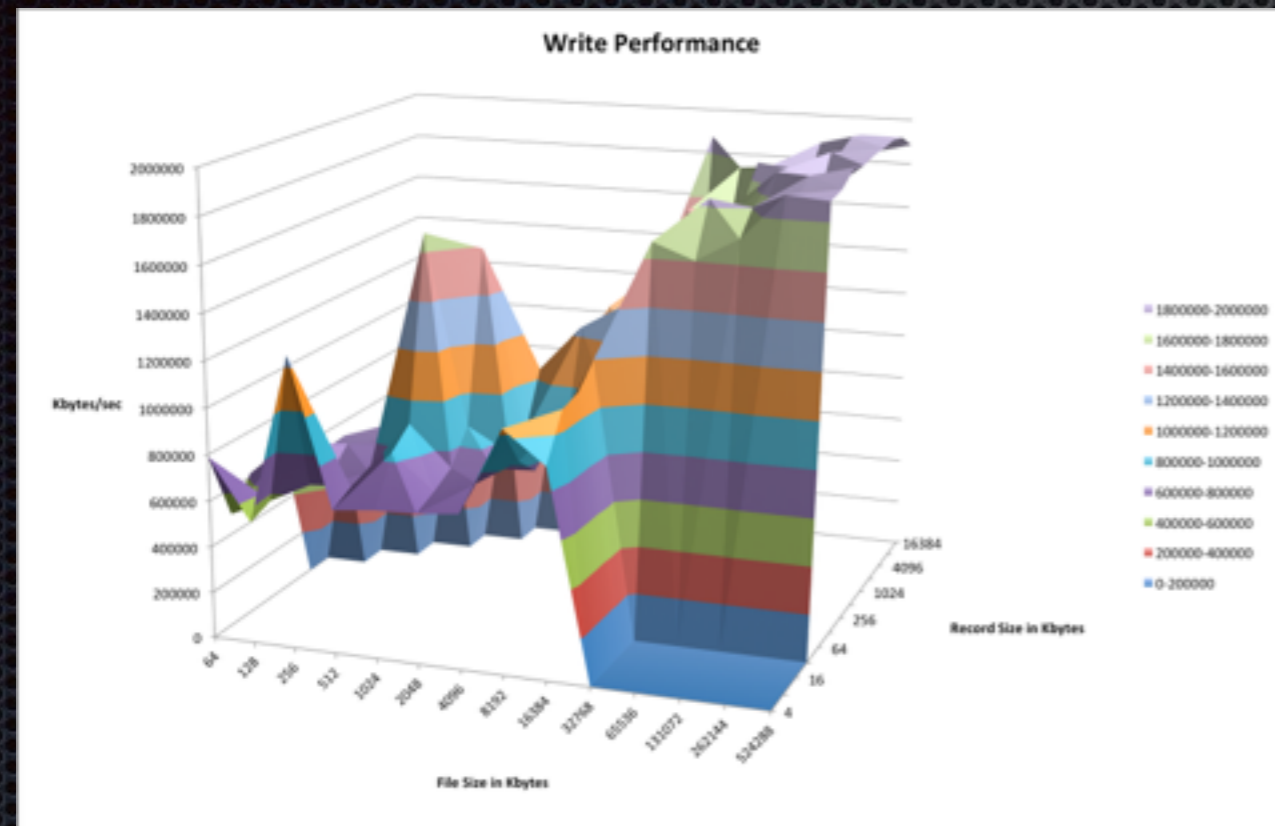


```
dd if=/dev/zero if=zero oflag=direct  
bs=4M count=250
```




IOZone



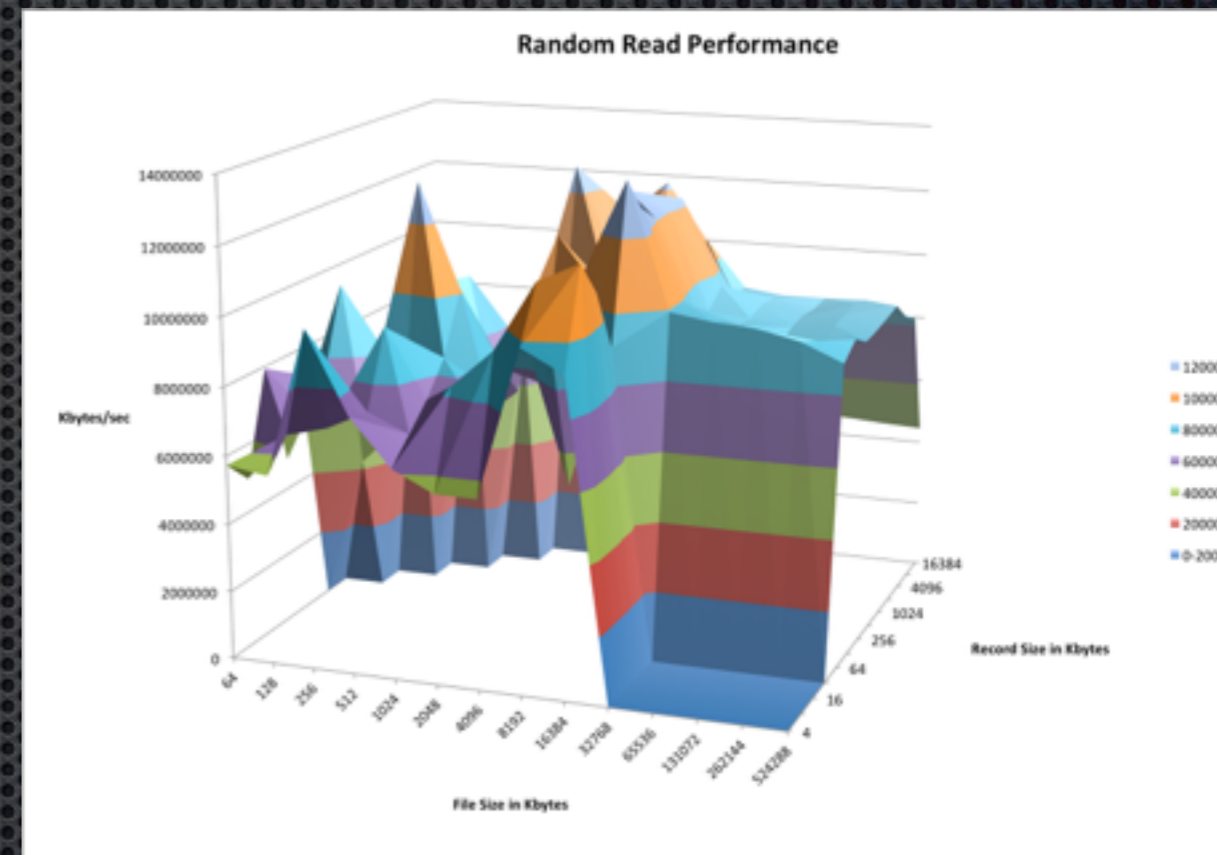
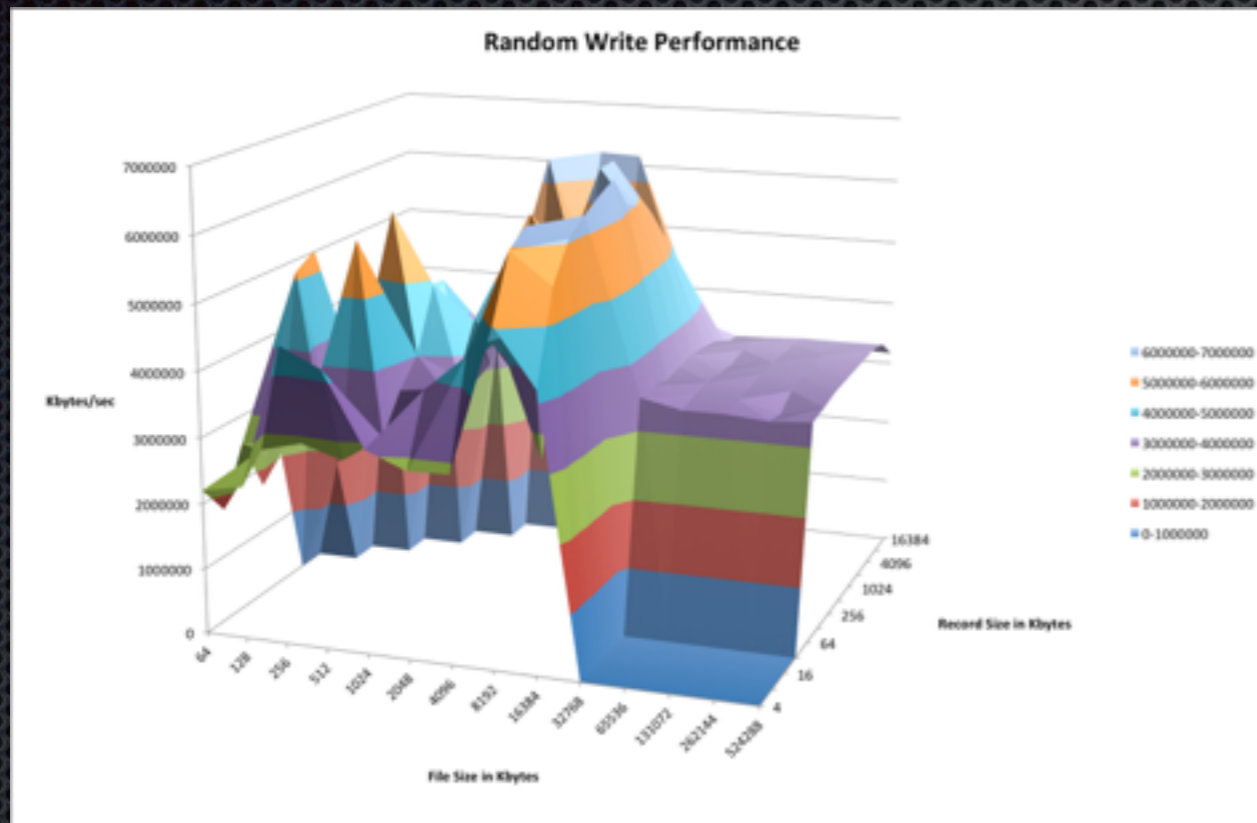
Sequential Read and Write



Reread Rewrite



Random Read and Write



Result Analysis

- DD for block size of 8M
- from level 0 to level 1 throughput drops by 2.5 times
- from level 1 to 2 drops by 1/2
- From level 0 to level 2 drops by 1/5
- DD for block size of 4M not stable, not consistent - take a bigger sample
- IOZONE



Lessons Learned

- Benchmarking Tools
- Its very tricky to get all the layers up and working
 - problems with Network, Storage space, Versions, Broken Packages, rebooting causes loss of connectivity - no console
 - Cloud Lab interface is constantly changing
 - Don't take random advice from the internet



FutureWork

- ✦ Multiple VM's
- ✦ NFS
- ✦ iSCSI



Thank you!!

- ✦ Ji Yong
- ✦ Weijia Song
- ✦ Zhiming Shen
- ✦ Prof Hakim
Weatherspoon :)



Appendix 1 : System Config

- Architecture: x86_64
- CPU op-mode(s): 32-bit, 64-bit
- Byte Order: Little Endian
- CPU(s): 32
- On-line CPU(s) list: 0-31
- Thread(s) per core: 2
- Core(s) per socket: 8
- Socket(s): 2
- NUMA node(s): 2
- Vendor ID: GenuineIntel
- CPU family: 6
- Model: 62
- Stepping: 4
- CPU MHz: 1200.000

Disk

Disk /dev/sda: 1000.2 GB, 1000204886016 bytes
255 heads, 63 sectors/track, 121601 cylinders, total
1953525168 sectors

Units = sectors of 1 * 512 = 512 bytes

Sector size (logical/physical): 512 bytes / 512 bytes

I/O size (minimum/optimal): 512 bytes / 512 bytes

Disk identifier: 0x90909090