

Cloud Computing through Virtualization and HPC technologies

William Lu, Ph.D.

Cloud Computing & HPC

A Case of HPC Implementation

Application Performance in VM

Summary

Cloud Computing & HPC

- **HPC users are always looking for more capacities**
- **Cloud Computing concept is very attractive for HPC sites with dynamic workload**

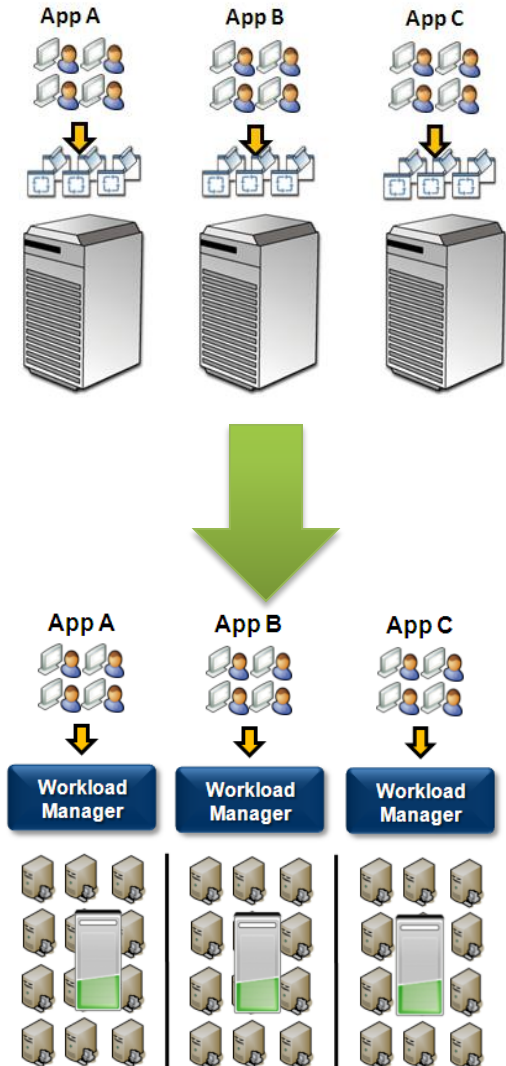
Questions

- **How cloud implementation look like in HPC?**
 - VM or not VM?
 - What is the VM performance cost for HPC applications?

A Case of HPC Implementation

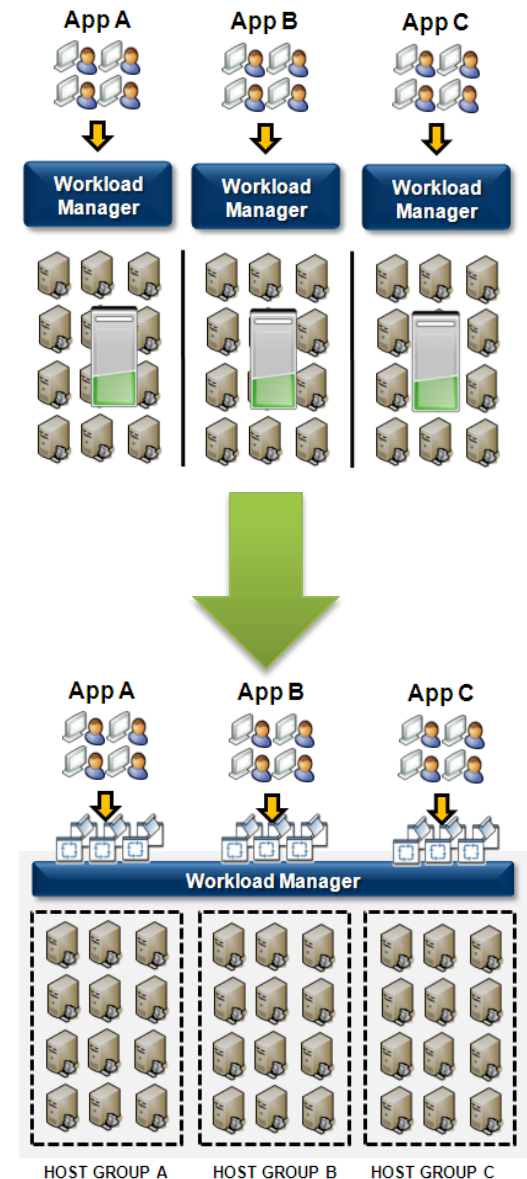
Stage 1: Commodity Cluster

- Problems:
 - Existing UNIX server does not have enough capacity to run CFD jobs
 - The maintenance cost and administration cost is about same as the original hardware
- Solution:
 - X86 Cluster
 - Open source cluster stack (Linux, MPI, open source scheduler)
- Result:
 - Capacity increases 5 times
 - Cost decrease 50%
 - Clusters are adopted in many departments



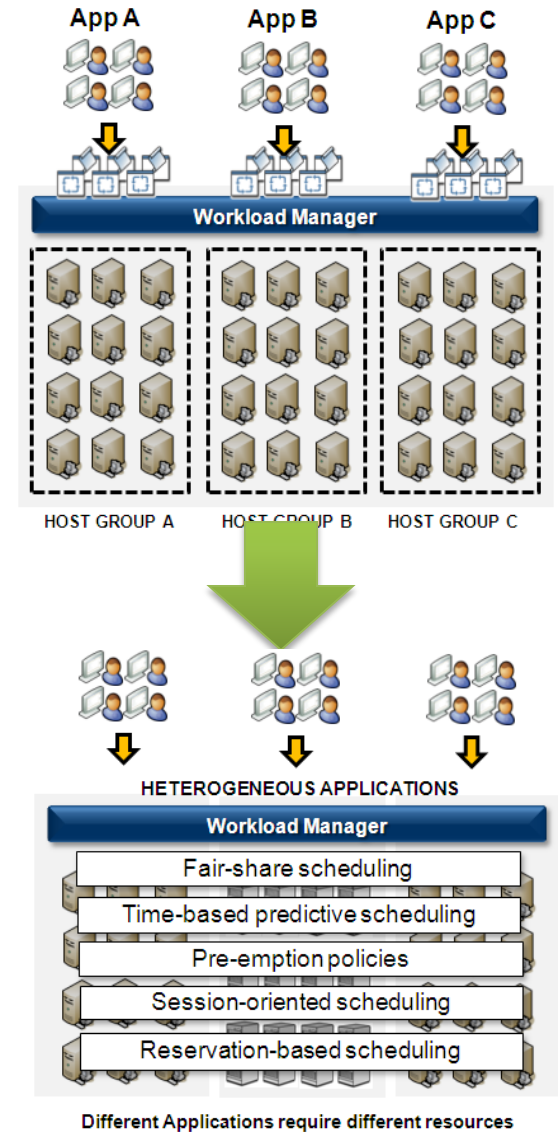
Stage 2: Centralized clusters

- Problems:
 - 50 clusters in an organization with different management practice
 - Each cluster needs budget to refresh
- Solution:
 - Consolidate 50 clusters into 2 HPC centers
 - Use robust workload scheduler (Platform LSF)
- Result:
 - Centralized budgeting and capacity planning
 - Cut hardware cost by 20%
 - Cut administration cost by 70%



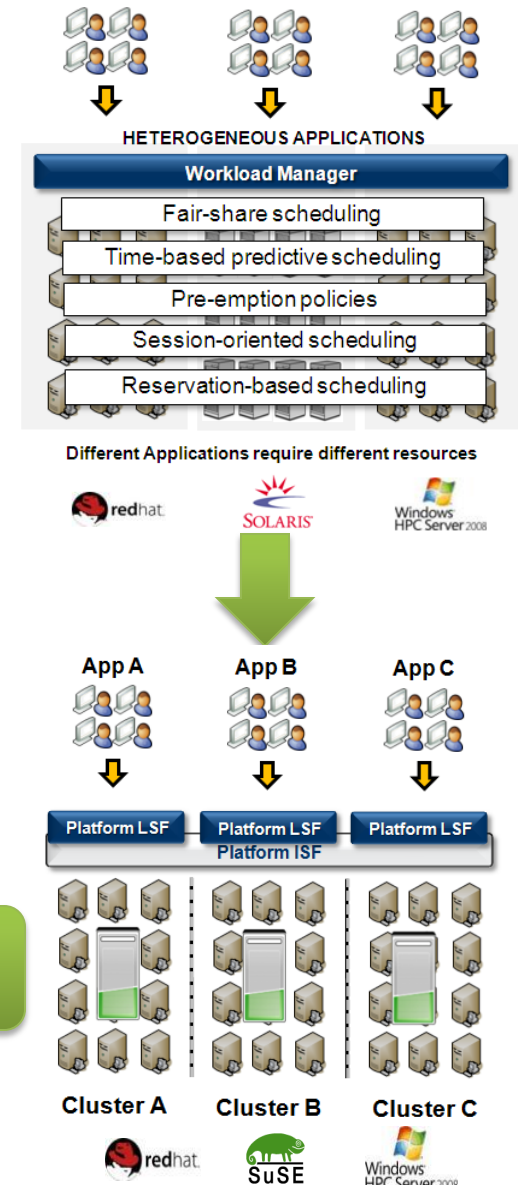
Stage 3: Breaking Group Silos

- Problems:
 - Each group has dedicated queue and a partition of a cluster
 - The utilization is low: 60%
 - No funding from management for increasing capacity before getting higher utilization
- Solution
 - Leverage advanced scheduling features in Platform LSF
- Result
 - Improved utilization from 60% to 75%



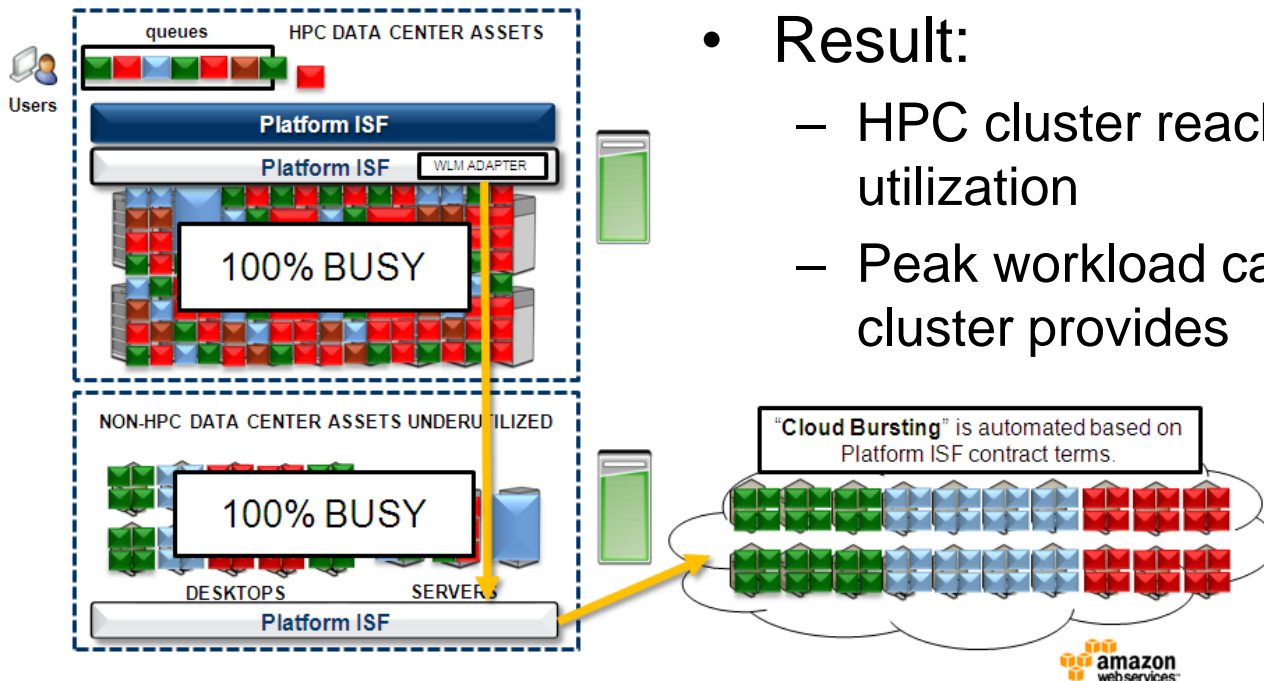
Stage 4: Make Users Happy

- Problems
 - Conflict scheduling policies among different groups
 - Conflict of OS/app environment across groups
 - Utilization is not optimal
- Solution
 - Multiple elastic logical clusters managed by Platform ISF
- Result
 - Utilization reaches to 85%
 - Users are much happier



Stage 5: A Shared Dynamic HPC Environment

- Problems:
 - Can we get near 100% utilization of HPC cluster?
 - How to get capacity without limit?
- Solution:
 - “Bursting”
- Result:
 - HPC cluster reaches near 100% utilization
 - Peak workload can go beyond what HPC cluster provides



Comparison of User Experiences

Public Cloud IaaS

Resource available on demand

Provisioning app stack manually

Pay per use

VM

A dynamic shared HPC infrastructure

Resource available on demand

On demand pre-defined stack

Shared cost with other user groups

Physical & VM

Private Cloud

Cloud Computing Performance Tests via HPC Advisory Council

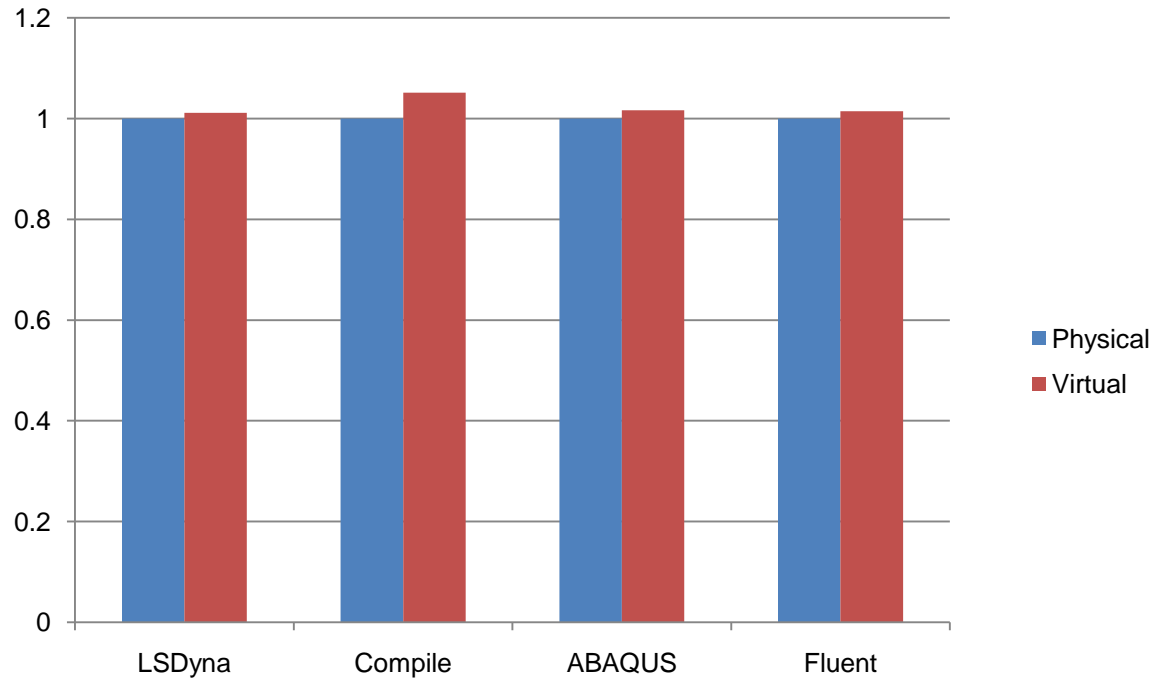


- **Application performance in a shared environment**
- **Application performance inside a VM on single node**
- **Application performance inside a VM using interconnects**

Environments

- **CPU: Dual socket Quad core 2.83 GHz Intel**
- **Memory: 8 GB RAM**
- **OS : CentOS Linux 5.3 x86_64**
- **Hypervisor: Citrix Xen 5.5**
- **VM Configuration**
 - 1x(8 GB VM with 8 Cores)
 - 8x(1 GB VMs with 1 core)


Serial Runs



Serial runs show the overhead of VM is small

Memory Bandwidth Test


Memory Bandwidth (Stream - 1.8 GB size)



VM adds small overhead to memory bandwidth


I/O Bandwidth Test

Disk I/O IOZone - 32 GB file




VM adds overhead to I/O bandwidth

LSDyna - MPP971 - Refined Neon (30ms)




Parallel LSDyna does not performance in VM

Fluent 12.1 - sedam_4m



Parallel Fluent performs well in VM

ABAQUS - Explicit - e2.inp



	serial	parallel-2	parallel-4	parallel-6
Physical	12768	6693	3958	3384
Virtual	12953	6836	4131	3397
Virtual - different VMs	12953	7175	4813	4370

Parallel Abaqus performs well in VM

Conclusion of the test

- **For common CAE applications running in serial mode, the VM overhead is small (< 4%)**
- **For common CAE applications running in parallel mode, the VM overhead depends**
- **CPU intensive or memory intensive application could be good candidates to run in VM**
- **I/O intensive application performance in VM depends**
- **We expect the VM overhead is unacceptable for MPI jobs over interconnects**

Summary

- **Technology is ready for implementation of private HPC cloud without VM**
- **A full HPC cloud solution running on VM is very application dependent**
- **We will continue to do application performance test and share the results in HPC Advisory Council**

Thank you