Fast and Efficient VM Migrations for Improving SLA and ROI in Cloud Infrastructures

DC CAVES 2010

Michael Kagan, Chief Technology Officer

Mellanox TECHNOLOGIES

Cloud Computing - Definition

Cloud computing is a pay-per-use model for enabling available, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.

NIST - *National Institute of Standards and Technology*

© 2010 MELLANOX TECHNOLOGIES

Service-Oriented View

- Server-centric view

- Service-centric view

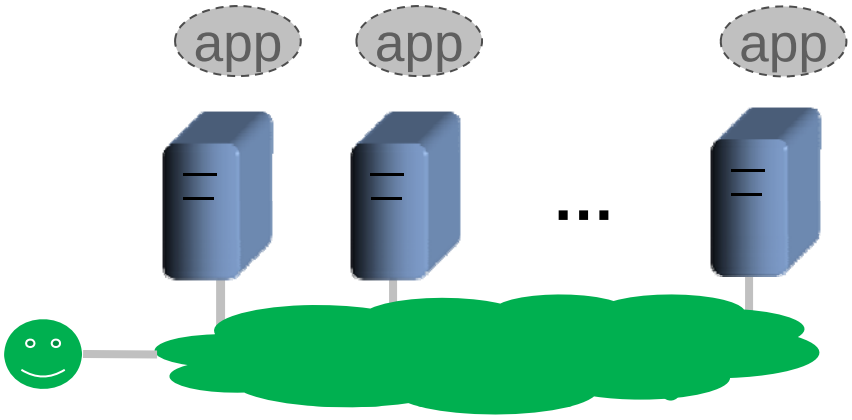

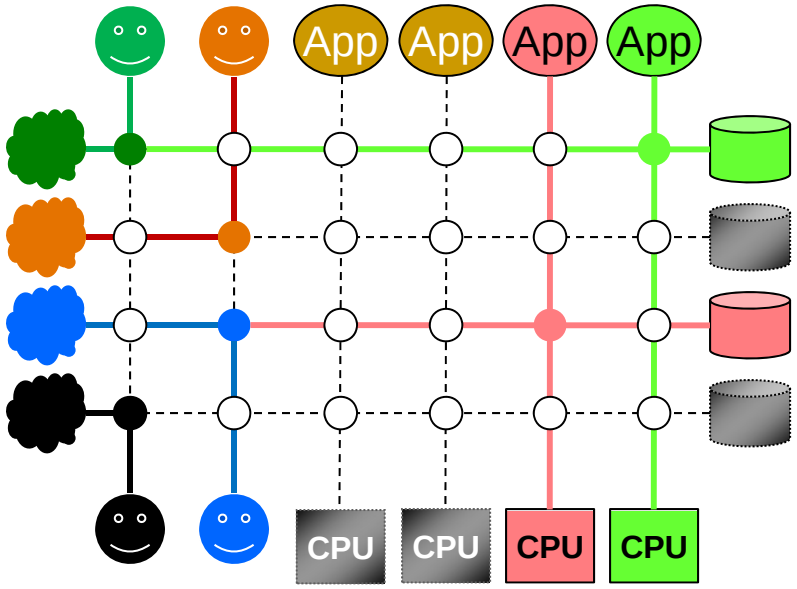

© 2010 MELLANOX TECHNOLOGIES

Virtual Machine – an Application Container

- exists in a virtual address space
- communicates with other containers
- does not own the physical address space
- must be transportable to other physical address spaces
- independent of physical infrastructure

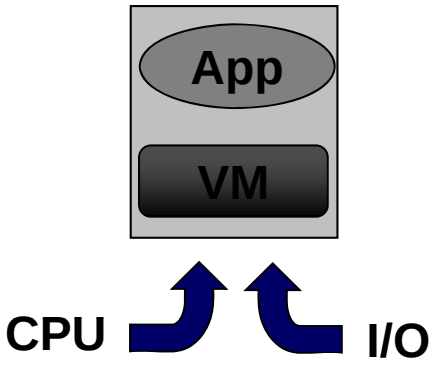

© 2010 MELLANOX TECHNOLOGIES

Importance of Cloud SLAs

Q: Rate the challenges/issues of the 'cloud'/on-demand model

(Scale: 1 = Not at all concerned 5 = Very concerned)

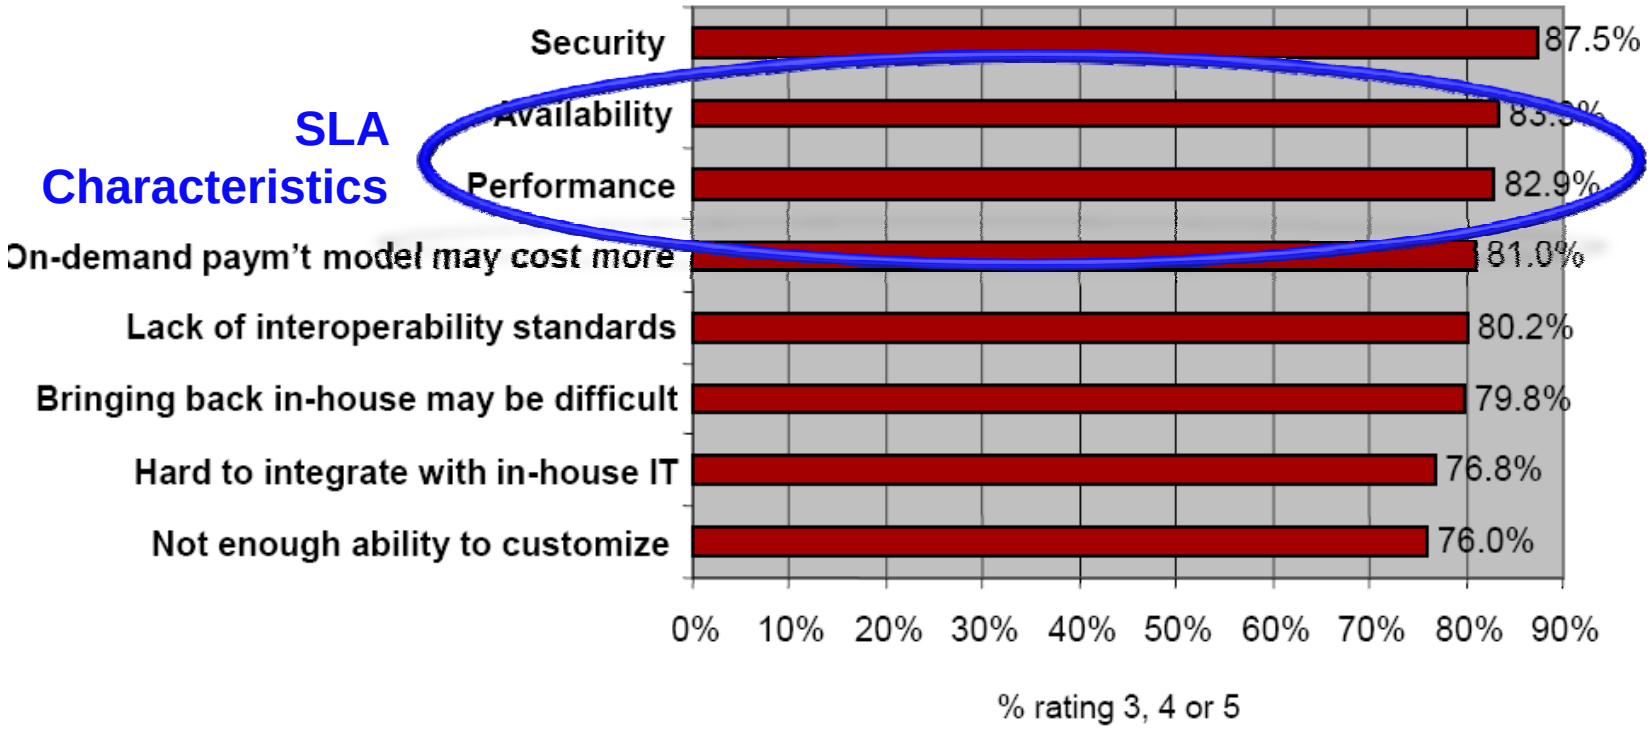

IDC Enterprise Panel, 3Q09, n = 263

Poor SLAs is a significant deterrent to cloud deployments

© 2010 MELLANOX TECHNOLOGIES

Improved SLA with Elastic VM

- Compute unit purchase in cloud is based on VMs
- Today VMs in cloud are not elastic
  - Fixed CPU, memory, storage, I/O per VM instance
- Tomorrow elastic VMs
  - Purchase a VM instance
  - Dynamic resources allocation
    - CPU, Memory, Storage, I/O
  - Pay based on resources usage
    - Like cell phone payment plans
- Challenge
  - Elastic VM requires fast VM Migration
    - Minimal VM “inactive” time

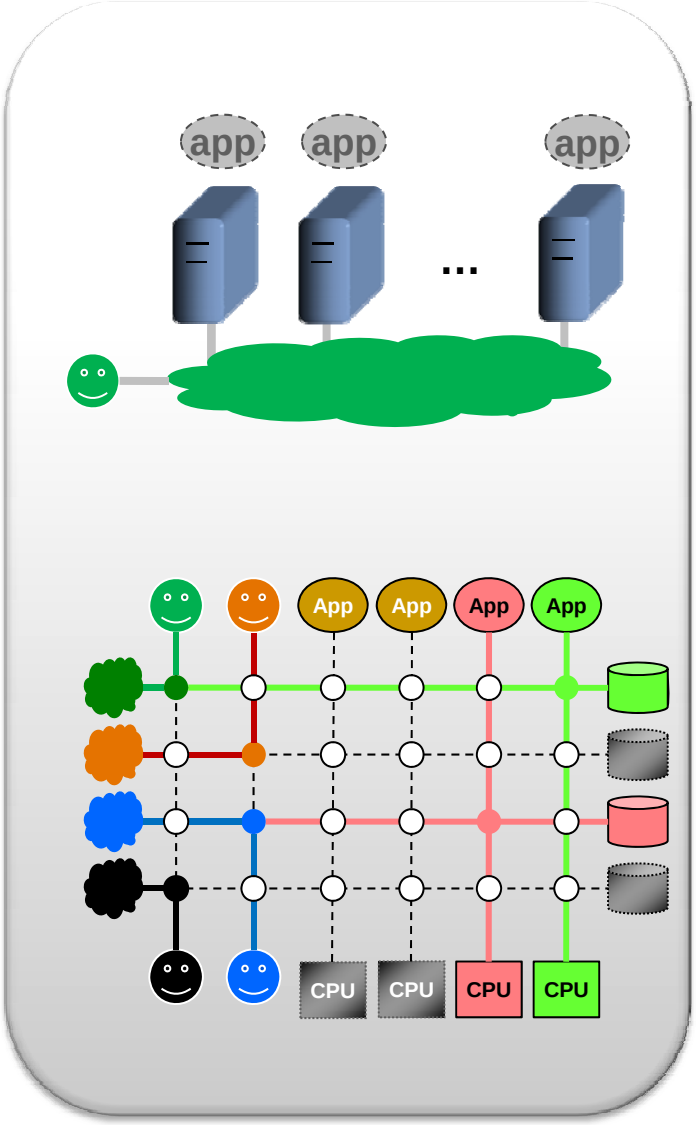

© 2010 MELLANOX TECHNOLOGIES

Case Study: VMware VMotion over 40Gb/s

- Servers
  - CPU: 2xNehalem Intel® Xeon® X5550 (8M Cache, 2.66 GHz)
  - Disks: 2*146GB (6 disks slots)
  - RAM: 24GB
- NICs
  - ConnectX-2 EN 1/10/40Gb/s
    - PCIe 2.0
  - Firmware version: 2.7.700
  - Driver Version: mlx4_en
- Storage
  - FC - connected via FC switch to the EMC storage
    - 4Gb/s HBA
    - EMC CX3-40F
- SW
  - ESX4.1 vMotion
  - VI Java API 2.0 hits Beta 2 milestone
  - Guest OS: RH 5.4

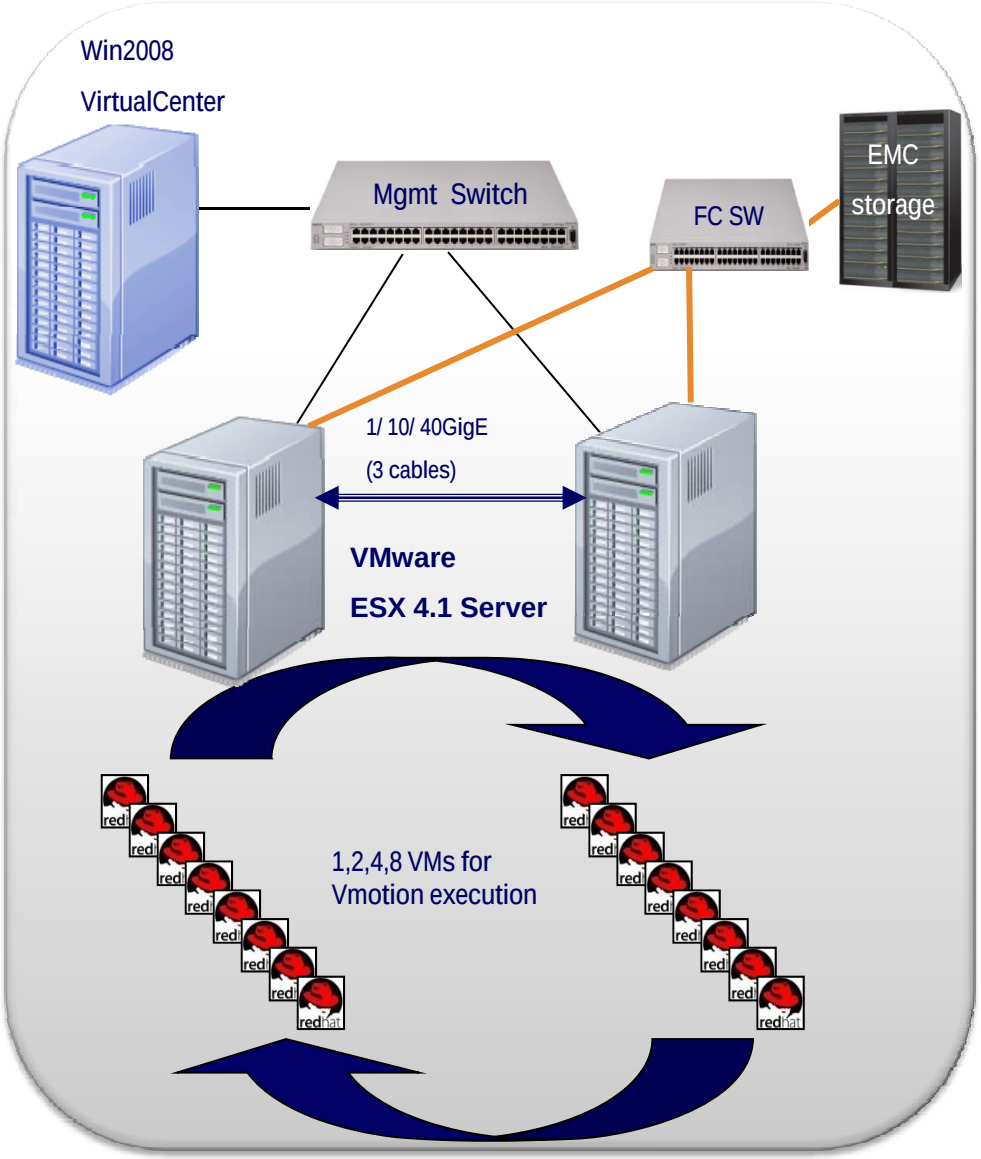

VM Migration Acceleration

Faster VM Migration with VMware vMotion over 40Gb/s

2.5X Boost

4X Boost

Migration of Passive VM

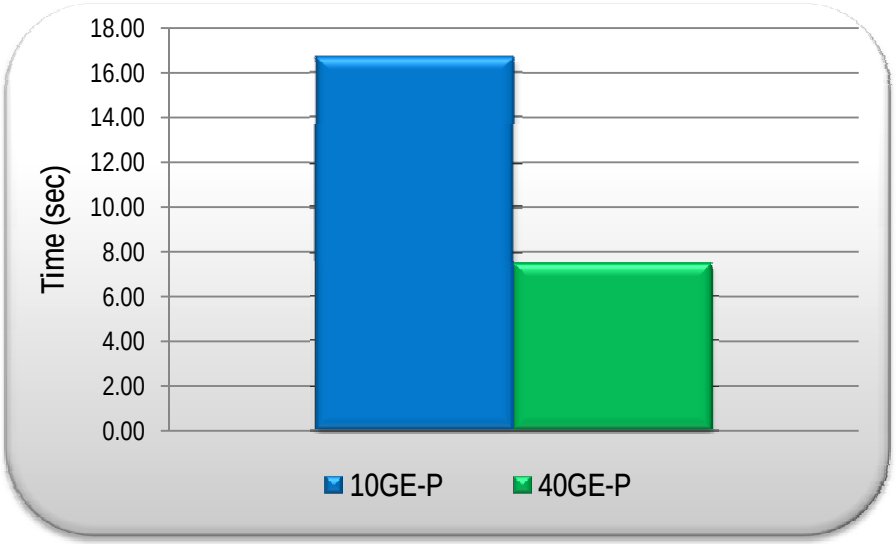

Migration of Active VM

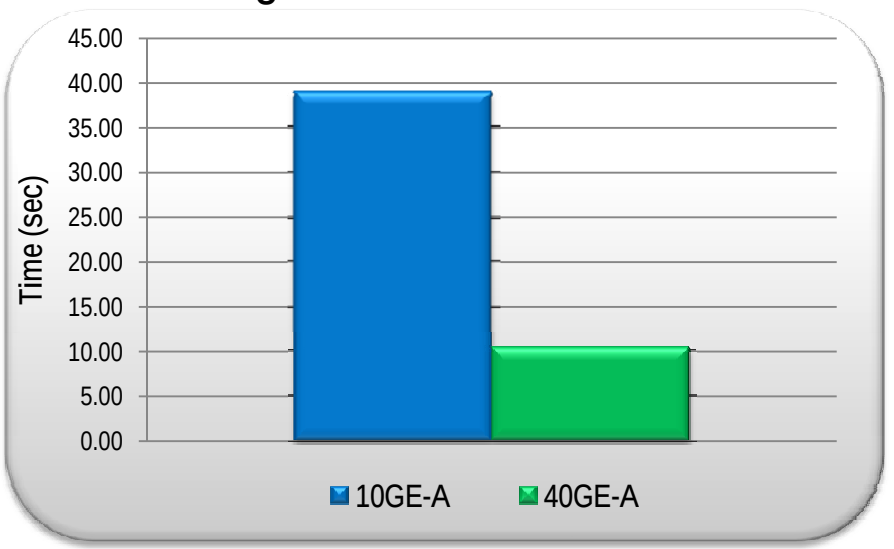

4X boost in VM migration over 40Gb/s

© 2010 MELLANOX TECHNOLOGIES

Fast Fabric Accelerates VM Migration
Reduces Network Traffic

- Passive VM (P) - Dirty the VM memory* once
- Active VM (A) - VM continue to "dirt" the VM memory*

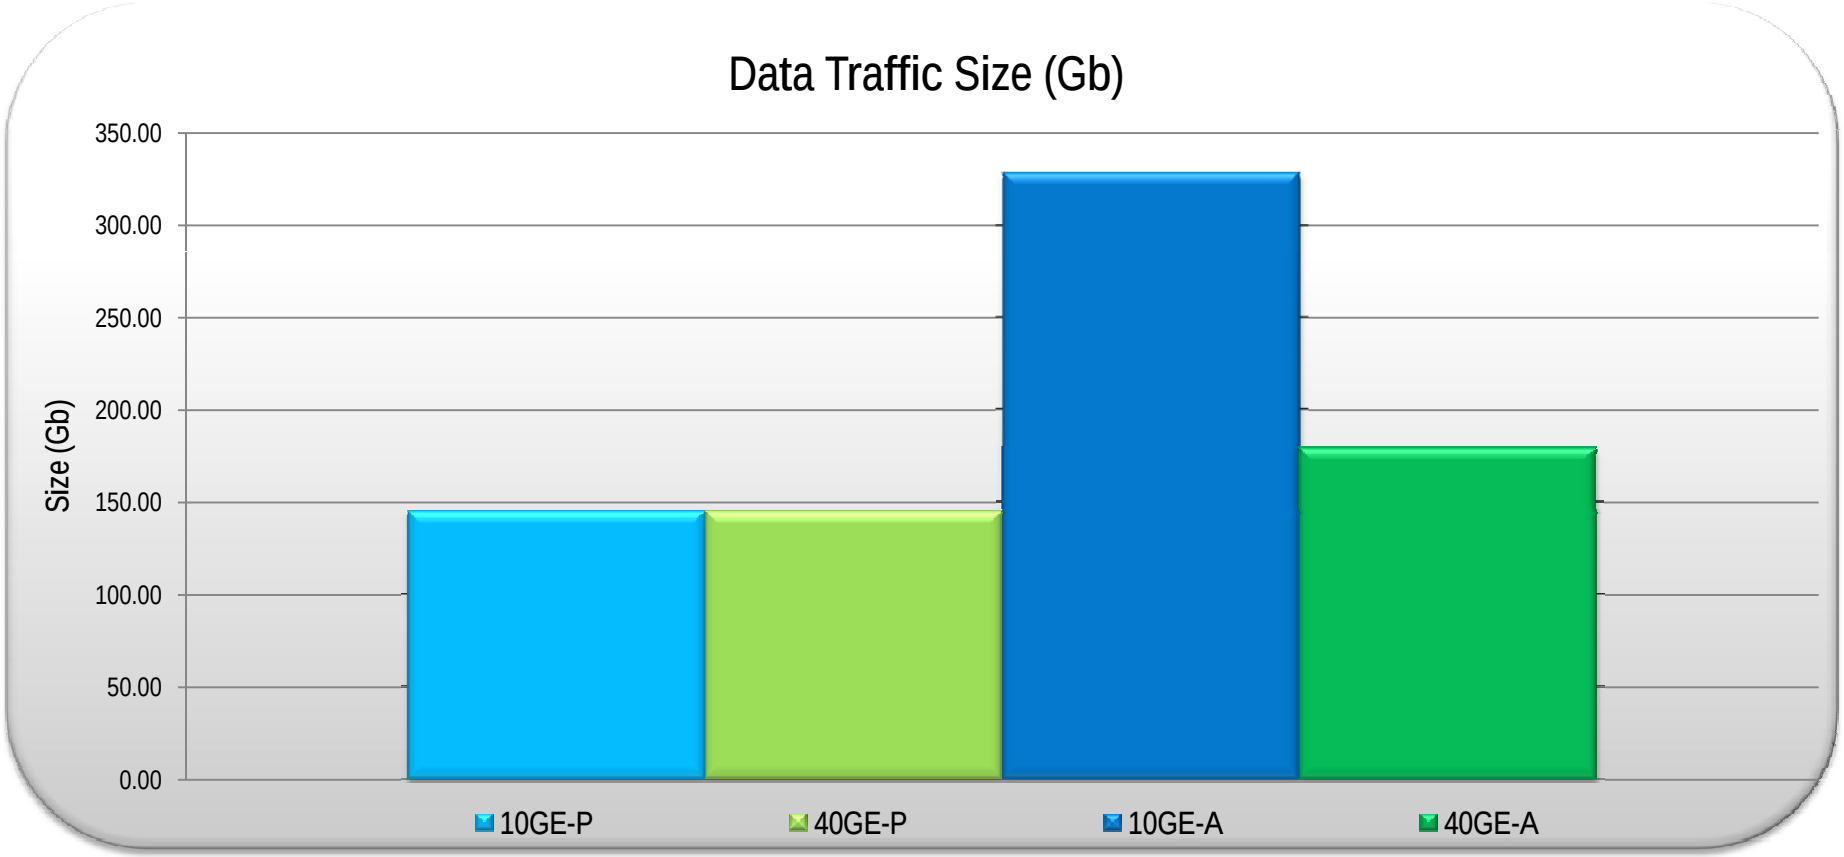

*** 16GB per VM**

Faster networking enables less memory to migrate

© 2010 MELLANOX TECHNOLOGIES

Multiple VM Migration Acceleration

Multiple VM Migration with VMware vMotion over 1/10/40Gb/s Adapters

vMotion Average Migration Duration*

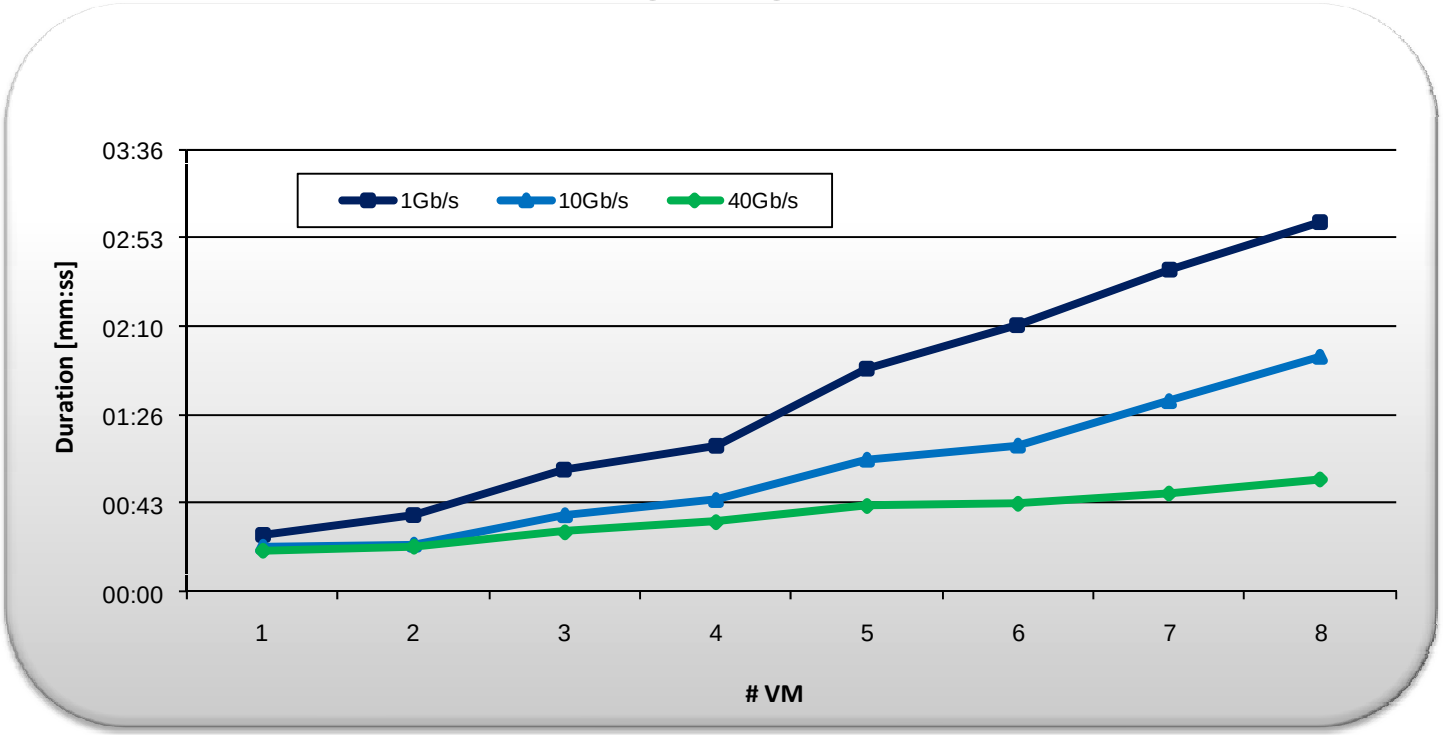

* Preliminary results.1GB per VM. Dirty once

Faster Fabric Enables Concurrent Migration of More VMs

© 2010 MELLANOX TECHNOLOGIES

SLA/ROI* Benefits with VM Migration Acceleration

Boosted SLA at a Lower TCO

32 VMs
4 x (8) Cores Servers
10Gb/s Connectivity

4 x 8 Cores Server
4 x 10Gb/s Ports
4 x 10Gb/s Switch Ports
4 x Cables
$15,800 Total Capital Cost
$987.50 Capital Cost per VM
$528.60 Power Cost per VM (3Y)

32 VMs
2 x (16 Cores) Servers
40Gb/s Connectivity

2 x 16 cores Server
2 x 40Gb/s Ports
2 x 40Gb/s Switch Ports
2 x Cables
$10,700 Total Capital Cost
$668.75 Capital Cost per VM
$295.65 Power Cost per VM (3Y)

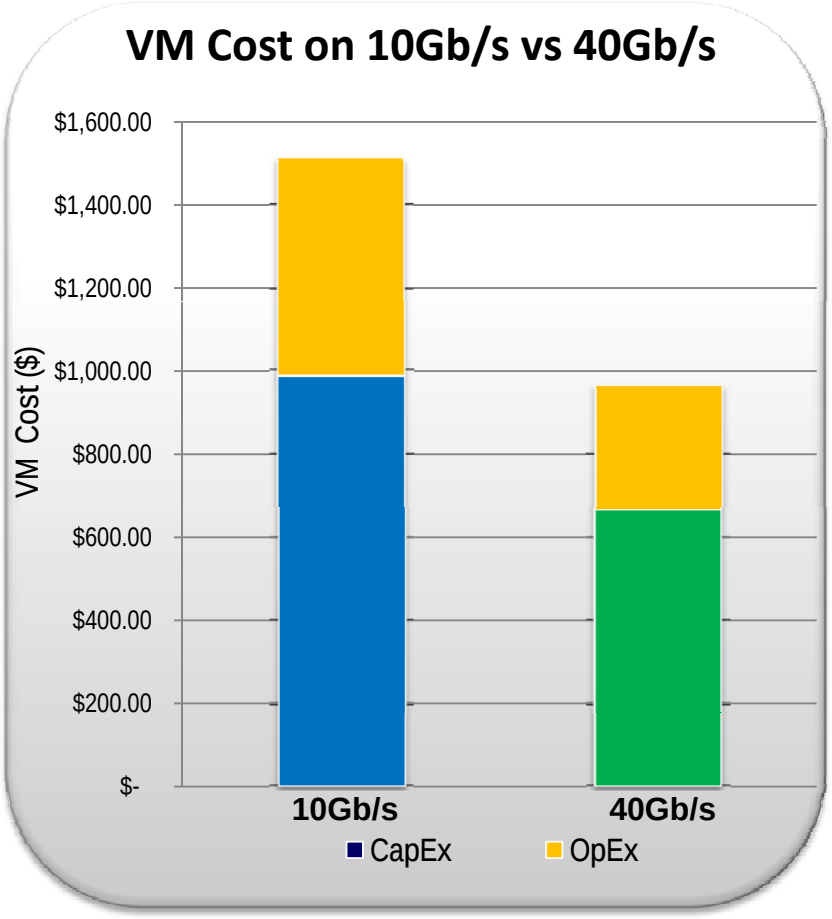

* Single VM Migration with 16GB memory

70%+ Operational cost saving over true 40Gb/s fabric

© 2010 MELLANOX TECHNOLOGIES

Efficient Fabric Accelerates VM Migration
Reduces CPU Utilization

VM Migration with VMware vMotion over InfiniBand*

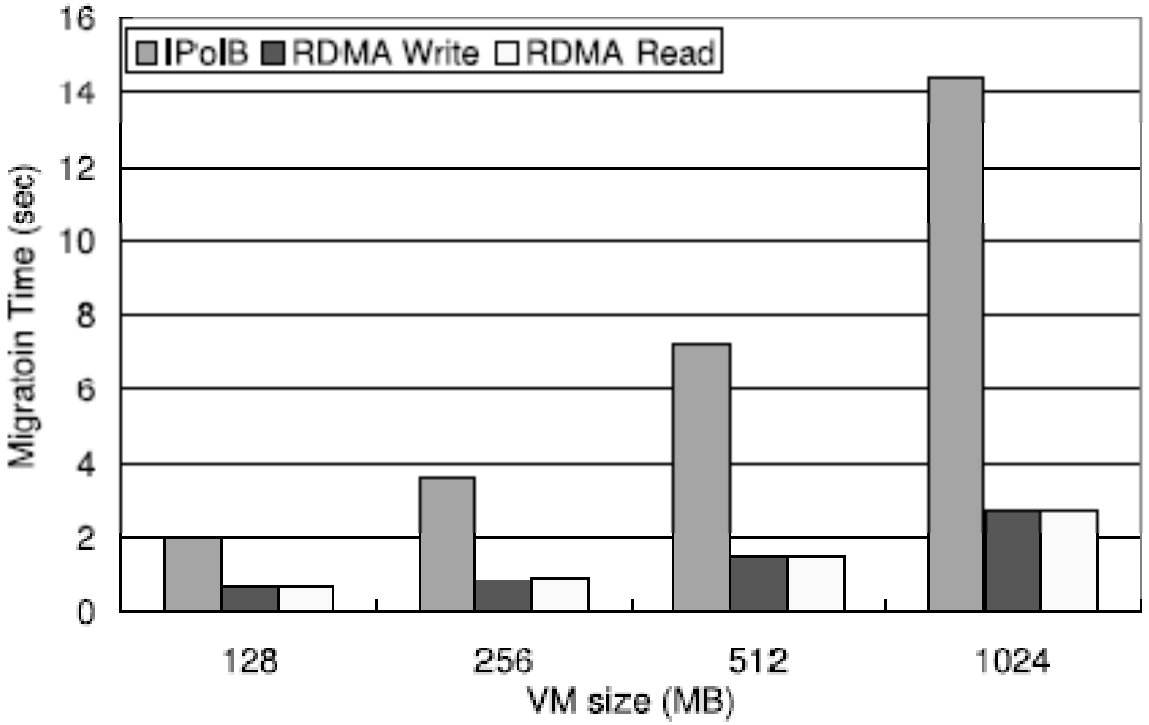

*
High Performance Virtual Machine Migration with RDMA over Modern Interconnects,
Wei Huang, Qi Gao, Jiuxing Liu, Dhabaleswar K. Panda,
Computer Science and Engineering, The Ohio State University and
IBM T. J. Watson Research Center

RDMA accelerates VM migration by 5X

© 2010 MELLANOX TECHNOLOGIES

The Network Determines the Cloud ROI

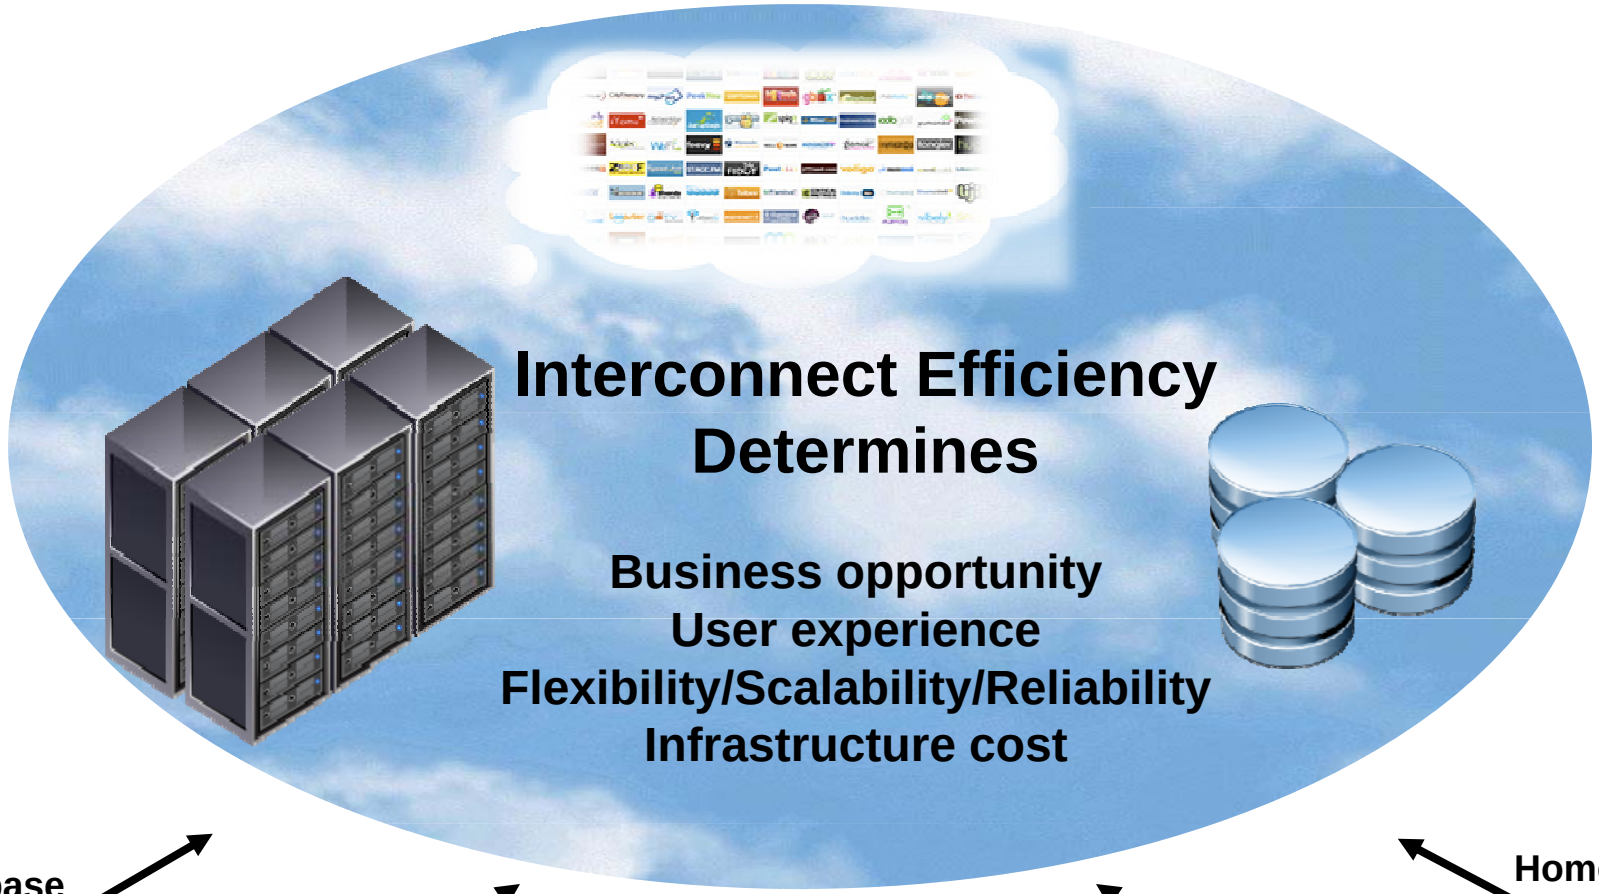

© 2010 MELLANOX TECHNOLOGIES

Summary

- Lower operational cost
  - Lower CapEx
    - Fewer cards, switch port, cables
  - Lower OpEx
    - Lower power, cooling, space
- SLA boost
  - Faster VM migration
  - Maximized VM performance
- Higher ROI
  - Improved resources utilization

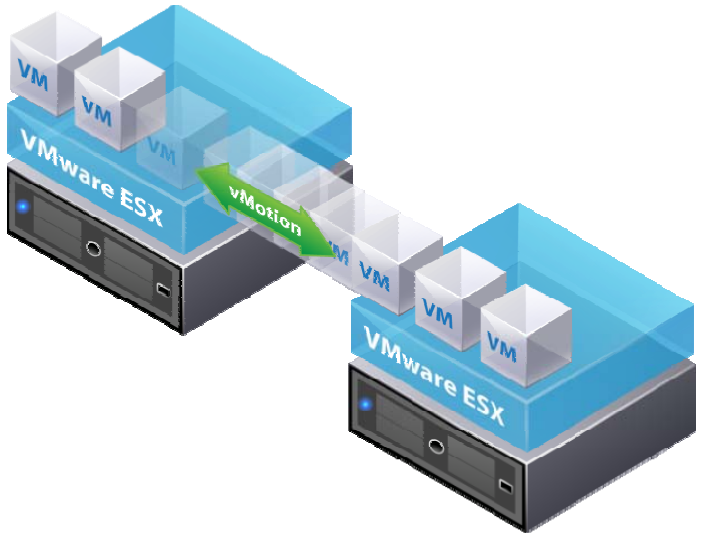

© 2010 MELLANOX TECHNOLOGIES