

Virtualization in Data Center

The Network Perspective

Michael Kagan

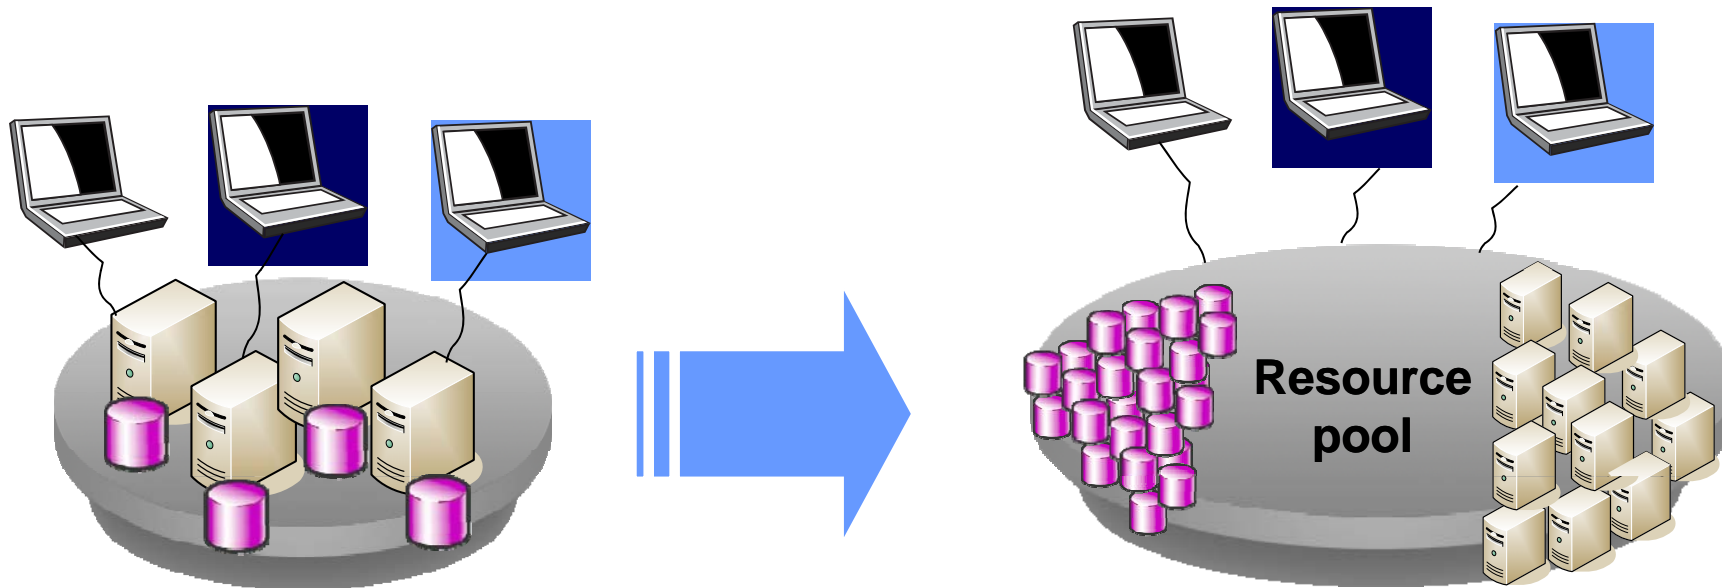
CTO, Mellanox Technologies

michael@mellanox.com




- Data Center Transition
- Servers as a Service
- Network as a Service
- IO as a Service
- Building Blocks
- Life Cases

Enterprise Datacenter Evolution



From equipment warehouse to service provider

Servers' Consolidation




News – nothing more than well-forgotten history

Service-Oriented View

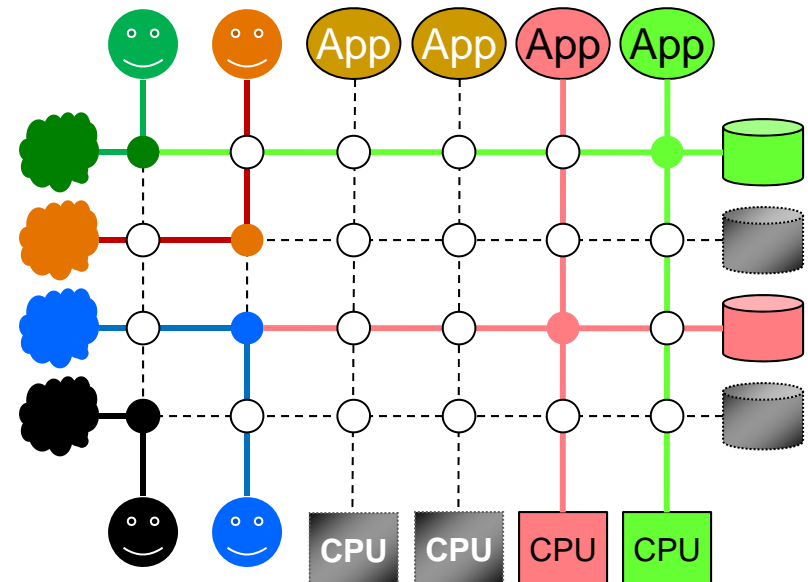
- **Hardware-centric view**


- Assign HW to consumer



- **Application-centric view**

- Deliver service to consumer

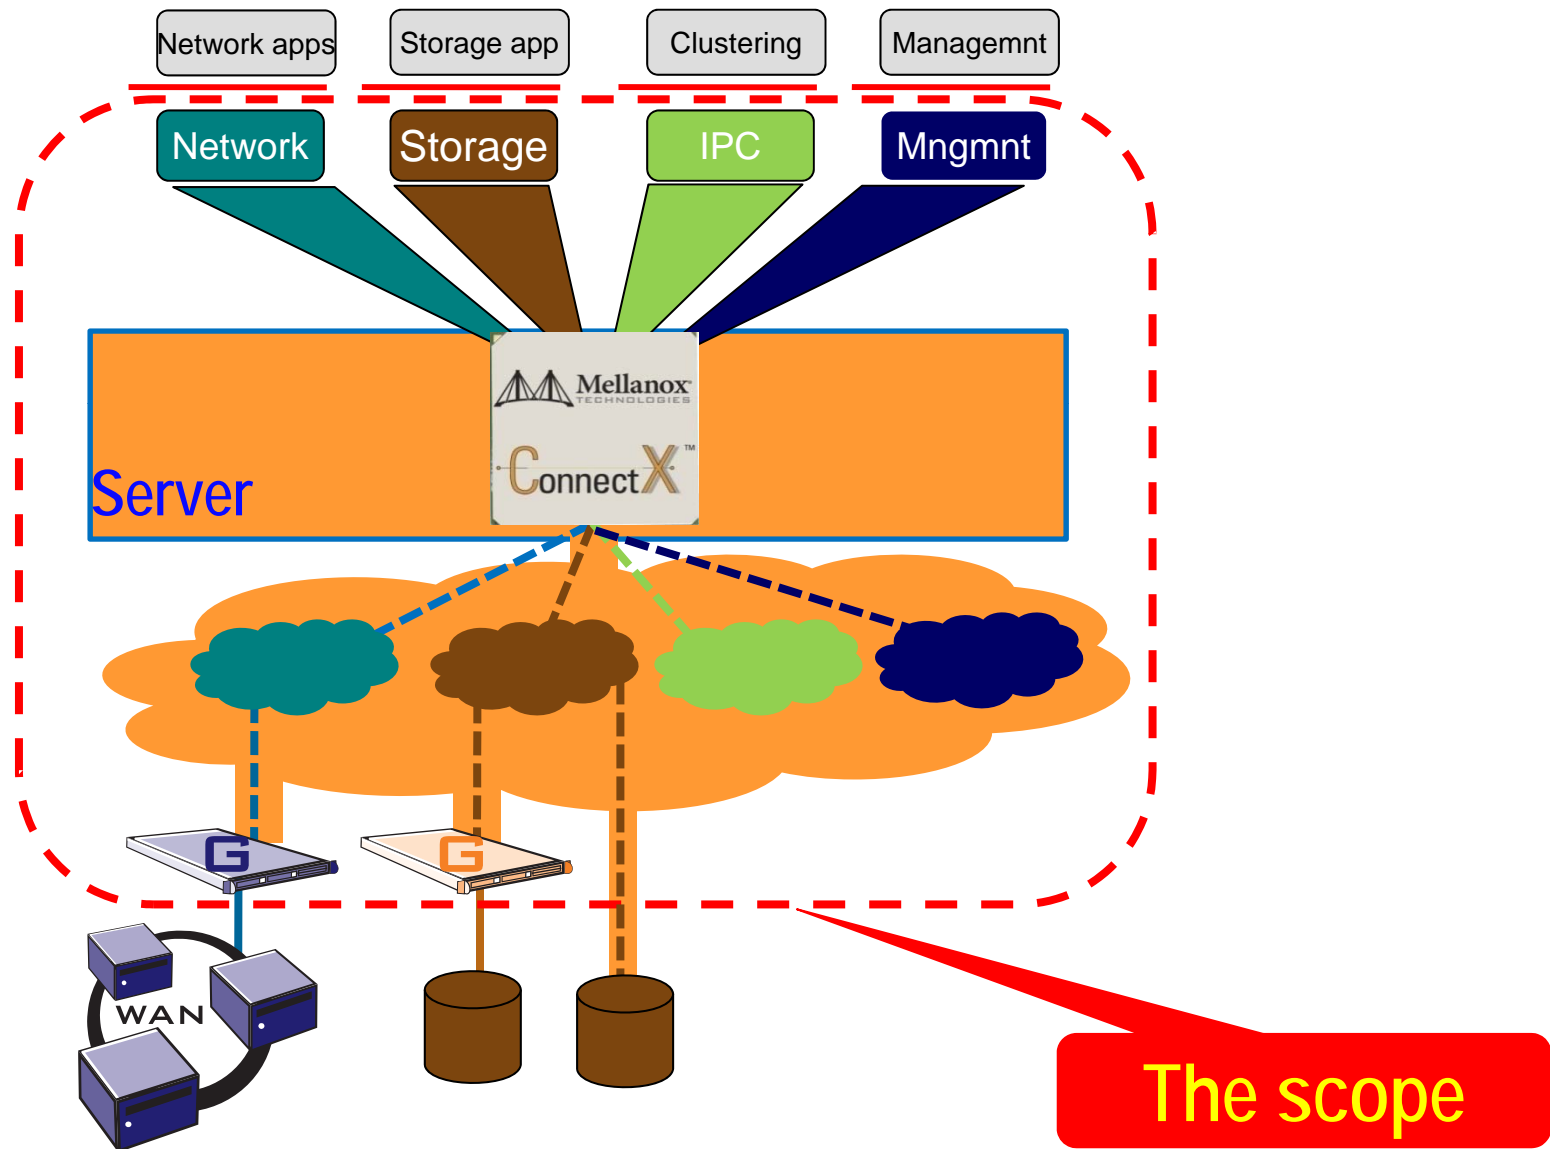





- **Highly-capable Hardware**
 - Standard multi-core CPUs, fast and large memory
- **Hardware Abstraction Layer**
 - Server Hypervisor
- **Management Infrastructure & tools**
 - Monitoring and provisioning

Computing as a Service

IO Services' Delivery – the Scope




Interconnect Consolidation




Find a Difference

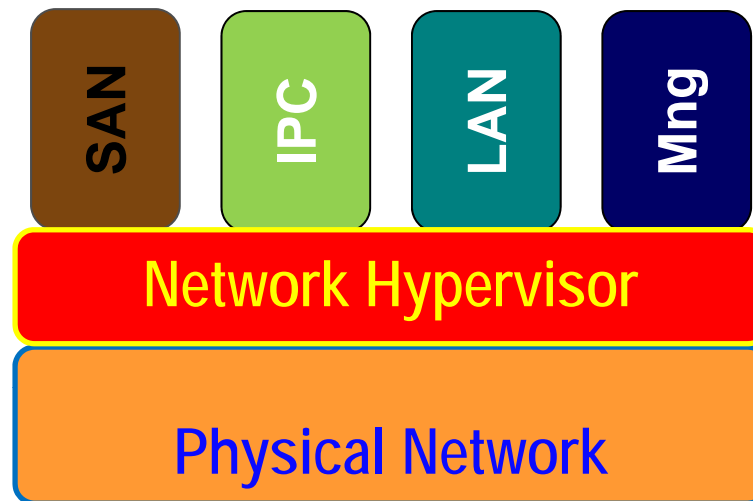


Converged Networks



Unified Network





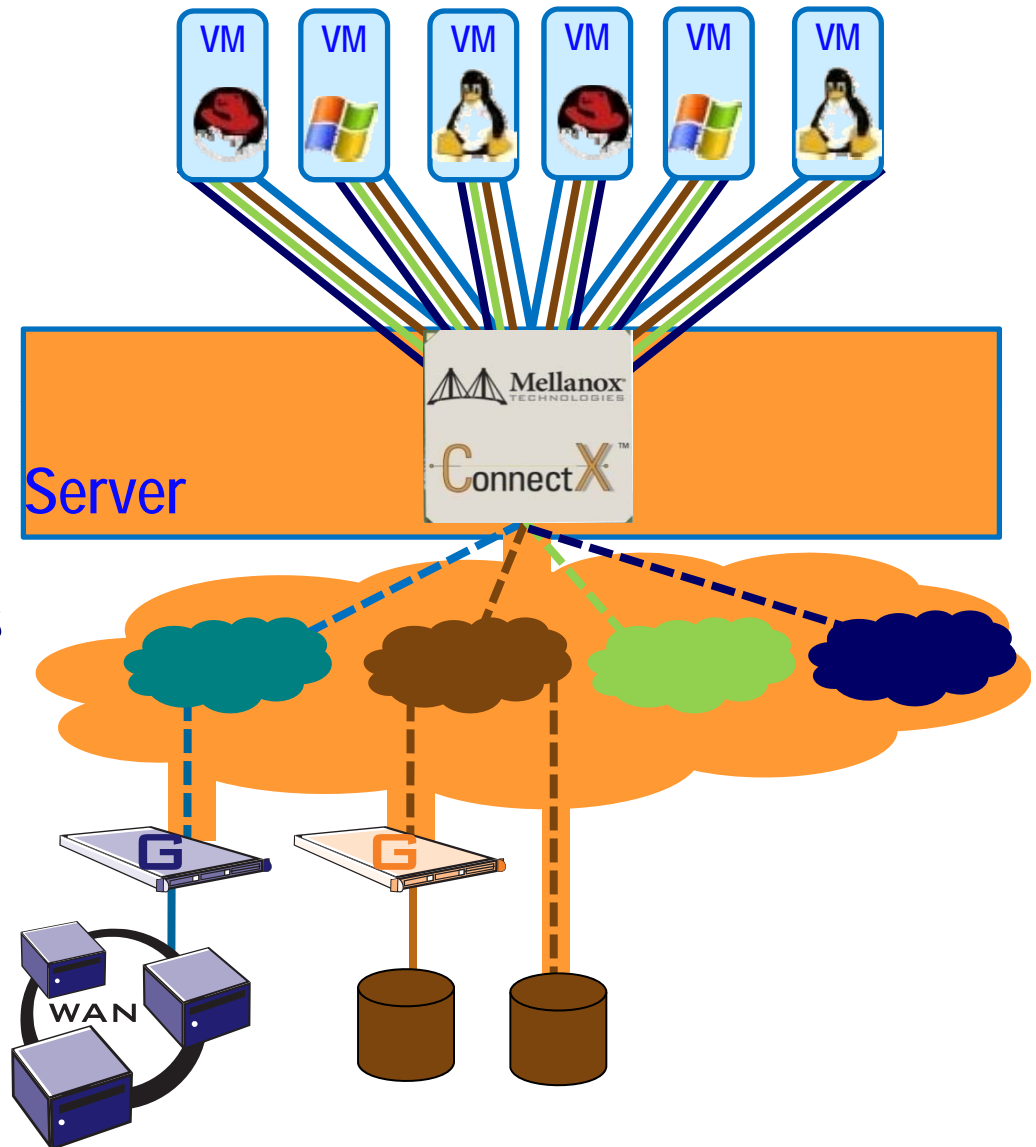
- **Highly-capable Hardware**
 - Standard high-performance NICs and switches
- **Hardware Abstraction Layer**
 - Network Hypervisor
- **Management Infrastructure & tools**
 - Monitoring and provisioning

Network as a Service

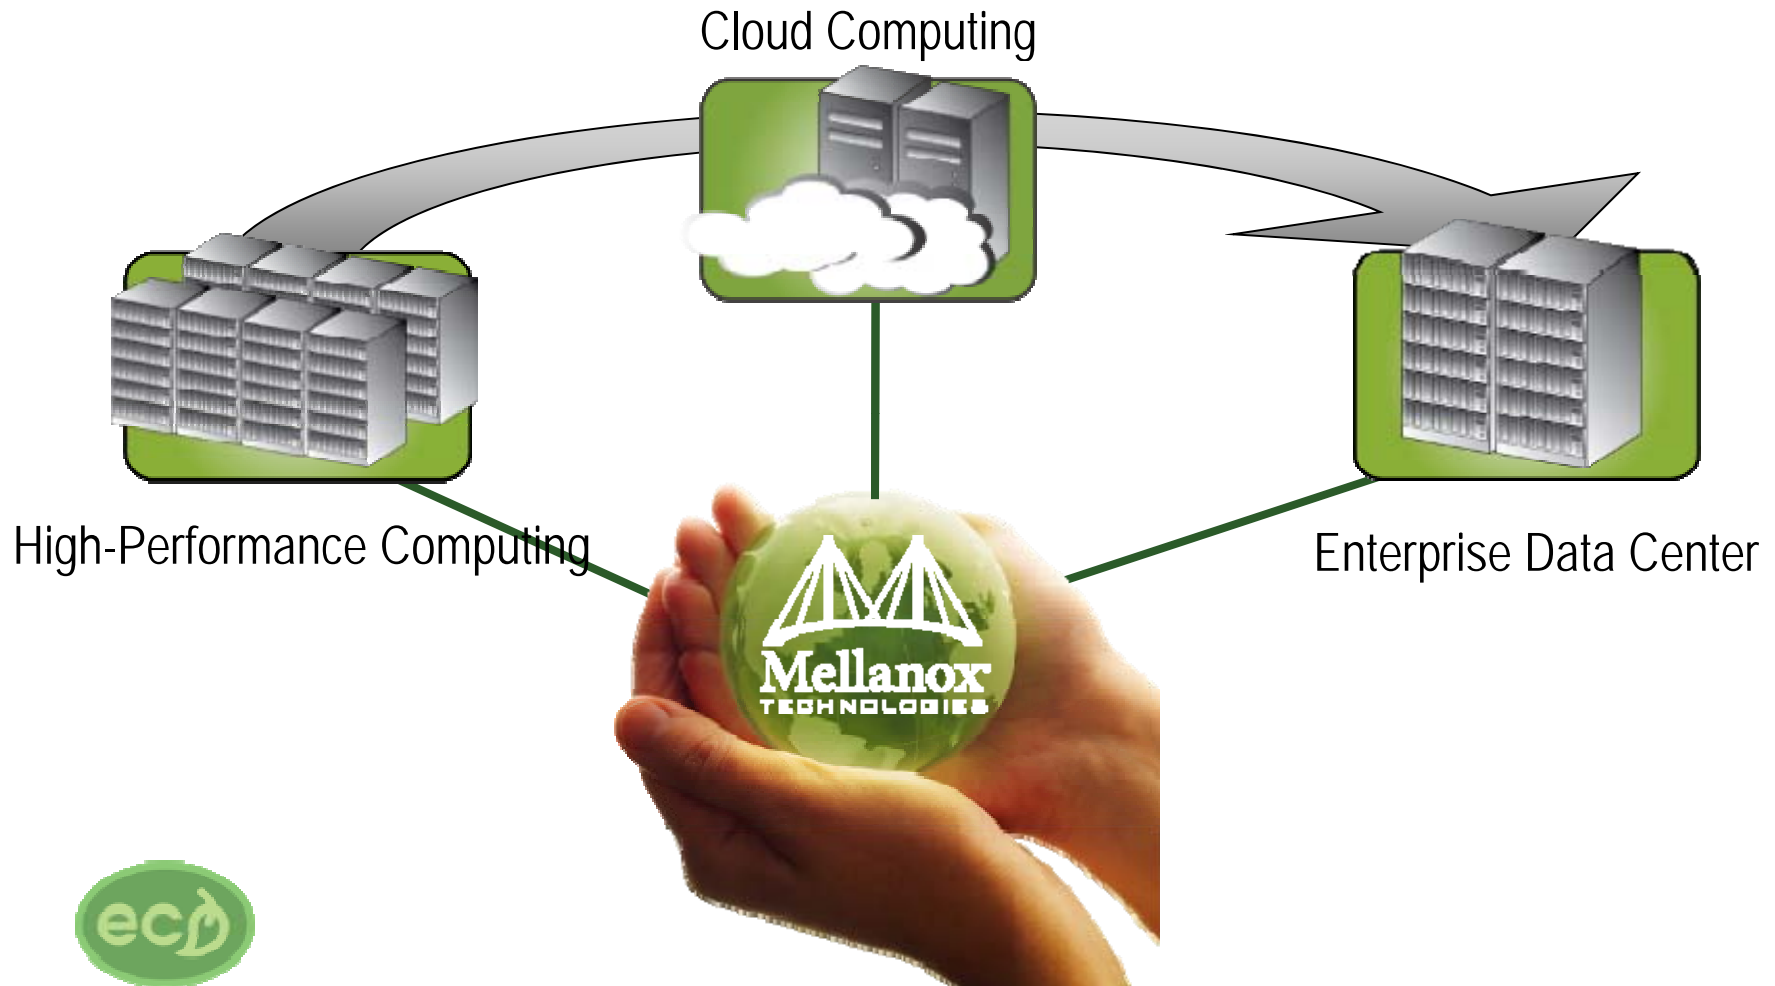
IO as a Service



- **Highly-capable IO adapter**
 - High performance
 - All IO services
 - Standard IO interfaces
- **Highly-capable Network**
 - High Performance
 - Virtual Networks
- **Highly-capable legacy bridges**
 - High Performance
 - High Flexibility
- **Hardware Abstraction Layer**
 - Server Hypervisor
 - Network Hypervisor
 - Integrated management



Mellanox – Network IO Solutions




Leading Connectivity Solution Provider For Servers and Storage

End-to-end Datacenter Networking Solutions




Adapters

10/20/40Gb/s InfiniBand and 10GE Ethernet



Switches

10/20/40Gb/s InfiniBand
From 8 to 648 ports in a single chassis




Converged Network



Bridges



InfiniBand to Ethernet and Fibre Channel



Ethernet to Fibre Channel

FCoE Gateway
12 ports of 2/4/8G FC

Software and Management



“Cloud” Data Center – Implementation



Intalio|Cloud Appliance



Compute Blade: 16 x HP ProLiant BL490c G6
2 x Quad Core Intel Xeon X5570 2.93GHz (8 cores total)
144GB (18 x 8GB) 1333MHz PC3-10600 DDR3 Registered DIMMs
No embedded storage



Database Storage: HP StorageWorks SB40c Storage Blade
6 x OCZ Vertex Series SATA II 2.5" SSD 250GB in RAID 0 configuration (striping)
20,000 IO/s, 1.5GB/s read, 1GB/s write, 100 microseconds latency
Used for database storage



File Storage: 4 x HP StorageWorks 600 Modular Disk System
70 LFF SAS or SATA Drives, 3Gb/s SAS connectivity
140TB in RAID 6 configuration (stripped set with dual distributed parity)
Used for file storage and database backup



Network Switch: 2 x HP 4X DDR InfiniBand Switch Module
20Gbps per port, 1.3 microseconds latency
16 Internal 4X DDR downlinks, 16 external 4X DDR QSFP uplinks
16 Internal 4X DDR downlinks, 8 external 4X DDR CX4 uplinks

System Requirements:

Customer Accounts: 2,304
Number of Users: 62,208 (27 per Account)
Memory per Account: 1GB
Database Storage per Account: 25GB
File Storage per Account: 250GB

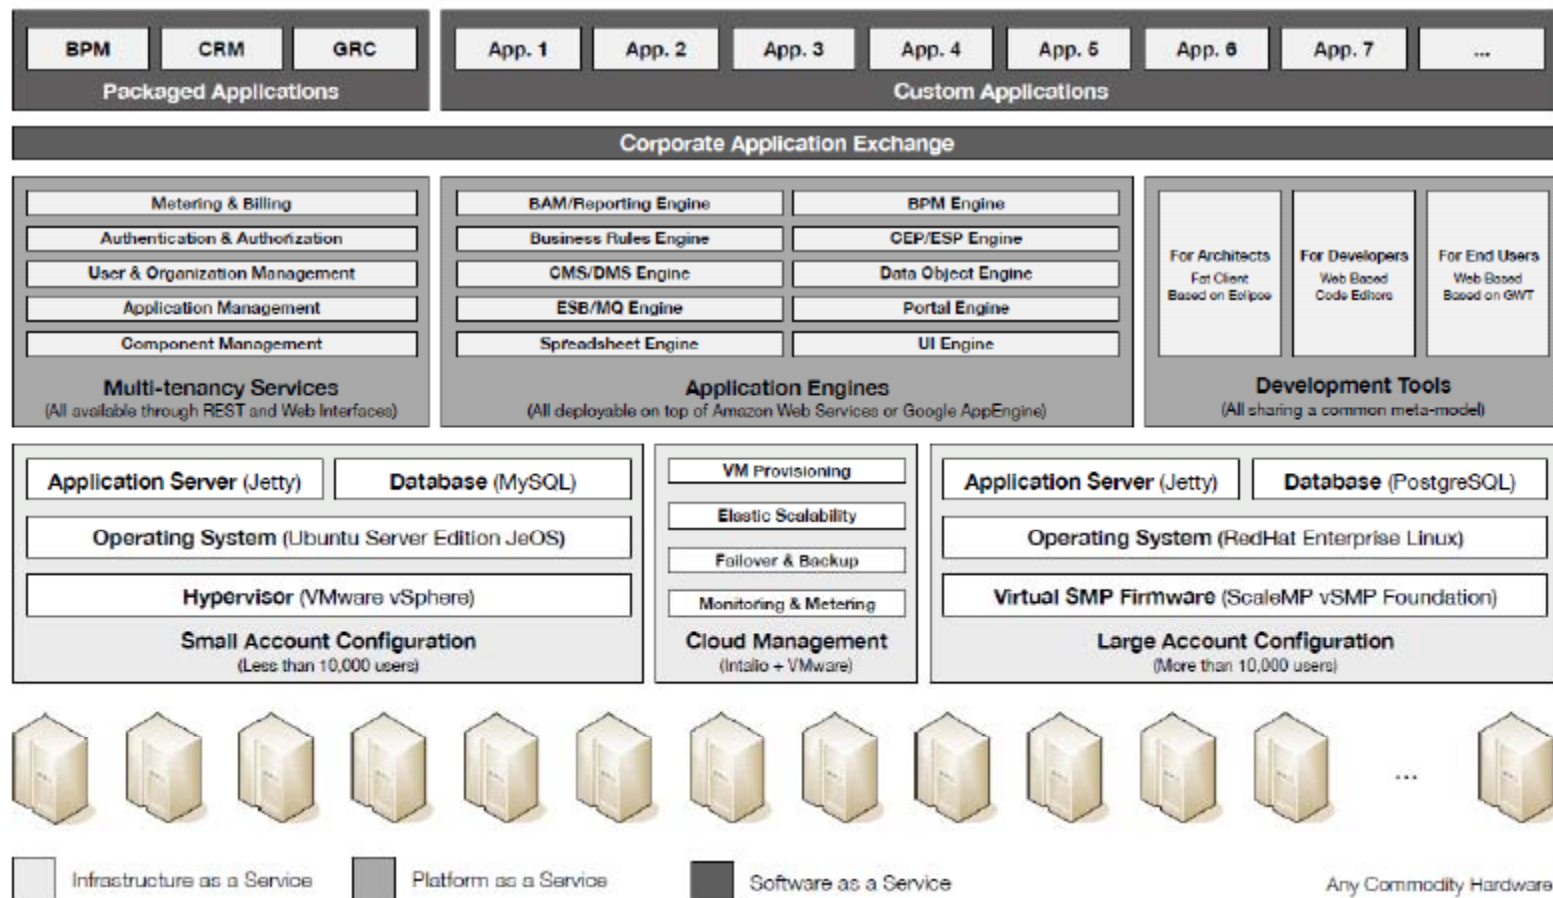
Configuration Overview:

32 CPUs, 128 CPU Cores
2.3TB Memory
24TB SSD Database Storage
628TB HDD File Storage
Hardware Cost: \$678K (\$19K/Month)



Source: Intalio

“Cloud” Architecture Integration of server and network hypervisors





Source: Intalio

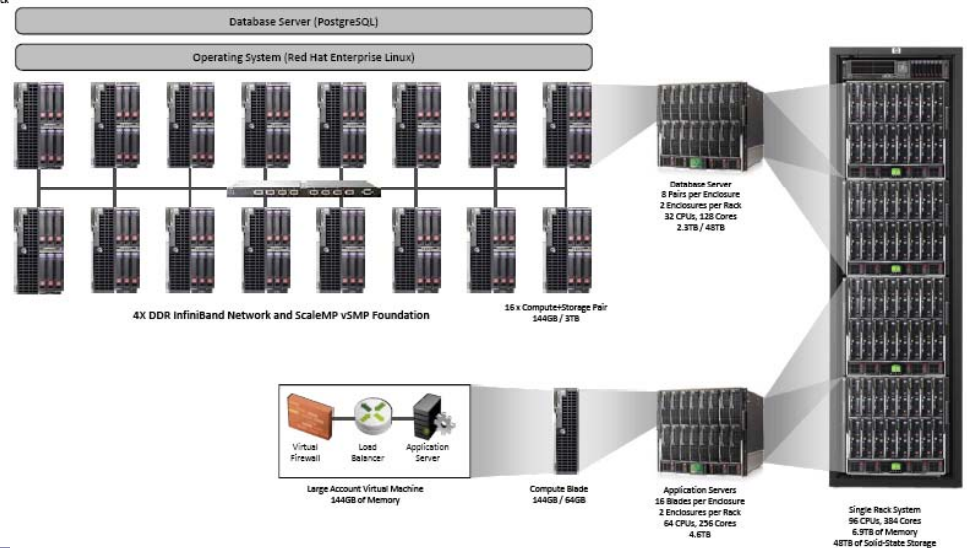
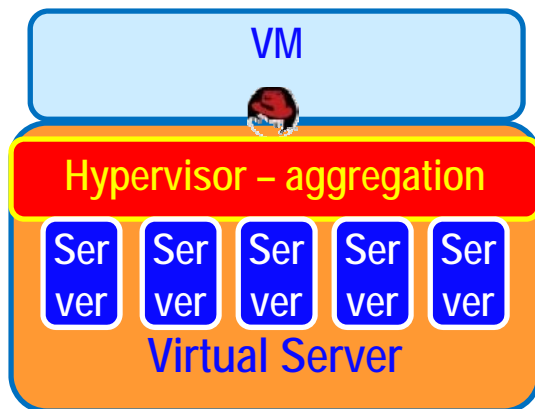
“Cloud” services



Small Account Configuration



Large Account Configuration



Source: Intalio

EDC Verticals & Case Studies



- Data Warehousing / Business Intelligence



Data Tier

- Managed Hosting/ Cloud Services



Logic and Data Tier

- Financial Services



Logic and Data Tier

Major Investments To Back Up Your Effort



- **HP blades**
 - Compute and storage
- **InfiniBand interconnect**
 - 20Gbit DDR
- **Oracle software**
 - Optimized for InfiniBand

**Runs Oracle
10x Faster***

The World's Fastest Database Machine

- Hardware by HP
- Software by Oracle

***But you have to be willing to spend 50% less on hardware.**

ORACLE®

10x faster based on comparing Oracle data warehouses on customer systems vs. HP Oracle Database Machines. Potential savings based on cost/terabyte comparison with selected competitor systems; Oracle Database and options licenses not included. Actual results and savings may vary.

Copyright © 2009, Oracle. All rights reserved. Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

#1 Price/Performance TPC-H over 11g Benchmark

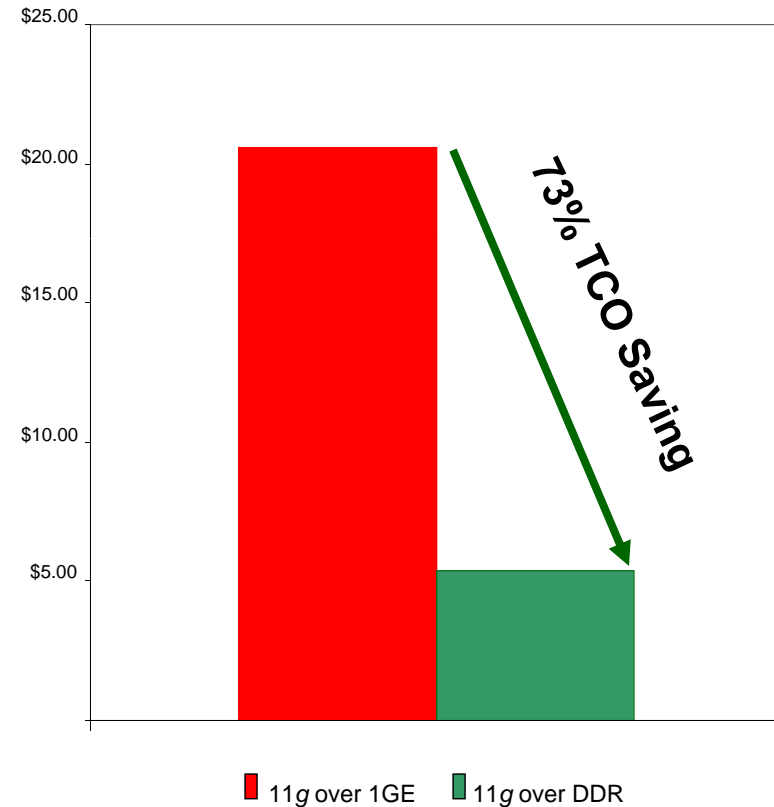


■ 11g over DDR

- Servers: 64 x ProLiant BL460c
 - CPU: 2 x Intel Xeon X5450
 - Quad-Core
- Fabric: Mellanox DDR InfiniBand
- Storage:
 - Native InfiniBand Storage
 - 6 x HP Oracle Exadata

	HP BladeSystem c-Class 128P RAC with HP Oracle Exadata Storage Servers	TPC-H Rev. 2.8.0 Report Date: June 3, 2009
	Total System Cost \$6,320,001USD	Composite Query per Hour Metric 1,166,976.6 QphH@1000GB

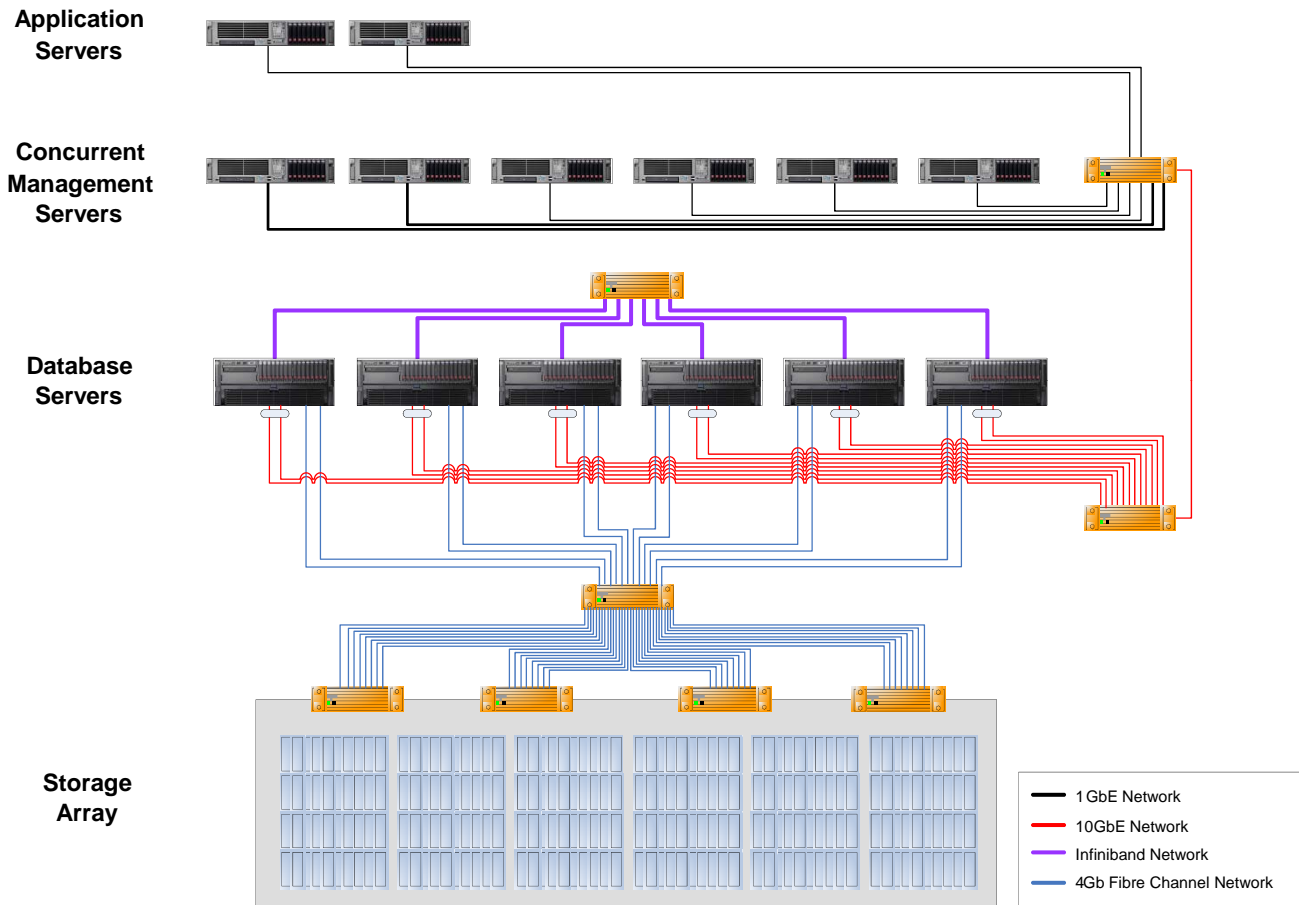
Price / QphH* @1000GB DB



World Record clustered TPC-H Performance and Price/Performance

* Query per Hour

Oracle Data Warehousing Case Study




- Application Servers**
 2x HP BL480C
 2 Processors / 8 core X560 3.16GHz
 64GB RAM
 4x 72GB 15K drives
 NIC: HP NC373i 1GB NIC
- Concurrent Manager Servers**
 6x HP BL480C
 2 Processors / 8 core X560 3.16GHz
 64GB RAM
 4x 72GB 15K drives
 NIC: HP NC373i 1GB NIC
- Database Servers**
 6x HP DL580 G5
 4 processors / 24 cores X7460 2.67GHz
 256GB RAM
 8x 72GB 15K drives
 NIC: Intel 10GbE XF SR 2 port PCIe NIC
 Interconnect: Mellanox 4x PCIe InfiniBand
- Storage Array**
 HP XP24000
 64GB cache / 20GB shared memory
 60 Array Groups of 4 spindles
 240 spindles total
 146GB 15K fibre channel disk drives

Datacenter hardware configuration

InfiniBand deliver 63% more TPS vs. 10GE

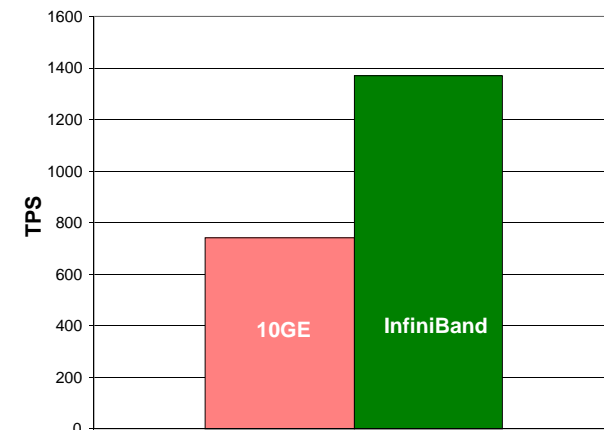


- TPS Rates for invoice load use case



	Activity	Start Time	End Time	Duration	Records	TPS
InfiniBand Interconnect						
1	Invoice Load - Load File	6/17/09 7:48	6/17/09 7:54	0:06:01	9,899,635	27,422.81
2	Invoice Load - Auto Invoice	6/17/09 8:00	6/17/09 9:54	1:54:21	9,899,635	1,442.89
3	Invoice Load - Total	N/A	N/A	2:00:22	9,899,635	1,370.76
10 GigE interconnect						
1	Invoice Load - Load File	6/25/09 17:15	6/25/09 17:20	0:05:21	7,196,171	22,417.98
2	Invoice Load - Auto Invoice	6/25/09 18:22	6/25/09 20:39	2:17:05	7,196,171	874.91
3	Invoice Load - Total	N/A	N/A	2:22:26	7,196,171	842.05

- Work Load
 - Nodes 1 through 4: Batch processing
 - Node 5: Extra Node not used
 - Node 6: EBS Other Activity
- Database size (2 TB)
 - ASM
 - 5 LUNS @ 400 GB




InfiniBand needs only 4 servers vs. 10 Servers needed by 10GE

Managed/Cloud Service Provider Case Study




■ VMware-based server virtualization application

Without Mellanox InfiniBand



256 VMs
4 racks, 4U servers
4X proc, 16GB
16 edge switches
192 I/O cards
\$744K capital cost

With Mellanox



256 VMs
1 rack, 1U servers
4X proc, 16GB
2 I/O Directors
32 I/O cards
\$347K capital

Savings

Cap Ex: **\$397,000**

Op Ex:

Floor space: **\$54,000**

Power: **\$20,000**

MAC*: **\$104,000**

TOTAL \$575,000

SAVINGS

(32 servers over 3 years)

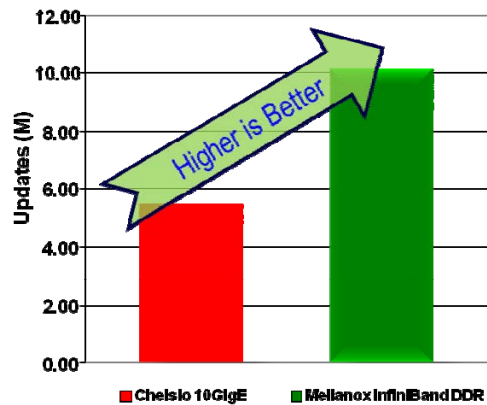
*2 moves/adds/changes per server per year

Source: Xsigo

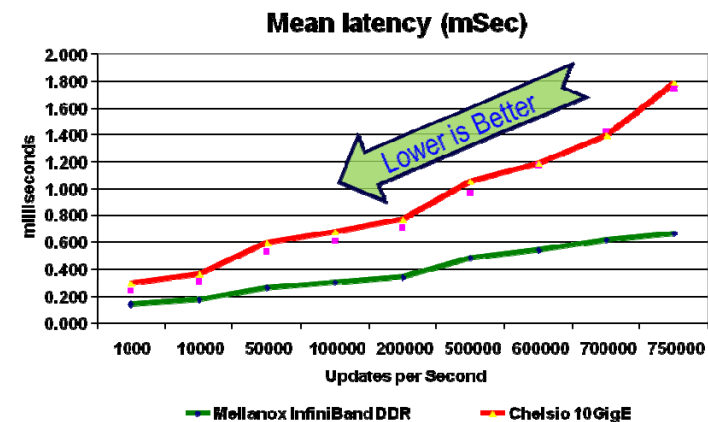
Financial Services: InfiniBand Delivers Performance & ROI Metrics



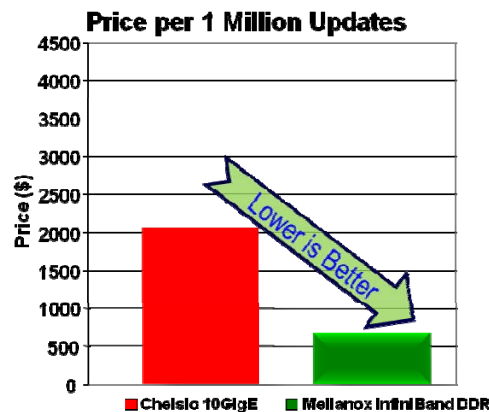
- 82% higher updates/sec



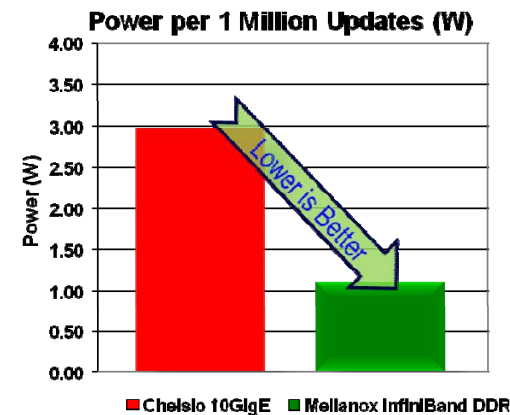
- 62% lower mean latency



- Costs 70% lower



- 3X less power consumption



Source: STAC

- **Virtualization to save Data Center TCO**
 - These are old news, but worth to remember

- **Network Virtualization – a mean to deliver services**
 - Highly-capable interconnect & management infrastructure

- **InfiniBand – efficient network virtualization backbone**
 - Available today, higher ROI than 10GE

- **InfiniBand deployment in data centers today**
 - Data and Logic tiers