

# 私有云平台上优化SQL Server

*Optimize SQL Server for Private Cloud*

殷皓 Howard Yin  
商务战略总监  
微软服务器与开发工具事业部  
中国云计算创新中心

# 课程目的及要点

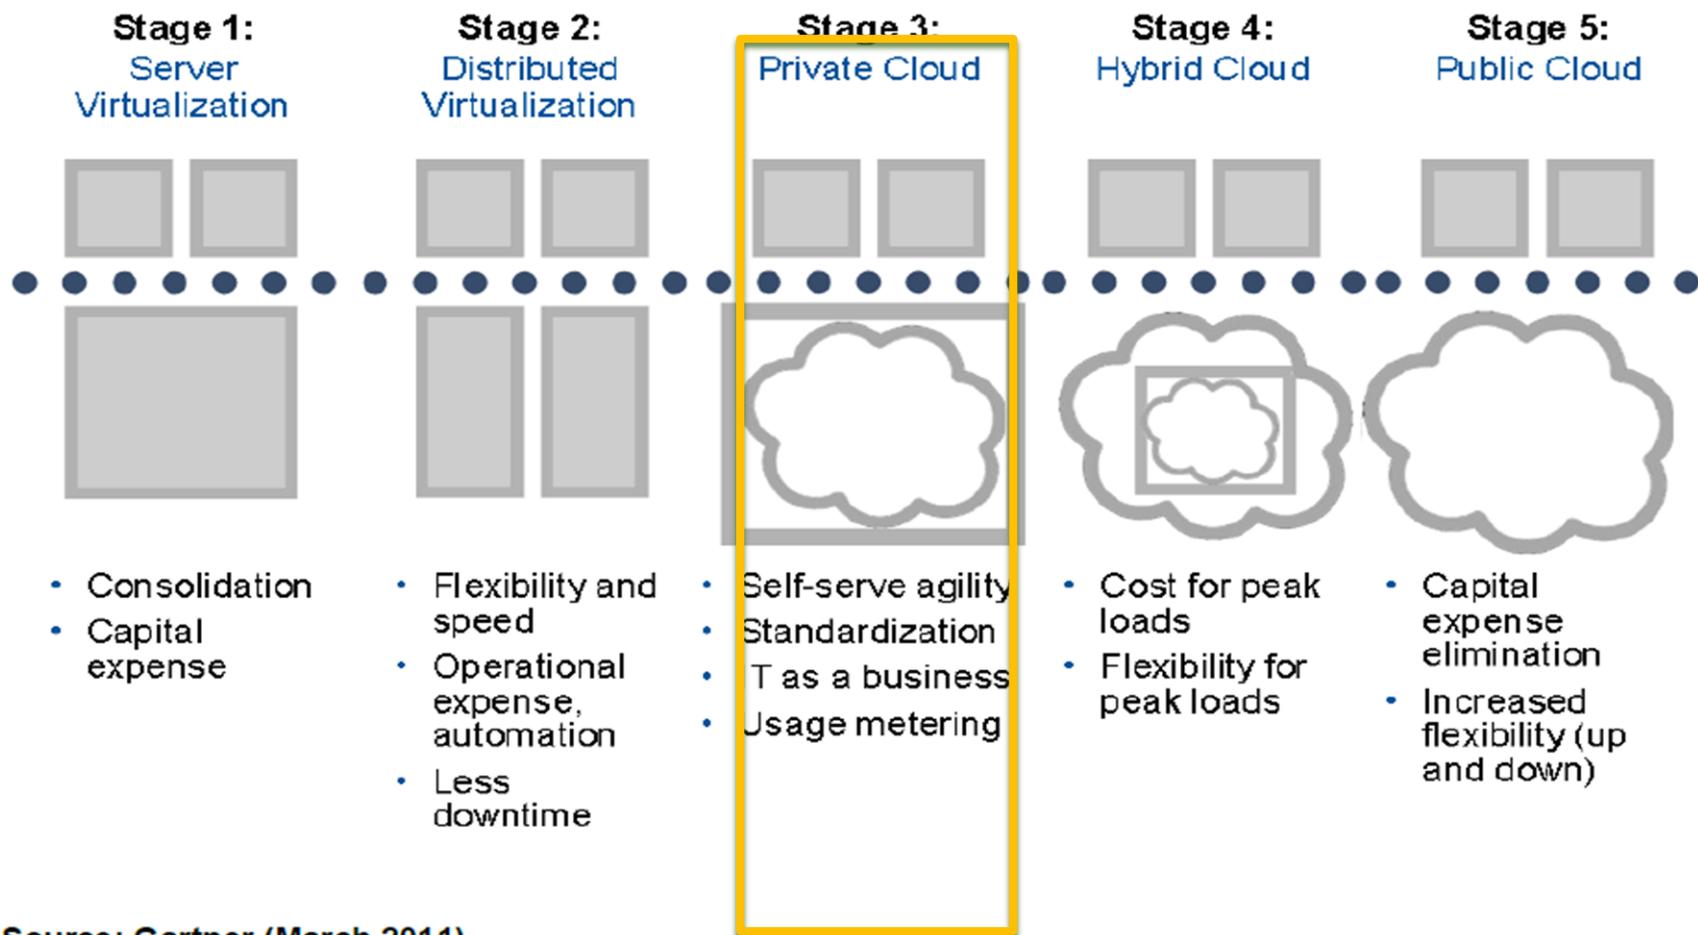
- 介绍SQL@私有云现状、路线图及关键技术
- 私有云的关键因素和优化SQL的最佳做法

- 什么是私有云？
- 怎样为私有云优化SQL Server？
- 怎样实现更高的密度？
- 真实案例

- 什么是私有云？

# 云征途中的五大阶段

Figure 1. Gartner Road Map: From Virtualization to Cloud Computing



Source: Gartner (March 2011)

Source: Gartner (March 2011)

# 私有云正在落地！

- 70% 的客户正在部署或计划私有云<sup>1</sup>
- 美国一半以上的政府部门会在12月内部署私有云<sup>2</sup>
- 虚拟化和云计算的普及在未来12月内将会显著提升<sup>3</sup>
- 自助服务正在迅速赶上虚拟化<sup>4</sup>
- 超过50%的客户 (Forrester)

<sup>1</sup> [AFCOM 2011 study](#)

<sup>2</sup> [Information Week 2011 survey](#)

<sup>3</sup> [Worldwide Executive Council 2011 survey on 100 CIOs](#)

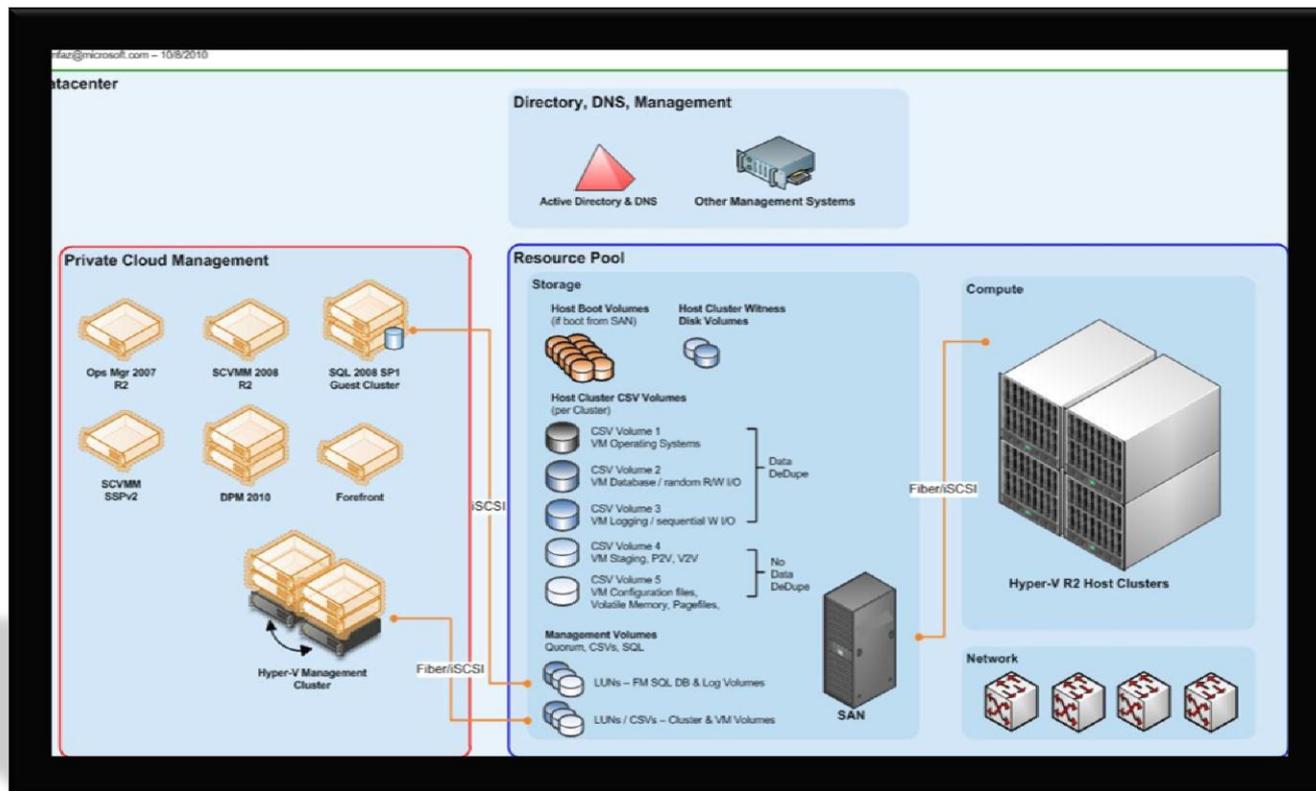
<sup>4</sup> [CompTIA 2011 cloud computing survey](#)

Top 3 Database Initiatives	Currently doing or planning in next 12 months
Database Virtualization	58%
Database Standardization	56%
Database Consolidation	51%

Source: Forrester

# 逻辑性架构

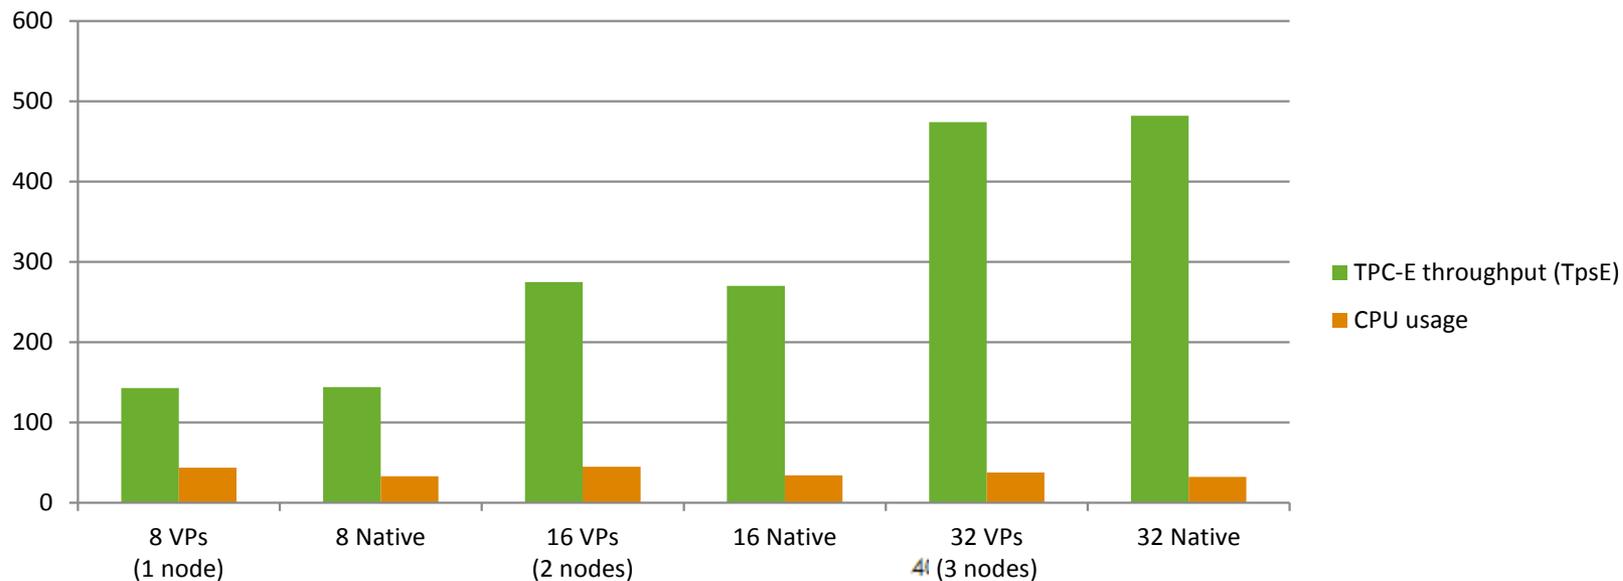
- 计算  
Compute
- 网络 Network
- 存储 Storage
- 管理  
Management



- 什么是私有云？
- 怎样为私有云优化SQL Server？

## 需要多大的服务器？

	Windows 2008 R2	Windows Server 8
CPU	4 CPUs	32 CPUs
Memory	64GB	512GB
VHD	2TB	>2TB
Host	64 LP	160 LPs



# 性能优化

	Win 2008 R1/R2	Windows Server 8
Enlightened VM	X	X
Better IO with Pass-through or Fixed VHDs	X	X
Minimize Roles in Host Partition	X	X
Virtualization Capability in New Processors	X	X
<u>Hyper-V on Server Core</u>	X	X
Guest Virtual Fiber Channel		X
Hyper-V over SMB2		X
VHDX		X
Guest networking (SR-IOV, VMQ)		X

# Server Core Installation 益处及选项

- **Benefits**

- Reduced maintenance.
- Reduced attack surface.
- Reduced management.
- Less disk space required.

- **Roles Supported (by Edition):**

KEY: ○ = Not Available    ● = Partial/Limited    ☑ = Full

Server Role	Enterprise	Datacenter	Standard	Web	Itanium	Foundation
Active Directory Certificate Services	☑	☑	☑	○	○	○
Active Directory Domain Services	☑	☑	☑	○	○	○
Active Directory Lightweight Directory Services	☑	☑	☑	○	○	○
BranchCache Hosted Cache	☑	☑	○	○	○	○
DHCP Server	☑	☑	☑	○	○	○
DNS Server	☑	☑	☑	☑	○	○
File Services	☑	☑	● <sup>1</sup>	○	○	○
Hyper-V	☑	☑	☑	○	○	○
Media Services*	☑	☑	☑	☑	○	○
Print Services	☑	☑	☑	○	○	○
Web Services (IIS)	☑	☑	☑	☑	○	○

<sup>1</sup> Limited to one standalone DFS root.

\* Must be downloaded separately.

- Install Hyper-V on Server Core
- `start /w ocsetup Microsoft-Hyper-V`
- Using Deployment Image Servicing and Management (DISM)
  - Discover available features:
  - `Dism /online /get-features /format:table`
  - To Install:
  - `Dism /online /enable-feature /featurename:<featurename>`
  - To Remove:
  - `Dism /online /enable-feature /featurename:<featurename>`
  - Beware, it is case sensitive
- Managing
  - Many options available, both locally and remotely
- Servicing

# 性能优化

	Win 2008 R1/R2	Windows Server 8
Enlightened VM	X	X
Better IO with Pass-through or Fixed VHDs	X	X
Minimize Roles in Host Partition	X	X
Virtualization Capability in New Processors	X	X
Hyper-V on Server Core	X	X
Guest Virtual Fiber Channel		X
Hyper-V over SMB2		X
VHDX		X
Guest networking (SR-IOV, VMQ)		X

# 大型的虚拟机的网络附属存储(NAS)

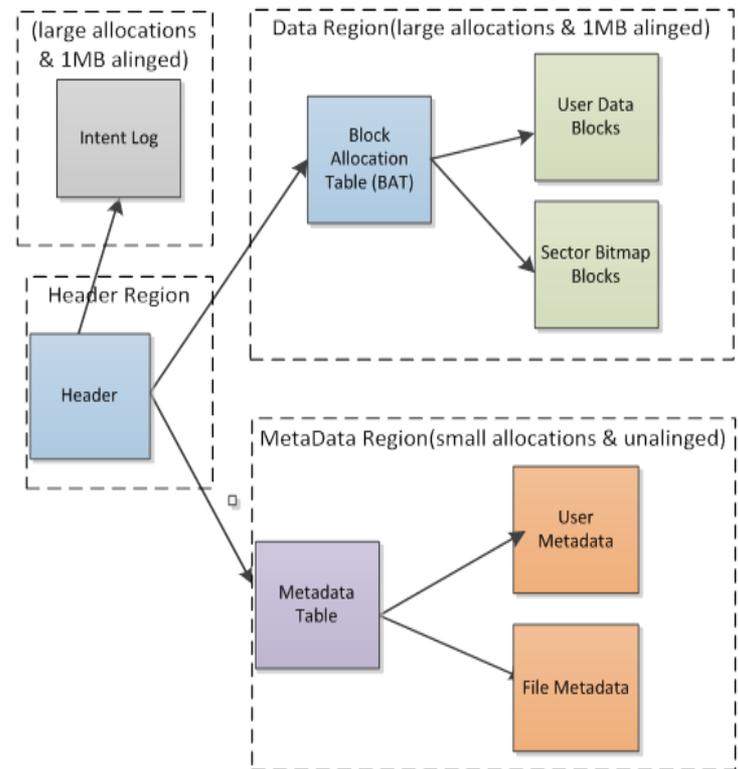
- Hyper-V Supported on SMB
- SMB comparable to FC
- Reliable (Transparent failover & multichannel)
- High performance & Low latency (RDMA)
- Works with existing infrastructure
  - Windows clustering (interchange CSV with SMB)
  - Support for remote VSS capabilities

# 最大限度提高性能

	Win 2008 R1/R2	Windows Server 8
Enlightened VM	X	X
Better IO with Pass-through or Fixed VHDs	X	X
Minimize Roles in Host Partition	X	X
Virtualization Capability in New Processors	X	X
Hyper-V on Server Core	X	X
Guest Virtual Fiber Channel		X
Hyper-V over SMB2		X
<u>VHDX</u>		X
Guest networking (SR-IOV, VMQ)		X

# VHDX

- >2TB disks
- **Better** performance (eliminate alignment issues)
- Resilient to corruption
- Embed user defined metadata
- Larger block sizes to adapt to workload requirements



# Maximize Performance

	Win 2008 R1/R2	Windows Server 8
Enlightened VM	X	X
Better IO with Pass-through or Fixed VHDs	X	X
Minimize Roles in Host Partition	X	X
Virtualization Capability in New Processors	X	X
Hyper-V on Server Core	X	X
Guest Virtual Fiber Channel		X
Hyper-V over SMB2		X
VHDX		X
Guest networking (SR-IOV, VMQ)		X

# SR-IOV (Single-Root I/O Virtualization)

## PCI-SIG STANDARD

- Expands virtualization concepts to PCI Express Devices
- Is not specific to an I/O class
- Allows a devices hardware surface to be 'assigned' to a VM
- Requires device, firmware, chipset and driver support

## FEATURE GOALS

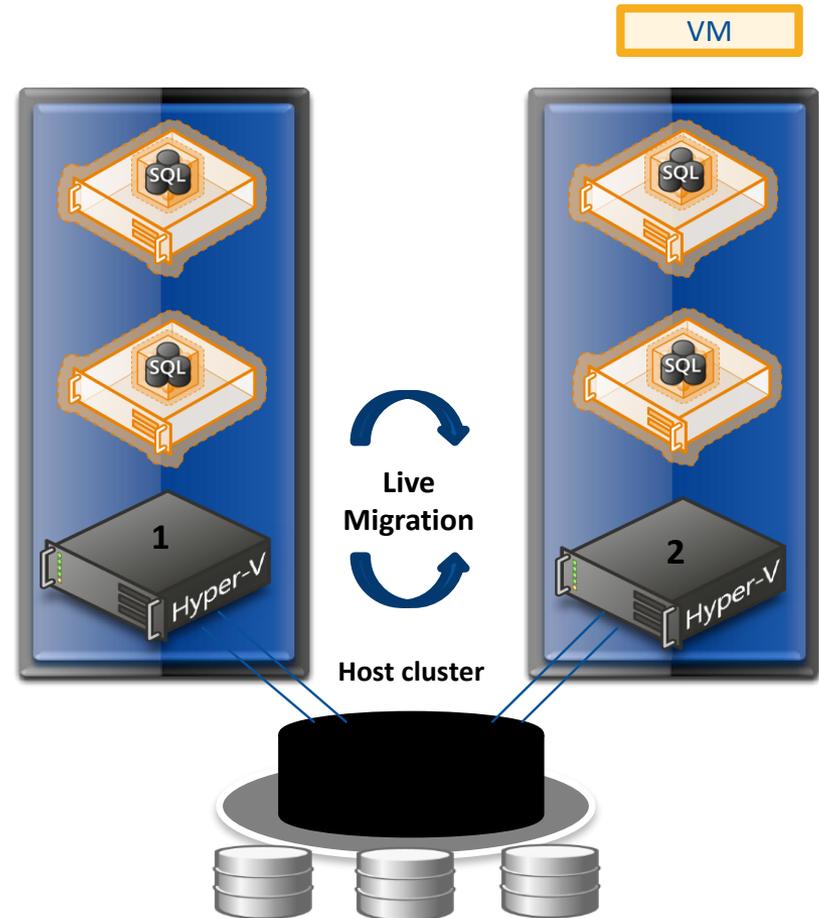
- Reduce latency of network path
- Reduce CPU utilization for processing network traffic
- Increased throughput
- Support Live Migration

# High Availability and Disaster Recovery

SQL 2005, 2008, 2008 R2, Denali	Windows 2008 R2	Windows Server8
Denali AlwaysOn	X	X
Database Mirroring	X	X
Guest Clustering	X	X
Database Replication	X	X
Log Shipping	X	X
Database Backup Restore	X	X
<u>Hyper-V Live Migration</u>	X	X
Hyper-V Replica		X
Hyper-V Live Storage Migration		X

# SQL Server with Live Migration

- Scenario description
- Virtualization benefits
  - No loss of service during failover with Live Migration.
  - Improve availability with less complexity
  - Better server utilization due to consolidation
  - Easier setup and management through System Center Virtual Machine Manager

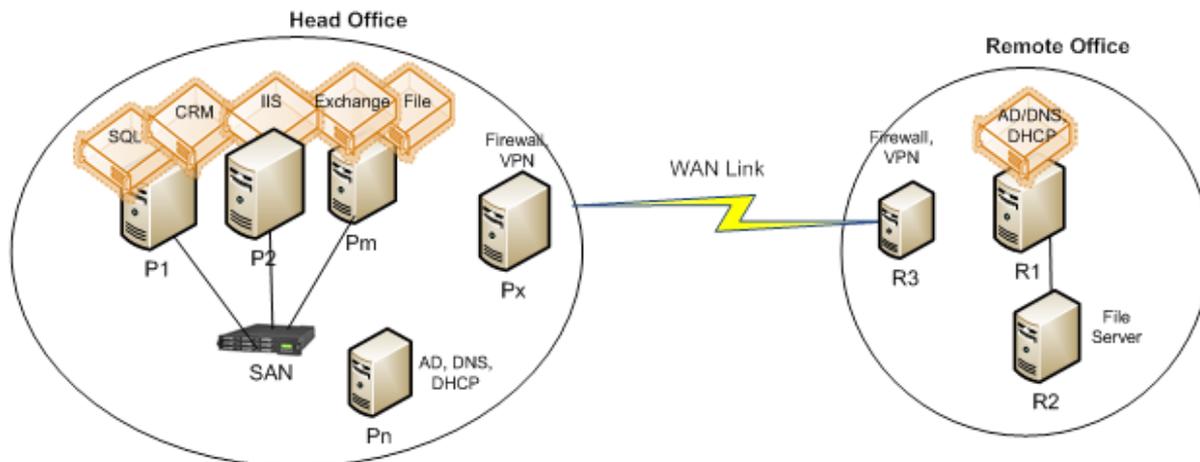


# High Availability and Disaster Recovery

SQL 2005, 2008, 2008 R2, Denali	Windows 2008 R2	Windows Server8
Denali AlwaysOn	X	X
Database Mirroring	X	X
Guest Clustering	X	X
Database Replication	X	X
Log Shipping	X	X
Database Backup Restore	X	X
Hyper-V Live Migration	X	X
Hyper-V Replica		X
Hyper-V Live Storage Migration		X

# Hyper-V Replica

- Part of Hyper-V install
- Hyper-V UI, PS and WMI support for all operations
- Integrated with Cluster Manager for all operations and status
- Asynchronous replication
- Support for asymmetric storage
- Works seamlessly across migrations
- Best Practices Analyzer (BPA) for getting set up correctly
- Status, perfmon counters and event logs for troubleshooting



# High Availability and Disaster Recovery

SQL 2005, 2008, 2008 R2, Denali	Windows 2008 R2	Windows Server8
Denali AlwaysOn	X	X
Database Mirroring	X	X
Guest Clustering	X	X
Database Replication	X	X
Log Shipping	X	X
Database Backup Restore	X	X
Hyper-V Live Migration	X	X
Hyper-V Replica		X
Hyper-V Live Storage Migration		X

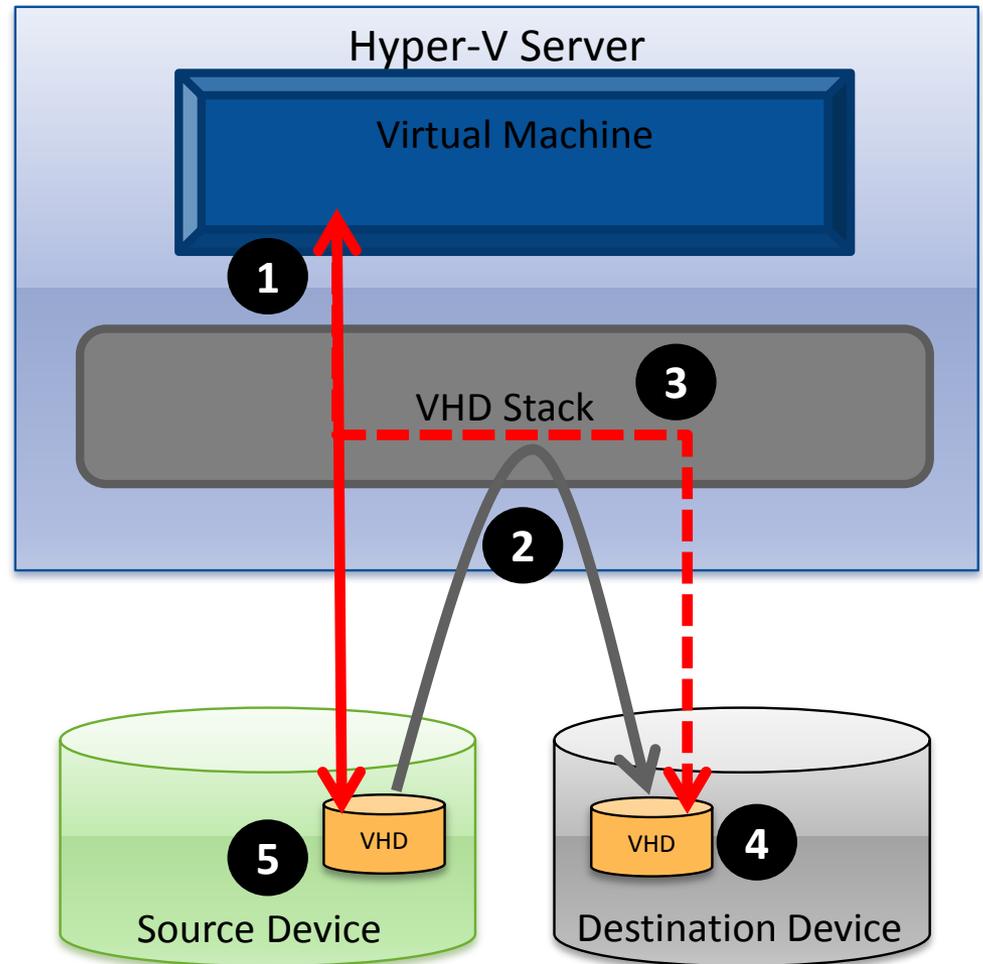
# Storage Migration

- ▶ The ability to move all “on disk” components of a virtual machine – while it is running
  - ▶ Virtual Hard Disks
  - ▶ Configuration Files
  - ▶ .BIN / .VSV / .SLP Files

# Storage Migration

## How does this all work?

- The ability to move all “on disk” components of a virtual machine – while it is running
  - Virtual Hard Disks
  - Configuration Files
  - .BIN / .VSV / .SLP Files



# Agenda

- What is Private Cloud?
- How to Optimize SQL for Private Cloud
- How to Achieve Higher Density?

# CPU Overcommit

- Can deliver higher virtualization density

	Windows Server 2008 R2	Windows Server 8
Virtual Processor – Host	64	160
Virtual Processor – Guest	4	32

- May introduce noticeable performance overhead when all workloads are busy
- Hyper-V benefits from newer generation of processor architecture (SLAT)

# CPU Over-commit

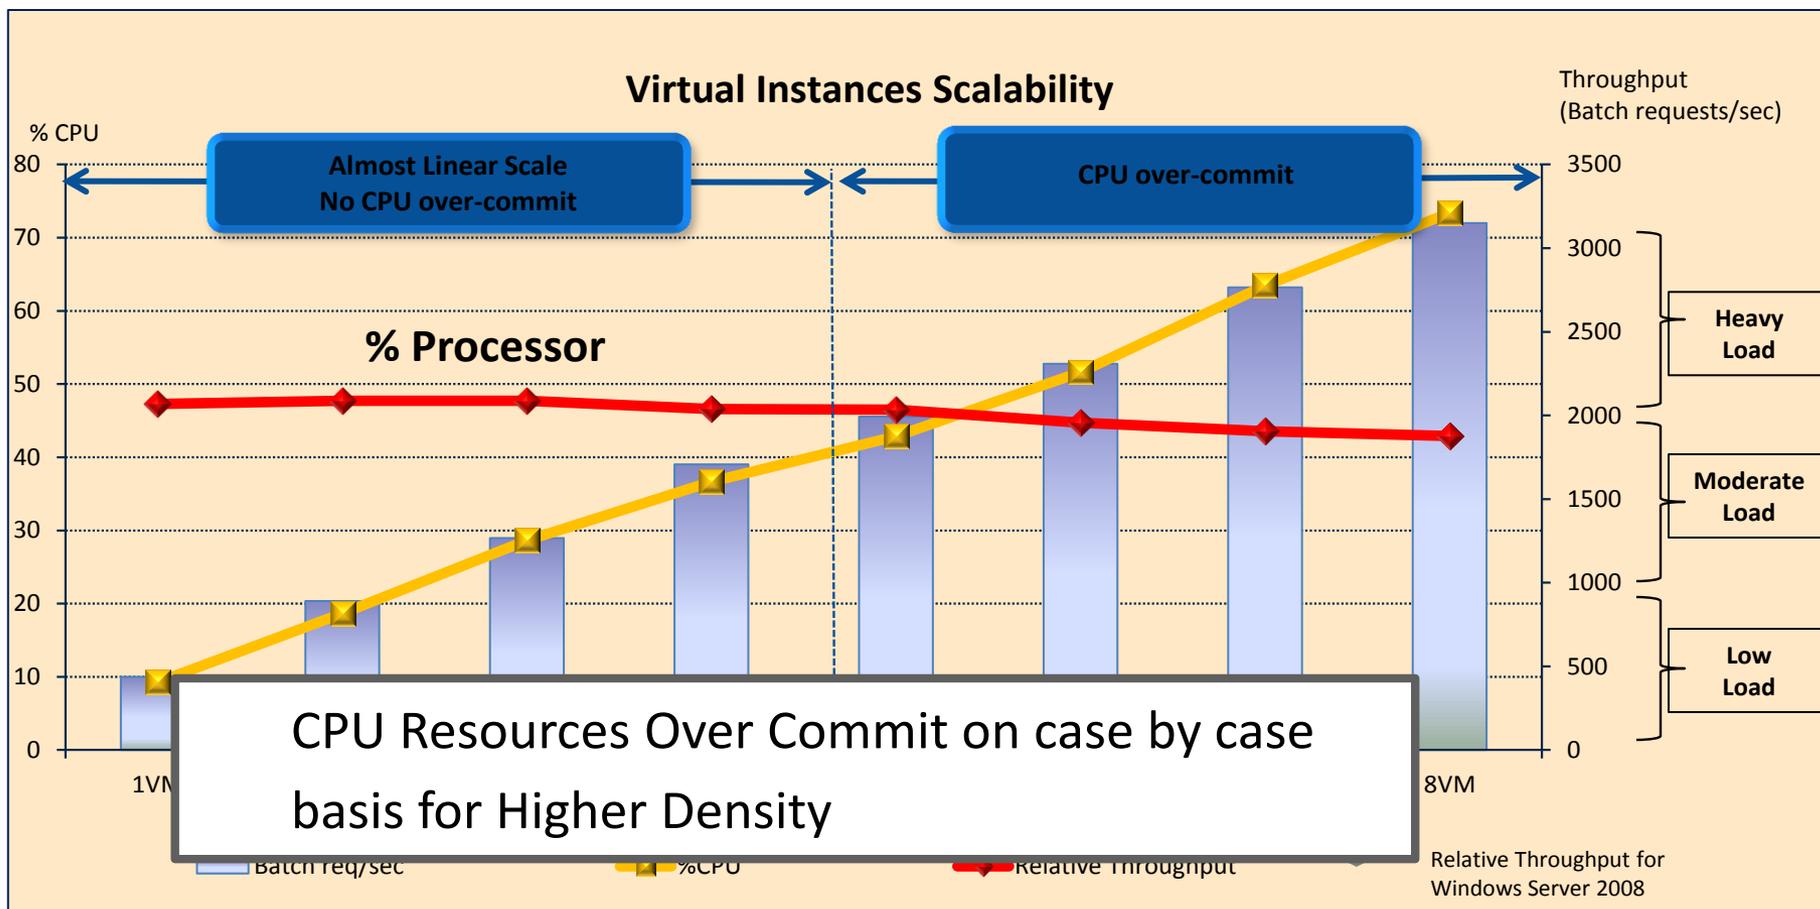
## Configuration:

- **OS:** Microsoft® Windows Server® 2008 R2 Hyper-V™
- **Hardware:**  
 HP DL585 (16 core) with SLAT  
 HP EVA 8000 storage
- **Virtual Machines:** 4 virtual processors and 7 GB RAM per virtual machine; Fixed size VHD



## Results:

- Increased throughput with consolidation
- Near linear scale in throughput with no CPU over-commit
- Improved performance with Windows Server 2008 R2 and SLAT processor architecture



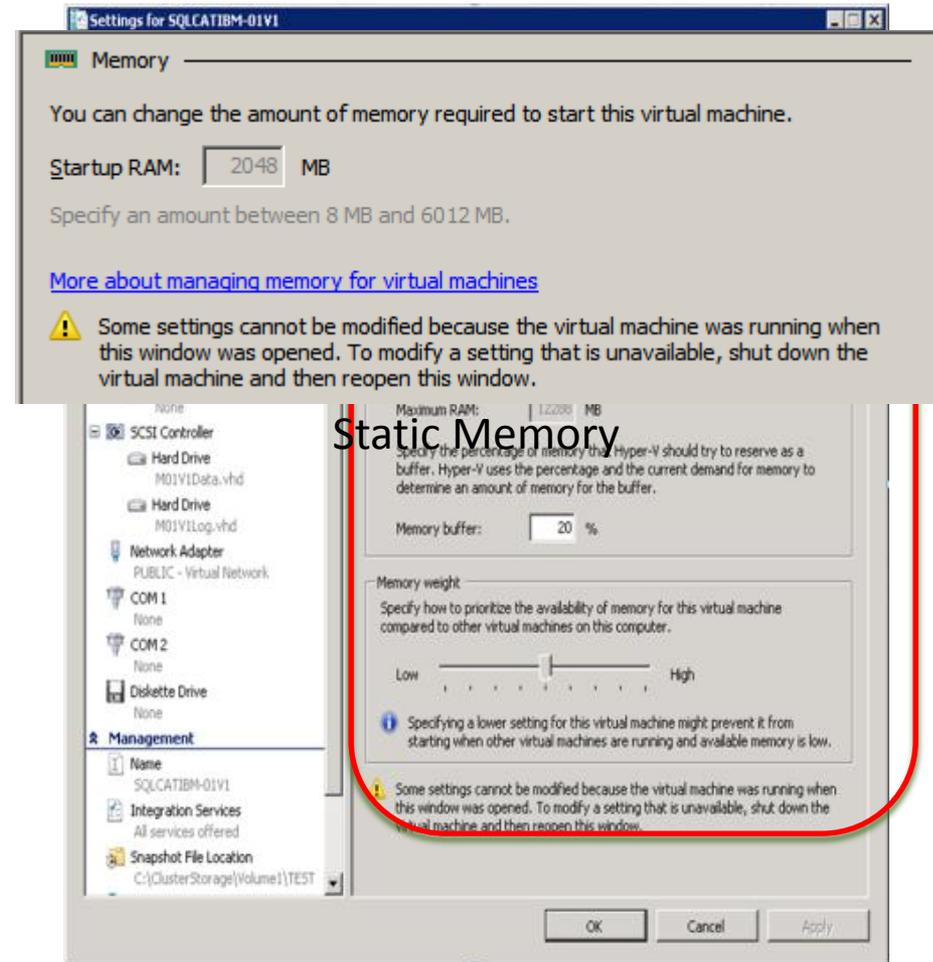
# SQL Server and Dynamic Memory

- What is dynamic memory?
- Primary goals
  - Maximum density
  - Minimal perf impact
- Primary scenarios
  - Unplanned failover
  - Planned maintenance (Live migrations)

Dynamic Memory is supported for SQL Server

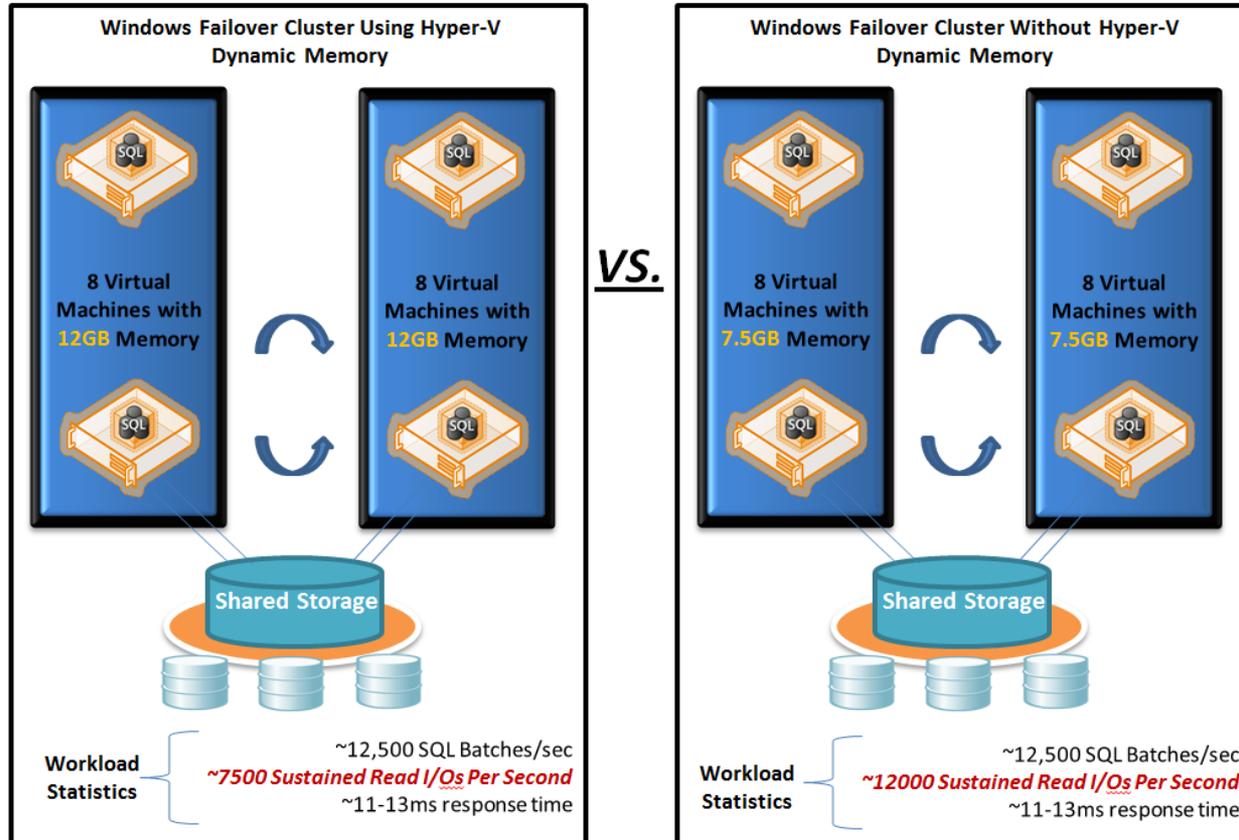
(<http://support.microsoft.com/?id=956893>)

Best practice guidance is [published](#)



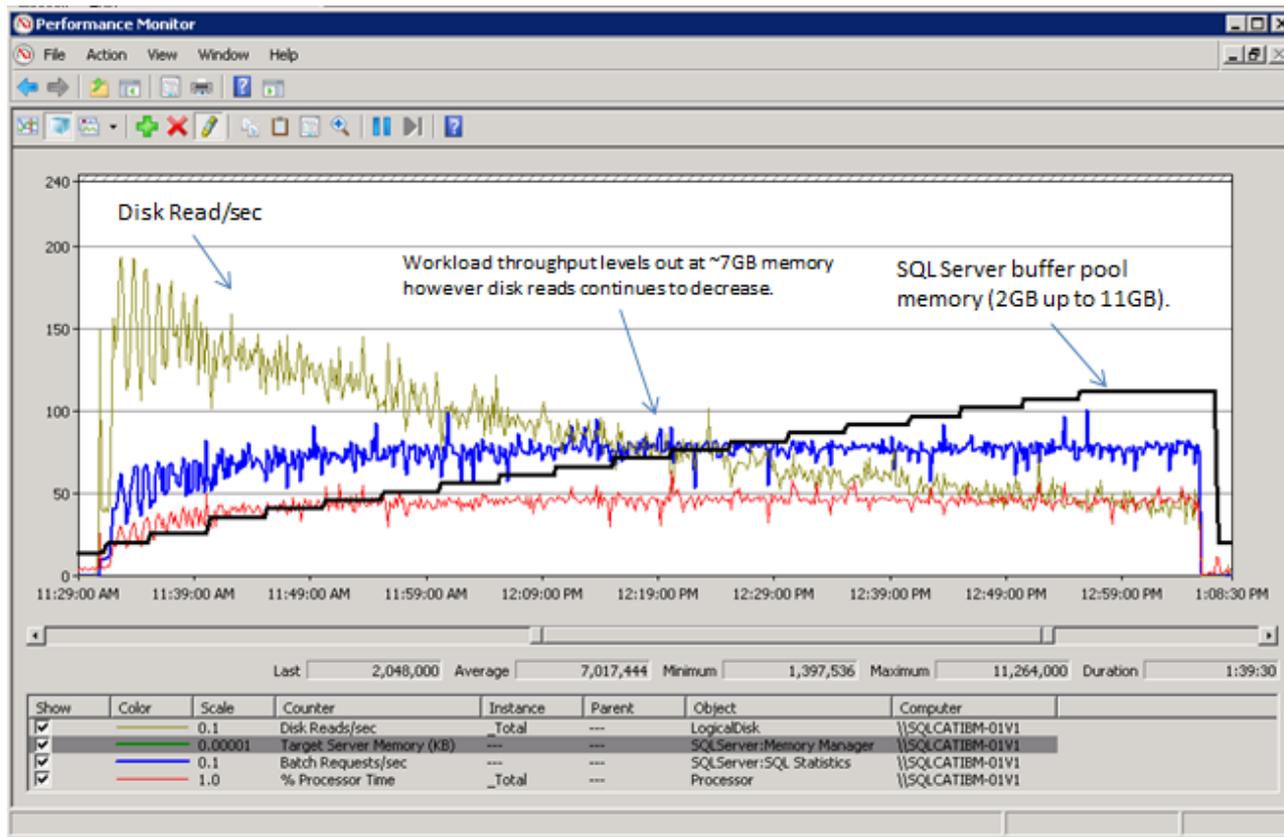
Dynamic Memory

# How does Dynamic Memory Benefit SQL Server?



\* Each cluster node has 128GB RAM

# How does Dynamic Memory Benefit SQL Server?



# Dynamic Memory Findings and Best Practices

- Reduce SQL Server memory using 'max server memory' sp\_configure setting prior to Live Migration
- Adjusting memory weight of VMs for Live Migration to an overcommitted server
- Grant SQL Server service accounts 'Lock Pages in Memory'

These settings combined avoid performance drop offs for SQL Server

# Agenda

- What is Private Cloud?
- How to Optimize SQL for Private Cloud
- How to Achieve Higher Density?
- A Real World Story

Case Study: Commerce Platform

# Virtualization and High Availability

# Commerce Platform

- ▶ Online payment and subscription processing service.
- ▶ Supporting 13 payment methods, 56 currencies , 74 markets
- ▶ 50 million active Microsoft customers and 2.5 billion dollars in transactions for FY11.
- ▶ Part of Server and Tools Business (STB).

# What is Commerce Platform



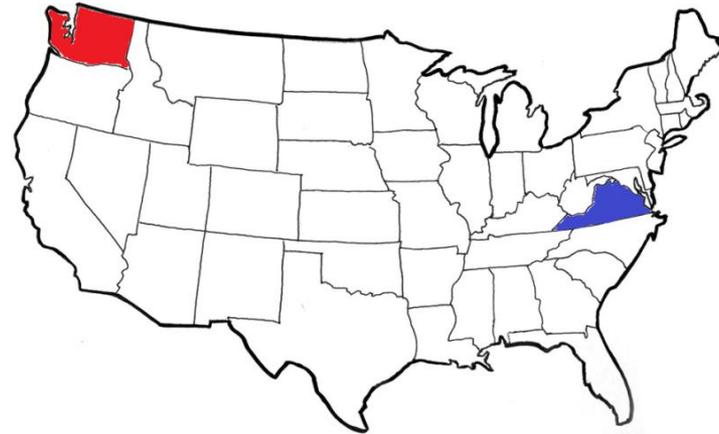
- ✓ Subscription Billing
- ✓ Payments
- ✓ Risk Management
- ✓ Mobile Operator Billing
- ✓ Point of Sale Activation
- ✓ Payouts



# SQL Infrastructure

## PROD & BCP

<b>2</b>	Datacenters
<b>210</b>	SQL 2008 SP2 physical user databases
<b>980</b>	<b>TB</b> active datafiles
<b>121</b>	<b>TB</b> storage
<b>420</b>	<b>TB</b> monthly growth
<b>12</b>	DB Mirror Pairs
<b>110</b>	Log Shipping pairs
<b>98</b>	Replication streams
<b>400</b>	Replicas / Subscribers

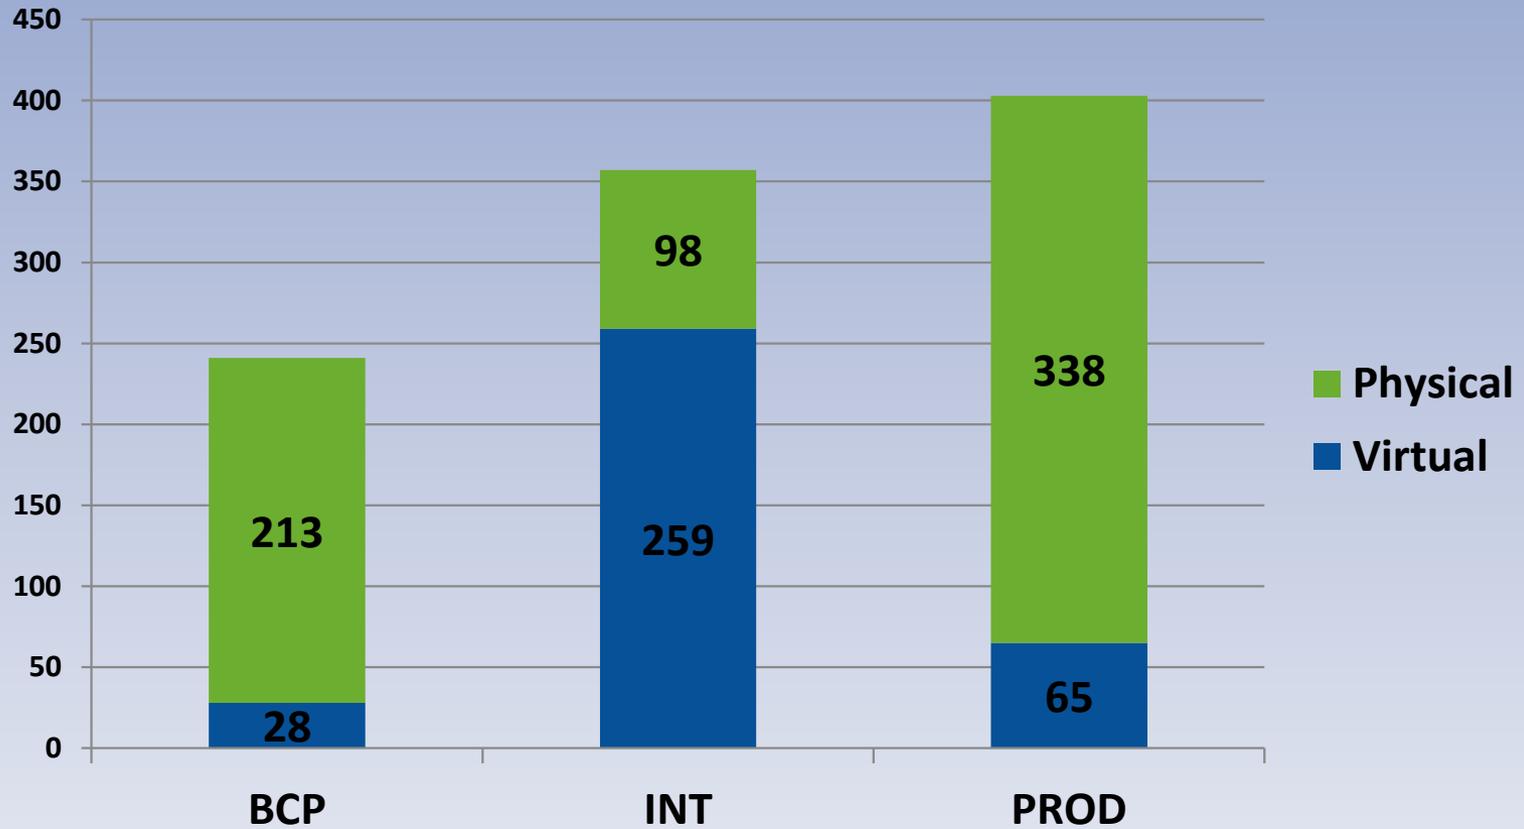


## INT

<b>1</b>	Datacenter
<b>132</b>	SQL 2008 SP2 Virtual Machines
<b>980</b>	user databases
<b>10</b>	<b>TB</b> active datafiles
<b>100</b>	<b>TB</b> storage
<b>110</b>	DB Mirror Pairs
<b>400</b>	Replication streams
<b>15</b>	Replicas / Subscribers

# VM/Physical

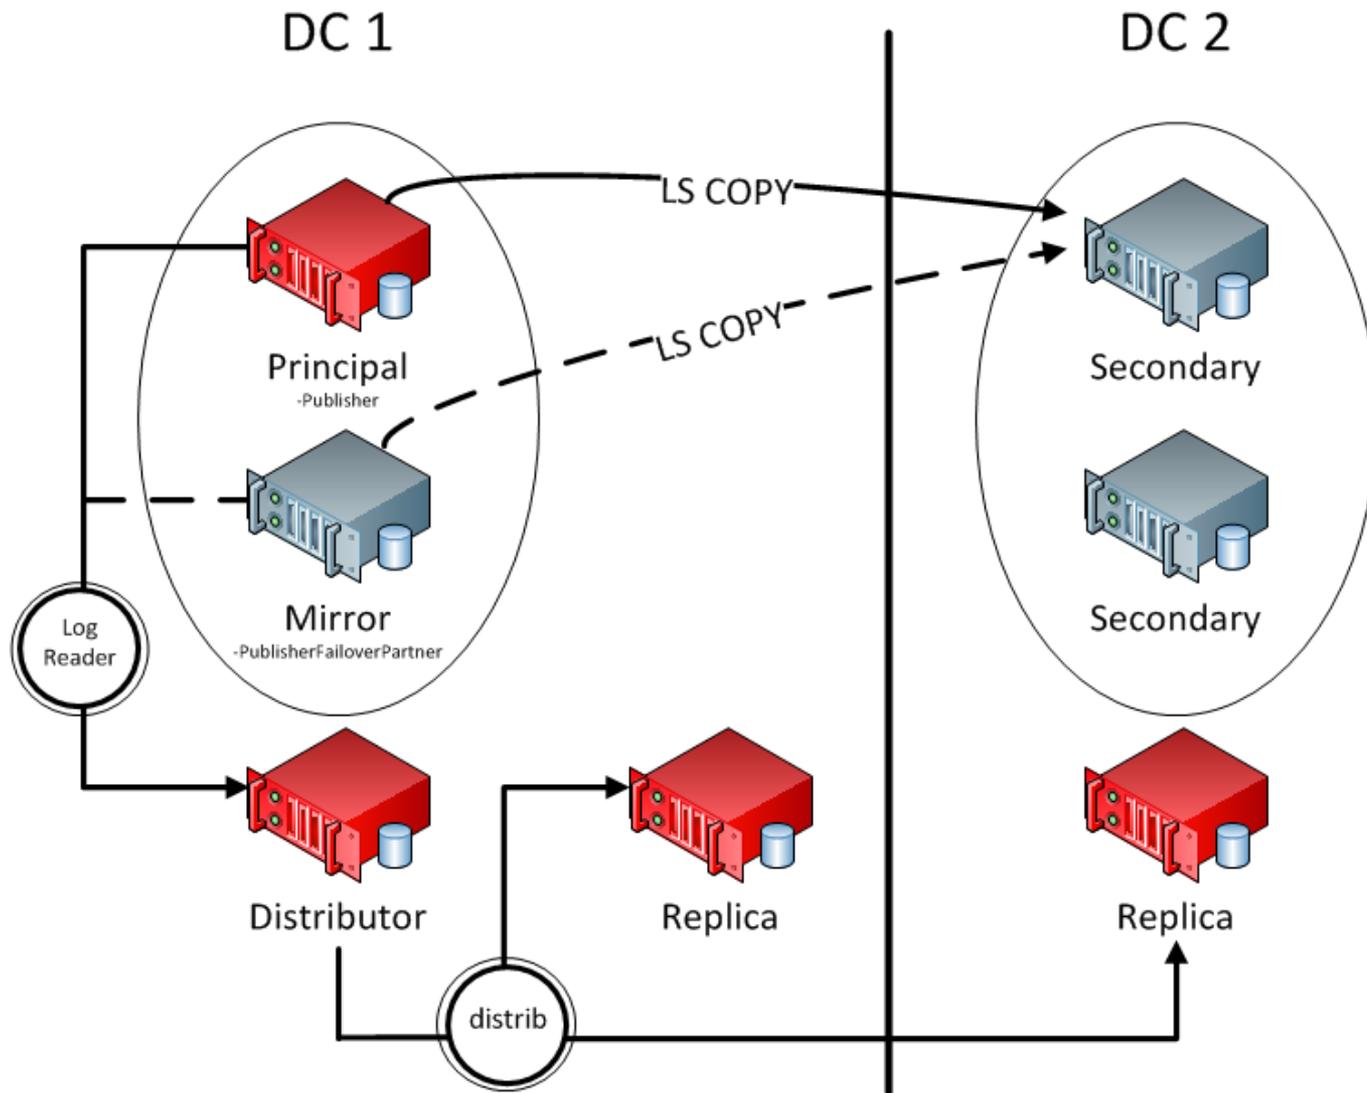
## Virtualization by Environment



# Virtualized Environment

- 100% VM in INT for over 2 years
- Elasticity allows 1:1 server mapping with INT & PROD
- Density of 4 – 12 VMs per host.
- Dynamic Disks, fixed 4GB memory, 4 processors.
- Standard organizational base image and template
- Scripted PowerShell VM creation
- 30 minutes to build VM, 1 hour to patch and fully deployed in 24 hours
- Utilizing VM Host Clusters for elasticity, HA and fast/live migration

# High Availability Architecture



```
C:\>TowerAdmin.exe mirroring plannedFailover /prisvr:TK2PRINCIPAL /pridb:MyDatabase
```

```
(07:10:39) *****
(07:10:39) PLANNEDFAILOVER [PRINCIPAL: TK2PRINCIPAL:MyDatabase; MIRROR: TK2MIRROR]
(07:10:39) *****
(07:10:39) Disable log shipping backup jobs
(07:10:50) Set mirroring partner safety to full
(07:10:50) Sync mirroring
(07:10:51) Mirroring state for database MyDatabase on server TK2PRINCIPAL is SYNCHRONIZED.
(07:10:51) Mirroring state for database MyDatabase on server TK2MIRROR is SYNCHRONIZED.
(07:10:52) Mirroring Failover LSN for database MyDatabase on server TK2PRINCIPAL is xxx.
(07:10:52) Last Log Backup LSN for database MyDatabase on server TK2PRINCIPAL is xxx.
(07:10:52) Sync mirroring Succeeded (Elapse: 00:00:02.0508200)
(07:10:54) Sync mirroring again Succeeded (Elapse: 00:00:02.0168064)
(07:10:54) *****
(07:10:54) DOWNTIME BEGINS
(07:10:54) *****
(07:10:55) Set partner to failover on principal Succeeded (Elapse: 00:00:01.1014404)
(07:10:55) *****
(07:10:55) DOWNTIME ENDS
(07:10:55) *****
(07:10:55) Set mirroring partner safety to off
(07:10:55) Setup partner safety to off for database MyDatabase on server TK2MIRROR.
(07:10:56) Set mirroring partner safety to off Succeeded (Elapse: 00:00:01.3495396)
(07:10:56) Enable Log Shipping Primary Agent Jobs TK2PRINCIPAL:MyDatabase
(07:10:56) Enable Job LSBak_MyDatabase@TK2PRINCIPAL
(07:10:56) Enable log shipping backup jobs Succeeded (Elapse: 00:00:00.0800320)
(07:10:56) Planned Failover ends successfully! (Elapse: 00:00:17.1408536)
```

# Future

- To the Cloud
- More Compression and Encryption
- More Tracing and Auditing
- SQL Server “Denali”

# Takeaways

- Best practices to maximize performance in virtualized environment
- Achieve higher density with CPU overcommit and dynamic memory
- Win8 is the ultimate private cloud OS
  - Better performance, scalability
  - Better availability and DR

- Start from smallest workload
- Continue to larger workload over time with experience
- Microsoft supports SQL Server virtualization:  
<http://support.microsoft.com/?id=956893>

Start small, learn, and continue to virtualize as you test and learn how to manage

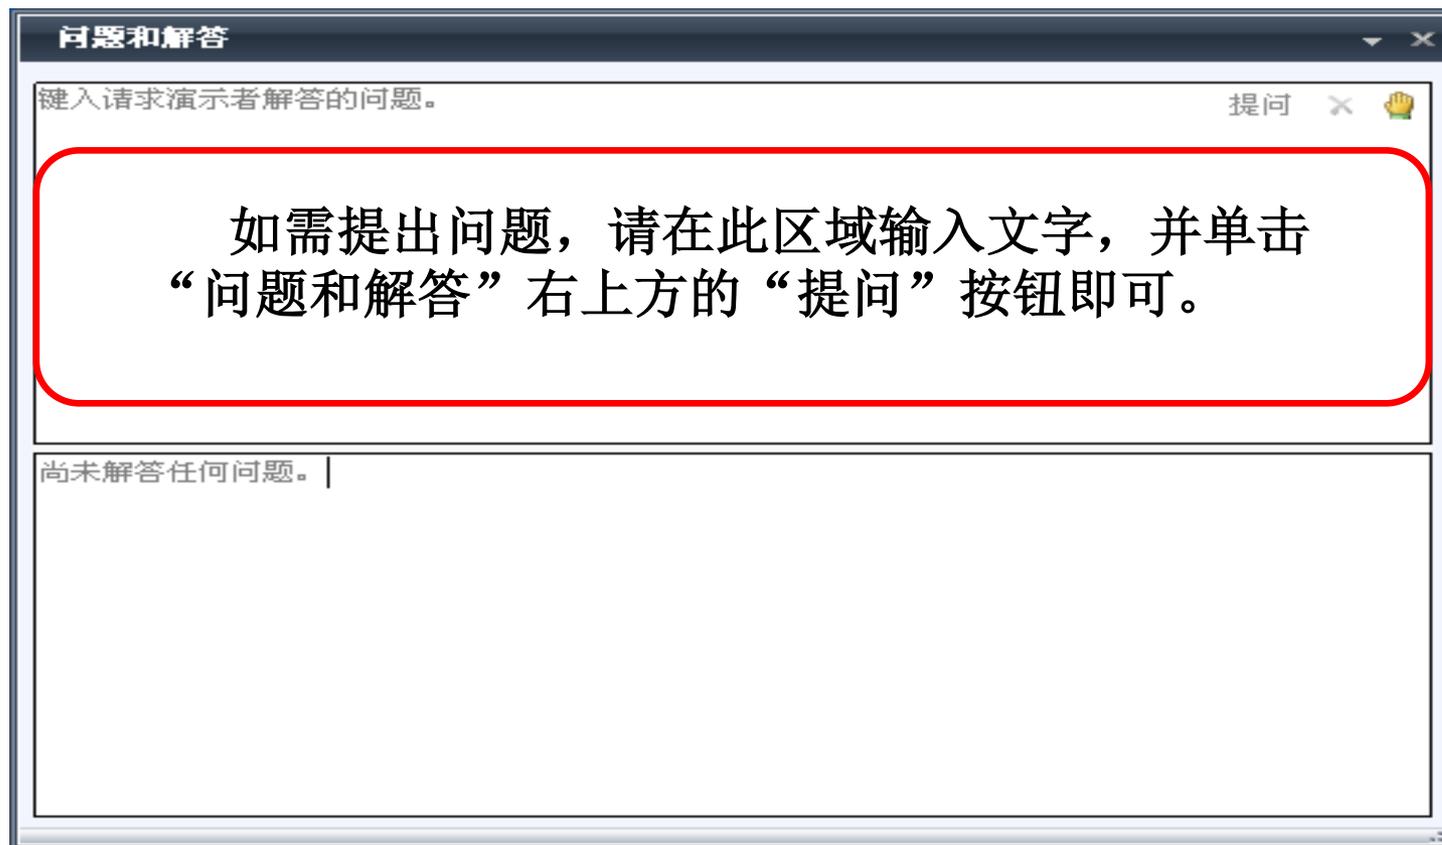
## Virtualization - Platform

### – SQL Server & Hyper-V

- Integration with Windows clustering: a key high availability (HA) technology
- Retain all Hyper-V features in failover clustering

	Windows Server 2008 R2	Windows Server 8
CPU	<=4 CPU	<=32 CPUs
Memory	<=64GB	<=512GB
Disk	<=2TB	<=16TB

# Question & Answer



Microsoft®  
tech.ed  
中国 | 2011

**Microsoft®**

您的潜力，我们的动力

© 2008 Microsoft Corporation. All rights reserved. Microsoft, Windows, Windows Vista and other product names are or may be registered trademarks and/or trademarks in the U.S. and/or other countries. The information herein is for informational purposes only and represents the current view of Microsoft Corporation as of the date of this presentation. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information provided after the date of this presentation. MICROSOFT MAKES NO WARRANTIES, EXPRESS, IMPLIED OR STATUTORY, AS TO THE INFORMATION IN THIS PRESENTATION.