

Microsoft

# tech·days

Hong Kong | 2013

You make  
the difference





# Exchange Server 2013 Virtualization Best Practices

Scott Schnoll  
Principal Technical Writer  
Microsoft Corporation



# Agenda

- Why virtualize? Why not?
- Exchange 2013 virtualization supportability
- Windows Server 2012 supportability
- Problem areas, and how to avoid them

Why and Why Not?

# Why Virtualize?

- Server consolidation
- Homogenous infrastructure
  - Deployment optimizations
  - Management optimizations
  - Monitoring optimizations
- Hardware utilization
- Cost

# Why Not Virtualize?

- Complexity
  - Additional deployment steps
  - Additional management layer
  - Additional monitoring layer
- Performance impact
- Workload incompatible / unsupported
- Cost

# Exchange Team's Take

- Customers should pick the simple solution
- Physical is often the simple solution, but won't be for every customer
- We support virtualized Exchange because it absolutely makes sense for a subset of our customers
- Customers that virtualize Exchange should be able to get something out of virtualizing

# Exchange 2013 Support



# Supported

## Hypervisors

Windows Server 2008 R2

Hyper-V Server 2008 R2

Windows Server 2012

Hyper-V Server 2012

Third-party hypervisors (SVVP)

## Exchange roles

Both Exchange 2013 roles supported

## Storage

Block-level

Same requirements as Exchange 2010 \*

## Host-based clustering

All Exchange 2013 roles supported

## Migration

All Exchange 2013 roles supported

## Jetstress testing in guests

Yes, on supported Windows hypervisors or ESX 4.1 or newer

# Explanation of Asterisk

- SMB 3.0 provides the ability to survive hardware failures that would otherwise impact file access
- Exchange 2013 supports VHD storage on SMB 3.0 file shares
  - Shares presented by Windows Server 2012 or other implementations of SMB 3.0
  - Specific to VHD storage – no direct access to shares from Exchange
  - Design for HA & failure handling
  - No change to our support of downlevel SMB & NFS

# Not Supported

- Dynamic memory, memory overcommit, or memory reclamation
- Significant processor oversubscription
  - Limited to 2:1, best practice is 1:1
- Hypervisor snapshots
- Differencing/delta disks
- Apps on the root
  - Only deploy management, monitoring, AV, etc.

# Not Supported

- Exchange on IaaS (e.g., Windows Azure)
  - No benefits when compared to Office 365
  - Limits on network configurations that prevent best practices for isolation from being implemented
  - Endpoint load balancing doesn't provide enough flexibility to meet Exchange load balancing requirements
  - Shared storage infrastructure doesn't provide IOPS and latency guarantees

Windows Server 2012

# Windows Server 2012

- Windows Server 2012 Hyper-V adds new features
  - Many deployment-blocking limits removed
- Customers who virtualize Exchange Server 2013 will have a great experience on Server 2012 Hyper-V
- Important to be aware of what does & doesn't work (and supportability limits)
  - In general, features that are transparent to Exchange will work
  - Supportability issues come up with features that require significant testing and don't align well with Exchange architecture or best practices

# Hyper-V Replica

- Not supported for Exchange
- Provides DR for VM via log shipping to remote host
- Makes sense for apps without built-in capabilities
- Use DAGs with Exchange 2010 and Exchange 2013

# Virtual Machine Storage Migration

- VM storage migration is a.k.a. VM mobility
- Not supported for Exchange
  - Not tested with current or downlevel Exchange
- Exchange servers typically have very large data storage requirements
  - Migrating data across network can be done more effectively via DAG seeding



# Noteworthy Areas

# Failure Domains

- Understanding failure domains is critical Exchange design
  - This is especially true with virtualized designs
- Embrace failure and prepare for it with redundancy and multiple paths to infrastructure
- Placing multiple copies of the same database in the same failure domain will lower availability
- Placing dependencies of Exchange on the same infrastructure will lower availability


# Host-Based Clustering for High Availability

- Not an Exchange-aware solution
- Only protects against hardware/network failures
- No HA for storage failures or data corruption
- Requires a shared storage deployment
- If you deploy this, combine with DAG when possible


# Hypervisor Memory Management

- Hyper-V's Dynamic Memory and VMware's Ballooning are fantastic for lab environments
- Not supported for production Exchange servers
- Exchange doesn't deal well with disappearing memory
  - Cache is served from the pagefile or you end up with undersized database cache

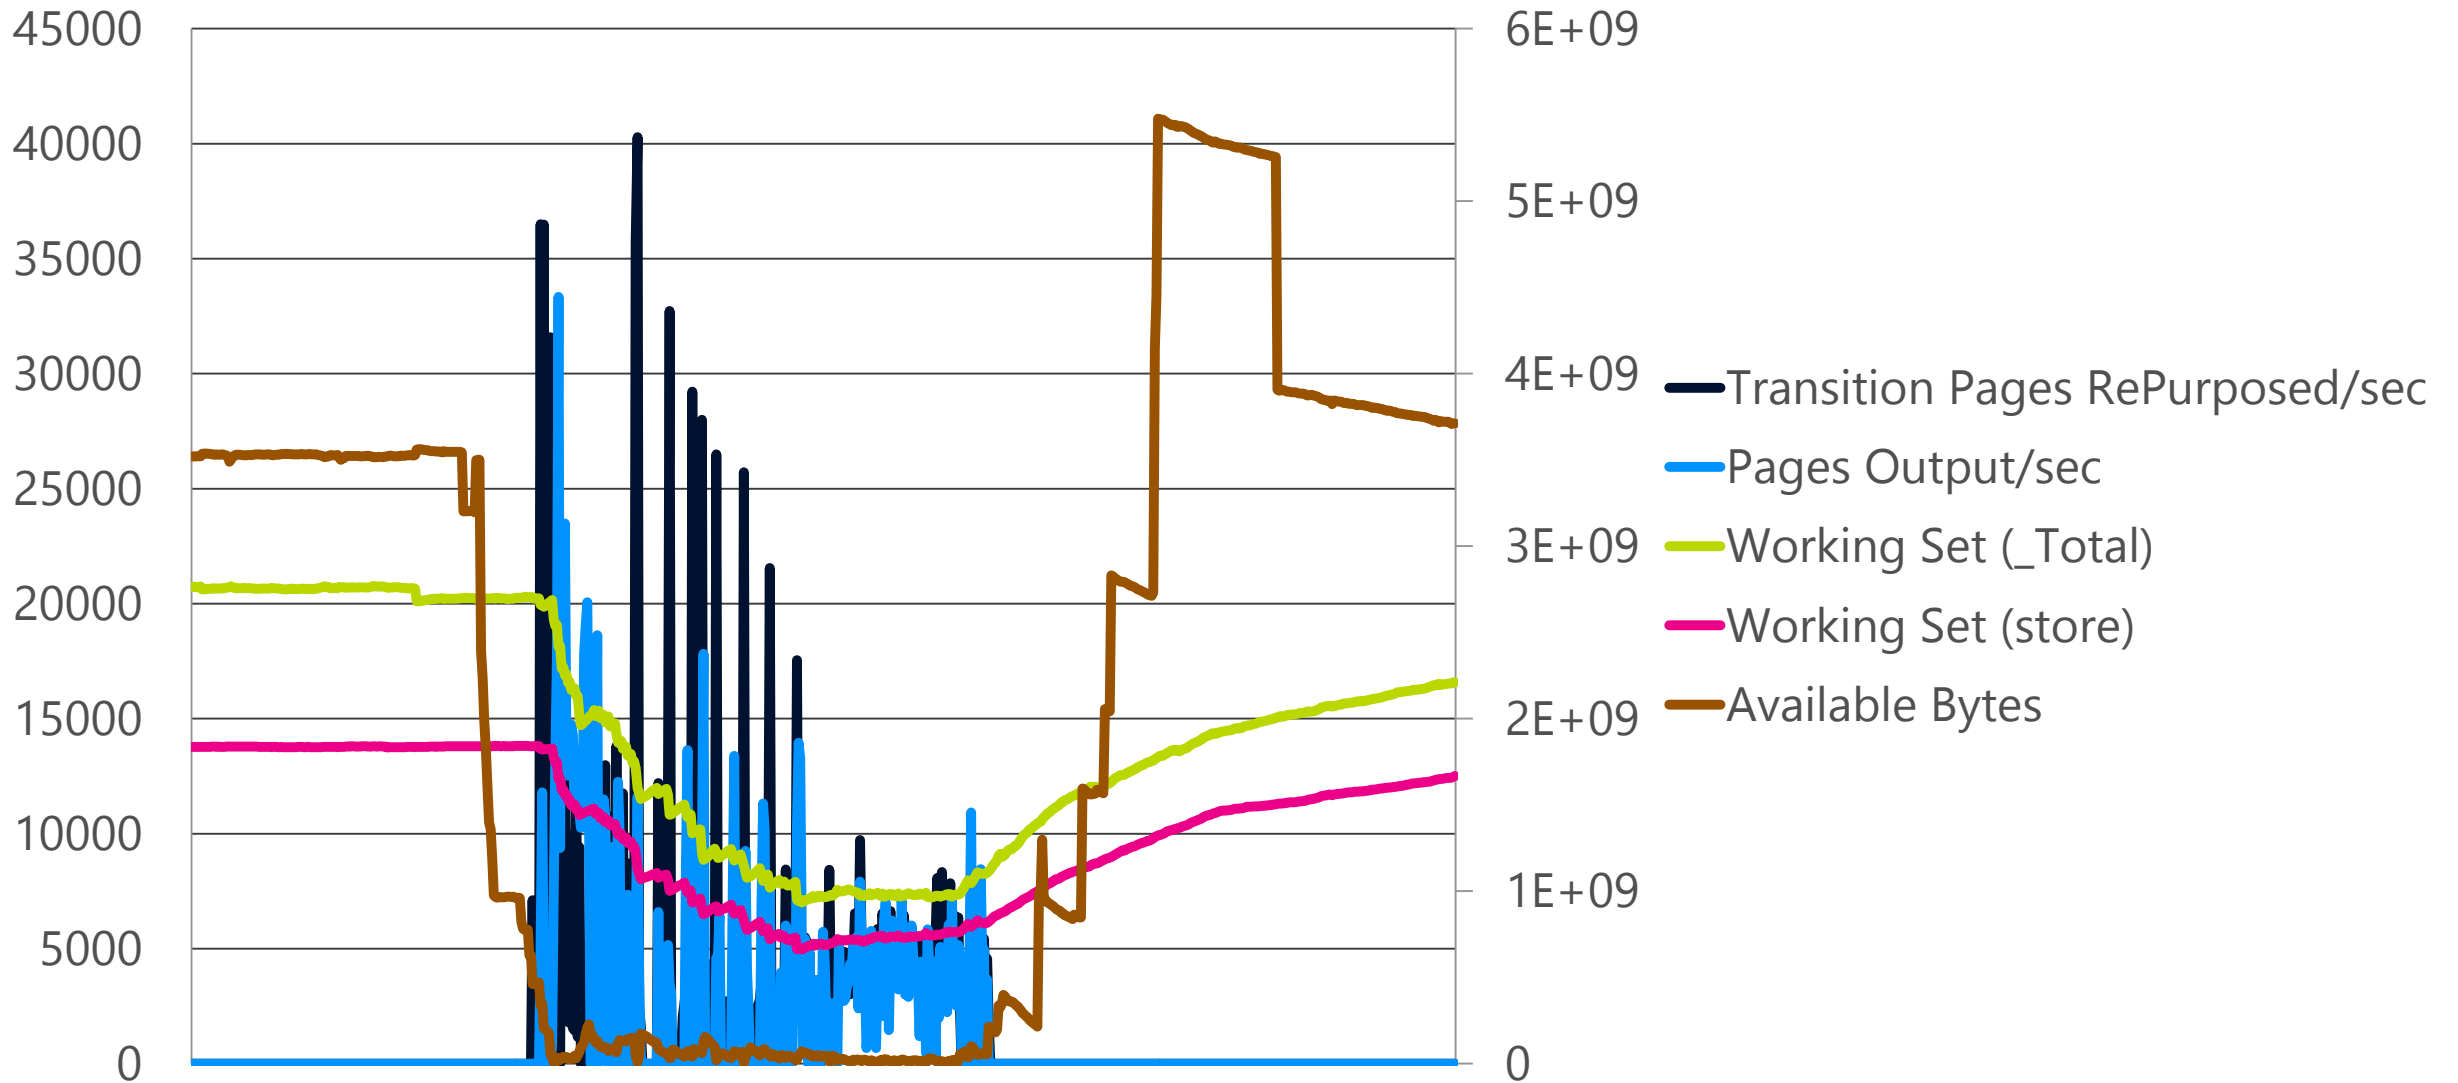
# Hyper-V Dynamic Memory



# Hyper-V Dynamic Memory



# Hyper-V Dynamic Memory



# Follow Hypervisor Vendor Guidance

- Memory Example

- VMWare - "Do not over-commit memory on ESXi hosts running Exchange workloads." - Page 11, [Microsoft Exchange 2010 on VMware Best Practices Guide](#)
- Reserve memory for Exchange VMs – don't overcommit

- CPU Example

- VMWare - "For performance-critical Exchange virtual machines (production systems), the total number of vCPUs assigned to all the virtual machines should be equal to or less than the total number of cores on the ESXi host machine." - Page 10, [Microsoft Exchange 2010 on VMware Best Practices Guide](#)
- Reserve CPU for Exchange VMs – don't oversubscribe



# Oversubscription

- Hypervisors don't make CPU resources out of thin air
  - Oversubscription can help with hardware consolidation, but it doesn't provide reliable high-performance Exchange services
  - Proper Exchange sizing ensures that resources are available on-demand, so don't allow hypervisors to take resources away
- CPU-constrained Exchange servers perform poorly
  - Delivery throughput reduction = queue growth
  - Content indexing throughput reduction = increased IOPS
  - Store ROP processing throughput reduction = RPC latency & end-user pain

# Oversubscription

- Exchange 2013 has Workload Management (WLM)
  - Dynamically adjusts background tasks to ensure that resources are being consumed efficiently
  - Monitors resource consumption and makes decisions based on resource availability
  - Inconsistent resource assignment results in bad WLM decisions

# Hypervisor Snapshots

- Snapshots make lab testing much easier
- Requires entire topology to be reverted back
- Not supported for production Exchange orgs
- Exchange must always move forward, and cannot travel back in time

# Summary

# Summary

- Don't oversubscribe anything
  - Exchange uses lots of memory – size properly and don't rely on hypervisor "magic"
  - Processor resources are critical for Exchange client experience – ensure that Exchange always gets what it needs
- Not all hypervisor features make sense for Exchange
- Follow hypervisor vendor's best practices

Thank You!



Scott Schnoll  
Principal Technical Writer  
[scott.schnoll@microsoft.com](mailto:scott.schnoll@microsoft.com)  
<http://aka.ms/schnoll>

