



# Windows Server 2008 R2 - Best Practices og Performance Tooling

*Brian Lauge Pedersen  
Technical Solutions Professional - Datacenter  
[blauge@microsoft.com](mailto:blauge@microsoft.com)*



Windows Server® 2008 R2

# Agenda

- Overview
  - New Hardware Support
  - New Performance Features
  - Performance Improvements
- Selected Deep Dive
  - What changed?
  - How?
  - Results?
- Questions and Answers

# In a Nutshell...

- R2 is more than a “patch” or “service pack”
- The most scalable release of Server
- Performance gains in a number of areas
- New features to address pain points
- Great today, ready for tomorrow


# Hardware Inflection Points



Multi-core



Virtualization



> 64 Logical  
Processors



Power



Windows Server® 2008 R2

# Overview

## *New Hardware Support*

- CPUs
  - > 64 Logical Processors (LP)
  - Power
    - Improved C State Support
    - Support for Core Parking
- Virtualization
  - Hardware SLAT (EPT/NPT)



# Overview

## *New Features – Part 1 of 2*

- ~~32bit~~
- > 64LP enlightenments
  - Network stack, storage stack, tools
- Cooperative Scheduling
  - User Mode Scheduling (UMS)
- NUMA enhancements
  - Topology APIs
  - NUMA-aware RSS

# Overview

## *New Features – Part 2 of 2*

- Hot Lock “removal”
  - Scheduler Dispatcher
  - Memory Manager PFN
  - Cache Manager VACB
  - Object Type

# Overview

## *Improvements - Part 1 of 3*

- Virtualization
  - 64 LP!
  - Scalability Improvements
  - NUMA Scheduling
  - Dynamic VHDs
- Boot Optimizations
  - Device initialization parallelization
- Power
  - Processor Power Management Policies
  - Timer Coalescing and Tick Skipping



# Overview

## *Improvements - Part 2 of 3*

- Minimization Work
  - Re-factored DLLs
  - Trigger start of Drivers and Services
  - Memory, disk and CPU optimizations
- File Server
  - Improvements in the WAN scenario by reducing network round trips
  - Scalability improvements

# Overview

## *Improvements - Part 3 of 3*


- Storage
  - Clustered Shared Volumes (Centipede)
  - Native VHD
  - NUMA I/O
- Terminal Server
  - Dynamic Fair Share Scheduling (DFSS)
  - Video Performance
  - WAN bandwidth reductions

# Agenda

- Overview
  - New Hardware Support
  - New Performance Features
  - Performance Improvements
- Selected Deep Dive
  - What changed?
  - How?
  - Results?
- Questions and Answers

# Details – > 64 LP

- What changed?
  - R2 supports 256 LPs
- How?
  - Groups
  - Static set of 1 to 64 LPs
  - Determined at boot time
  - Minimal number of groups
  - LP spatial locality important



# Details – > 64 LPs

*An example - 2 Group, 128LPs*



# Details – > 64 LPs

## • Results?

- Processes assigned round-robin to groups
  - Can be overridden to inherit parent process group
- Threads inherit group affinity
  - Thread can be affinitized to only a single group
- Most applications not affected
  - For “Legacy” APIs, group is implied
  - “Legacy” drivers are Group 0
  - New APIs to expose group information
  - Applications that use or store per-processor information for the entire system must be modified



# Details – > 64 LPs

## *Code Impacts – Minimized*

```
#define MAXIMUM_PROC_PER_GROUP    64  
#define MAXIMUM_PROCESSORS  
    MAXIMUM_PROC_PER_GROUP
```

// Examples of new APIs

GetMaxiumProcссорGroupCount(...)

GetMaxiumProcссорCount(...)

CreateRemoteThreadEx(...)

GetActiveProcessorGroupCount(...)

GetCurrentProcessorNumberEx(...)

GetLogicalProcessorInformationEx(...)

GetMaxiumProcessorCount(...)

// and many more...

*video*

SQL on 256 LPs



# Details – Lock “Removal”

## *Terminology*



- What's a lock
  - A spinlock is a locking primitive associate with global data structures - prevents multiple threads from simultaneously modifying important data
  - Waiting threads “spin” doing nothing waiting for access to the lock
- Contention
  - Threads being stalled waiting for their turn to access a lock

# Details – Lock “Removal”


## *Details*

- What’s changed?
  - Four key Kernel locks have been “removed”
    - Scheduler Dispatcher, Memory Manager PFN, Cache Manger VACB, and Object Manager Type
- How?
  - Decomposition into “smaller” locks
- Results?
  - Less contention, less waiting ⇒ better scalability
  - No detectable change for user applications

# Details – Lock “Removal”

## *Scaling without the Dispatcher Lock*

- 1.7x scaling going from 64 to 128 LPs



# Details – Power Savings


- What changed?
  - Processor Power Management algorithms & settings
  - Increased processor idle state usage
  - Intelligent Tick Distribution, Tick Skipping, Timer Coalescing, and Core Parking
  - New metering and budgeting features
- How?
  - Lots of tuning on diverse workloads (with Intel/AMD assistance and validation)
  - Refactoring of the OS & kernel to minimize idle activity
  - Support for onboard metering
- Results?
  - WS08R2 improves greatly from W2k3





# Details – Power Saving


## OS Comparison – Out of the Box Settings



— WS2k3 SP2 - High Perf — WS08 RTM - Balanced — WS08 R2 RC - Balanced

# Details – Power Saving

## Power vs. Performance – R2 Power Plans



# Details – Power Saving

## *Balanced vs. High Performance*




# *Demo*

Core parking



# Details – Virtualization

## *Hyper-V Power*



# Details – Virtualization

## *Scalability*

Category	WS08 (v1)	WS08 R2 (v2)
Max LP	24	<b>64</b>
VP:LP	4:1	<b>8:1</b>
Max Active VMs	96	<b>384</b>
Max VPs per VM	4	4
Max VPs	96	<b>512</b>
Max Mem per VM	64GB	64GB
Max Mem (Host)	1TB	1TB



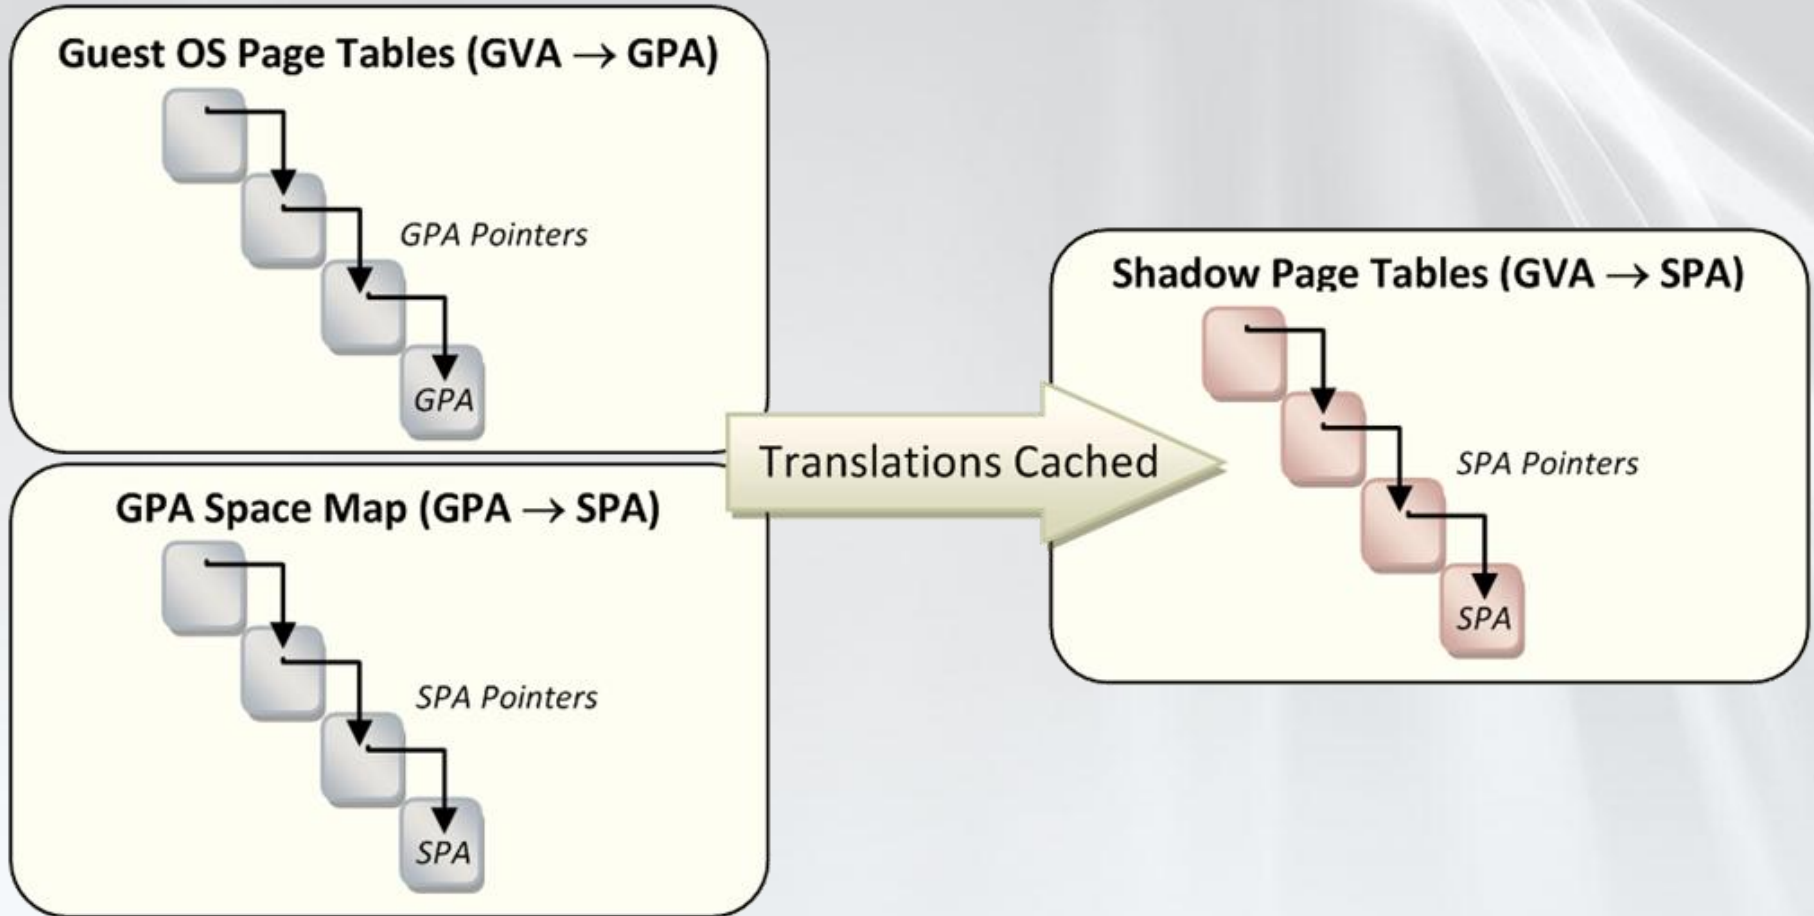
# Details – Virtualization

## *Memory Mapping – Terminology*

- Memory Management Units (MMU)
  - Guest Virtual Address (GVA)
  - Guest Physical Address (GPA)
  - System Physical Address (SPA)
- 
- Optimization – Shadow Page Tables

# Details – Virtualization

## Shadow Page Tables



# Details – Virtualization

## *Hardware SLAT*

- What's new?
  - Shadow Page Tables stored in hardware
  - Second Level Address Translation (SLAT)
- How?
  - Support for hardware SLAT
  - Replaces Multiple Shadow Address Space (MSAS)
  - EPT on Intel, NPT on AMD
- Results?
  - Same performance with less overhead


# Details – Virtualization

## *Live Migration*

- What's new?
  - VM moving between Hosts without noticeable interruption
- How?
  - Incremental copy and restore of VM
    - Quick Migration without the downtime
  - Clustered Shared Volumes (Centipede)
    - Required SAN for seamless VM movement between clusters

# Details – Virtualization

## *Live Migration*



# Details – Virtualization

## *Live Migration*

- Results?
  - Improved Reliability
    - Load balancing across VM Hosts
  - Migration times can be longer than Quick Migration
  - Down times are very short



# Details – Virtualization


## *Virtual Hard Drivers (VHD)*

- What's changed?
  - A number of performance improvements
- How?
  - Rewritten VHD stack
- Results?
  - Random Write much improved
  - Over all considerably closer to native

# Details – Virtualization

## *Virtual Hard Drivers (VHD)*

Higher is Better



# Details – Footprint

## *Terminology*


- Memory Working Set
  - The amount of RAM the memory manager assigns to process or kernel memory type
  - Shown in Task Manager
- Memory Reference Set
  - The amount of Working Set that is actively used by a process or the kernel
- Paged Pool
- Non-paged Pool

# Details – Footprint

- What changed?
  - Memory usage
  - Disks usage
- How?
  - Demand start of Drivers and Services (UBPM)
  - Memory Manager changes
  - Targeted testing and analysis of usage
- Results
  - Memory
    - Better utilization
    - Smaller Reference Set
    - Reduced Non-paged pool memory usage
  - Enterprise Full installation is now 2GB smaller on disk



# Details – Footprint



# Details – File Copy Changes

- What changed?
  - Reduced per-file network round trips
  - Reduced L2 cache misses and CPU churn.
  - Reduced spurious I/O activity on the system
  - Greater Parallelism
- How?
  - Core copy engine optimizations in CopyFileEx and Shell
  - Kernel optimizations in Cache Manager and Memory Manager
  - Improved SMB2 request compounding






# Details – File Copy Changes

- Results

- Improvements in local file copy
- Faster WAN transfer for large file sets
- Robocopy now supports multithreaded copy
  - Multi-threaded copy (“/MT” switch)  
Improvement up to 8x on high-latency networks

# Details – WAN Office File Open

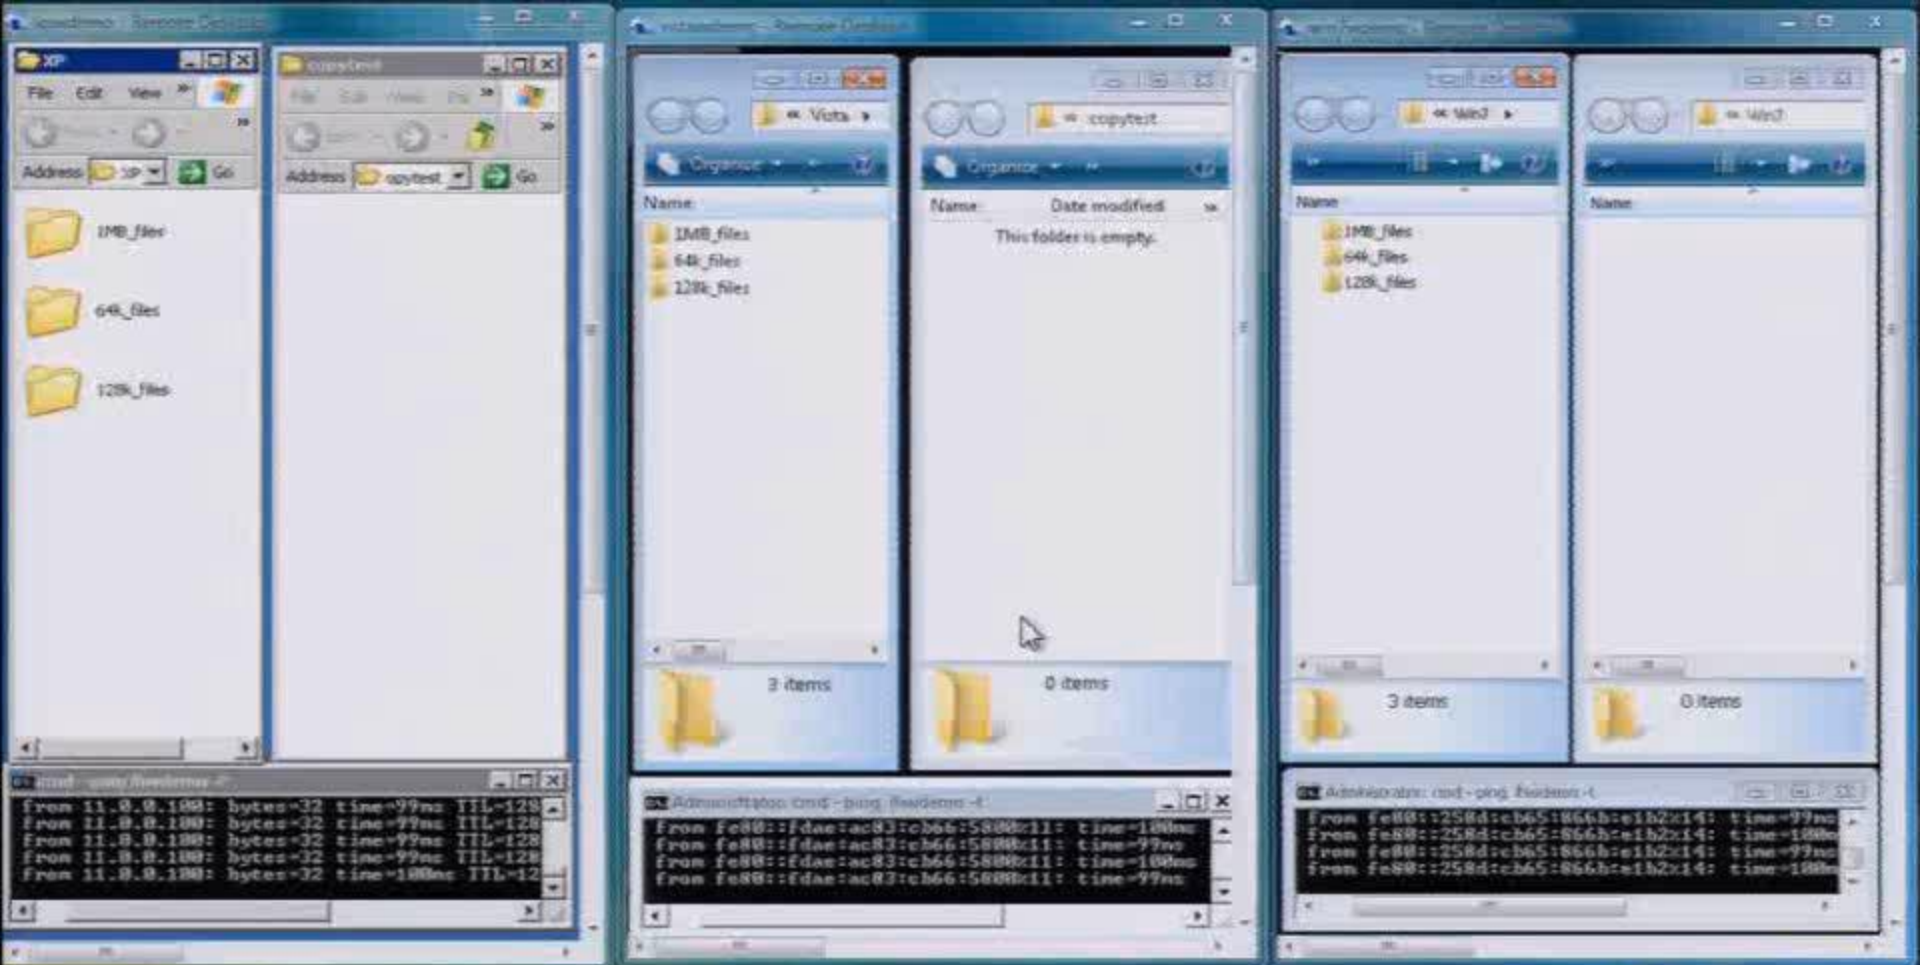
## *Better Together!*



■ Vista SP2 - WS08 ■ Vista SP2 - WS08R2 ■ Win7 - WS08 ■ Win7 - WS08R2



Windows Server 2008 R2



# Performance Tuning

- Windows Hardware Developer Central (WHDC)

  - <http://www.microsoft.com/whdc>

  - PC Fundamentals

    - Performance Page

  - Turning Guides

    - Freshly updated for R2

  - [http://www.microsoft.com/whdc/system/sysperf/Perf\\_tun\\_srv-R2.msp](http://www.microsoft.com/whdc/system/sysperf/Perf_tun_srv-R2.msp)

  - Virtualization and Partitioning

- Windows Server Performance Team Blog

  - <http://blogs.technet.com/winserverperformance/>

# Performance Tuning

## *Performance Analyzer*

- XPerf
  - Sampling based profiler
  - Built on top of the Event Tracing for Windows (ETW) infrastructure
- What it allows:
  - Driver delays analysis
  - CPU sampling analysis
  - Disk I/O analysis
  - Network analysis

<http://msdn.microsoft.com/en-us/performance/cc752957.aspx>



# Performance Turning XPerfView





# Summary

- R2 is more than a “patch” or “service pack”
- The most scalable release of Server
- Performance gains in a number of areas
- New features to address pain points
- Great today, ready for tomorrow

# Additional Resources

- > 64 Logical Processors  
<http://code.msdn.microsoft.com/64plusLP>
- Channel 9  
<http://channel9.msdn.com/tags/w2k8r2>
- Power savings and Management
  - Plug and Play Power Management  
<http://www.microsoft.com/whdc/system/pnppwr/default.mspx>
  - Power In, Dollars Out: How to Stem the Flow in the Data Center  
[http://www.microsoft.com/whdc/system/pnppwr/powermgmt/Svr\\_Pwr\\_ITA\\_dmin.mspx](http://www.microsoft.com/whdc/system/pnppwr/powermgmt/Svr_Pwr_ITA_dmin.mspx)
- Windows Performance Analysis Developer Center:  
<http://msdn.microsoft.com/en-us/performance/default.aspx>

Campus Days  
14.-16. januar 2010



**THANK YOU!**



Windows Server® 2008 **R2**

**Microsoft** TechNet