



SQL Server Virtualization 499

Advanced Troubleshooting and Performance Topics

SQL PASS Virtualization Virtual Chapter

2014.10.08

About David Klee



@kleegeek
davidklee.net
gplus.to/kleegeek
linked.com/a/davidaklee

Specialties / Focus Areas / Passions:

- Performance Tuning & Troubleshooting
- Virtualization
- Cloud Enablement
- Infrastructure Architecture
- Health Monitoring
- High Availability
- Disaster Recovery
- Capacity Management
- **SQL PASS Virtual Chapters**
 - **Virtualization**
 - **Performance**
 - **HA & DR**



SQL Server Virtualization 101 - 301

- Refresher for previous session topics at:
www.davidklee.net/education/videos/
- Focused on:
 - What is virtualization?
 - What does it do for DBAs?
 - Private cloud infrastructure concepts
 - Resources and queues
 - Data management benefits
 - HA & DR features
 - Performance topics

PASS VIRTUALIZATION
VIRTUAL CHAPTER



State of Virtualization Today


- Most SQL Servers are virtualized
- A handful of the top performers are not
- Today's hypervisors are functionally transparent to performance
- Virtualization limits today – 64 vCPU & 1TB RAM
- Done right, the only risks are perceived risks
- How many are greater than 80% virtual?
- Are these issues holding your back?



Troubleshooting Scenarios

- “Just virtualized and it’s slow”
- Noisy neighbors
- Memory pressure
- Storage latency
- CPU pressure
- Need to scale up
- Need to scale out

No silver bullet solutions



(Image source bit.ly/ZsrvG9)


HFX
TECH

© 2014 Heraflux Technologies

5

“Just Virtualized and Now It’s Slow”

- Virtualized and it just “feels” slower
- Where are your benchmarks & baselines?
- Compare these items before and after
- Infrastructure bottlenecks must exist in the VM environment and not in the previous physical one



HFX
TECH

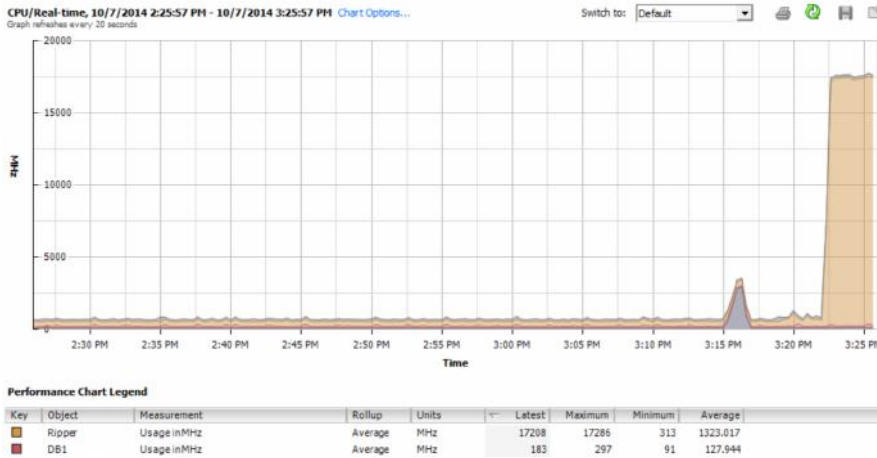
© 2014 Heraflux Technologies

6



“Noisy Neighbors”

- Random & not sustained performance slowdowns
- Are other VMs consuming the resources you need?



© 2014 Heraflux Technologies

7

How to Battle “Noisy Neighbors”

- Consider resource reservations
 - Awkward and tedious to manage
- Tiered resource pools
- Determine the resource bottleneck

DB0 - Edit Settings

Virtual Hardware | VM Options | SDRS Rules | vApp Options

CPU: 4
Cores per Socket: 2 | Sockets: 2
CPU Hot Plug: Enable CPU Hot Add
Reservation: 0 MHz
Limit: Unlimited MHz
Shares: Normal | 4000



© 2014 Heraflux Technologies

8



Memory Pressure


- Hosts generally designed for maximum consolidation
- Host memory is usually the first compute resource to max out
- Memory overcommit a feature, not error
- Not to be used with SQL Server VMs
- Memory pressure? Hosts >93% memory consumed

Name	S...	% CPU	% Memory
esxi1.180overe.net		1	34
esxi2.180overe.net		1	34
esxi3.180overe.net		3	41



Host Memory Pressure Protection

- MS Hyper-V: Dynamic memory?
(Image source: bit.ly/1CTqaHf)



- VMware vSphere: Memory ballooning?

Performance Chart Legend

Key	Object	Measurement	Rollup	Units	Latest	Maximum	Minimum	Average
■	Gateway	Granted	Average	Kilobytes	655564	655564	652604	654567.68
■	Gateway	Consumed	Average	Kilobytes	507428	517600	496256	504695.06
■	Gateway	Balloon	Average	Kilobytes	2016648	2016648	2016648	2016648
■	Gateway	Active	Average	Kilobytes	31456	314572	0	29709.067

- **Answer – no host memory overcommitment!**



Host Memory Pressure Protection

- VMware specific: reserve all guest memory
- SQL Server: Lock pages in memory
 - Last line of defense
 - Pros?
 - Cons?

DB0 - Edit Settings

Virtual Hardware | VM Options | SDRS Rules | vApp Options

GPU: 4

***Memory**

RAM: 6144 MB

Reservation (*): 6144 MB

Reserve all guest memory (All locked)

Limit: Unlimited MB

Shares: Normal 61440

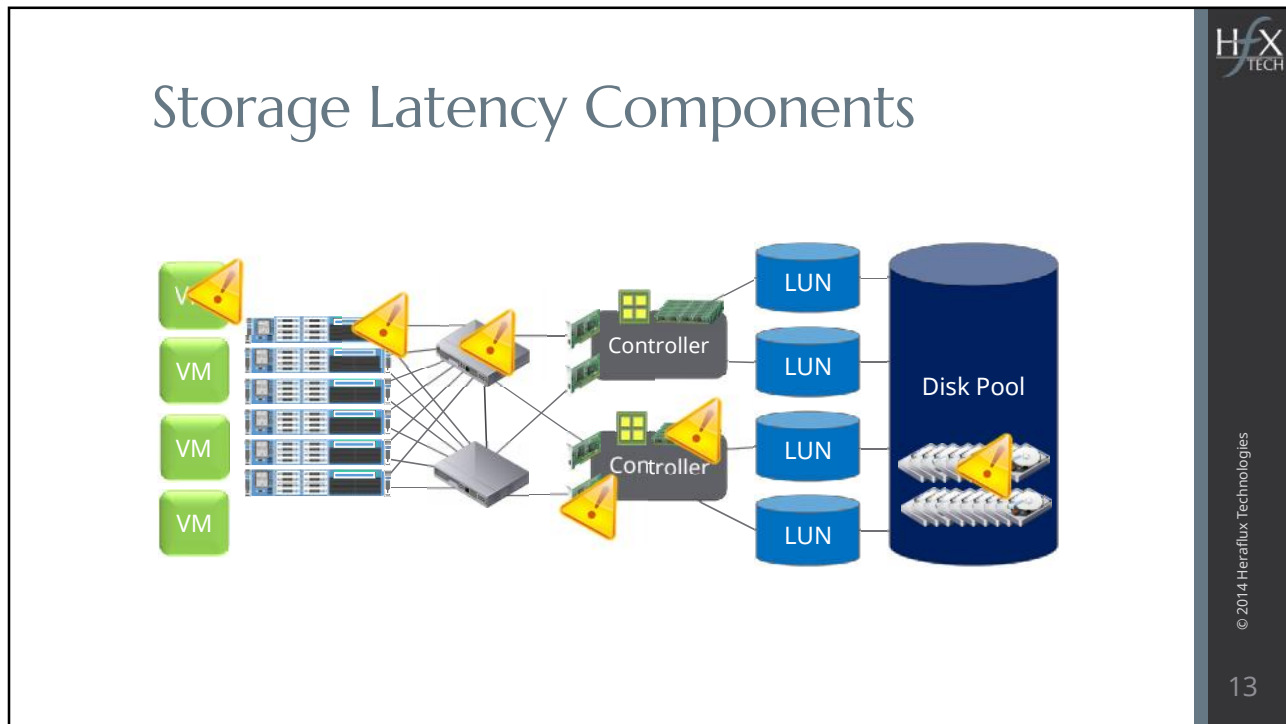
Memory Hot Plug: Enable

HX
TECH
© 2014 Heraflux Technologies
11

Storage Latency Components


The diagram illustrates storage latency components. On the left, two blue cylindrical icons labeled 'LUN' are connected by lines to a central dark blue rounded square labeled 'SAN'. Inside the SAN box, there is an illustration of a server rack with multiple server units.

HX
TECH
© 2014 Heraflux Technologies
12



Situation: Storage Performance

- Poor storage performance reported
- Get stats that you can see!
- Perfmon data – ongoing metrics
- SQLIO Batch – maximums load test
 - PoSH @ bit.ly/1vV2eih
- Symptoms?
 - High read/write latency
 - Low maximum sustained IOps
 - Low maximum throughput




HX
TECH

© 2014 Heraflux Technologies

14

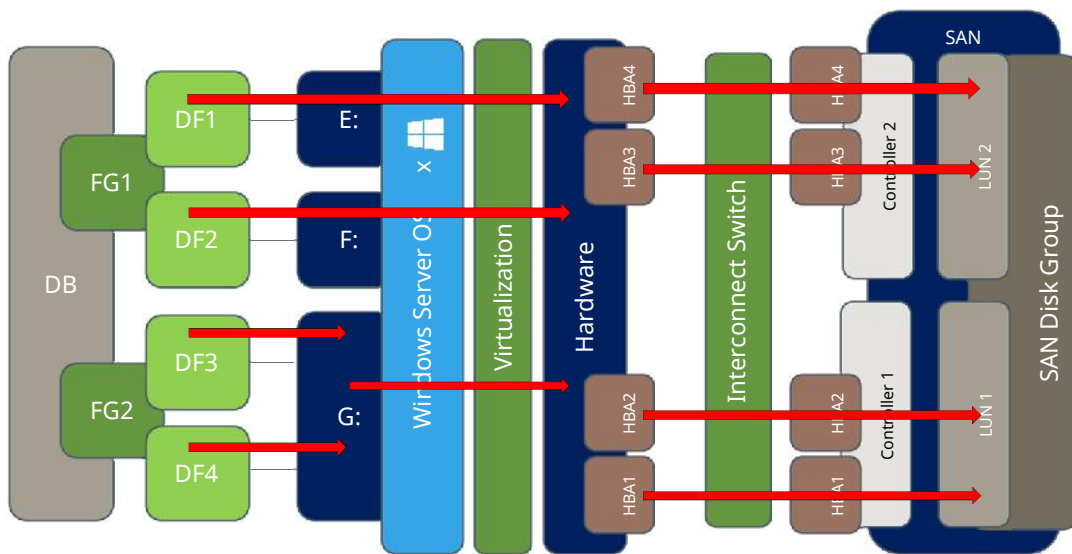


Storage Performance Measurements

- Perfmon data extractor
 - By  Video: bit.ly/1CTt52w Code: bit.ly/1tyl4JB
- SQL Server storage latency collector
 - Scripts to collect ongoing stats @ bit.ly/1xnqhYq
- Take the results back to storage group
- Work together the infrastructure bottleneck
- What can you do?



Spread out the Workload



Things You Can Do for Storage

- Spread out the workload
 - More file groups / data files
 - More virtual disks
 - More active SAN LUNs & paths
- Archive old data
- More memory (more I/O buffer)
- Query tuning
- Better indexing strategy
- Better table data model
- Table row/page compression
- Appropriate data types
- Index / statistic maintenance
- In-memory constructs
- Etc. etc. etc...




CPU Pressure

- Not as simple as 1:1 vCPU:pCPU ratio
 - Usually a major waste of pCPU resources
- First, "right-size" the vCPU count
 - One size does not fit all
 - Too few is BAD for performance (obvious)
 - Too many is also BAD (not obvious)
 - Review your utilization and workload profile
 - Pick "right" amount of vCPUs for the job
 - Stay tuned for announcements on this topic
 - Goal: 50-60% average utilization
 - Very workload specific



CPU NUMA

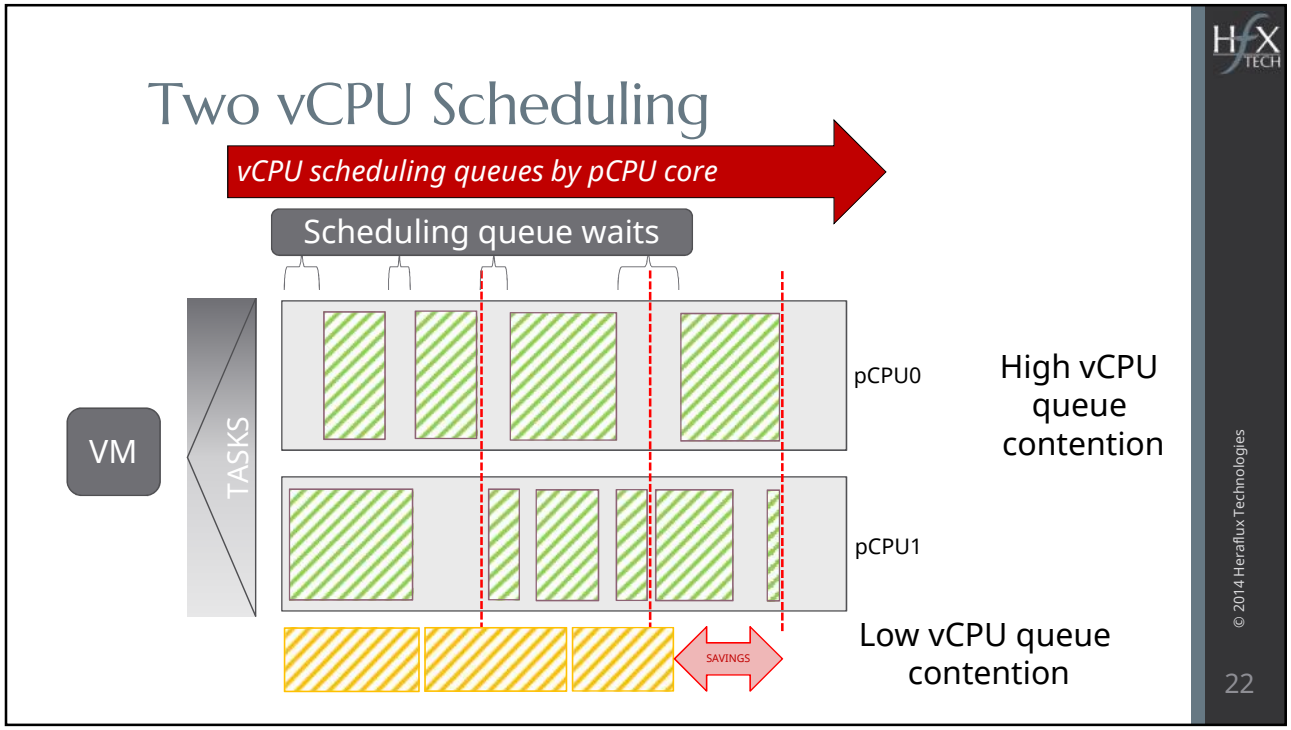
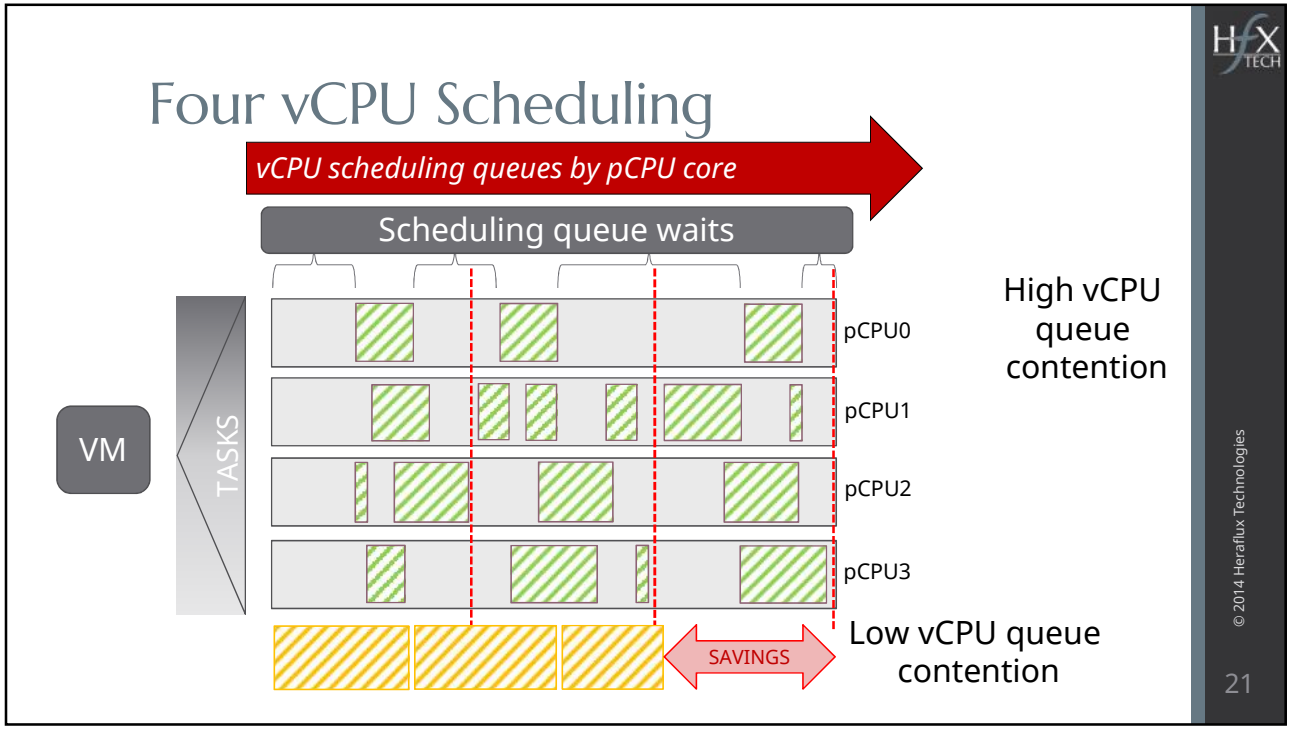
- Balance across physical NUMA nodes
- Or keep VM within one NUMA node
- Queue contention?



CPU Scheduling Contention

- All vCPU requests get scheduled onto CPU scheduler
- Queues are not linear
- More background activity, more waits
- Waits = time stolen from VM performance
- CPU, memory, storage, network all have this challenge
- CPU impacts the most common





Metrics to Collect

- SQL Server
 - Raw CPU / mem / disk usage
 - NUMA memory usage
 - Storage latency by DB file
 - Wait statistics
- Windows
 - CPU & memory consumption
 - Storage IOPs / latency / throughput
 - Processes (SQL Server vs other)
 - Perfmon how-to @ bit.ly/1sqSVns



Metrics to Collect

- Virtualization
 - Resource consumption by VM
 - Resource utilization by host
 - CPU scheduling queue wait
 - Overcommitment metrics
 - VMware vSphere: CPU Ready
 - MS Hyper-V: CPU Wait Time per Dispatch
- Work with your VM admin to get these



CPU Pressure Remedies

- Schedule less on resource schedulers
- Reduce MaxDOP / Increase Cost Threshold for Parallelism
 - Less widely parallelized queries
- Large memory pages
 - Less pointers
- Spread out workload
 - More & smaller footprint VMs
- Tune high CPU consuming queries



Trouble Scaling Up?

- Continue to look for performance gains within your scope
 - Windows OS
 - SQL Server
 - Application data handling
- Faster / larger host hardware
 - Increase the maximum footprint of your VM
- More host hardware
 - Spread out the workload
- Faster SAN
 - Flash & hybrid arrays
 - Faster & more storage interconnects



Trouble Scaling Out?

- Dedicated SQL Server virtual infrastructure cluster
- Segregate workloads
- Resource pools & resource reservations
- Ongoing “right-sizing”
- Individual VM efficiency gains add up
- More host hardware
 - Usually memory bound rather than CPU
- Faster SAN
 - Flash & hybrid arrays
 - Faster & more storage interconnects



Conclusions

- Go forth and virtualize your remaining physical servers
- Double check the performance of your virtual servers
- Apply these tips to all of your virtual SQL Servers
 - Ongoing, not one-time
- Efficiency goes up
- Performance goes up
- Consolidation ratio goes up
- Licensing costs go down
- Customer satisfaction goes up
- Your stress levels go down



Questions?

-  @kleegeek
-  davidklee.net
-  gplus.to/kleegeek
-  linked.com/a/davidaklee

© 2014 Heraflux Technologies

29



Thanks for attending!

www.heraflux.com

© 2014 Heraflux Technologies