

# Monitoring a Virtualized Database Server

Thomas LaRock  
Head Geek  
SolarWinds



# Why Are You Here?

- » You have virtualized database servers
- » You are going to have virtualized database servers

# Why Are You Here?

- » You have virtualized database servers
- » You are going to have virtualized database servers
- » You want/need to troubleshoot
- » You want to avoid issues

# A Little About Me...



SQL Server 2008




# What Is Virtualization?

- » Virtualization turns hardware into a resource queue
- » Started with mainframes, fell out of favor when PC became popular, but making a comeback

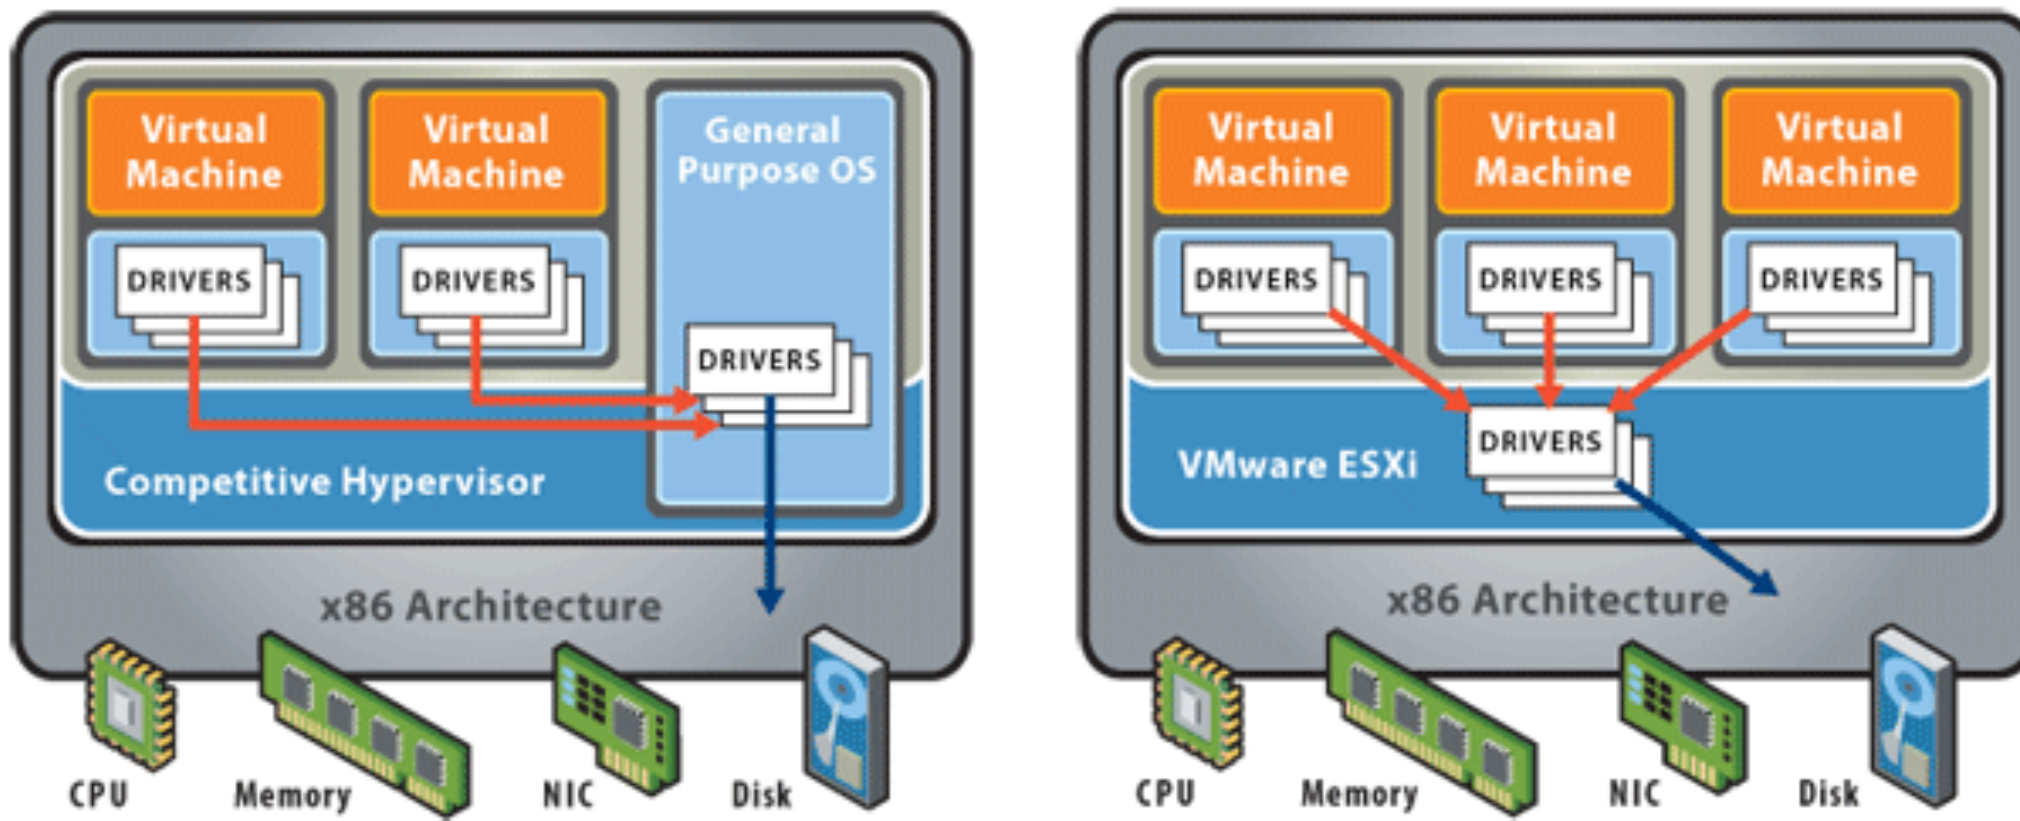


<http://www.dan-dare.org/FreeFun/Images/TheMatrixWallpaper800.jpg>

# Virtualization Vendors



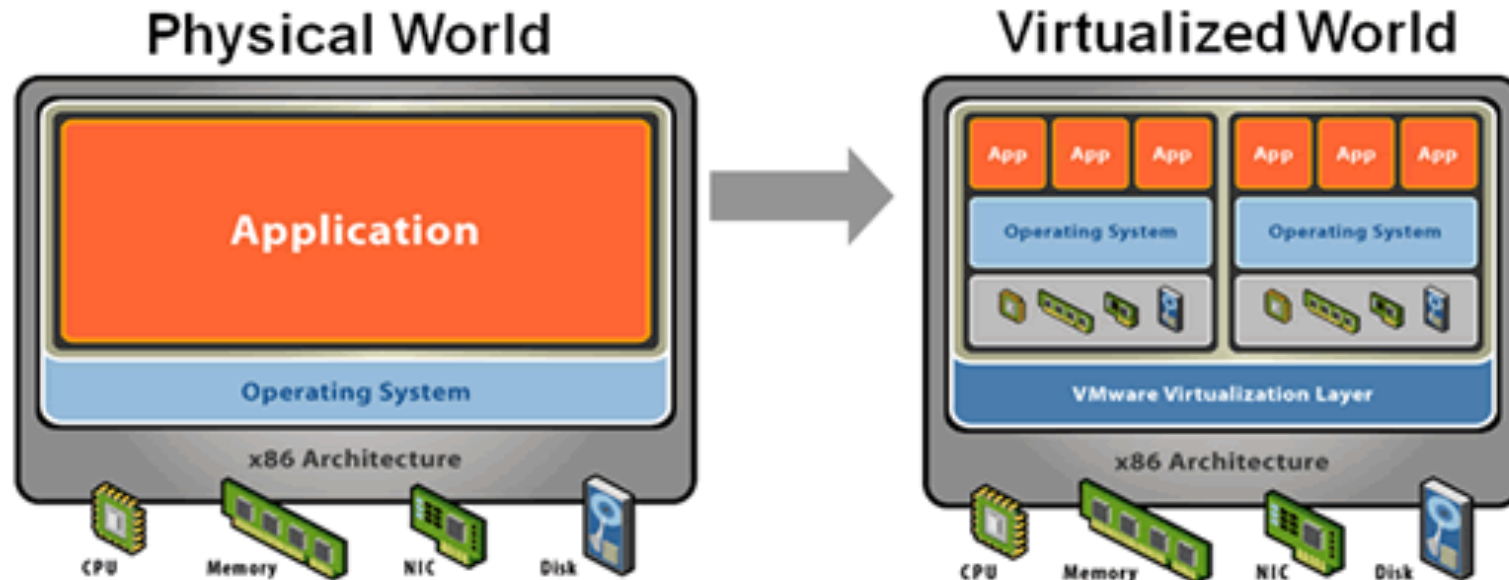
# Hypervisor Layer



<http://www.vmware.com/files/images/thumbnails/vmw-dgrm-vsphere-most-secure-hypervisor-lg.gif>



# Physical vs. Virtual



## Traditional x86 Architecture

Single OS image per machine  
Underutilized resources

## Virtualization

OS and application contained in a single file  
Applications are isolated from one another  
Hardware independence and flexibility

<http://cdn.ttgtmedia.com/ITKE/uploads/blogs.dir/28/files/2009/10/whatisvirt21.jpg>



# Why Virtualize?

- » Maximize resource usage
  - Finance loves you

# Why Virtualize?

- » Maximize resource usage
  - Finance loves you
- » Biggest strength – shared resources
- » Biggest weakness – shared resources!

# Why Virtualize?

- » Maximize resource usage
  - Finance loves you
- » Biggest strength – shared resources
- » Biggest weakness – shared resources!
- » Database servers often last to go

# Resource Bottlenecks



<http://www.econtech.com/newsletter/img/delays.jpg>

» They don't change!

- CPU
- Memory
- Disk I/O
- Network
- Locking/blocking

# Database Monitoring

» Standard tools may include:

- Perfmon
- MDW
- SQL Server DMVs

# Database Monitoring

» Standard tools may include:

- Perfmon
- MDW
- SQL Server DMVs

» **Waits and queues are vital**

# Database Monitoring

## » Standard tools may include:

- Perfmon
- MDW
- SQL Server DMVs

## » **Waits and queues are vital**

## » Must examine hosts

- vSphere, vCenter
- System Center VMM



# Database Monitoring

## » Standard tools may include:

- Perfmon
- MDW
- SQL Server DMVs

## » Waits and queues are vital

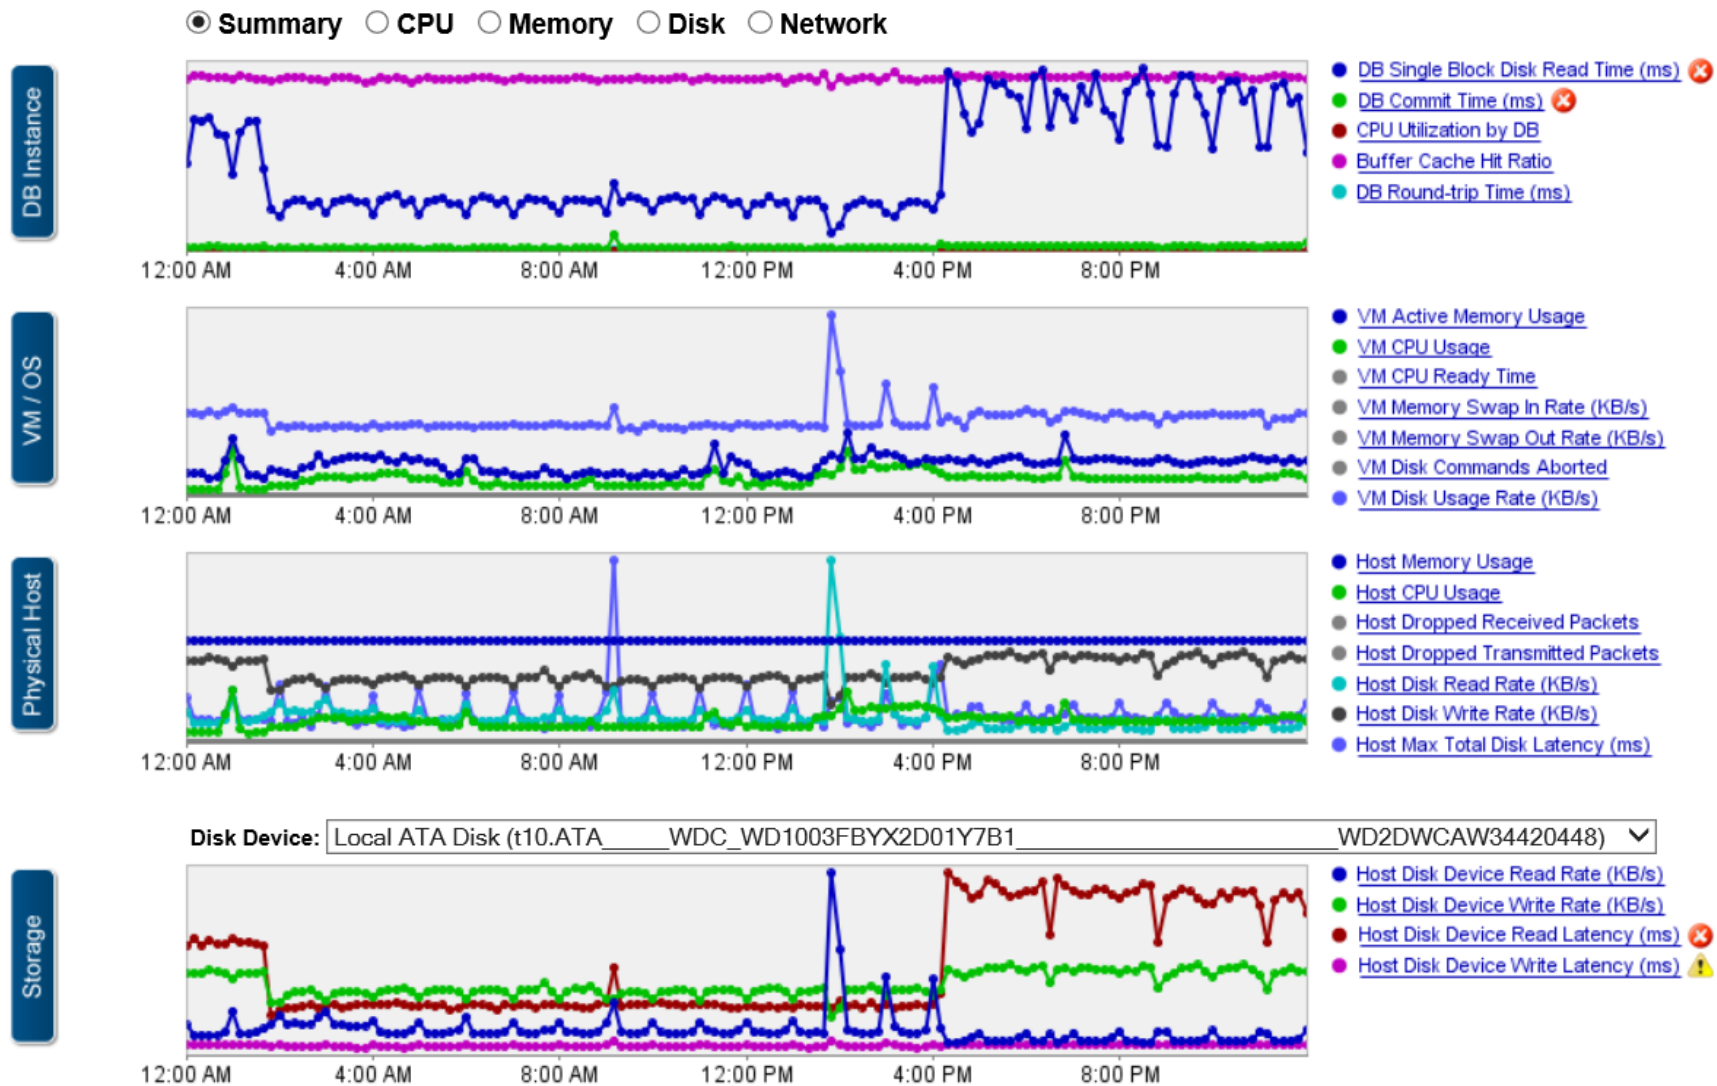
## » Must examine hosts

- vSphere, vCenter
- System Center VMM

## » I want to eliminate host/guest as culprit

- **BEFORE** trying to tune a query!

# Which Layer of Cake?



# Agenda

- » CPU
- » Memory
- » Disk
- » Network
- » Tips/tricks
- » Questions/comments/concerns

# CPU Configuration

## » vCPU

- Start at 1.5:1 ratio of vCPU to logical cores
- vNUMA and hot-plug
- Idle vCPUs can actually hinder performance

# CPU Configuration

## » vCPU

- Start at 1.5:1 ratio of vCPU to logical cores
- vNUMA and hot-plug
- Idle vCPUs can actually hinder performance

## » Know how many vCPUs you need?

- Probably not, so start with 1 and go up as needed
- Remember vCPU \* GHz formula

# CPU Configuration

## » vCPU

- Start at 1.5:1 ratio of vCPU to logical cores
- vNUMA and hot-plug
- Idle vCPUs can actually hinder performance

## » Know how many vCPUs you need?

- Probably not, so start with 1 and go up as needed
- Remember vCPU \* GHz formula

## » CPU and NUMA

- Set MAXDOP = vNUMA core count

# CPU Monitoring

» Databases are not **typically** CPU bound



# CPU Monitoring

» Databases are not **typically** CPU bound

» BMFF metric – VM Ready Time

- Similar to 'runnable' queue
- Amount of time VM was 'ready' to run, but needed to wait for CPU resource
- Hyper-V Hypervisor Virtual Processor\CPU Wait Time Per Dispatch



[http://2.bp.blogspot.com/\\_jIPNvBuH41s/TLgMSVTkyQI/AAAAAAAAAFu/K\\_atz9IM5QQ/s1600/Intel-718028.jpg](http://2.bp.blogspot.com/_jIPNvBuH41s/TLgMSVTkyQI/AAAAAAAAAFu/K_atz9IM5QQ/s1600/Intel-718028.jpg)

# CPU Monitoring

- » Databases are not **typically** CPU bound
- » BMFF metric – VM Ready Time
  - Similar to 'runnable' queue
  - Amount of time VM was 'ready' to run, but needed to wait for CPU resource
  - Hyper-V Hypervisor Virtual Processor\CPU Wait Time Per Dispatch
- » Other metrics include:
  - VM CPU utilization
  - Host CPU utilization



[http://2.bp.blogspot.com/\\_jIPNvBuH41s/TLgMSVTKyQI/AAAAAAAAAFu/K\\_atz9IM5QQ/s1600/Intel-718028.jpg](http://2.bp.blogspot.com/_jIPNvBuH41s/TLgMSVTKyQI/AAAAAAAAAFu/K_atz9IM5QQ/s1600/Intel-718028.jpg)

# Is There a CPU Problem?

## » Where does it hurt?

- DB performance is “slow”, for no apparent reason
- Host server has high CPU utilization

# Is There a CPU Problem?

## » Where does it hurt?

- DB performance is “slow”, for no apparent reason
- Host server has high CPU utilization

## » Possible causes

- CPU over committed
- Additional VM guests (planned or unplanned)

# Is There a CPU Problem?

## » Where does it hurt?

- DB performance is “slow”, for no apparent reason
- Host server has high CPU utilization

## » Possible causes

- CPU over committed
- Additional VM guests (planned or unplanned)

## » Potential remedies

- Move VM guests back to where they belong
- Reconfigure vCPU settings
- Purchase more CPUs

# CPU Scenario

- » Measure host CPU usage
  - AVG > 75% -or- Peak > 90%
- » Check guest VM ready time
  - If any vCPU > 1000ms (or 10%)
- » Host CPU saturation exists
  - Balance guest resources
  - Get more CPUs for the host

# Memory Configuration

## » vRAM assigned to VM guest

- Ballooning – unused memory that was given back
- Swapping – memory given back forcibly



# Memory Configuration

- » vRAM assigned to VM guest
  - Ballooning – unused memory that was given back
  - Swapping – memory given back forcibly
- » Plan for 80% host allocation

# Memory Configuration

- » vRAM assigned to VM guest
  - Ballooning – unused memory that was given back
  - Swapping – memory given back forcibly
- » Plan for 80% host allocation
- » Set Memory Reservation  $\geq$  DB Memory

# Memory Configuration

- » vRAM assigned to VM guest
  - Ballooning – unused memory that was given back
  - Swapping – memory given back forcibly
- » Plan for 80% host allocation
- » Set Memory Reservation  $\geq$  DB Memory
- » Configure min/max SQL memory

# Memory Configuration


- » vRAM assigned to VM guest
  - Ballooning – unused memory that was given back
  - Swapping – memory given back forcibly
- » Plan for 80% host allocation
- » Set Memory Reservation  $\geq$  DB Memory
- » Configure min/max SQL memory
- » Lock Pages in Memory granted

# Memory Monitoring

- » Databases are **often** memory bound
  - Often due to bad configurations

# Memory Monitoring


- » Databases are **often** memory bound
  - Often due to bad configurations
- » BMFF – Swapping, Ballooning
  - If you got the memory, you want to keep it



<http://www.gadgetimages.org/wp-content/uploads/2013/07/ram.jpg>

# Memory Monitoring

- » Databases are **often** memory bound
  - Often due to bad configurations
- » BMFF – Swapping, Ballooning
  - If you got the memory, you want to keep it
- » Other Metrics include
  - Guest and Host Memory Utilization
  - Available Mbytes



<http://www.gadgetimages.org/wp-content/uploads/2013/07/ram.jpg>



# Is There a Memory Problem?

## » Where does it hurt?

- Low memory counters
- Host ballooning > 0

# Is There a Memory Problem?

## » Where does it hurt?

- Low memory counters
- Host ballooning > 0

## » Possible causes

- Host/VM/SQL memory misconfigured

# Is There a Memory Problem?

## » Where does it hurt?

- Low memory counters
- Host ballooning > 0

## » Possible causes

- Host/VM/SQL memory misconfigured

## » Potential remedies

- Set VM memory reservation = memory provisioned
- Check memory allowed per NUMA node

# Memory Scenario

- » Measure Available Mbytes
  - Is it  $< 300$ , or  $< 10\%$  (which do you prefer?)
- » Check VM Ballooning
  - Is it  $> 0$ ?
  - Set min/max to non-default values
- » Assume min/max is set, check for 'lock pages in memory' granted to service account
  - If not, then DO IT!
  - Use of 'large pages'

# Disk Configuration

## » You are sharing your storage!

- Storage issues are usually related to configuration and not capabilities of ESX
- Follow best practices from storage vendor

# Disk Configuration

## » You are sharing your storage!

- Storage issues are usually related to configuration and not capabilities of ESX
- Follow best practices from storage vendor

## » Create dedicated datastores for databases

# Disk Configuration

## » You are sharing your storage!

- Storage issues are usually related to configuration and not capabilities of ESX
- Follow best practices from storage vendor

## » Create dedicated datastores for databases

## » Isolate data and log files

# Disk Monitoring

## » BMFF – Latency

- Host maxTotalLatency
- Host Device Latency (by device)
- VM Command Latency (for all VMs)



<http://tinyurl.com/m5oxpdj>



# Disk Monitoring

## » BMFF – Latency

- Host maxTotalLatency
- Host Device Latency (by device)
- VM Command Latency (for all VMs)

## » Other metrics include

- Disk Read Rate
- Disk Write Rate



<http://tinyurl.com/m5oxpdj>

# Is There a Disk Problem?

## » Where does it hurt?

- High latencies

# Is There a Disk Problem?

## » Where does it hurt?

- High latencies

## » Possible causes

- Overloaded storage
- Overloaded network
- I/O intensive queries
- Memory/CPU issues

# Is There a Disk Problem?

## » Where does it hurt?

- High latencies

## » Possible causes

- Overloaded storage
- Overloaded network
- I/O intensive queries
- Memory/CPU issues

## » Potential remedies

- Reconfigure storage and/or network
- Rewrite queries

# Disk Scenario

- » Monitor disk latency
  - Is Host maxTotalLatency > 30ms ?
- » Is VM Command Latency  $\geq$  30ms for your VM?
  - Look for PAGEIOLATCH\_XX waits
  - Tune Disk I/O intensive processes on database
  - Are Memory / CPU issues causing I/O problems
- » Review device latency
  - Review Disk Read/Write rates for that device

# Network Configuration

## » How big is that pipe?

- vSwitch – software switch inside of VMKernel
- vSwitch can be tied to 1 or more NICs
- VMWare claims to handle over 30GB/sec
- Most db's don't need that much

# Network Configuration

## » How big is that pipe?

- vSwitch – software switch inside of VMKernel
- vSwitch can be tied to 1 or more NICs
- VMWare claims to handle over 30GB/sec
- Most db's don't need that much

## » VMXNET paravirtualized network adapter

- Installed into guest O/S capable of 1Gbps
- Minimizes overhead between VM and Host
- Requires VMware Tools

# Network Configuration

## » How big is that pipe?

- vSwitch – software switch inside of VMKernel
- vSwitch can be tied to 1 or more NICs
- VMWare claims to handle over 30GB/sec
- Most db's don't need that much

## » VMXNET paravirtualized network adapter

- Installed into guest O/S capable of 1Gbps
- Minimizes overhead between VM and Host
- Requires VMware Tools

## » Co-locate VMs when possible



# Network Monitoring

- » BMFF – Dropped Packets
- » Other Metrics Include
  - Network Rate
  - Bytes Total/sec
  - Output Queue Length



<http://tinyurl.com/n8nqwuu>

# Network Monitoring

- » BMFF – Dropped Packets
- » Other Metrics Include
  - Network Rate
  - Bytes Total/sec
  - Output Queue Length
- » Look for dropped packets
  - Any errors on host NICs?
  - One NIC getting more traffic than another?



<http://tinyurl.com/n8nqwuu>

# Network Monitoring

- » BMFF – Dropped Packets
- » Other Metrics Include
  - Network Rate
  - Bytes Total/sec
  - Output Queue Length
- » Look for dropped packets
  - Any errors on host NICs?
  - One NIC getting more traffic than another?
- » Is Network Rate is getting close to max?



<http://tinyurl.com/n8nqwuu>

# Is There a Network Problem?

## » Where does it hurt?

- Network performance is slow or intermittent

# Is There a Network Problem?

## » Where does it hurt?

- Network performance is slow or intermittent

## » Possible causes

- Overloaded host CPU
- Single NIC
- Wrong network adapter drivers

# Is There a Network Problem?

## » Where does it hurt?

- Network performance is slow or intermittent

## » Possible causes

- Overloaded host CPU
- Single NIC
- Wrong network adapter drivers

## » Potential remedies

- Reduce CPU pressure on host
- NIC teaming

# Network Scenario

- » Monitor droppedTx, droppedRx for all vmnic objects
  - Are they > 0?
- » Check for overloaded host CPU
  - Add additional vCPU
  - Add additional vNIC
- » Check VM driver configuration for network devices
  - Bad config options lead to bad network perf

# Tips and Tricks

1. Don't build your own host from spare parts
  - Unless your name is "MacGyver"



<http://www.unclebobs.com/getstorganized/wp-content/uploads/2012/11/mcgyver.jpeg>



# Tips and Tricks

## 2. Baseline/benchmark for performance

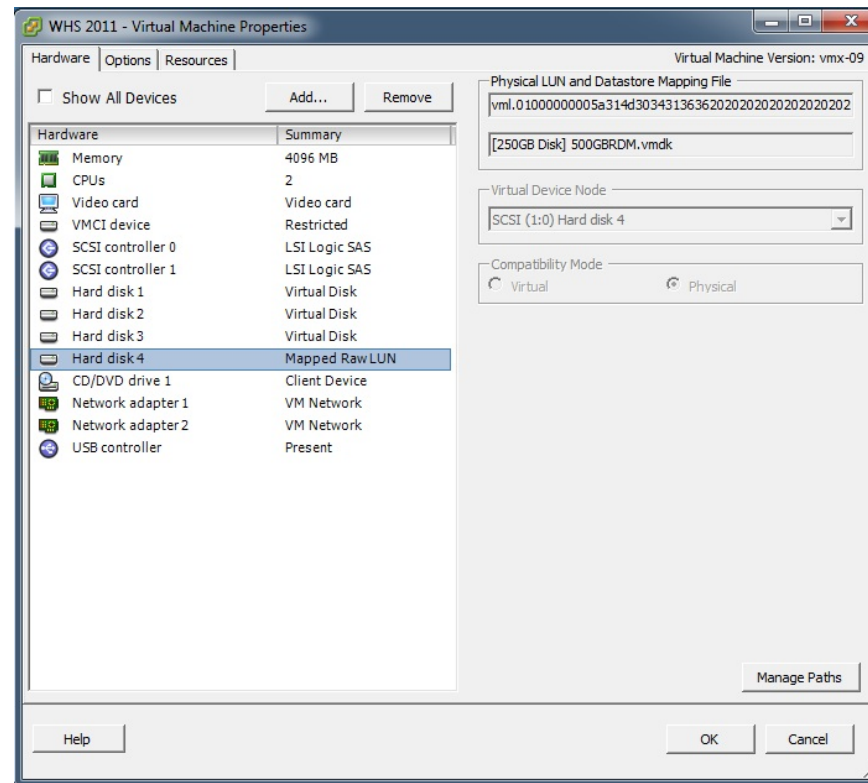
- Otherwise you have no idea what is “good” or “bad”



<http://www.gulfcoastnews.com/2010-Images/KatrinaFloodMarkerRodenburg.jpg>

### 3. Know datastore options

- VMFS versus RDM, which one is right for you

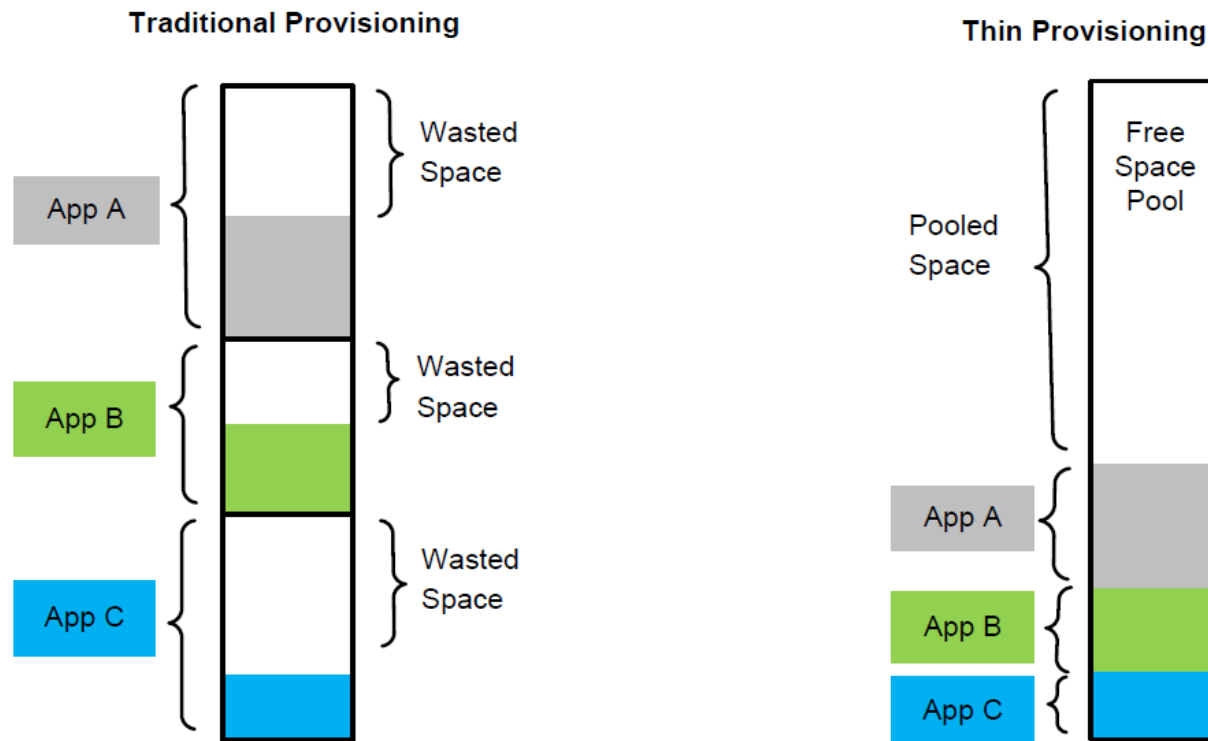


<http://forza-it.co.uk/wp-content/uploads/Properties-of-Disk-4-RDM.jpg>

# Tips and Tricks

## 4. Avoid thin provisioning

- Unless you enjoy headaches



<https://communities.netapp.com/servlet/JiveServlet/showImage/38-8422-15789/ThinProvisioning+1.png>

# Tips and Tricks

## 5. Avoid over-allocation of CPU and memory

- Over-allocation leads to over-commit



[http://www.endlessssimmer.com/wp-content/uploads/2009/10/3590888947\\_7e80ca52df.jpg](http://www.endlessssimmer.com/wp-content/uploads/2009/10/3590888947_7e80ca52df.jpg)

# Tips and Tricks

## 6. Don't trust O/S counters

- O/S doesn't know it is virtualized



<http://www.bryaneisenberg.com/wp-content/uploads/2013/11/bigstock-Mousetrap-35879422.jpg>



# Tips and Tricks

## 7. Running it all at once

- Know your workloads



<http://epicswag.net/wp-content/uploads/2012/02/816.jpg>

# Tips and Tricks

## 8. Capacity planning

- Leave room for growth, failovers



<http://www.paleycenter.org/assets/public-programs/PP-spring-2013/DrWhoFridge-web.jpg>

# Quick BMFF Sheet

Resource	Metric	Host / VM	Description
CPU	Ready	VM	CPU time spent in ready state
	Usage	Both	CPU usage as a percentage during a defined interval
Memory	Swapin, Swapout	Both	Memory the host swaps in/out from/to disk (per VM, or cumulative over host)
	Vmmemctl	Both	Amount of memory reclaimed from resource pool by way of ballooning
Disk	maxtotallatency	Host	Highest latency value across all disks used by the host.
	deviceLatency	Host	Average time to complete a command from the physical device.
	totalLatency	Host	Average latency in all guests.
Network	droppedTx, dropped Rx	Both	Drop packets per second
	usage	Both	Sum of data transmitted and received



# Questions?



## For More Information

- » <http://tinyurl.com/vm-perf-counters>
- » <http://tinyurl.com/common-vm-network-issues>
- » <http://tinyurl.com/perf-whitepaper>
- » <http://tinyurl.com/vmware-sql-bp-guide>
- » <http://tinyurl.com/vmdk-or-rdm>
- » <http://tinyurl.com/large-pages>
- » <http://tinyurl.com/klee-vnuma>
- » <http://tinyurl.com/measure-hyper-v-perf>