

« Upgrade your Virtual Hardware in a few minutes, with a twist.

Performance Troubleshooting VMware vSphere – Network »

Search

Search

Recent Comments


- > B. Riley on Single Throat to Choke
- > John Call on Single Throat to Choke
- > Cg on Understanding ESX/ESXi Equivalency...Are we there yet?
- > Jeff Snively on Single Throat to Choke
- > MrLoftcraft on Upgrade your Virtual Hardware in a few minutes, with a twist.

Archives

- > July 2010
- > June 2010
- > May 2010
- > April 2010
- > March 2010
- > February 2010
- > January 2010
- > December 2009
- > November 2009
- > October 2009
- > September 2009
- > August 2009
- > June 2009
- > May 2009
- > April 2009
- > March 2009
- > February 2009
- > January 2009
- > December 2008
- > November 2008
- > October 2008
- > August 2008
- > July 2008
- > June 2008
- > May 2008

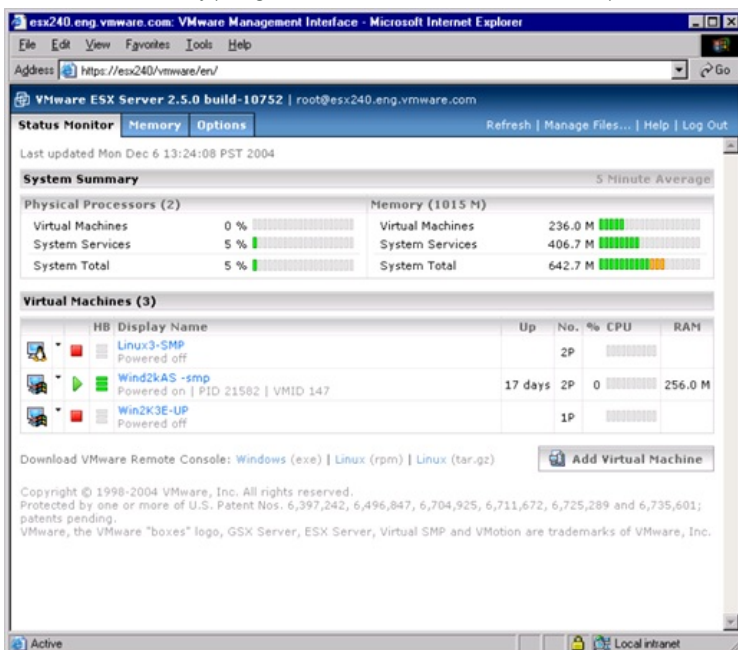
Performance Troubleshooting VMware vSphere – Storage

March 16th, 2010 | Author: Scott Sauer



Introduction

Personally one of the most interesting components of the VMware architecture I/O stack is storage. There are a plethora of diverse storage solutions in the industry today that offer unique different ways of addressing storage performance, as well as the increase in capacity demands. Storage problems are the most common mis-configuration effecting performance that exists in VMware today. An oversaturated LUN will effect all virtual machines that share that same data store. Take this concept up a level, a group of disks (RAID group) that are saturated with I/O will negatively impact all LUNS that share those same physical spindles. Storage traditionally has been the "red headed step child" in VMware and hasn't gotten a lot of visibility. Storage I/O bottlenecks can create serious virtual machine problems and yet it wasn't until ESX 3.x that graphic visibility was even displayed to VMware administrators, see 2.x MUI reflects CPU and memory (Management User Interface for those newer to VMware).




This is still evident today, consider DRS, it has the ability to migrate based on CPU and memory, but storage is an after thought. Speculation and rumor suggest that VMware has been working on storage DRS and will introduce this sometime in the future. With all of the different solutions and approaches from different vendors, I will try to focus in a more general overview of storage concepts rather than specific array types and how they work. Having said that, my background is with EMC storage products, so excuse any "biased" terminology or methodologies.

Fundamentals

Before getting into too much detail on troubleshooting let's cover some of the basics with storage to ensure we are on the same page. From a VMware perspective, there are different storage types that the server can utilize for its VMFS (Virtual Machine File System) partitions. Local storage which can range from anything from IDE drives to locally attached SCSI drives, or shared storage. Shared storage connectivity offers a distinct advantage over locally attached storage as it allows you to leverage VMotion which can actively move a virtual machine from one physical host to another.

Three shared storage options that are supported in VMware vSphere are fiber channel (SAN), iSCSI (hardware or software) and NFS. NFS and iSCSI use standard ethernet networks rather than a devoted fiber channel switched network. Traditionally with VMware, fiber channel SANs have been the more widely adopted and deployed configurations since historically it has been around the longest and offers fast, reliable connectivity. iSCSI and NFS have quickly caught up with fiber channel SANs as more vendors offer outstanding products while avoiding some of the expenses that are tied to a traditional SAN. The following chart depicts testing that VMs are conducted on the various options, (lower is better).





Raw device mappings (RDM) is yet another option that VMware allows an administrator to utilize. A raw device mapping is simply the act of giving a virtual machine full and complete access to a LUN, there is no underlying VMFS structure. A great use case for this might be a database that needs guaranteed disk performance and can't suffer contention.

Across these different shared storage options one can now choose a variety of configuration methods. SATA disks, fiber channel disks, flash disks are all now common drive types that one can implement in an enterprise storage array. RAID types can vary from RAID 1, 10 and 5, and there are technologies that will allow you to stack multiple RAID types together for performance (Metaluns). Storage vendors typically leverage cache to assist with workloads by having frequent requests sit in memory since traditional hard disks take longer to access data (seek time). If you're not a storage administrator this might sound overwhelming, and sometimes it is. You need to step through the entire stack when looking at storage related problems. Check the array, then check the switch, then check the host. Make sure you are analyzing the entire gamut, not just the part you are responsible for (which may just be the ESX server).

What to look for


Before I jump into my bulleted list of issues to be aware of, I wanted to stress some key concepts that will help prevent storage related problems down the road. Physical and virtual workloads share the same characteristics when it comes to storage, and that is **planning**. The more requirements you can get up front from your customer or your own environment, the better this will aid in your storage configuration design work. Don't let virtualization add a layer of confusion, try and keep things simple. If you had a high I/O SQL database on a physical server, would you put this on NFS using a RAID 3+1 SATA configuration? Then don't do it in a virtual environment! Consider your workloads and give them the resources they need to thrive as an application or you will have users coining phrases like "VMware performance is slow".




- ✦ Check for excessive demands placed on a LUN. This is the number one cause of storage performance problems and sometimes the hardest to identify. Consider this, one physical fiber channel disk drive will average 125 IOPS per disk. Have you given your LUN enough physical spindles to accommodate the workload? Are you using the right RAID type?
- ✦ Check for excessive demand on the disk group or RAID group. Is your VMFS LUN sharing physical spindles with other LUNS? A LUN that might be used by another application that is not performing correctly can put a serious load on your VMware environment.
- ✦ Check for front end saturation on array. Are you oversubscribing your fiber or ethernet connections of the array? Make sure you are giving enough bandwidth to the hosts that are connecting to the storage device.
- ✦ Check the ESX host ports for saturation. Are you giving the host enough bandwidth to the array? Are you saturating your connectivity and need to add more bandwidth?
- ✦ Check the array cache performance. Is it time to upgrade the cache on the storage processors? Maybe you have completely outgrown your array and it's time to consider moving towards Solid State drives for heavier workloads.
- ✦ Know your **configuration maximums** on ESX storage. Don't give your ESX clusters an excessive number of LUNS, this can create excessive SCSI chatter between the hosts and create an additional load on the storage.
- ✦ Leverage Paravirtual SCSI drivers for heavy workloads inside the virtual machine.
- ✦ Consider backups. Backups can put a tremendous load on your infrastructure from the network all the way down to the storage array.
- ✦ Try to isolate the problem and re-create in a test environment if possible.

Monitoring with Virtual Center

The first place I would start with checking storage configurations is Virtual Center. Virtual Center provides excellent reporting and gives you granular control over which metrics you would like to report against. VMware vSphere now includes a nice graphical summary in the performance tab of the physical host. This gives you a quick dashboard type view of the overall health of the system over a 24 hour period. Here are some storage samples:



Can you take a guess when the backup occurred on the host above?





Monitoring with ESXTOP


Esxtop is another excellent way to monitor performance metrics on an ESX host. Similar to the Unix/Linux "Top" command, this is designed to give an administrator a snapshot of how the system is performing. SSH to one of your ESX servers and execute the command "esxtop". The default screen that you should see is the CPU screen, if you need to monitor the disk adapters select the "d" key. If you would like to monitor the LUNS specifically select the "u" key. Esxtop gives you great real-time information and can even be set to log data over a longer time period, try "esxtop -a -b > performance.csv". Check your I/O on your physical adapters here. Examine what your virtual machines are doing, if you want to isolate the display to the virtual machine worlds hit the "v" key. Track the following metrics:

GuestDAVG/cmd -Average latency (ms) from the Device
(LUN)KAVG/cmd -Average latency (ms) in the
VMKernelGAVG/cmd -Average latency (ms) in the Guest


10:14:26am up 63 days 20:00, 261 worlds; CPU load average: 0.28, 0.29, 0.30															
ADAPTER	CID	TID	LID	NCHNS	NTGTS	NLUNS	CMDS/s	READS/s	WRITES/s	MBREAD/s	MBWRN/s	DAVG/cmd	KAVG/cmd	GAVG/cmd	QAVG/cmd
vmhba0	-	-	-	1	1	1	39.88	0.00	39.88	0.00	2.10	5.63	0.01	5.64	0.03
vmhba1	-	-	-	1	4	100	142.18	2.60	139.58	0.01	2.87	1.00	0.01	1.01	0.00
vmhba2	-	-	-	1	4	100	47.16	19.42	27.74	0.44	0.17	2.01	0.40	2.41	0.00

Monitor at the Storage Array

Each storage array typically has some level of monitoring and reporting that you can utilize to assist in your troubleshooting efforts. These tools might not be accessible to you depending on your organization and/or customer. I can't speak to all array types but here is an illustration of a EMC Clarion array with LUNS that are being used for VMs and are trying to find the "hot spots".



Look for your LUN outliers, make sure your storage processors are evenly balanced in terms of workloads. The chart below is displaying an unbalanced storage processor configuration. Try to correct this by balancing your LUNS when created, even numbers on SPA, odd numbers on SPB.



Conclusion

If your using VMWare vSphere, there are many different ways to monitor for storage related problems. The Virtual Center database is the first place you should start. Check your physical storage adapters, then work your way down the stack to the LUNS, then the virtual machine(s) that might be indicating a problem. Take a look at esxtop, check some of the key metrics that we discussed above. Storage, unlike some of the other VMWare components needs to be closely examined. You might need to use third party tools to assist your troubleshooting efforts. Check all components of the architecture like the adapters, the switch and then the array. They all need to work in conjunction together to provide quality performance to your virtual machines. If all else fails, engage VMWare support and open a service request. Support contracts exist for a reason and I have opened many SRs that were new technical problems that have never been discovered by VMWare support.

Related Posts:

- Performance Troubleshooting VMWare vSphere – Network
- Performance Troubleshooting VMWare vSphere – CPU
- Performance Troubleshooting VMWare vSphere – Memory
- More Bang for Your Buck with PVSCSI (Part 1)
- Performance Troubleshooting VMWare vSphere – The Tetralogy
- Powered by Contextual Related Posts

DISQUS ▾

Sort by Popular now [Subscribe by email](#) [Subscribe by RSS](#)

have you considered looking at solutions such as Atlantis Computing Vscaler to decrease the load on Storage?

Like Reply

Great article.

Like Reply

Excellent Article

Like Reply

Thanks! Storage is a hard topic to cover since there are so many solutions out there. Glad you liked it.

Like Reply

Post as ...

Trackback URL

blog comments powered by DISQUS



Copyright © 2010 virtual insanity. All