

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 – Memory

February 19th, 2010 | Author: [Scott Sauer](#)



Introduction

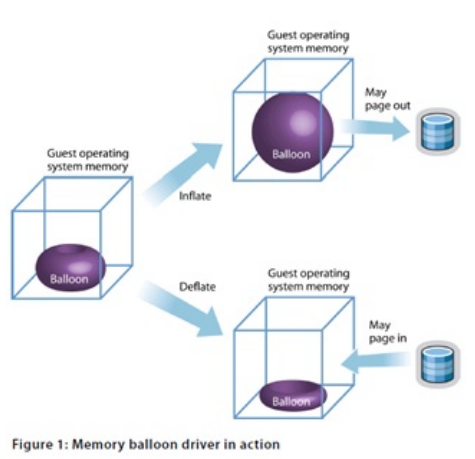
As memory prices continue to drop and the x64 bit architecture is embraced and adopted more in the industry, we continue to see a rise in memory demands. Only a few years ago, 1-2 GB virtual machines were the norm, 95% of these being 32 bit operating systems. From my personal experience I have seen this trend change to 2-4 GB as a norm, with the more high performing virtual machines consuming anywhere from 4-16 GB of memory. VMware has answered this demand with vSphere now delivering up to 1TB of addressable memory per physical host, and up to 255GB per virtual machine.

With processors now more powerful than ever, the general shift of virtual machine limitations is changing from compute to memory. This is reflected in our industry today as we see an increase in the memory footprint on traditional servers (Intel Nehalem), and vendors such as Cisco introducing **extended memory technology** which can more than double the standard memory configuration. I recently had the opportunity to sit in on a Cisco Unified Computing System architectural overview class, and was impressed with what I saw. The extended memory technology is quite unique because it not only allows you to scale out on your memory configuration, it uses a special ASIC to virtualize the memory so there is no reduction in bus speed. A financial advantage to having this many DIMM sockets is you can use lower capacity DIMMs (2 GB or 4GB) to achieve the same memory configuration in a standard server where you would have to use 8GB DIMMs.

Memory Technologies in VMware vSphere

There are some major benefits of virtualization when it comes to memory. VMware implements some sophisticated and unique ways of maximizing physical memory workloads within an ESX host. All of these features work out of the box with no advanced configuration necessary. To understand problems that might occur in your environment you need to be familiar with these basic memory concepts.

- ❏ **Transparent Page Sharing** – The VMkernel will compare physical memory pages to find duplicates, then free up this redundant space and replaces it with a pointer. If multiple operating systems are running on one physical host, why should you load the same files multiple times? Think of this as the data de-duplication process we are seeing in a majority of backup solutions in the industry.
- ❏ **Memory Overcommitment** – The act of assigning more memory to powered on virtual machines than the physical server has available. This allows for virtual machines that have heavier memory demands to utilize the memory that is not actively being used on under utilized machines.
- ❏ **Memory Overhead** – Once a virtual machine is powered on the ESX host reserves memory for the the normal operations of VMware infrastructure. This memory can't be used for swapping or ballooning, and is reserved for the system.
- ❏ **Memory Balloon Driver** – When VMware tools are installed on a virtual machine they provide device drivers into the host virtualization layer, from within the guest operating system. Part of this package that is installed is the balloon driver or "vmtoolsd" which can be observed inside the guest. The balloon driver communicates to the hypervisor to reclaim memory inside the guest when it's no longer valuable to the operating system. If the Physical ESX server begins to run low on memory it will grow the balloon driver to reclaim memory from the guest. This process reduces the chance that the physical ESX host will begin to swap, which you will cause performance degradation. Here is an illustration of ballooning in ESX:



What to look for

- ❏ Check ESX host swapping. If you are overcommitting memory on the physical ESX host you can run into a situation when each virtual machine is in need of the total amount of what is granted. When the host is out of memory it will begin to page out. Keep an eye on your oversubscription

rates of physical hosts, or ensure you have enough memory resources across your DFS clusters so it can balance the load more effectively. Swapping will occur when the following formula is met:

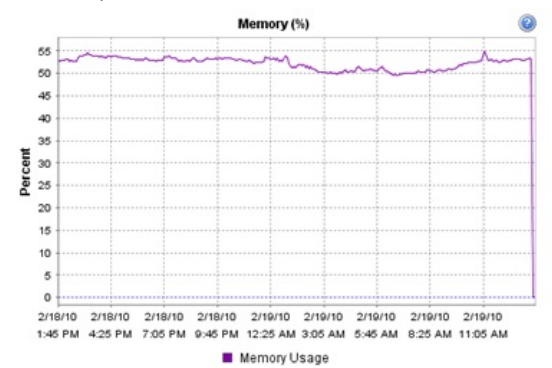
$$\text{Total_active_memory} > (\text{Memory_Capacity} - \text{Memory_Overhead}) + \text{Total_balloonable_memory} + \text{Page_sharing_savings}$$

- Check for Virtual machine swapping. Make sure your virtual machines have enough memory for the application workload that they are supporting. If virtual machine swapping starts to occur this can put a strain on the disk subsystem.
- Check to ensure VMware tools are installed and updated. VMware tools not only provides drivers from the guest to the hypervisor, but the balloon driver also gets installed with VMware tools. For proper memory management the ESX host relies on the balloon driver to manage memory.
- Check memory reservation settings. By default VMware ESX dynamically tries to reclaim memory when not needed. There are situations when you might choose to utilize memory reservations. If you set memory reservations in your environment be aware that this memory is permanently assigned to the host and can not be reallocated when it's not being used. Don't sell the balloon driver short, many third part application vendors over spec their configurations for personal safety, and ballooning can help counteract some of that wasted "fluff factor".

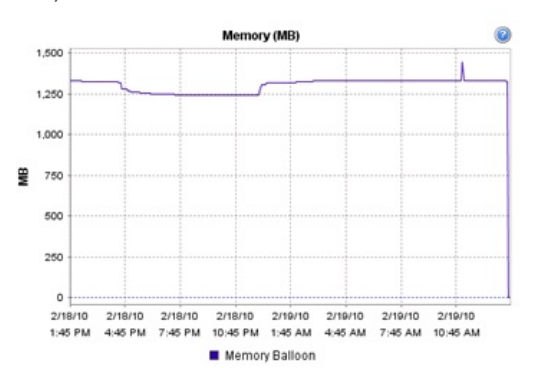
Monitoring with Virtual Center

The first place I would start with checking memory 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 memory samples:

Check your over all % usage (lower is better)

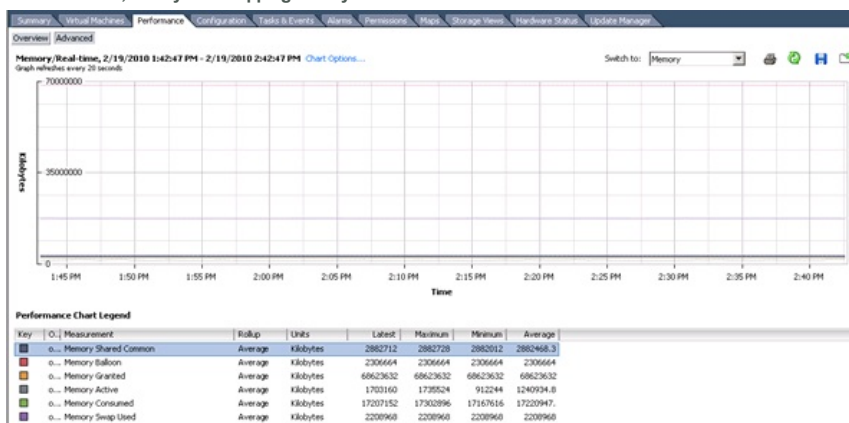


Check your Ballooning (lower is better)



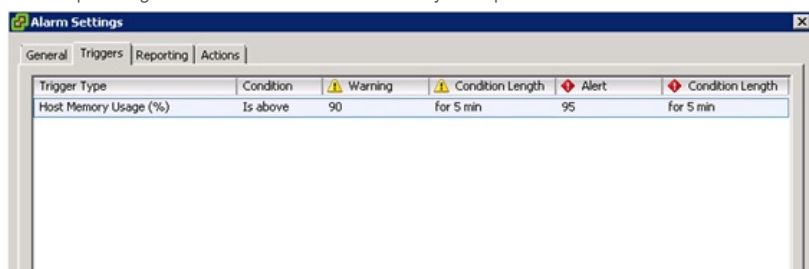
Selecting the advance tab gives you a much more granular way of viewing performance data. At first glance this might look like overkill, but with a little bit of fine tuning, you can make it report on some great historical information. Here is a snapshot of memory utilization with many of the variables we just discussed above, great snapshot of what's going on (looks healthy below):

Check your various metrics, mainly for swapping activity



The virtual center performance statistics by default display the past hour of statistics, and show a more detailed analysis of what's currently happening on your host. Select the option "Chart Options" to change values such as time/date range and which counters you would like to display.

Virtual Center Alarms are an excellent tool that can sometimes be overlooked and forgotten about. While this is more of a proactive tool than a reactive or troubleshooting tool, I thought it was worth mentioning. Setup Memory alerts so you will be notified via e-mail if a problem starts to manifest itself. Here is an alarm configured to trigger if physical host Memory usage is above 90% for 5 minutes or greater. A lot of these alerts are built into Virtual Center so you don't have to do a lot of pre-configuration work. You do need to make sure you setup the e-mail notifications under the "Actions" tab.





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 memory select the "m" 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 total Physical memory here, make sure you aren't over committing and causing swapping. Examine what your virtual machines are doing, if you want to just display the virtual machine worlds hit the "V" key.

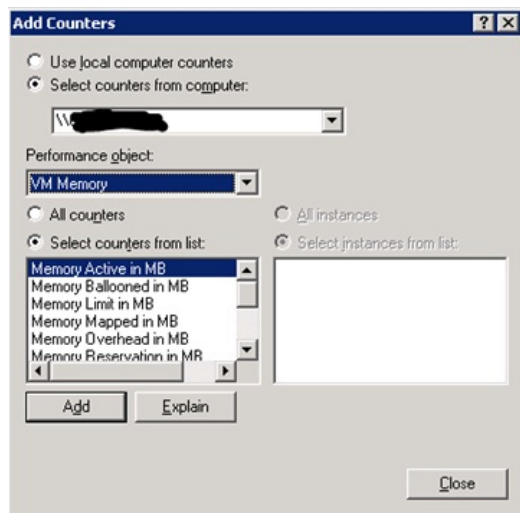
```

2:57:36pm up 140 days 5:54, 169 worlds; MEM overcommit avg: 1.85, 1.85, 1.85
PMEM /MB: 40958 total: 800 cos, 747 vmk, 28274 other, 11136 free
VMKMEM/MB: 39618 managed: 2377 minfree, 7225 rsvd, 31937 ursvd, high state
COSMEM/MB: 74 Free: 1066 swap_t, 1066 swap_f: 0.00 r/s, 0.00 w/s
PSHARE/MB: 109471 shared, 7143 common: 102328 saving 0.00 r/s, 0.00 w/s
SWAP /MB: 276 curr, 30 target: 0.00 r/s, 0.00 w/s
MEMCTL/MB: 0 curr, 0 target, 82527 max
  
```

GID	NAME	MEM/SZ	GRANT	SZTGT	TCHD	SACTV	SACTVS	SACTVP	SACTVN	OVHDUR	OVHD	OVHDMAX
18	vobd.4231	20.63	1.44	20.63	1.44	0	0	0	0	0.00	0.00	0.00
19	vmware-vmkauthd	5.93	2.13	5.93	2.13	0	0	0	0	0.00	0.00	0.00
65		4096.00	4095.50	1195.98	40.96	2	0	1	1	53.55	181.23	215.75
67		4096.00	4095.45	1175.85	40.96	0	0	0	0	53.55	182.02	215.70
69		2048.00	2047.89	960.60	102.40	6	1	5	2	57.54	190.62	237.67
71		4096.00	4095.75	1348.05	163.84	2	2	2	2	53.55	181.56	212.79
73		4096.00	40959.58	3034.05	0.00	0	0	0	0	667.32	906.76	1084.53
86		4096.00	4095.73	2327.55	122.88	3	2	3	1	89.63	232.01	279.55
90		4096.00	4095.73	1941.50	122.88	4	1	3	2	53.55	184.54	215.32
97		2048.00	2047.83	1453.70	102.40	6	3	5	4	37.46	146.88	178.50
103		2048.00	2047.85	2018.95	450.56	12	14	12	12	57.40	190.62	238.01
104		4096.00	40951.35	5471.68	0.00	0	0	0	0	667.32	882.29	1059.04
106		4096.00	4095.65	1665.11	122.88	4	1	3	0	53.55	151.94	183.90
108		4096.00	3962.84	1767.62	245.76	8	2	6	2	53.55	138.02	170.20
113		4096.00	4095.57	1765.90	0.00	1	0	0	0	53.55	141.75	172.94
116		4096.00	3960.60	1152.54	0.00	1	0	0	0	53.55	139.68	170.92
130		2048.00	2047.85	1838.47	143.36	4	4	4	4	57.54	138.41	186.95

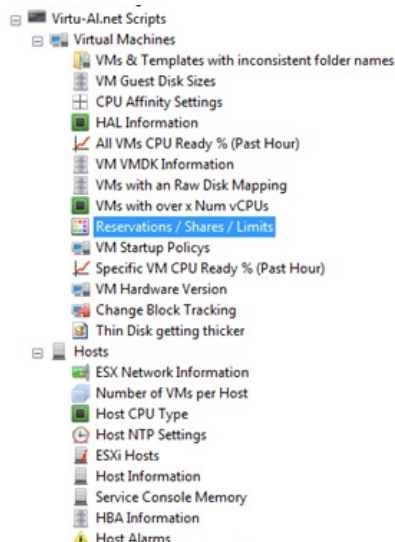
Monitor inside the Virtual Machine

A great feature VMware introduced for Windows virtual machines was integrating VMware performance counters right into the Performance Monitor or "perfmon" tool. If your running vSphere 4 update 1 make sure you [read this post](#) first as there is a bug with the vmttools that will prevent them from showing up. You can monitor the same metrics found in Virtual Center and esxtop here. Just another way of getting at the data especially if you have a background in Microsoft Windows and are familiar with perfmon.



Monitoring with PowerCLI

Another great place to go to for finding potential memory problems and bottlenecks is PowerCLI. I have been using PowerGUI from Quest, accompanied by a powerpack from [Alan Renouf](#). If your not a command line guru don't let this discourage you. PowerGUI is a windows application that allows you to run pre-defined PowerCLI commands against your Virtual Center server or your physical ESX hosts. Want to find out what your ESX host service console memory is set to? How about virtual machines that have memory reservations, shares or limits configured? You can pull all of this information using Alan's powerpack.



Conclusion

If your using VMWare vSphere, there are many different ways to monitor for memory problems. The Virtual Center database is the first place you should start. Check your physical host memory conditions, then work your way down the stack to 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.

Look for the outliers in your environment. If something doesn't look right, that's probably the case. Scratch away at the surface and see if something pops up. Use all possible tools available to you like PowerCLI. Approaching problems from a different perspective will sometimes bring light to a situation you weren't aware of. 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 – CPU
- Performance Troubleshooting VMWare vSphere – Storage
- Performance Troubleshooting VMWare vSphere – Network
- VMWare Windows Perfmon counters missing in vSphere 4u1
- Performance Troubleshooting VMWare vSphere – The Tetralogy
- Powered by Contextual Related Posts



Showing 4 comments

Sort by Popular now Subscribe by email Subscribe by RSS

scottdsauer ★ 5 months ago

Thanks Hany! That is going to be a difficult one especially since there are so many diverse storage solutions out there.



Hany Michael 5 months ago

Another great post Scott! I'm really enjoying this series, and i can't wait for the storage one.



scottdsauer ★ 5 months ago

Thanks Eric, used to rounding up. Fixed.



Eric Sloof 5 months ago

and up to 256GB per virtual machine. Should be 255GB



Add New Comment

Post as ...

Trackback URL

blog comments powered by DISQUS



Content © 2010 virtual insanity. All rights reserved. Privacy Policy