

Memory Management Metrics for VMware Environments

How to effectively manage memory by examining
VMware performance metrics

WHITEPAPER BY JONATHAN KLICK



Introduction

The number one capacity constraint in most virtualized environments is memory. If memory isn't managed properly, an environment's performance and possibly the data center's ROI will suffer. Unfortunately, some virtualization administrators are unfamiliar with which memory metrics should be used, and how to use them to correctly size virtual machines and troubleshoot performance problems. Additionally, in scenarios where the "large memory pages" setting is active, memory management becomes even more complicated. The purpose of this document is to enhance knowledge for virtualized environment memory management and to recommend memory metrics to use for VM (virtual machine) memory allocation sizing, performance issue monitoring and furthermore, how to make these assessments when large memory pages are in use.

Capacity Assessing Memory Metrics at the VM and Hardware Level

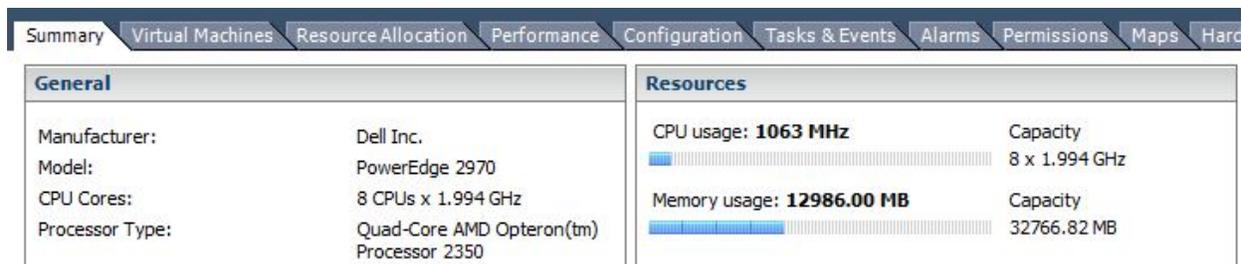
Capacity management takes place on two different levels: the hardware and the virtual machine. For the virtual machine, managing capacity involves making sure that VMs have been allocated all the resources needed without significant overprovisioning. At the hardware (host and cluster) level, managing capacity is a matter of maximizing VM density without sacrificing performance. In both cases, memory management becomes a balancing act between performance and cost efficiency.

To manage memory at both the virtual machine and hardware level, the following key memory metrics must be collected from VMware vCenter:

- Memory Consumed
- Memory Active
- Memory Balloon
- Memory Swapped

Hardware Level Memory Management with Large Memory Pages

For hosts, the primary performance metric is Memory Consumed, which indicates the total amount of virtual machine memory that has physical memory mapped to it. Memory consumed is what is shown in "Memory usage" for a host in either the summary view or the performance charts in vCenter. Memory Consumed and Memory Usage can be used interchangeably.



The Large Memory Pages Setting and its Effects to Memory Metrics

As organizations migrate to vSphere and at the same time install hosts with newer, high-capacity processors, data centers may begin to encounter issues with “large memory pages”. Large memory pages exist to increase performance but can vastly skew the results shown by some memory metrics. The large memory pages setting was introduced in ESX 3.5 but importantly, is enabled by default within vSphere with the newer AMD and Intel processors. Some virtualization administrators may not be aware that this setting is enabled.

Important for hardware-level memory management, large memory pages will skew the results shown by the host Memory Usage metric. With large memory pages enabled, Memory Usage will show utilization at or around 100%.

Memory Consumed Cannot be Relied on with Large Memory Pages

Managing the memory capacity in a virtual environment that doesn’t accurately indicate the host’s memory usage with Memory Consumed can be exceedingly difficult. The VMware Knowledge Base attempts to provide some guidance, but does not offer any recommendations on how exactly to calculate an optimal allocation. The VMware Knowledge Base states:

Capacity planning requires a thorough understanding of the average and the peak memory demands of each virtual machine and resource pool as well as the implications of their reservation, limit, and share settings. The risk and reward of performance and memory overcommitment also must be closely considered.

You should base your capacity planning decisions on the examination of multiple memory performance metrics and not the host memory usage metric alone.

SOURCE: FROM VMWARE.COM

([HTTP://KB.VMWARE.COM/SELFERVICE/MICROSITES/SEARCH.DO?CMD=DISPLAYKC&DOCTYPE=KC&EXTERNALID=1021896](http://kb.vmware.com/selfservice/microsites/search.do?cmd=displaykc&doctype=kc&externalid=1021896))

It is important to understand that Memory Consumed is the metric that is being skewed by large memory pages. Host memory usage is simply a percentage value calculated by dividing Memory Consumed by the total host memory. Before the large memory pages setting existed, Memory Consumed was much more useful at indicating what a VM or host was using for memory resources. The differences in calculation caused by this change may cause confusion with veteran VMware administrators, as many are surprised to find that they cannot fit nearly as many VMs on their hardware as with ESX 3.x.

So, if host Memory Consumed cannot be relied on for planning capacity and alerting when large memory pages are in use, what measure will yield this insight?

Memory Active Becomes the Alternative Metric to Use

The best alternative to the Memory Consumed metric to judge memory usage when large memory pages are enabled is Memory Active. This metric indicates the amount of physical memory that has recently been “touched” - i.e. used. By looking at the peak Memory Active value for the past week or

month, a virtualization administrator can get a good idea of the memory workload on their hosts and clusters. Keeping peak memory active values at or around 70% will ensure VM performance without overallocating memory resources.

Additional Memory Management Strategies

Over-committing memory on a host is a normal thing to do with or without the large memory pages setting active - in fact, it is encouraged. However, memory over-utilization is something to be wary of because it can cause performance problems. As a host's memory usage (i.e. Memory Consumed) starts to reach ~94%, it will begin to take steps to reclaim and conserve memory resources. A full explanation of these memory reclamation techniques can be found in section 3 of VMware's document, [Understanding Memory Resource Management](#).

Some of the memory reclamation techniques mentioned can be considered good or normal, while some of them should be avoided at all costs. For example, Transparent Page Sharing (TPS) is a very useful technology for reclaiming redundant memory pages. Furthermore, small amounts of ballooning, indicated by the Memory Balloon metric, are not harmful and can be a normal way for hosts to shift around memory resources from VMs that don't need them to those that do. However, any amount of host swapping, indicated by the Memory Swapped metric, should be an immediate cause for concern. Section 5 of the aforementioned document also illustrates in further detail how each of these technologies can hurt or help memory performance.

VM Level Memory Management

The micro level of memory management is about making sure that VMs are allocated the correct amount of resources without significantly over-allocating. Over-allocations will also cause unnecessary memory overhead commitments which can affect a data center's ability to raise VM consolidation ratios.

Memory Metrics for the VM Level

Similar to how memory is managed at the hardware level, depending on whether or not an environment has large memory pages enabled, a virtualization administrator would want to use Memory Consumed or Memory Active to monitor the memory usage of VMs. This is how to determine what VMs need for the purpose of sizing their memory allocations.

Avoiding Unnecessary Memory Overhead Commitments

The more memory (and CPU) that is allocated to a VM, the more memory overhead the VM consumes. A table of memory overhead values based on allocations is provided by VMware (shown in the table on the next page).

Table 3-2. Overhead Memory on Virtual Machines

Memory (MB)	1 VCPU	2 VCPUs	3 VCPUs	4 VCPUs	5 VCPUs	6 VCPUs	7 VCPUs	8 VCPUs
256	113.17	159.43	200.53	241.62	293.15	334.27	375.38	416.50
512	116.68	164.96	206.07	247.17	302.75	343.88	385.02	426.15
1024	123.73	176.05	217.18	258.30	322.00	363.17	404.34	445.52
2048	137.81	198.20	239.37	280.53	360.46	401.70	442.94	484.18
4096	165.98	242.51	283.75	324.99	437.37	478.75	520.14	561.52
8192	222.30	331.12	372.52	413.91	591.20	632.86	674.53	716.19
16384	334.96	508.34	550.05	591.76	900.44	942.98	985.52	1028.07
32768	560.27	863.41	906.06	948.71	1515.75	1559.42	1603.09	1646.76
65536	1011.21	1572.29	1616.19	1660.09	2746.38	2792.30	2838.22	2884.14
131072	1912.48	2990.05	3036.46	3082.88	5220.24	5273.18	5326.11	5379.05
262144	3714.99	5830.60	5884.53	5938.46	10142.83	10204.79	10266.74	10328.69

SOURCE: FROM VMWARE.COM

([HTTP://PUBS.VMWARE.COM/VSP40_1/WWHELP/WWHIMPL/COMMON/HTML/WWHELP.HTM#HREF=RESMGMT/R_OVERHEAD_MEMORY_ON_VIRTUAL_MACHINES.HTML#1_7_9_9_10_1&SINGLE=TRU](http://pubs.vmware.com/vsp40_1/wwhelp/wwhimpl/common/html/wwhelp.htm#href=resmgmt/r_overhead_memory_on_virtual_machines.html#1_7_9_9_10_1&single=tru))

The overhead amounts listed above might not seem like much at an initial glance. However, as memory is typically the number one capacity constraint, in an environment where over-allocations run rampant, it may be possible to achieve up to a 10% to 15% increase in VM consolidation ratios by reducing CPU and memory allocations where necessary.

Memory Management as Environments Scale

As an environment grows to hundreds or thousands of VMs, keeping track of performance throughout the virtual environment (not just for memory), can become a difficult task with a vastly increasing number of VMs, nodes, and metrics to assess as often as possible. Installing a dedicated capacity management system to help manage this complexity will make this task easier and more automated. Without such a system in place, to manage memory performance and improve memory usage efficiency, the following steps should be undertaken:

1. Extract the Memory Consumed, Memory Active, Memory Balloon and Memory Swapped metrics for all VMs and all hosts at least on a daily basis
2. Monitor for memory usage peaks to indicate if a VM is getting the resources it needs
3. Using those same memory usage peaks, decrease the allocated size of VMs that have been overprovisioned. As mentioned above, the target value should be around 70% allocation.
4. Monitor for any hypervisor swapping on both the VMs and hosts to ensure performance isn't degraded by overutilization of hardware.

Conclusion

Memory management is an essential aspect of monitoring application performance and working to increase ROI for the data center. This is a task that many virtualization administrators may not be fully versed in and is only made more difficult by the changes to memory metrics that occur when large memory pages are being used. Managing memory becomes more time-consuming and complicated as environments scale. A capacity management system will help to automate the tasks required to effectively manage memory. Without such a system, memory metric observations should be taken and analyzed as often as possible to gain insights which will help to avoid performance problems and allow an environment to use its memory resource as efficiently as possible.