

# Memory optimization ...

---



Tasara 2 posts since

Mar 6, 2007

I have a few questions ESX memory optimization techniques.

1. Will ESX reclaim memory from a VM that has been configured with a full memory reservation i.e. reserved memory equals VM memory ?

In the resource management guide it is stated that it will not reclaim memory even if the VM has idle memory but I have seen posts that states the opposite.

2. If you configure a VM with a memory reservation that equals the VM memory will transparent page sharing automatically be disabled for this VM ?



weinstein5 6,331 posts since

Nov 19, 2005 1. **Re: Memory optimization** May 3, 2008 8:01 PM

1) My understanding, if the VM is not using the memory the VMKernel will assign it to other vms that need it - even if the reservation is set to the vms available memory - do you have the section in the resource management guide where it says that?

2) Transparent Page Sharing is always active for all VMs on an ESX Host even if configured the way you describe - you can only disable it for the entire host - which is something I would not recommend -



Tasara 2 posts since

Mar 6, 2007 2. **Re: Memory optimization** May 4, 2008 6:02 AM

👤 in response to: [weinstein5](#)

@1 : It is stated on page 17 in the Resource Management guide

[http://www.vmware.com/pdf/vi3\\_35/esx\\_3/r35/vi3\\_35\\_25\\_resource\\_mgmt.pdf](http://www.vmware.com/pdf/vi3_35/esx_3/r35/vi3_35_25_resource_mgmt.pdf) (see quote below). Actually it is not tied to the situation where the VM has all its memory reserved, but valid for all memory reservations. As I understand the sentence it means that once a VM has accessed its full memory reservation that memory is allocated to the VM from that point on and will not be reclaimed even if the memory is idle. Consequently, it will require a reboot to free up the memory. This behaviour will have a significant consequence if you allocate a high memory reservation to workloads that are peaky by nature (Test and development VM's for example). This way a substantial amount of memory can be wasted in the cluster.

"After a virtual machine has accessed its full reservation, ESX Server allows the virtual machine to retain this much memory, and will not reclaim it, even if the virtual machine becomes idle and stops accessing memory."

@2 : This is also my understanding but I have just seen some posts in the forum that states that TPS is disabled when a VM has a memory reservation that is equal to configured VM memory.

I can imagine situations where available memory is low where ESX would need to page to satisfy reservations in a situation where the TPS level is high. Otherwise the VM would have to wait until ballooning takes effect to have its reservation honored.

Unfortunately, the advanced memory techniques are not that well described in the product documentation.

\_\_\_\_\_ {size}{font}



[kitcolbert](#) 310 posts since

Feb 15, 2006 **3. Re: Memory optimization** May 7, 2008 9:48 AM


1. In answer to your specific question, no, the VM's memory will not be reclaimed through swapping or ballooning (but could be reclaiming through transparent page sharing). However, if the VM has not used its full memory reservation yet, then the unused region will be allocated to other VMs until that VM decides it wants to use it.

2. No, transparent page sharing is not disabled by default.



[Ken.Cline](#) 5,146 posts since

Jul 7, 2004 **4. Re: Memory optimization** May 22, 2008 10:28 AM

 in response to: [kitcolbert](#) Kit, thanks for taking the time to answer. A couple questions for clarification, if I could



kitcolbert wrote:

1. In answer to your specific question, no, the VM's memory will not be reclaimed through swapping or ballooning (but could be reclaiming through transparent page sharing). However, if the VM has not used its full memory reservation yet, then the unused region will be allocated to other VMs until that VM decides it wants to use it.

**Scenario 1:** I create a VM and allocate 2GB vRAM and set a Reservation of 1GB. I install Windows as the Guest OS. Since this is a Windows VM, on startup, it will touch all allocated memory, which will indicate that the Reservation has been claimed.

- - ESX will not, via swapping/ballooning, reclaim the Reserved vRAM (1GB).
- - vRAM from Reservation to allocation could be reclaimed via swapping/ballooning
- - All pages of vRAM are eligible for collapse via TPS

**Scenario 2:** I create a VM and allocate 2GB vRAM and set a Reservation of 1GB. I install Linux as the Guest OS. Since this is a Linux VM, on startup, it will not gratuitously touch all allocated memory.

- - ESX will not, via swapping/ballooning, reclaim Reserved vRAM (up to 1GB) that has been touched by the Guest OS.
- - vRAM from Reservation to allocation could be reclaimed.
- - All pages of vRAM are eligible for collapse via TPS

Questions:

- - Under what circumstances will the "touched" Reserved vRAM be released by ESX? (Guest OS shutdown/VM power cycle/etc.)
- - Under which counter(s) do these values show up in VC? (mem.granted.average / mem.consumed.average / etc.)

2. No, transparent page sharing is not disabled by default. Questions:

- - Does setting Reservation == Allocation in the .vmx file disable TPS for a VM?
- - Disabling TPS has been suggested for Citrix/TS workloads, does disabling TPS impact performance?
- - Can you explain the purpose of each of the following advanced configuration options?
- --- Mem.ShareScanGHZ
- --- Mem.ShareScanTime

Memory optimization ...

- --- Mem.ShareRateMax
- --- ShareScanTotal
- --- ShareScanVM

Thanks!

Ken Cline

Technical Director, Virtualization

[Wells Landers](#)

VMware Communities User Moderator



[esiebert7625](#) 6,794 posts since

Oct 23, 2006 5. **Re: Memory optimization** May 22, 2008 12:13 PM

👤 in response to: [kitcolbert](#) Also another question, from your VMworld presentation on guest & host memory it sounded like that a host that was not memory over-committed would never reclaim memory from a VM. So if a host is not memory over-committed then features like swapping, ballooning and TPS would not be necessary, why waste the time and resources doing them if they are not needed. Does TPS kick in at all on a host that is not memory over-committed?

Eric Siebert

VMware Communities User Moderator

-----

Check out my website: [VMware-land](#)

Read my virtualization blog: [SSV Blog](#)

-----



[kitcolbert](#) 310 posts since

Feb 15, 2006 6. **Re: Memory optimization** May 22, 2008 12:33 PM

👤 in response to: [Ken.Cline](#) Hi Ken,

I have answers for some of your questions, but for some of the VC stuff I'll have to find other people to answer them.

**Scenario 1:** I create a VM and allocate 2GB vRAM and set a Reservation of 1GB. I install Windows as the Guest OS. Since this is a Windows VM, on startup, it will touch all allocated memory, which will indicate that the Reservation has been claimed.

- - ESX will not, via swapping/ballooning, reclaim the Reserved vRAM (1GB).
- - vRAM from Reservation to allocation could be reclaimed via swapping/ballooning
- - All pages of vRAM are eligible for collapse via TPS

Yep. So initially the VM may be using very close to its full memory, but over time TPS will bring that down (most likely even below the reservation).

**Scenario 2:** I create a VM and allocate 2GB vRAM and set a Reservation of 1GB. I install Linux as the Guest OS. Since this is a Linux VM, on startup, it will not gratuitously touch all allocated memory.

- - ESX will not, via swapping/ballooning, reclaim Reserved vRAM (up to 1GB) that has been touched by the Guest OS.

Memory optimization ...

- - vRAM from Reservation to allocation could be reclaimed.
- - All pages of vRAM are eligible for collapse via TPS

Yep.

Questions:

- - Under what circumstances will the "touched" Reserved vRAM be released by ESX? (Guest OS shutdown/VM power cycle/etc.)

- Transparent page sharing

- "Hard" VM reset. In a hard reset, the underlying hypervisor code does a reset of the VM, thus releasing the VM's memory back to the free memory pool. If it's just a "soft" reset, initiated internally by the VM, then the hypervisor code doesn't even realize anything's happening, so the VM keeps all the memory it's allocated.

\* - Under which counter(s) do these values show up in VC? (mem.granted.average / mem.consumed.average / etc.)

Unfortunately my knowledge is more ESX-specific, so I'm not sure about VC. But what values are you looking for in particular?

2. No, transparent page sharing is not disabled by default. Questions:

- - Does setting Reservation == Allocation in the .vmx file disable TPS for a VM?

Nope. TPS will run unless you explicitly disable it. (There's some VM config file option for it, but I can't think of it off the top of my head.)

-- Disabling TPS has been suggested for Citrix/TS workloads, does disabling TPS impact performance?

I've heard of this rumor as well, but I've yet to see any hard data that this actually helps. It's certainly possible that it does, but in general I'd recommend leaving TPS enabled unless you have specific data showing that disabling it will help improve performance. I'll ask some other engineers in the resource mgmt group and get their feedback as well.

-- Can you explain the purpose of each of the following advanced configuration options?

- --- Mem.ShareScanGHZ

Basically how aggressive do you want TPS to scan for shared pages? Obviously TPS has a cost to running in terms of CPU, but has benefits in terms of reduced memory usage. So there's a tradeoff. We have a default that we think is a good balance, but we let the user modify that if they want.

--- Mem.ShareScanTime

- --- Mem.ShareRateMax

I think both of these are just adjusting how aggressive you want to scan a particular VM. Scanning too quickly could share pages that are in the write working set, meaning that the sharing will be broken very quickly. Thus scanning too quickly may not be the best. But scanning too slowly means that memory that could be shared won't be for longer, meaning a higher memory usage. Another tradeoff, we set some defaults.

--- ShareScanTotal

- --- ShareScanVM

Hmm, I don't see these... Where are you seeing them?

Thanks,

Memory optimization ...

Kit



[kitcolbert](#) 310 posts since

Feb 15, 2006 7. **Re: Memory optimization** May 22, 2008 12:38 PM

👤 in response to: [esiebert7625](#) Hi Eric,

Yes, that slide was slightly incorrect. I meant to say that TPS is always running, even if not overcommitted.

One reason that we continuously do TPS is that it makes the system much more amenable to new VMs starting. If we didn't do TPS, then when you started a new VM, the ESX host would have to swap/balloon like crazy to make room for the new VM (in addition to start TPS). Obviously everything would slow down for awhile in this scenario. Another idea might be to delay the start of the new VM while we do TPS and balloon/swap whatever else we can. This also isn't great since you would have to wait awhile for your new VM to start.

Thus the best solution is just to keep TPS running. It has a small overhead, so it shouldn't hurt to keep it enabled. And it allows you to easier add/remove VMs from the system without as bad of an impact as they could otherwise have.

Thanks,

Kit



[oreeh](#) 9,872 posts since

Nov 30, 2005 8. **Re: Memory optimization** May 22, 2008 12:42 PM

👤 in response to: [kitcolbert](#) There's some VM config file option for it, but I can't think of it off the top of my head.

Could it be sched.mem.pshare.enable ?



[esiebert7625](#) 6,794 posts since

Oct 23, 2006 9. **Re: Memory optimization** May 22, 2008 12:43 PM

👤 in response to: [kitcolbert](#)

Thanks, I suspected that might be the case, especially in a DRS/HA environment where one host may end up with more VM's then intended.

Great presentation by the way, I thought it was one of the best technical ones of VMworld and I'm sure it helped everyone understand the often confusing area of how memory works in ESX.

Eric Siebert

VMware Communities User Moderator

-----

Check out my website: [VMware-land](#)

Read my virtualization blog: [SSV Blog](#)


-----



[kitcolbert](#) 310 posts since

Feb 15, 2006 10. **Re: Memory optimization** May 22, 2008 2:04 PM

Memory optimization ...


 in response to: [oreeh](#) Could it be sched.mem.pshare.enable ?

Yeah, that sounds like it.



[kitcolbert](#) 310 posts since

Feb 15, 2006 11. **Re: Memory optimization** May 22, 2008 2:17 PM

 in response to: [esiebert7625](#) Great presentation by the way, I thought it was one of the best technical ones of VMworld and I'm sure it helped everyone understand the often confusing area of how memory works in ESX.

Thanks! Definitely let me know if you think anything could be improved/clarified.


Thanks,

Kit



[Ken.Cline](#) 5,146 posts since


Jul 7, 2004 12. **Re: Memory optimization** May 22, 2008 10:08 PM

 in response to: [kitcolbert](#) Kit - thanks for taking the time to answer these questions. You're definitely helping clear up some muddy water!

Unfortunately my knowledge is more ESX-specific, so I'm not sure about VC. But what values are you looking for in particular?

No worries. I'm not really looking for anything other than an understanding of how the various memory states are reported. Seems that every time I get on a customer site, they ask "What, exactly, does *that* value represent?" Just want to have a better answer than "I'm not 100% sure".

Hmm, I don't see these... Where are you seeing them?

Legacy stuff, I guess. Dates back to some notes I had from ESX 2 days 

Thanks again - your help is greatly appreciated!

Ken Cline

Technical Director, Virtualization

[Wells Landers](#)

VMware Communities User Moderator