

## Vmware vSphere 4.1 Memory Compression Understanding

04/08/2010 – 9:41 am

Posted in [Vmware Technology](#), [vmware vSphere 4.1](#)


Tagged [Vmware](#), [vSphere](#), [Understanding](#), [4.1](#), [Memory compression](#)

There is a new feature both Microsoft Hyper-v and Vmware present in their latest version of product, that's Memory compression. How does Memory compression work for Hyper-v is still an unknown but for Vmware vSphere 4.1, it's pretty easy to understand. This post is dedicated for Memory Compression and I hope I can give you a easy explanation.

There are always challenge for Vmware to reclaim free host memory from VMs because Microsoft doesn't share internal index file to tell Vmware which memory pages are free. All Memory used by VMs are always aggregates and eventually it will bust up if you overcommitted memory on your host. So Vmware develop following technologies to work around this issue.


- Transparent Page Sharing
- Ballooning
- Swapping
- Memory Compression

### Transparent Page Sharing




One of biggest advantage comparing is Transparent Page Sharing. Basically, vmKernel scan most memory pages of VMs sitting on the same host. It divide memory page as small piece (4kB) and generate a hash signature for each 4kB block. If there are more than one identical blocks on the host physical memory, Vmware keeps one block and free rest of blocks to host. The exception is if you use hardware-assisted memory virtualization with large index table, Vmware doesn't do page sharing till host memory is overcommitted. The block size and scan schedule are adjustable.

### Ballooning



Ballooning is used via VMware Driver installed in the OS. This driver is still in the Kernel level and will try to claim memory when your host memory is short. It pins memory pages of VM and tell Host to use those pinned memory as free memory. Unlike Sharing pages, it's not proactive action.

### Swapping



Swapping is a bad bad choice. If your VM start to swap memory to physical disk, your VM performance will have dramatic impact. This is very last result of VMware try to make VM still alive.


### Memory Compression

Here it is. New technology of vSphere 4.1. It's not something you will use everyday, it's a technology which gives your VM another last breath before it sinks with swapping. Let's see how it works.

#### What is Memory Compression?


The idea of memory compression is very straightforward: if the swapped out pages can be compressed and stored in a compression cache located in the main memory, the next access to the page only causes a page decompression which can be an order of magnitude faster than the disk access. With memory compression, only a few uncompressible pages need to be swapped out if the compression cache is not full.

Basically, when Swapping is about to happen, VMcompress part of memory, which will be swapped out to physical disk, to its self memory space (not host memory). It divide VM swapping memory to 4KB and tried to compress them to 2KB. If compress success, then, you can save 50% space, if it fails, VM still swap the original 4KB to physical disk.



Therefore, it may or may not work 100% for all swapping memory pages. Even it can compress everything, please be aware it's using its own VM memory space not host memory, which means it can't use all VM memory. In default, that's 10% of VM memory can be acted as Memory compression Cache. No other options will happen to those compressed memory. Anything happens, VM needs to uncompress them first and do other action next.

Let's see a real case.



If you notice those two lines, the gap between 28GB and 24GB memory, the memory compression kept throughput with very 6% performance difference. If we look the two bars with swapping read, you can clearly see even with the worst case, memory compression still manage to take more than half of load of Swapping and make VM live longer and operate-able.

#### Summary:

Memory compression give your VMs a last chance to struggle on operation level. It's better than nothing, isn't it?

#### One Trackback/Pingback

By Tweets that mention VMware vSphere 4.1 Memory Compression Understanding « GeekSilver's Blog – Topsy.com on 04-Aug-2010 at 11:15 am

[...] This post was mentioned on Twitter by VM Digest, Silver Chen. Silver Chen said: VMware vSphere 4.1 Memory Compression Understanding: <http://wp.me/pVbEv-6B> [...]

#### Post a Comment

Name \*

Email \*

Website

Comment

Post Comment »

- ☐ Notify me of follow-up comments via email.  
☐ Notify me of site updates

#### GeekSilver Email Subscription

Enter your email address to subscribe to this blog and receive notifications of new posts by email.

Sign me up!

#### GeekSilver RSS Link



RSS - Posts

#### GeekSilver Twitter

- VMware vSphere 4.1 vs ESX 3.x storage identifier Understanding: <http://wp.me/pVbEv-6R> 16 hours ago
- VMware vSphere 4.1 Memory Compression Understanding: <http://wp.me/pVbEv-6B> 5 days ago
- VMware vSphere 4.1 vStorage APIs for Array Integration (VAI) understanding: <http://wp.me/pVbEv-61> 1 week ago
- Difference between vSphere Hypervisor 4.1 and ESXi 4.1: <http://wp.me/pVbEv-6t> 1 week ago
- VMware vSphere 4.1 Storage IO Control (SIOC) understanding: <http://wp.me/pVbEv-65> 1 week ago

#### Search

Find »

#### Tags

4.1 101 Architecture asco Configuration Control  
Courses difference EMC ESX esxi fast hot add How to hp cim  
agent hypervisor Inform2010 Install IO logs Memory compression  
Microsoft Netioc Network Control Network Distributed Switch  
pvscsi qos SAN SGU Stargate Storage Tools Training  
Troubleshooting Understanding vaai vcp vds VIEW  
vMA VMware vmworld 2010 vSphere  
Windows 2008 sp2 winsxs

#### Search

Find »

#### Blogroll

- Development Blog
- Documentation
- Plugins
- Suggest Ideas
- Support Forum
- Themes
- WordPress Planet

#### RSS Feeds

- All posts
- All comments

#### Meta

- Register
- Log in