

# Understanding Memory Resource Management in VMware ESX 4.1

VMware ESX 4.1

---

VMware® ESX™ is a hypervisor designed to efficiently manage hardware resources including CPU, memory, storage, and network among multiple, concurrent virtual machines. This paper describes the basic memory management concepts in ESX, the configuration options available, and provides results to show the performance impact of these options. The focus of this paper is in presenting the fundamental concepts of these options.<sup>1</sup>

This paper is organized as follows:

- Section 1. “Introduction” on page 1
- Section 2. “ESX Memory Management Overview” on page 2
- Section 3. “Memory Reclamation in ESX” on page 5
- Section 4. “ESX Memory Allocation Management for Multiple Virtual Machines” on page 13
- Section 5. “Performance Evaluation” on page 15
- Section 6. “Best Practices” on page 24
- Section 7. “References” on page 25

Memory compression is a new feature for VMware ESX 4.1 and is covered in Section 3.5, “Memory Compression” on page 10 and Section 5.4, “Memory Compression Performance” on page 22.

## 1. Introduction

ESX uses high-level resource management policies to compute a target memory allocation for each virtual machine (VM) based on the current system load and parameter settings for the virtual machine (shares, reservation, and limit [2]). The computed target allocation is used to guide the dynamic adjustment of the memory allocation for each virtual machine. In the cases where host memory is overcommitted, the target allocations are still achieved by invoking several lower-level mechanisms to reclaim memory from virtual machines.

This paper assumes a pure virtualization environment in which the guest operating system running inside the virtual machine is not modified to facilitate virtualization (often referred to as paravirtualization). Knowledge of ESX architecture will help you understand the concepts presented in this paper.


---

<sup>1</sup> More details can be found in *Memory Resource Management in VMware ESX Server* [1].

The VMware vSphere™ Client exposes several memory statistics in the performance charts. Among them are charts for the following memory types: consumed, active, shared, granted, overhead, balloon, swapped, and compressed.. A complete discussion about these metrics can be found in “Memory Performance Chart Metrics in the vSphere Client” [3] and “VirtualCenter Memory Statistics Definitions” [4].

Two important memory statistics are Consumed Memory and Active Memory. You can use the charts for these statistics to quickly monitor the host memory and virtual machine memory usage.

**Figure 1.** Host and Active Memory usage in vSphere Client Performance Charts



Consumed Memory usage is defined as the amount of host memory that is allocated to the virtual machine, Active Memory is defined as the amount of guest memory that is currently being used by the guest operating system and its applications. These two statistics are quite useful for analyzing the memory status of the virtual machine and providing hints to address potential performance issues.

This paper helps answer these questions:

- Why is the Consumed Memory so high?
- Why is the Consumed Memory usage sometimes much larger than the Active Memory?
- Why is the Active Memory different from what is seen inside the guest operating system?

These questions cannot be easily answered without understanding the basic memory management concepts in ESX. Understanding how ESX manages memory will also make the performance implications of changing ESX memory management parameters clearer.

## 2. ESX Memory Management Overview

### 2.1 Terminology

The following terminology is used throughout this paper.

- Host physical memory<sup>2</sup> refers to the memory that is visible to the hypervisor as available on the system.

<sup>2</sup> The terms host physical memory and host memory are used interchangeably in this paper. They are also equivalent to the term machine memory used in “Memory Resource Management in VMware ESX Server” [1].

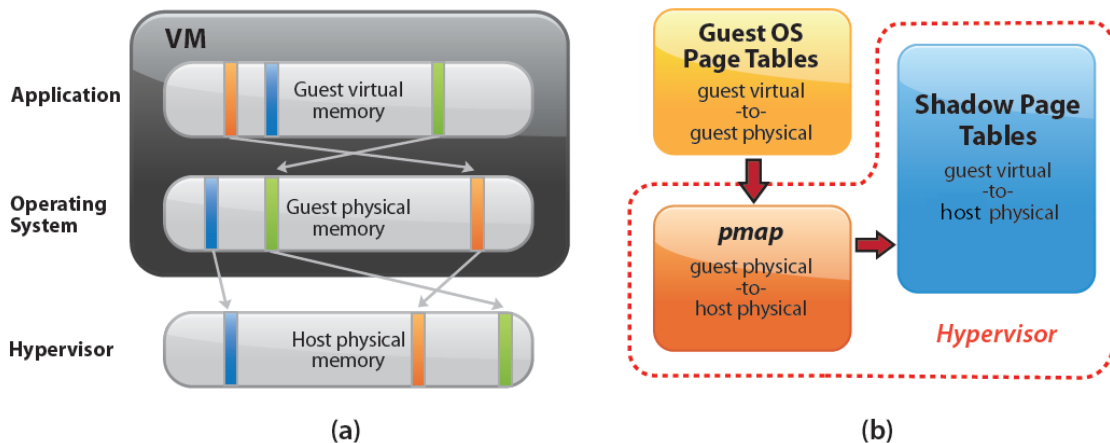
- Guest physical memory refers to the memory that is visible to the guest operating system running in the virtual machine.
- Guest virtual memory refers to a continuous virtual address space presented by the guest operating system to applications. It is the memory that is visible to the applications running inside the virtual machine.
- Guest physical memory is backed by host physical memory, which means the hypervisor provides a mapping from the guest to the host memory.
- The memory transfer between the guest physical memory and the guest swap device is referred to as guest level paging and is driven by the guest operating system. The memory transfer between guest physical memory and the host swap device is referred to as hypervisor swapping, which is driven by the hypervisor.

## 2.2 Memory Virtualization Basics

Virtual memory is a well-known technique used in most general-purpose operating systems, and almost all modern processors have hardware to support it. Virtual memory creates a uniform virtual address space for applications and allows the operating system and hardware to handle the address translation between the virtual address space and the physical address space. This technique not only simplifies the programmer's work, but also adapts the execution environment to support large address spaces, process protection, file mapping, and swapping in modern computer systems.

When running a virtual machine, the hypervisor creates a contiguous addressable memory space for the virtual machine. This memory space has the same properties as the virtual address space presented to the applications by the guest operating system. This allows the hypervisor to run multiple virtual machines simultaneously while protecting the memory of each virtual machine from being accessed by others. Therefore, from the view of the application running inside the virtual machine, the hypervisor adds an extra level of address translation that maps the guest physical address to the host physical address. As a result, there are three virtual memory layers in ESX: guest virtual memory, guest physical memory, and host physical memory. Their relationships are illustrated in Figure 2 (a).

**Figure 2.** Virtual memory levels (a) and memory address translation (b) in ESX



As shown in Figure 2 (b), in ESX, the address translation between guest physical memory and host physical memory is maintained by the hypervisor using a physical memory mapping data structure, or *pmap*, for each virtual machine. The hypervisor intercepts all virtual machine instructions that manipulate the hardware translation lookaside buffer (TLB) contents or guest operating system page tables, which contain the virtual to physical address mapping. The actual hardware TLB state is updated based on the separate shadow page tables, which contain the guest virtual to host physical address mapping. The shadow page

tables maintain consistency with the guest virtual to guest physical address mapping in the guest page tables and the guest physical to host physical address mapping in the pmap data structure. This approach removes the virtualization overhead for the virtual machine's normal memory accesses because the hardware TLB will cache the direct guest virtual to host physical memory address translations read from the shadow page tables. Note that the extra level of guest physical to host physical memory indirection is extremely powerful in the virtualization environment. For example, ESX can easily remap a virtual machine's host physical memory to files or other devices in a manner that is completely transparent to the virtual machine.

Recently, some new generation CPUs, such as third generation AMD Opteron and Intel Xeon 5500 series processors, have provided hardware support for memory virtualization by using two layers of page tables in hardware. One layer stores the guest virtual to guest physical memory address translation, and the other layer stores the guest physical to host physical memory address translation. These two page tables are synchronized using processor hardware. Support for hardware memory virtualization eliminates the overhead required to keep shadow page tables in synchronization with guest page tables in software memory virtualization. For more information about hardware-assisted memory virtualization, see "Performance Evaluation of Intel EPT Hardware Assist" [5] and "Performance Evaluation of AMD RVI Hardware Assist." [6]

## 2.3 Memory Management Basics in ESX

Prior to talking about how ESX manages memory for virtual machines, it is useful to first understand how the application, guest operating system, hypervisor, and virtual machine manage memory at their respective layers.

- An application starts and uses the interfaces provided by the operating system to explicitly allocate or deallocate virtual memory during its execution.
- In a non-virtual environment, the operating system assumes it owns all physical memory in the system. The hardware does not provide interfaces for the operating system to explicitly "allocate" or "free" physical memory. The operating system establishes the definitions of "allocated" or "free" physical memory. Different operating systems have different implementations to realize this abstraction. One example is that the operating system maintains an "allocated" list and a "free" list, so whether or not a physical page is free depends on which list the page currently resides in.
- Because a virtual machine runs an operating system and several applications, the virtual machine memory management properties combine both application and operating system memory management properties. Like an application, when a virtual machine first starts, it has no pre-allocated physical memory. Like an operating system, the virtual machine cannot explicitly allocate host physical memory through any standard interface. The hypervisor also creates the definitions of "allocated" and "free" host memory in its own data structures. The hypervisor intercepts the virtual machine's memory accesses and allocates host physical memory for the virtual machine on its first access to the memory. In order to avoid information leaking among virtual machines, the hypervisor always writes zeroes to the host physical memory before assigning it to a virtual machine.
- Virtual machine memory deallocation acts just like an operating system, such that the guest operating system frees a piece of physical memory by adding these memory page numbers to the guest free list, but the data of the "freed" memory may not be modified at all. As a result, when a particular piece of guest physical memory is freed, the mapped host physical memory will usually not change its state and only the guest free list will be changed.

The hypervisor knows when to allocate host physical memory for a virtual machine because the first memory access from the virtual machine to a host physical memory will cause a page fault that can be easily captured by the hypervisor. However, it is difficult for the hypervisor to know when to free host physical memory upon virtual machine memory deallocation because the guest operating system free list is generally

not publicly accessible. Hence, the hypervisor cannot easily find out the location of the free list and monitor its changes.

Although the hypervisor cannot reclaim host memory when the operating system frees guest physical memory, this does not mean that the host memory, no matter how large it is, will be used up by a virtual machine when the virtual machine repeatedly allocates and frees memory. This is because the hypervisor does not allocate host physical memory on every virtual machine's memory allocation. It only allocates host physical memory when the virtual machine touches the physical memory that it has never touched before. If a virtual machine frequently allocates and frees memory, presumably the same guest physical memory is being allocated and freed again and again. Therefore, the hypervisor just allocates host physical memory for the first memory allocation and then the guest reuses the same host physical memory for the rest of the allocations. That is, if a virtual machine's entire guest physical memory (configured memory) has been backed by the host physical memory, the hypervisor does not need to allocate any host physical memory for this virtual machine any more. This means that Equation 1 always holds true:

**Equation 1.**

VM's host memory usage  $\leq$  VM's guest memory size + VM's overhead memory

Here, the virtual machine's overhead memory is the extra host memory needed by the hypervisor for various virtualization data structures besides the memory allocated to the virtual machine. Its size depends on the number of virtual CPUs and the configured virtual machine memory size. For more information, see the *vSphere Resource Management Guide* [2].

### 3. Memory Reclamation in ESX

ESX uses several innovative techniques to reclaim virtual machine memory, which are:

- Transparent page sharing (TPS)—reclaims memory by removing redundant pages with identical content
- Ballooning—reclaims memory by artificially increasing the memory pressure inside the guest
- Hypervisor swapping—reclaims memory by having ESX directly swap out the virtual machine's memory
- Memory compression—reclaims memory by compressing the pages that need to be swapped out

The following sections describe these techniques and the motivation behind memory reclamation.


#### 3.1 Motivation

According to Equation 1 above, if the hypervisor cannot reclaim host physical memory upon virtual machine memory deallocation, it must reserve enough host physical memory to back all virtual machine's guest physical memory (plus their overhead memory) in order to prevent any virtual machine from running out of host physical memory. This means that memory overcommitment cannot be supported. The concept of memory overcommitment is fairly simple: host memory is overcommitted when the total amount of guest physical memory of the running virtual machines is larger than the amount of actual host memory. ESX supports memory overcommitment from the very first version, due to two important benefits it provides:

- Higher memory utilization: With memory overcommitment, ESX ensures that host memory is consumed by active guest memory as much as possible. Typically, some virtual machines may be lightly loaded compared to others. Their memory may be used infrequently, so for much of the time their memory will sit idle. Memory overcommitment allows the hypervisor to use memory reclamation techniques to take the inactive or unused host physical memory away from the idle virtual machines and give it to other virtual machines that will actively use it.

- Higher consolidation ratio: With memory overcommitment, each virtual machine has a smaller footprint in host memory usage, making it possible to fit more virtual machines on the host while still achieving good performance for all virtual machines. For example, as shown in Figure 3, you can enable a host with 4G host physical memory to run three virtual machines with 2G guest physical memory each. Without memory overcommitment, only one virtual machine can be run because the hypervisor cannot reserve host memory for more than one virtual machine, considering that each virtual machine has overhead memory.

**Figure 3.** Memory overcommitment in ESX




In order to effectively support memory overcommitment, the hypervisor must provide efficient host memory reclamation techniques. ESX leverages several innovative techniques to support virtual machine memory reclamation. These techniques are transparent page sharing, ballooning, and host swapping. In ESX 4.1, before host swapping, ESX applies a new technique called memory compression in order to reduce the amount of pages that need to be swapped out, while reclaiming the same amount of host memory.

### 3.2 Transparent Page Sharing (TPS)

When multiple virtual machines are running, some of them may have identical sets of memory content. This presents opportunities for sharing memory across virtual machines (as well as sharing within a single virtual machine). For example, several virtual machines may be running the same guest operating system, have the same applications, or contain the same user data. With page sharing, the hypervisor can reclaim the redundant copies and keep only one copy, which is shared by multiple virtual machines in the host physical memory. As a result, the total virtual machine host memory consumption is reduced and a higher level of memory overcommitment is possible.

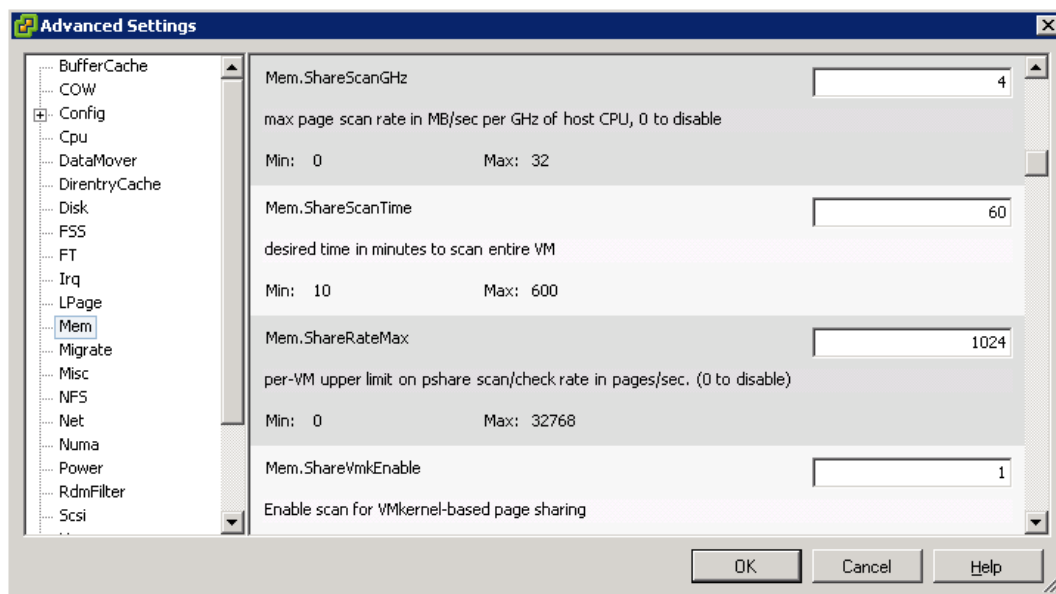
In ESX, the redundant page copies are identified by their contents. This means that pages with identical content can be shared regardless of when, where, and how those contents are generated. ESX scans the content of guest physical memory for sharing opportunities. Instead of comparing each byte of a candidate guest physical page to other pages, an action that is prohibitively expensive, ESX uses hashing to identify potentially identical pages. The detailed algorithm is illustrated in Figure 4.

**Figure 4.** Content-based page sharing in ESX

A hash value is generated based on the candidate guest physical page's content. The hash value is then used as a key to look up a global hash table, in which each entry records a hash value and the physical page number of a shared page. If the hash value of the candidate guest physical page matches an existing entry, a full comparison of the page contents is performed to exclude a false match. Once the candidate guest physical page's content is confirmed to match the content of an existing shared host physical page, the guest physical to host physical mapping of the candidate guest physical page is changed to the shared host physical page, and the redundant host memory copy (the page pointed to by the dashed arrow in Figure 4) is reclaimed. This remapping is invisible to the virtual machine and inaccessible to the guest operating system. Because of this invisibility, sensitive information cannot be leaked from one virtual machine to another.

A standard copy-on-write (CoW) technique is used to handle writes to the shared host physical pages. Any attempt to write to the shared pages will generate a minor page fault. In the page fault handler, the hypervisor will transparently create a private copy of the page for the virtual machine and remap the affected guest physical page to this private copy. In this way, virtual machines can safely modify the shared pages without disrupting other virtual machines sharing that memory. Note that writing to a shared page does incur overhead compared to writing to non-shared pages due to the extra work performed in the page fault handler.

In VMware ESX, the hypervisor scans the guest physical pages randomly with a base scan rate specified by `Mem.ShareScanTime`, which specifies the desired time to scan the virtual machine's entire guest memory. The maximum number of scanned pages per second in the host and the maximum number of per-virtual machine scanned pages, (that is, `Mem.ShareScanGHz` and `Mem.ShareRateMax` respectively) can also be specified in ESX advanced settings. An example is shown in Figure 5.

**Figure 5.** Configure page sharing in vSphere Client

The default values of these three parameters are carefully chosen to provide sufficient sharing opportunities while keeping the CPU overhead negligible. In fact, ESX intelligently adjusts the page scan rate based on the amount of current shared pages. If the virtual machine's page sharing opportunity seems to be low, the page scan rate will be reduced accordingly and vice versa. This optimization further mitigates the overhead of page sharing.

In hardware-assisted memory virtualization (for example, Intel EPT Hardware Assist and AMD RVI Hardware Assist [6]) systems, ESX will automatically back guest physical pages with large host physical pages (2MB contiguous memory region instead of 4KB for regular pages) for better performance due to less TLB misses. In such systems, ESX will not share those large pages because: 1) the probability of finding two large pages having identical contents is low, and 2) the overhead of doing a bit-by-bit comparison for a 2MB page is much larger than for a 4KB page. However, ESX still generates hashes for the 4KB pages within each large page. Since ESX will not swap out large pages, during host swapping, the large page will be broken into small pages so that these pre-generated hashes can be used to share the small pages before they are swapped out. In short, we may not observe any page sharing for hardware-assisted memory virtualization systems until host memory is overcommitted.

### 3.3 Ballooning

Ballooning is a completely different memory reclamation technique compared to transparent page sharing. Before describing the technique, it is helpful to review why the hypervisor needs to reclaim memory from virtual machines. Due to the virtual machine's isolation, the guest operating system is not aware that it is running inside a virtual machine and is not aware of the states of other virtual machines on the same host. When the hypervisor runs multiple virtual machines and the total amount of the free host memory becomes low, none of the virtual machines will free guest physical memory because the guest operating system cannot detect the host's memory shortage. Ballooning makes the guest operating system aware of the low memory status of the host.


In ESX, a balloon driver is loaded into the guest operating system as a pseudo-device driver.<sup>3</sup> It has no external interfaces to the guest operating system and communicates with the hypervisor through a private

<sup>3</sup> VMware Tools must be installed in order to enable ballooning. This is recommended for all workloads.

channel. The balloon driver polls the hypervisor to obtain a target balloon size. If the hypervisor needs to reclaim virtual machine memory, it sets a proper target balloon size for the balloon driver, making it “inflate” by allocating guest physical pages within the virtual machine. Figure 6 illustrates the process of the balloon inflating.

In Figure 6 (a), four guest physical pages are mapped in the host physical memory. Two of the pages are used by the guest application and the other two pages (marked by stars) are in the guest operating system free list. Note that since the hypervisor cannot identify the two pages in the guest free list, it cannot reclaim the host physical pages that are backing them. Assuming the hypervisor needs to reclaim two pages from the virtual machine, it will set the target balloon size to two pages. After obtaining the target balloon size, the balloon driver allocates two guest physical pages inside the virtual machine and pins them, as shown in Figure 6 b. Here, “pinning” is achieved through the guest operating system interface, which ensures that the pinned pages cannot be paged out to disk under any circumstances. Once the memory is allocated, the balloon driver notifies the hypervisor about the page numbers of the pinned guest physical memory so that the hypervisor can reclaim the host physical pages that are backing them. In Figure 6 (b), dashed arrows point at these pages. The hypervisor can safely reclaim this host physical memory because neither the balloon driver nor the guest operating system relies on the contents of these pages. This means that no processes in the virtual machine will intentionally access those pages to read/write any values. Thus, the hypervisor does not need to allocate host physical memory to store the page contents. If any of these pages are re-accessed by the virtual machine for some reason, the hypervisor will treat it as a normal virtual machine memory allocation and allocate a new host physical page for the virtual machine. When the hypervisor decides to deflate the balloon—by setting a smaller target balloon size—the balloon driver deallocates the pinned guest physical memory, which releases it for the guest’s applications.

**Figure 6.** Inflating the balloon in a virtual machine



Typically, the hypervisor inflates the virtual machine balloon when it is under memory pressure. By inflating the balloon, a virtual machine consumes less physical memory on the host, but more physical memory inside the guest. As a result, the hypervisor offloads some of its memory overload to the guest operating system while slightly loading the virtual machine. That is, the hypervisor transfers the memory pressure from the host to the virtual machine. Ballooning induces guest memory pressure. In response, the balloon driver allocates and pins guest physical memory. The guest operating system determines if it needs to page out guest physical memory to satisfy the balloon driver’s allocation requests. If the virtual machine has plenty of free guest physical memory, inflating the balloon will induce no paging and will not impact guest performance. In this case, as illustrated in Figure 6, the balloon driver allocates the free guest physical memory from the guest free list. Hence, guest-level paging is not necessary. However, if the guest is already under memory pressure, the guest operating system decides which guest physical pages to be paged out to the virtual swap device in order to satisfy the balloon driver’s allocation requests. The genius of ballooning is that it allows the guest operating system to intelligently make the hard decision about which pages to be paged out without the hypervisor’s involvement.

For ballooning to work as intended, the guest operating system must install and enable the balloon driver, which is included in VMware Tools. The guest operating system must have sufficient virtual swap space configured for guest paging to be possible. Ballooning might not reclaim memory quickly enough to satisfy host memory demands. In addition, the upper bound of the target balloon size may be imposed by various guest operating system limitations.

### 3.4 Hypervisor Swapping

In the cases where ballooning and transparent page sharing are not sufficient to reclaim memory, ESX employs hypervisor swapping to reclaim memory. At virtual machine startup, the hypervisor creates a separate swap file for the virtual machine. Then, if necessary, the hypervisor can directly swap out guest physical memory to the swap file, which frees host physical memory for other virtual machines.

Besides the limitation on the reclaimed memory size, both page sharing and ballooning take time to reclaim memory. The page-sharing speed depends on the page scan rate and the sharing opportunity. Ballooning speed relies on the guest operating system's response time for memory allocation.

In contrast, hypervisor swapping is a guaranteed technique to reclaim a specific amount of memory within a specific amount of time. However, hypervisor swapping is used as a last resort to reclaim memory from the virtual machine due to the following limitations on performance:

- **Page selection problems:** Under certain circumstances, hypervisor swapping may severely penalize guest performance. This occurs when the hypervisor has no knowledge about which guest physical pages should be swapped out, and the swapping may cause unintended interactions with the native memory management policies in the guest operating system.
- **Double paging problems:** Another known issue is the double paging problem. Assuming the hypervisor swaps out a guest physical page, it is possible that the guest operating system pages out the same physical page, if the guest is also under memory pressure. This causes the page to be swapped in from the hypervisor swap device and immediately to be paged out to the virtual machine's virtual swap device.

Page selection and double-paging problems exist because the information needed to avoid them is not available to the hypervisor.

- **High swap-in latency:** Swapping in pages is expensive for a VM. If the hypervisor swaps out a guest page and the guest subsequently accesses that page, the VM will get blocked until the page is swapped in from disk. High swap-in latency, which can be tens of milliseconds, can severely degrade guest performance.

ESX mitigates the impact of interacting with guest operating system memory management by randomly selecting the swapped guest physical pages.

### 3.5 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. This means the number of future synchronous swap-in operations will be reduced. Hence, it may improve application performance significantly when the host is in heavy memory pressure. In ESX 4.1, only the swap candidate pages will be compressed. This means ESX will **not** proactively compress guest pages when host swapping is not necessary. In other words, memory compression does not affect workload performance when host memory is undercommitted.


#### 3.5.1 Reclaiming Memory Through Compression

Figure 8 illustrates how memory compression reclaims host memory compared to host swapping.

Assuming ESX needs to reclaim two 4KB physical pages from a VM through host swapping, page A and B

are the selected pages (Figure 8a). With host swapping only, these two pages will be directly swapped to disk and two physical pages are reclaimed (Figure 8b). However, with memory compression, each swap candidate page will be compressed and stored using 2KB of space in a per-VM compression cache. Note that page compression would be much faster than the normal page swap out operation which involves a disk I/O. Page compression will fail if the compression ratio is less than 50% and the uncompressible pages will be swapped out. As a result, every successful page compression is accounted for reclaiming 2KB of physical memory. As illustrated in Figure 8c, pages A and B are compressed and stored as half-pages in the compression cache. Although both pages are removed from VM guest memory, the actual reclaimed memory size is one page.

**Figure 8.** Host swapping vs. memory compression in ESX



If any of the subsequent memory access misses in the VM guest memory, the compression cache will be checked first using the host physical page number. If the page is found in the compression cache, it will be decompressed and push back to the guest memory. This page is then removed from the compression cache. Otherwise, the memory request is sent to the host swap device and the VM is blocked.

### 3.5.2 Managing Per-VM Compression Cache

The per-VM compression cache is accounted for by the VM's guest memory usage, which means ESX will not allocate additional host physical memory to store the compressed pages. The compression cache is transparent to the guest OS. Its size starts with zero when host memory is undercommitted and grows when virtual machine memory starts to be swapped out.

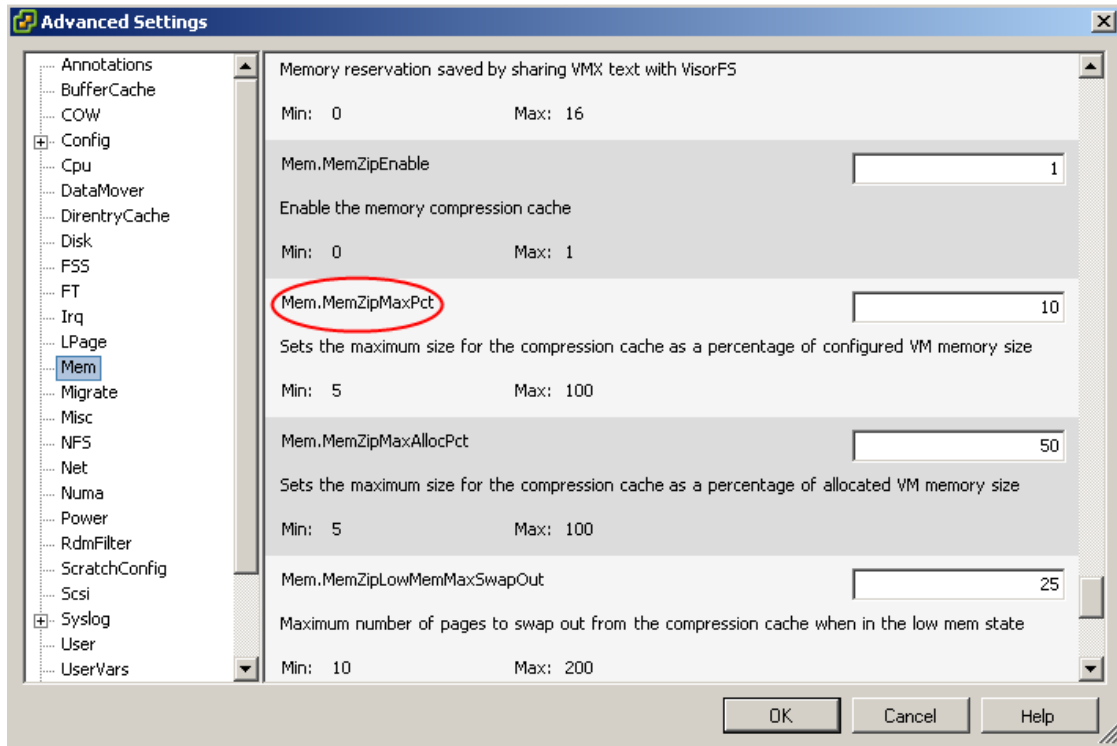
If the compression cache is full, one compressed page must be replaced in order to make room for a new compressed page. An age-based replacement policy is used to choose the target page. The target page will be decompressed and swapped out. ESX will not swap out compressed pages.

If the pages belonging to compression cache need to be swapped out under severe memory pressure, the compression cache size is reduced and the affected compressed pages are decompressed and swapped out.

The maximum compression cache size is important for maintaining good VM performance. If the upper bound is too small, a lot of replaced compressed pages must be decompressed and swapped out. Any following swap-ins of those pages will hurt VM performance. However, since compression cache is accounted for by the VM's guest memory usage, a very large compression cache may waste VM memory and unnecessarily create VM memory pressure especially when most compressed pages would not be touched in the future. In ESX 4.1, the default maximum compression cache size is conservatively set to 10%

of configured VM memory size. This value can be changed through the vSphere Client in Advanced Settings by changing the value for `Mem.MemZipMaxPct`.

**Figure 7.** Change the maximum compression cache size in the vSphere Client




### 3.6 When to Reclaim Host Memory<sup>4</sup>

ESX maintains four host free memory states: high, soft, hard, and low, which are reflected by four thresholds: 6%, 4%, 2%, and 1% of host memory respectively. Figure 8 shows how the host free memory state is reported in `esxtop`.

By default, ESX enables page sharing since it opportunistically “frees” host memory with little overhead. When to use ballooning or swapping (which activates memory compression) to reclaim host memory is largely determined by the current host free memory state.

<sup>4</sup> The discussions and conclusions made in this section may not be valid when the user specifies a resource pool for virtual machines. For example, if the resource pool that contains a virtual machine is specified as a small memory limit, ballooning or hypervisor swapping occur for the virtual machine even when the host free memory is in the high state. The detailed explanation of resource pool is beyond the scope of this paper. Most of the details can be found in the “Managing Resource Pools” section of the *vSphere Resource Management Guide* [2].

**Figure 8.** Host free memory state in esxtop


```

4:39:57am up 20:35, 141 worlds; MEM overcommit avg: 0.00, 0.00, 0.00
PMEM /MB: 73722 total: 800 cos, 1067 vmk, 8 other, 71846 free
VMKMEM/MB: 72053 managed: 4323 minfree, 5119 rsvd, 66933 ursvd, high state
COSMEM/MB: 77 free: 1239 swap_t, 1239 swap_f: 0.00 r/s, 0.00 w/s
NUMA /MB: 17365 (16739), 18357 (18224), 18357 (18269), 18357 (18250)
PSHARE/MB: 10 shared, 10 common: 0 saving
SWAP /MB: 0 curr, 0 rclmtgt: 0.00 r/s, 0.00 w/s
ZIP /MB: 0 zipped, 0 saved
MEMCTL/MB: 0 curr, 0 target, 0 max

```

GID	NAME	MEMSZ	GRANT	SZTGT	TCHD	TCHD W	SWCUR	SWTGT
53	vmkiscsid.4303	55.04	1.73	1.90	1.73	1.73	0.00	0.00
49	vobd.4299	25.43	5.00	5.50	5.00	5.00	0.00	0.00
52	net-lbt.4302	14.45	4.22	4.64	4.22	4.22	0.00	0.00
55	vmware-vmkauthd	5.88	1.73	1.91	1.73	1.73	0.00	0.00
51	net-cdp.4301	3.41	0.29	0.32	0.29	0.29	0.00	0.00

In the high state, the aggregate virtual machine guest memory usage is smaller than the host memory size. Whether or not host memory is overcommitted, the hypervisor will not reclaim memory through ballooning or swapping. (This is true only when the virtual machine memory limit is not set.)

If host free memory drops towards the soft threshold, the hypervisor starts to reclaim memory using ballooning. Ballooning happens before free memory actually reaches the soft threshold because it takes time for the balloon driver to allocate and pin guest physical memory. Usually, the balloon driver is able to reclaim memory in a timely fashion so that the host free memory stays above the soft threshold.

If ballooning is not sufficient to reclaim memory or the host free memory drops towards the hard threshold, the hypervisor starts to use swapping in addition to using ballooning. During swapping, memory compression is activated as well. With host swapping and memory compression, the hypervisor should be able to quickly reclaim memory and bring the host memory state back to the soft state.

In a rare case where host free memory drops below the low threshold, the hypervisor continues to reclaim memory through swapping and memory compression, and additionally blocks the execution of all virtual machines that consume more memory than their target memory allocations.


In certain scenarios, host memory reclamation happens regardless of the current host free memory state. For example, even if host free memory is in the high state, memory reclamation is still mandatory when a virtual machine's memory usage exceeds its specified memory limit. If this happens, the hypervisor will employ ballooning and, if necessary, swapping and memory compression to reclaim memory from the virtual machine until the virtual machine's host memory usage falls back to its specified limit.

## 4. ESX Memory Allocation Management for Multiple Virtual Machines

This section describes how ESX allocates host memory to multiple virtual machines, especially when the host memory is overcommitted.

ESX employs a share-based allocation algorithm to achieve efficient memory utilization for all virtual machines and to guarantee memory to those virtual machines which need it most. [1]

ESX provides three configurable parameters to control the host memory allocation for a virtual machine: Shares, Reservation, and Limit. The interface in the vSphere Client is shown in Figure 9.

**Figure 9.** Configure virtual machine memory allocation

Limit is the upper bound of the amount of host physical memory allocated for a virtual machine. By default, limit is set to unlimited, which means a virtual machine's maximum allocated host physical memory is its specified virtual machine memory size (according to Equation 1 on page 5). Reservation is a guaranteed lower bound on the amount of host physical memory the host reserves for a virtual machine even when host memory is overcommitted. Memory Shares entitle a virtual machine to a fraction of available host physical memory, based on a proportional-share allocation policy. For example, a virtual machine with twice as many shares as another is generally entitled to consume twice as much memory, subject to its limit and reservation constraints.

Periodically, ESX computes a memory allocation target for each virtual machine based on its share-based entitlement, its estimated working set size, and its limit and reservation. Here, a virtual machine's working set size is defined as the amount of guest physical memory that is actively being used. When host memory is undercommitted, a virtual machine's memory allocation target is the virtual machine's consumed host physical memory size with headroom. The maximum memory allocation target is:

#### Equation 2.

Maximum allocation target =  $\min\{\text{VM's memory size, VM's limit}\}$

When host memory is overcommitted, a virtual machine's allocation target is somewhere between its specified reservation and specified limit depending on the virtual machine's shares and the system load. If a virtual machine's host memory usage is larger than the computed allocation target, which typically happens in memory overcommitment cases, ESX employs a ballooning or swapping mechanism to reclaim memory from the virtual machine in order to reach the allocation target. Whether to use ballooning or to use swapping is determined by the current host free memory state as described in previous sections.

Shares play an important role in determining the allocation targets when memory is overcommitted. When the hypervisor needs memory, it reclaims memory from the virtual machine that owns the fewest shares-per-allocated page.

A significant limitation of the pure proportional-share algorithm is that it does not incorporate any information about the actual memory usage of the virtual machine. As a result, some idle virtual machines with high shares can retain idle memory unproductively, while some active virtual machines with fewer shares suffer from lack of memory.

ESX resolves this problem by estimating a virtual machine's working set size and charging a virtual machine more for the idle memory than for the actively used memory through an idle tax. [1] A virtual machine's shares-per-allocated page ratio is adjusted to be lower if a fraction of the virtual machine's memory is idle.

Hence, memory will be reclaimed preferentially from the virtual machines that are not fully utilizing their allocated memory. The detailed algorithm can be found in *Memory Resource Management in VMware ESX Server* [1], in section 5.2 “Reclaiming Idle Memory.” The effectiveness of this algorithm relies on the accurate estimation of the virtual machine’s working set size. ESX uses a statistical sampling approach to estimate the aggregate virtual machine working set size without any guest involvement. At the beginning of each sampling period, the hypervisor intentionally invalidates several randomly selected guest physical pages and starts to monitor the guest accesses to them. At the end of the sampling period, the fraction of actively used memory can be estimated as the fraction of the invalidated pages that are re-accessed by the guest during the epoch. The detailed algorithm can be found in *Memory Resource Management in VMware ESX Server* [1], in section 5.3 “Measuring Idle Memory.” By default, ESX samples 100 guest physical pages for each 60-second period. The sampling rate can be adjusted by changing `Mem.SamplePeriod` through the vSphere Client in Advanced Settings.

By overpricing the idle memory and effective working set estimation, ESX is able to efficiently allocate host memory under memory overcommitment while maintaining the proportional-share allocation.

## 5. Performance Evaluation

In this section, the performance of various ESX memory reclamation techniques is evaluated. The purpose is to help users understand how individual techniques impact the performance of various applications.

### 5.1 Experimental Environment

ESX 4.0 RC was installed on a Dell PowerEdge 6950 system and experiments were conducted against SPECjbb, Kernel Compile, Swingbench, Exchange, and SharePoint workloads to evaluate page sharing, ballooning and host swapping performance. ESX 4.1 RC was installed on a Dell PowerEdge R710 system to evaluate the memory compression feature using SharePoint and Swingbench workloads. The system hardware configurations and workload descriptions are summarized in Table 1 and Table 2 respectively.

**Table 1.** Server configurations

PowerEdge 6950	Processor:	4 socket dual core AMD Opteron 8222SE processors @ 3GHZ
	Memory:	64GB DDR2
PowerEdge R710	Processor:	2 quad-core Intel Xeon X5570 (Nehalem) processors @ 2.93GHz, with hyper-threading
	Memory:	96GB DDR3
SAN Storage	LUN Size:	2TB Fibre Channel

**Table 2.** Workload descriptions

SPECjbb2005	Heap Size:	2.5GB
	Number of Warehouse:	1
	Run Time:	10 minutes
	VM configuration:	1vcpu, 4GB memory

<b>Kernel Compile</b>	Linux Kernel Version:	2.6.17
	Command:	"make -j 1 bzimage > /dev/null"
	VM configuration:	1vcpu, 512MB memory
<b>Swingbench</b>	Database:	Oracle 11g
	Functional benchmark:	Order Entry
	Run time:	20 minutes
	VM configuration:	4 vcpus, 5GB memory
<b>Exchange 2007</b>	Server:	Microsoft Exchange 2007
	LoadGen client:	2000 heavy exchange users
	VM configuration:	4vcpus, 12GB memory
<b>SharePoint 2007</b>	SQL Server:	SQL server 2008 SP1 with 200GB Sharepoint user data
	Web Server:	Windows Server 2008 with Web Server and Query Server roles enabled
	Application Server:	Index server
	VM configurations:	3 Web Server VMs (2vcpu, 4GB memory) 1 Application Server VM (2vcpu, 4GB memory) 1 SQL Server VM (2vcpu, 16GB memory)

The guest operating system running inside the SPECjbb, kernel compile, and Swingbench virtual machines was 64-bit Red Hat Enterprise Linux 5.2 Server. The guest operating system running inside the Exchange virtual machine and SharePoint VMs was Windows Server 2008.

For SPECjbb2005 and Swingbench, the throughput was measured by calculating the number of transactions per second. For kernel compile, the performance was measured by calculating the inverse of the compilation time. For Exchange, the performance was measured using the average Remote Procedure Call (RPC) latency. For SharePoint, the performance was measured using the number of requests per second. In addition, for Swingbench, Exchange and SharePoint, the client applications were installed in a separate native machine.

## 5.2 Transparent Page Sharing Performance

In this experiment, two instances of workloads were run. The overall performance of workloads with page sharing enabled is compared to that with page sharing disabled. The focus is on evaluating the overhead of page scanning. Since the page scan rate (number of scanned pages per second) is largely determined by the

Mem.ShareScanTime<sup>5</sup>, in addition to the default 60 minutes, the minimal Mem.ShareScanTime of 10 minutes was tested, which potentially introduces the highest page scanning overhead.

**Figure 10.** Performance impact of transparent page sharing



Figure 10 confirms that enabling page sharing introduces a negligible performance overhead in the default setting and only <1 percent overhead when Mem.ShareScanTime is 10 minutes for all workloads.

Page sharing sometimes improves performance because the virtual machine's host memory footprint is reduced so that it fits the processor cache better.

Besides page scanning, breaking copy-on-right (CoW) pages is another source of page sharing overhead. Unfortunately, such overhead is highly application dependent and it is difficult to evaluate it through configurable options. Therefore, the overhead of breaking CoW pages was omitted in this experiment.

## 5.3 Ballooning vs. Host Swapping

In the following experiments, VM memory reclamation was enforced by changing each virtual machine's memory limit value from the default of unlimited, to values that are smaller than the configured virtual machine memory size. Page sharing was turned off to isolate the performance impact of ballooning or swapping. Since the host memory is much larger than the virtual machine memory size, the host free memory is always in the high state. Hence, by default, ESX only uses ballooning to reclaim memory. Both ballooning and memory compression were turned off to obtain the performance of using swapping only. The ballooned and swapped memory sizes were also collected when the virtual machine ran steadily.


### 5.3.1 Linux Kernel Compile

Figure 10 presents the throughput of the kernel compile workload with different memory limits when using ballooning or swapping. This experiment was contrived to use only ballooning or swapping, not both. While this case will not often occur in production environments, it shows the performance penalty due to

<sup>5</sup> This is an advanced setting, shown in Figure 5 on page 8.

either technology on its own. The throughput is normalized to the case where virtual machine memory is not reclaimed.

**Figure 11.** Performance of kernel compile when using the ballooning vs. the swapping



By using ballooning, the kernel compile virtual machine only suffers from 3 percent throughput loss even when the memory limit is as low as 128MB (1/4 of the configured virtual machine memory size). This is because the kernel compile workload has very little memory reuse and most of the guest physical memory is used as buffer caches for the kernel source files. With ballooning, the guest operating system reclaims guest physical memory upon the balloon driver's allocation request by dropping the buffer pages instead of paging them out to the guest virtual swap device. Because dropped buffer pages are not reused frequently, the performance impact of using ballooning is trivial.


However, with hypervisor swapping, the selected guest buffer pages are unnecessarily swapped out to the host swap device and some guest kernel pages are swapped out occasionally, making the performance of the virtual machine degrade when the memory limit decreases. When the memory limit is 128MB, the throughput loss is about 34 percent in the swapping case. Balloon inflation is a better approach to memory reclamation from a performance perspective.

Figure 11 illustrates that as the memory limit decreases, the ballooned and swapped memory sizes increase almost linearly. There is a difference between the ballooned memory size and the swapped memory size. In the ballooning cases, when virtual machine memory usage exceeds the specified limit, the balloon driver cannot force the guest operating system to page out guest physical pages immediately unless the balloon driver has used up most of the free guest physical memory, which puts the guest operating system under memory pressure. In the swapping cases, however, as long as the virtual machine memory usage exceeds the specified limit, the extra amount of pages will be swapped out immediately. Therefore, the ballooned memory size is roughly equal to the virtual machine memory size minus the specified limit, which means that the free physical memory is included. The swapped memory size is roughly equal to the virtual machine host memory usage minus the specified limit. In the kernel compile virtual machine, since most of the guest physical pages are used to buffer the workload files, the virtual machine's effective host memory usage is close to the virtual machine memory size. Hence, the swapped memory size is similar to the ballooned memory size.

### 5.3.2 Oracle/Swingbench

Figure 12 presents the throughput of an Oracle database tested by the Swingbench workload with different memory limits when using ballooning vs. swapping. The throughput is normalized to the case where virtual machine memory is not reclaimed.

**Figure 12.** Performance of Swingbench when using ballooning vs. swapping




As shown in the kernel compile test, using ballooning barely impacts the throughput of the Swingbench virtual machine until the memory limit decreases below 2048MB. This occurs when the guest operating system starts to page out the physical pages that are heavily reused by the Oracle database.

In contrast to ballooning, any amount of swapping causes significant throughput penalty. The throughput loss is already 17 percent when the memory limit is 3584MB. In hypervisor swapping, some guest buffer pages are unnecessarily swapped out and some guest kernel or performance-critical database pages are also unintentionally swapped out because of the random page selection policy. For the Swingbench virtual machine, the virtual machine host memory usage is very close to the virtual machine memory size, so the swapped memory size is very close to the ballooned memory size.

### 5.3.3 SPECjbb

Figure 13 presents the throughput of the SPECjbb workload with different memory limits when using ballooning vs. swapping.

The throughput is normalized to the case where virtual machine memory is not reclaimed.

**Figure 13.** Performance of SPECjbb when using ballooning vs. swapping

From this figure, it is seen that when the virtual machine memory limit decreases beyond 2816MB, the throughput of SPECjbb degrades significantly in both ballooning and swapping cases. When the memory limit is reduced to 2048MB, the throughput losses are 89 percent and 96 percent for ballooning and swapping respectively. Since the configured JVM heap size is 2.5GB, the actual virtual machine working set size is around 2.5GB plus guest operating system memory usage (about 300MB). When the virtual machine memory limit falls below 2816MB, the host memory cannot back the entire virtual machine's working set, so that virtual machine starts to suffer from guest-level paging in the ballooning cases or hypervisor swapping in the swapping cases.

Since SPECjbb is an extremely memory intensive workload, its throughput is largely determined by the virtual machine memory hit rate. In this instance, virtual machine memory hit rate is defined as the percentage of guest memory accesses that result in host physical memory hits. A higher memory hit rate means higher throughput for the SPECjbb workload. Since ballooning and host swapping similarly decrease memory hit rate, both guest level paging and hypervisor swapping largely hurt SPECjbb performance.

Surprisingly, when the memory limit is 2506MB or 2304MB, swapping yields higher throughput than ballooning does. This seems to be counterintuitive because hypervisor swapping typically causes a higher performance penalty. One reasonable explanation is that the random page selection policy used in hypervisor swapping largely favors the access patterns of the SPECjbb virtual machine. More specifically, with ballooning, when the guest operating system (Linux in this case) pages out guest physical pages to satisfy the balloon driver's allocation request, it chooses the pages using an LRU-approximated<sup>6</sup> policy. However, JVM often scans the allocated guest physical memory (Java heap) for garbage collection purpose. This behavior may fall into a well-known LRU pathological case in which the memory hit rate drops dramatically even when the memory size is slightly smaller than the working set size. In contrast,

<sup>6</sup> LRU refers to "least recently used." LRU is a caching policy that determines how pages are discarded when the cache is full and a new page must be accepted into the cache. LRU specifies that the least recently used page is discarded.

when using hypervisor swapping, the swapped physical pages are randomly selected by the hypervisor, which makes the memory hit rate reduce more gradually as the memory limit decreases. That is why using swapping achieves higher throughput when the memory limit is smaller than the virtual machine's working set size. However, when the memory limit drops to 2304MB, the virtual machine memory hit rate is equivalently low in both swapping and ballooning cases. Using swapping starts to cause worse performance compared to using ballooning. Note that the above two configurations where swapping outperforms ballooning are rare pathological cases for ballooning performance. In most cases, using ballooning achieves much better performance compared to using swapping.

Since the virtual machine's working set size (~2.8GB) is much smaller than the configured virtual machine memory size (4GB), the ballooned memory size is much higher than the swapped memory size.

### 5.3.4 Microsoft Exchange Server 2007

This section presents how ballooning and swapping impact the performance of an Exchange Server virtual machine. Exchange Server is a memory intensive workload that is optimized to use all the available physical memory to cache the transactions for fewer disk I/Os.

The Exchange Server performance was measured using the average Remote Procedure Call (RPC) latency during two hours of stable performance. The RPC latency gauges the server processing time for an RPC from LoadGen (the client application that drives the Exchange server). Therefore, lower RPC latency means better performance. The results are presented in Figure 14.

**Figure 14.** Average RPC latency of Exchange when using ballooning vs. using swapping




Figure 14 (a), illustrates that when the memory limit decreases from 12GB to 3GB, the average RPC latency is gradually increased from 4.6ms to 7.3ms with ballooning. However, as shown in Figure 14 (b), the RPC latency is dramatically increased from 4.6ms to 143ms when solely swapping out 2GB host memory. When the memory limit is reduced to 9GB, hypervisor swapping makes the RPC latency too high, which resulted in the failure of the LoadGen application (due to timeout).

Overall, this figure confirms that using ballooning to reclaim memory is much more efficient than using hypervisor swapping for the Exchange Server virtual machine.


## 5.4 Memory Compression Performance

In this experiment, host memory overcommitment is enforced by reducing the host physical memory size instead of setting the VM memory limit. The maximum compression cache size is fixed at 10% of configured guest memory size.

### 5.4.1 SharePoint

SharePoint Server is composed of five VMs with a total VM memory size of 32GB. The SQL server VM is given a full 16GB memory reservation because any memory reclamation through ballooning or swapping from this VM will make the SharePoint performance degrade significantly. No reservations are set for the other four VMs. In addition, all the memory reclamation techniques are enabled as default. Figure 15 presents how memory compression helps improve the SharePoint performance when host memory size is reduced. The performance result is normalized to the 36GB case where the host memory is undercommitted.

**Figure 15.** Throughput and Swap Read rate of SharePoint with different host memory sizes



As shown in Figure 15, when host memory size is reduced from 36GB to 28GB, there is only 6% performance degradation. This is because page sharing and ballooning efficiently reclaim most of VM memory with little performance penalty. There are a few swapping activities when host memory is reduced to 28GB and 26GB, but they do not impact performance much.

However, when host memory is reduced to 26GB, ballooning and page sharing are not enough to reclaim host memory. ESX has to swap out guest pages from VMs. Without memory compression, the performance loss is 95% because of the severe penalty of host swapping. Interestingly, by using memory compression, the performance loss is brought back from 95% to 6%. Such a huge improvement is mainly due to the significant reduction in the amount of swapped in pages. As we can see, applying memory compression reduces the swap in (swap read) rate by around 85% since most of the missing pages can be found in the compression caches.


When host memory size is reduced to 24GB, even with memory compression, the throughput drops to 25% of the throughput in the 36GB case because the amount of swapped in pages is considerably increased. By checking the compression cache statistics, we found that compression caches of those VMs reached their maximum limit size in this case.

### 5.4.2 Swingbench

In this experiment, 16 Swingbench VMs were used to overcommit the host memory and none of the VMs had any memory reservation. The total VM configured memory size was 80GB. Figure 16 shows the normalized overall throughput of Swingbench with memory compression versus without memory compression when host memory size is reduced from 96GB to 50GB. The performance result is normalized to the 96GB case. Each of the 16 VMs had the same copy of the database installed, which led to artificial page sharing and made the impact of memory compression less obvious. So for benchmarking purposes, we temporarily turned off page sharing in this experiment.

From Figure 16, we can see that when host memory is reduced from 96GB to 70GB, the performance loss of Swingbench is < 5% since ballooning efficiently reclaims the guest pages from the VMs without causing severe guest paging. However, when host memory is reduced to 60GB, the memory pressure is high enough to make the host start to swap out guest pages. Using memory compression can easily recover 14% lost throughput. Again, as shown in this figure, the improvement is mainly due to the significant reduction in host page swap in (swap read) operations. For the same reason, memory compression can bring the normalized throughput from 42% to 70% when host memory is only 50GB. In this case, even with memory compression, there still were considerable swap in operations. Upon close observation, we found that although the compression caches of those VMs were not full, there were quite a lot of swapped out pages that failed on compression because they could not be compressed into half pages. Hence, memory compression cannot completely eliminate host swapping.

**Figure 16.** Throughput and Swap Read rate of Swingbench with different host memory sizes



Memory compression itself introduces some CPU overhead when performing compression or decompression. In this experiment, we found that the physical CPU utilization is slightly increased by 1-2% when memory compression is enabled.

## 6. Best Practices

Based on the memory management concepts and performance evaluation results presented in the previous sections, the following are some best practices for host and guest memory usage.

- Do not disable page sharing or the balloon driver. As described, page sharing is a lightweight technique which opportunistically reclaims redundant host memory with trivial performance impact. In the cases where hosts are heavily overcommitted, using ballooning is generally more efficient and safer than using hypervisor swapping, based on the results presented in Section 5.3. These two techniques are enabled by default in ESX 4 and should not be disabled unless application testing shows that the benefits of doing so clearly outweigh the costs.
- Carefully specify memory limits and reservations. The virtual machine memory allocation target is subject to the VM's memory limit and reservation. If these two parameters are misconfigured, users may observe ballooning or swapping even when the host has plenty of free memory. For example, a virtual machine's memory may be reclaimed when the specified limit is too small or when other virtual machines reserve too much host memory, even though they may only use a small portion of the reserved memory. If a performance-critical virtual machine needs a guaranteed memory allocation, the reservation needs to be specified carefully because it may impact other virtual machines.
- Host memory size should be larger than guest memory usage. For example, it is unwise to run a virtual machine with a 2GB working set size in a host with only 1GB host memory. If this is the case, the hypervisor has to reclaim the virtual machine's active memory through ballooning or hypervisor swapping, which will lead to potentially serious virtual machine performance degradation. Although it is difficult to tell whether the host memory is large enough to hold all of the virtual machines' working sets, the bottom line is that the host memory should not be excessively overcommitted making the guests have to continuously page out guest physical memory.
- Use shares to adjust relative priorities when memory is overcommitted. If the host's memory is overcommitted and the virtual machine's allocated host memory is too small to achieve a reasonable performance, adjust the virtual machine's shares to escalate the relative priority of the virtual machine so that the hypervisor will allocate more host memory for that virtual machine.
- Set an appropriate virtual machine memory size. The virtual machine memory size should be slightly larger than the average guest memory usage. The extra memory will accommodate workload spikes in the virtual machine. Note that the guest operating system only recognizes the specified virtual machine memory size. If the virtual machine memory size is too small, guest-level paging is inevitable, even though the host may have plenty of free memory. If the virtual machine memory size is set to a very large value, virtual machine performance will be fine, but more virtual machine memory means that more overhead memory needs to be reserved for the virtual machine.

## 7. References

- [1] Carl A. Waldspurger. "Memory Resource Management in VMware ESX Server". Proceeding of the fifth Symposium on Operating System Design and Implementation, Boston, Dec 2002.  
<http://www.waldspurger.org/carl/papers/esx-mem-osdi02.pdf>
- [2] vSphere Resource Management Guide. VMware guide.  
[http://www.vmware.com/pdf/vsphere4/r40/vsp\\_40\\_upgrade\\_guide.pdf](http://www.vmware.com/pdf/vsphere4/r40/vsp_40_upgrade_guide.pdf)
- [3] Memory Performance Chart Statistics in the vSphere Client. VMware Community document.  
<http://communities.vmware.com/docs/DOC-10398>
- [4] VirtualCenter Memory Statistics Definitions. VMware Community document.  
<http://communities.vmware.com/docs/DOC-5230>
- [5] Performance Evaluation of Intel EPT Hardware Assist. VMware performance study.  
<http://www.vmware.com/resources/techresources/10006>
- [6] Performance Evaluation of AMD RVI Hardware Assist. VMware performance study.  
<http://www.vmware.com/resources/techresources/1079>
- [7] "The buffer cache," Linux System Administrators Guide. The Linux Documentation Project.  
<http://tldp.org/LDP/sag/html/buffer-cache.html>

---

If you have comments about this documentation, submit your feedback to: [docfeedback@vmware.com](mailto:docfeedback@vmware.com)

**VMware, Inc. 3401 Hillview Ave., Palo Alto, CA 94304 [www.vmware.com](http://www.vmware.com)**

Copyright © 2010 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

Item: EN-000411-00

---