

# vSphere Resource Management

Update 1  
VMware vSphere 6.0  
VMware ESXi 6.0  
vCenter Server 6.0

This document supports the version of each product listed and supports all subsequent versions until the document is replaced by a new edition. To check for more recent editions of this document, see <http://www.vmware.com/support/pubs>.

EN-001903-00

**vmware**<sup>®</sup>

You can find the most up-to-date technical documentation on the VMware Web site at:

<http://www.vmware.com/support/>

The VMware Web site also provides the latest product updates.

If you have comments about this documentation, submit your feedback to:

[docfeedback@vmware.com](mailto:docfeedback@vmware.com)

Copyright © 2006–2015 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

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

# Contents

About vSphere Resource Management	7
<b>1 Getting Started with Resource Management</b>	<b>9</b>
Resource Types	9
Resource Providers	9
Resource Consumers	10
Goals of Resource Management	10
<b>2 Configuring Resource Allocation Settings</b>	<b>11</b>
Resource Allocation Shares	11
Resource Allocation Reservation	12
Resource Allocation Limit	12
Resource Allocation Settings Suggestions	13
Edit Resource Settings	13
Changing Resource Allocation Settings—Example	14
Admission Control	15
<b>3 CPU Virtualization Basics</b>	<b>17</b>
Software-Based CPU Virtualization	17
Hardware-Assisted CPU Virtualization	18
Virtualization and Processor-Specific Behavior	18
Performance Implications of CPU Virtualization	18
<b>4 Administering CPU Resources</b>	<b>19</b>
View Processor Information	19
Specifying CPU Configuration	19
Multicore Processors	20
Hyperthreading	20
Using CPU Affinity	22
Host Power Management Policies	23
<b>5 Memory Virtualization Basics</b>	<b>27</b>
Virtual Machine Memory	27
Memory Overcommitment	28
Memory Sharing	28
Types of Memory Virtualization	29
<b>6 Administering Memory Resources</b>	<b>33</b>
Understanding Memory Overhead	33
How ESXi Hosts Allocate Memory	34

- Memory Reclamation 35
- Using Swap Files 36
- Sharing Memory Across Virtual Machines 40
- Memory Compression 41
- Measuring and Differentiating Types of Memory Usage 42
- Memory Reliability 43
- About System Swap 43
  
- 7 View Graphics Information 45**
  
- 8 Managing Storage I/O Resources 47**
  - Storage I/O Control Requirements 47
  - Storage I/O Control Resource Shares and Limits 48
  - Set Storage I/O Control Resource Shares and Limits 49
  - Enable Storage I/O Control 49
  - Set Storage I/O Control Threshold Value 50
  - Storage DRS Integration with Storage Profiles 51
  
- 9 Managing Resource Pools 53**
  - Why Use Resource Pools? 54
  - Create a Resource Pool 55
  - Edit a Resource Pool 56
  - Add a Virtual Machine to a Resource Pool 56
  - Remove a Virtual Machine from a Resource Pool 57
  - Remove a Resource Pool 58
  - Resource Pool Admission Control 58
  
- 10 Creating a DRS Cluster 61**
  - Admission Control and Initial Placement 62
  - Virtual Machine Migration 63
  - DRS Cluster Requirements 65
  - Configuring DRS with Virtual Flash 66
  - Create a Cluster 66
  - Edit a Cluster 67
  - Create a DRS Cluster 68
  - Set a Custom Automation Level for a Virtual Machine 69
  - Disable DRS 70
  - Restore a Resource Pool Tree 70
  
- 11 Using DRS Clusters to Manage Resources 71**
  - Adding Hosts to a Cluster 71
  - Adding Virtual Machines to a Cluster 73
  - Removing Virtual Machines from a Cluster 73
  - Removing a Host from a Cluster 74
  - DRS Cluster Validity 75
  - Managing Power Resources 80
  - Using DRS Affinity Rules 84

- 12 Creating a Datastore Cluster 89**
  - Initial Placement and Ongoing Balancing 90
  - Storage Migration Recommendations 90
  - Create a Datastore Cluster 90
  - Enable and Disable Storage DRS 91
  - Set the Automation Level for Datastore Clusters 91
  - Setting the Aggressiveness Level for Storage DRS 92
  - Datastore Cluster Requirements 93
  - Adding and Removing Datastores from a Datastore Cluster 94
  
- 13 Using Datastore Clusters to Manage Storage Resources 95**
  - Using Storage DRS Maintenance Mode 95
  - Applying Storage DRS Recommendations 97
  - Change Storage DRS Automation Level for a Virtual Machine 98
  - Set Up Off-Hours Scheduling for Storage DRS 98
  - Storage DRS Anti-Affinity Rules 99
  - Clear Storage DRS Statistics 102
  - Storage vMotion Compatibility with Datastore Clusters 103
  
- 14 Using NUMA Systems with ESXi 105**
  - What is NUMA? 105
  - How ESXi NUMA Scheduling Works 106
  - VMware NUMA Optimization Algorithms and Settings 107
  - Resource Management in NUMA Architectures 108
  - Using Virtual NUMA 108
  - Specifying NUMA Controls 110
  
- 15 Advanced**
  - Attributes 113
    - Set Advanced Host Attributes 113
    - Set Advanced Virtual Machine Attributes 116
  - Latency Sensitivity 118
  - About Reliable Memory 118
  
- 16 Fault Definitions 121**
  - Virtual Machine is Pinned 122
  - Virtual Machine not Compatible with any Host 122
  - VM/VM DRS Rule Violated when Moving to another Host 122
  - Host Incompatible with Virtual Machine 122
  - Host has Virtual Machine that Violates VM/VM DRS Rules 122
  - Host has Insufficient Capacity for Virtual Machine 122
  - Host in Incorrect State 122
  - Host has Insufficient Number of Physical CPUs for Virtual Machine 123
  - Host has Insufficient Capacity for Each Virtual Machine CPU 123
  - The Virtual Machine is in vMotion 123
  - No Active Host in Cluster 123
  - Insufficient Resources 123
  - Insufficient Resources to Satisfy Configured Failover Level for HA 123

	No Compatible Hard Affinity Host	123
	No Compatible Soft Affinity Host	123
	Soft Rule Violation Correction Disallowed	123
	Soft Rule Violation Correction Impact	124
<b>17</b>	<b>DRS Troubleshooting Information</b>	<b>125</b>
	Cluster Problems	125
	Host Problems	128
	Virtual Machine Problems	131
	Index	135

# About vSphere Resource Management

---

*vSphere Resource Management* describes resource management for VMware® ESXi and vCenter® Server environments.

This documentation focuses on the following topics.

- Resource allocation and resource management concepts
- Virtual machine attributes and admission control
- Resource pools and how to manage them
- Clusters, vSphere® Distributed Resource Scheduler (DRS), vSphere Distributed Power Management (DPM), and how to work with them
- Datastore clusters, Storage DRS, Storage I/O Control, and how to work with them
- Advanced resource management options
- Performance considerations

## Intended Audience

This information is for system administrators who want to understand how the system manages resources and how they can customize the default behavior. It's also essential for anyone who wants to understand and use resource pools, clusters, DRS, datastore clusters, Storage DRS, Storage I/O Control, or vSphere DPM.

This documentation assumes you have a working knowledge of VMware ESXi and of vCenter Server.





# Getting Started with Resource Management

---

# 1

To understand resource management, you must be aware of its components, its goals, and how best to implement it in a cluster setting.

Resource allocation settings for a virtual machine (shares, reservation, and limit) are discussed, including how to set them and how to view them. Also, admission control, the process whereby resource allocation settings are validated against existing resources is explained.

Resource management is the allocation of resources from resource providers to resource consumers.

The need for resource management arises from the overcommitment of resources—that is, more demand than capacity and from the fact that demand and capacity vary over time. Resource management allows you to dynamically reallocate resources, so that you can more efficiently use available capacity.

This chapter includes the following topics:

- [“Resource Types,”](#) on page 9
- [“Resource Providers,”](#) on page 9
- [“Resource Consumers,”](#) on page 10
- [“Goals of Resource Management,”](#) on page 10

## Resource Types

Resources include CPU, memory, power, storage, and network resources.

---

**NOTE** ESXi manages network bandwidth and disk resources on a per-host basis, using network traffic shaping and a proportional share mechanism, respectively.

---

## Resource Providers

Hosts and clusters, including datastore clusters, are providers of physical resources.

For hosts, available resources are the host’s hardware specification, minus the resources used by the virtualization software.

A cluster is a group of hosts. You can create a cluster using vSphere Web Client, and add multiple hosts to the cluster. vCenter Server manages these hosts’ resources jointly: the cluster owns all of the CPU and memory of all hosts. You can enable the cluster for joint load balancing or failover. See [Chapter 10, “Creating a DRS Cluster,”](#) on page 61 for more information.

A datastore cluster is a group of datastores. Like DRS clusters, you can create a datastore cluster using the vSphere Web Client, and add multiple datastores to the cluster. vCenter Server manages the datastore resources jointly. You can enable Storage DRS to balance I/O load and space utilization. See [Chapter 12, “Creating a Datastore Cluster,”](#) on page 89.

## Resource Consumers

Virtual machines are resource consumers.

The default resource settings assigned during creation work well for most machines. You can later edit the virtual machine settings to allocate a share-based percentage of the total CPU, memory, and storage I/O of the resource provider or a guaranteed reservation of CPU and memory. When you power on that virtual machine, the server checks whether enough unreserved resources are available and allows power on only if there are enough resources. This process is called admission control.

A resource pool is a logical abstraction for flexible management of resources. Resource pools can be grouped into hierarchies and used to hierarchically partition available CPU and memory resources. Accordingly, resource pools can be considered both resource providers and consumers. They provide resources to child resource pools and virtual machines, but are also resource consumers because they consume their parents' resources. See [Chapter 9, "Managing Resource Pools,"](#) on page 53.

ESXi hosts allocate each virtual machine a portion of the underlying hardware resources based on a number of factors:

- Resource limits defined by the user.
- Total available resources for the ESXi host (or the cluster).
- Number of virtual machines powered on and resource usage by those virtual machines.
- Overhead required to manage the virtualization.

## Goals of Resource Management

When managing your resources, you should be aware of what your goals are.

In addition to resolving resource overcommitment, resource management can help you accomplish the following:

- Performance Isolation—prevent virtual machines from monopolizing resources and guarantee predictable service rates.
- Efficient Utilization—exploit undercommitted resources and overcommit with graceful degradation.
- Easy Administration—control the relative importance of virtual machines, provide flexible dynamic partitioning, and meet absolute service-level agreements.

# Configuring Resource Allocation Settings

---

# 2

When available resource capacity does not meet the demands of the resource consumers (and virtualization overhead), administrators might need to customize the amount of resources that are allocated to virtual machines or to the resource pools in which they reside.

Use the resource allocation settings (shares, reservation, and limit) to determine the amount of CPU, memory, and storage resources provided for a virtual machine. In particular, administrators have several options for allocating resources.

- Reserve the physical resources of the host or cluster.
- Set an upper bound on the resources that can be allocated to a virtual machine.
- Guarantee that a particular virtual machine is always allocated a higher percentage of the physical resources than other virtual machines.

This chapter includes the following topics:

- [“Resource Allocation Shares,”](#) on page 11
- [“Resource Allocation Reservation,”](#) on page 12
- [“Resource Allocation Limit,”](#) on page 12
- [“Resource Allocation Settings Suggestions,”](#) on page 13
- [“Edit Resource Settings,”](#) on page 13
- [“Changing Resource Allocation Settings—Example,”](#) on page 14
- [“Admission Control,”](#) on page 15

## Resource Allocation Shares

Shares specify the relative importance of a virtual machine (or resource pool). If a virtual machine has twice as many shares of a resource as another virtual machine, it is entitled to consume twice as much of that resource when these two virtual machines are competing for resources.

Shares are typically specified as **High**, **Normal**, or **Low** and these values specify share values with a 4:2:1 ratio, respectively. You can also select **Custom** to assign a specific number of shares (which expresses a proportional weight) to each virtual machine.

Specifying shares makes sense only with regard to sibling virtual machines or resource pools, that is, virtual machines or resource pools with the same parent in the resource pool hierarchy. Siblings share resources according to their relative share values, bounded by the reservation and limit. When you assign shares to a virtual machine, you always specify the priority for that virtual machine relative to other powered-on virtual machines.

The following table shows the default CPU and memory share values for a virtual machine. For resource pools, the default CPU and memory share values are the same, but must be multiplied as if the resource pool were a virtual machine with four virtual CPUs and 16 GB of memory.

**Table 2-1.** Share Values

Setting	CPU share values	Memory share values
High	2000 shares per virtual CPU	20 shares per megabyte of configured virtual machine memory.
Normal	1000 shares per virtual CPU	10 shares per megabyte of configured virtual machine memory.
Low	500 shares per virtual CPU	5 shares per megabyte of configured virtual machine memory.

For example, an SMP virtual machine with two virtual CPUs and 1GB RAM with CPU and memory shares set to **Normal** has  $2 \times 1000 = 2000$  shares of CPU and  $10 \times 1024 = 10240$  shares of memory.

**NOTE** Virtual machines with more than one virtual CPU are called SMP (symmetric multiprocessing) virtual machines. ESXi supports up to 128 virtual CPUs per virtual machine.

The relative priority represented by each share changes when a new virtual machine is powered on. This affects all virtual machines in the same resource pool. All of the virtual machines have the same number of virtual CPUs. Consider the following examples.

- Two CPU-bound virtual machines run on a host with 8GHz of aggregate CPU capacity. Their CPU shares are set to **Normal** and get 4GHz each.
- A third CPU-bound virtual machine is powered on. Its CPU shares value is set to **High**, which means it should have twice as many shares as the machines set to **Normal**. The new virtual machine receives 4GHz and the two other machines get only 2GHz each. The same result occurs if the user specifies a custom share value of 2000 for the third virtual machine.

## Resource Allocation Reservation

A reservation specifies the guaranteed minimum allocation for a virtual machine.

vCenter Server or ESXi allows you to power on a virtual machine only if there are enough unreserved resources to satisfy the reservation of the virtual machine. The server guarantees that amount even when the physical server is heavily loaded. The reservation is expressed in concrete units (megahertz or megabytes).

For example, assume you have 2GHz available and specify a reservation of 1GHz for VM1 and 1GHz for VM2. Now each virtual machine is guaranteed to get 1GHz if it needs it. However, if VM1 is using only 500MHz, VM2 can use 1.5GHz.

Reservation defaults to 0. You can specify a reservation if you need to guarantee that the minimum required amounts of CPU or memory are always available for the virtual machine.

## Resource Allocation Limit

Limit specifies an upper bound for CPU, memory, or storage I/O resources that can be allocated to a virtual machine.

A server can allocate more than the reservation to a virtual machine, but never allocates more than the limit, even if there are unused resources on the system. The limit is expressed in concrete units (megahertz, megabytes, or I/O operations per second).

CPU, memory, and storage I/O resource limits default to unlimited. When the memory limit is unlimited, the amount of memory configured for the virtual machine when it was created becomes its effective limit.

In most cases, it is not necessary to specify a limit. There are benefits and drawbacks:

- **Benefits** — Assigning a limit is useful if you start with a small number of virtual machines and want to manage user expectations. Performance deteriorates as you add more virtual machines. You can simulate having fewer resources available by specifying a limit.
- **Drawbacks** — You might waste idle resources if you specify a limit. The system does not allow virtual machines to use more resources than the limit, even when the system is underutilized and idle resources are available. Specify the limit only if you have good reasons for doing so.

## Resource Allocation Settings Suggestions

Select resource allocation settings (reservation, limit and shares) that are appropriate for your ESXi environment.

The following guidelines can help you achieve better performance for your virtual machines.

- Use **Reservation** to specify the minimum acceptable amount of CPU or memory, not the amount you want to have available. The amount of concrete resources represented by a reservation does not change when you change the environment, such as by adding or removing virtual machines. The host assigns additional resources as available based on the limit for your virtual machine, the number of shares and estimated demand.
- When specifying the reservations for virtual machines, do not commit all resources (plan to leave at least 10% unreserved). As you move closer to fully reserving all capacity in the system, it becomes increasingly difficult to make changes to reservations and to the resource pool hierarchy without violating admission control. In a DRS-enabled cluster, reservations that fully commit the capacity of the cluster or of individual hosts in the cluster can prevent DRS from migrating virtual machines between hosts.
- If you expect frequent changes to the total available resources, use **Shares** to allocate resources fairly across virtual machines. If you use **Shares**, and you upgrade the host, for example, each virtual machine stays at the same priority (keeps the same number of shares) even though each share represents a larger amount of memory, CPU, or storage I/O resources.

## Edit Resource Settings

Use the Edit Resource Settings dialog box to change allocations for memory and CPU resources.

### Procedure

- 1 Browse to the virtual machine in the vSphere Web Client navigator.
- 2 Right-click and select **Edit Resource Settings**.
- 3 Edit the CPU Resources.

Option	Description
<b>Shares</b>	CPU shares for this resource pool with respect to the parent's total. Sibling resource pools share resources according to their relative share values bounded by the reservation and limit. Select <b>Low</b> , <b>Normal</b> , or <b>High</b> , which specify share values respectively in a 1:2:4 ratio. Select <b>Custom</b> to give each virtual machine a specific number of shares, which expresses a proportional weight.
<b>Reservation</b>	Guaranteed CPU allocation for this resource pool.
<b>Limit</b>	Upper limit for this resource pool's CPU allocation. Select <b>Unlimited</b> to specify no upper limit.

- 4 Edit the Memory Resources.

Option	Description
<b>Shares</b>	Memory shares for this resource pool with respect to the parent's total. Sibling resource pools share resources according to their relative share values bounded by the reservation and limit. Select <b>Low</b> , <b>Normal</b> , or <b>High</b> , which specify share values respectively in a 1:2:4 ratio. Select <b>Custom</b> to give each virtual machine a specific number of shares, which expresses a proportional weight.
<b>Reservation</b>	Guaranteed memory allocation for this resource pool.
<b>Limit</b>	Upper limit for this resource pool's memory allocation. Select <b>Unlimited</b> to specify no upper limit.


- 5 Click **OK**.

## Changing Resource Allocation Settings—Example

The following example illustrates how you can change resource allocation settings to improve virtual machine performance.

Assume that on an ESXi host, you have created two new virtual machines—one each for your QA (VM-QA) and Marketing (VM-Marketing) departments.

**Figure 2-1.** Single Host with Two Virtual Machines



In the following example, assume that VM-QA is memory intensive and accordingly you want to change the resource allocation settings for the two virtual machines to:

- Specify that, when system memory is overcommitted, VM-QA can use twice as much CPU and memory resources as the Marketing virtual machine. Set the CPU shares and memory shares for VM-QA to **High** and for VM-Marketing set them to **Normal**.
- Ensure that the Marketing virtual machine has a certain amount of guaranteed CPU resources. You can do so using a reservation setting.

### Procedure

- 1 Browse to the virtual machines in the vSphere Web Client navigator.
- 2 Right-click **VM-QA**, the virtual machine for which you want to change shares, and select **Edit Settings**.
- 3 Under **Virtual Hardware**, expand CPU and select **High** from the **Shares** drop-down menu.
- 4 Under **Virtual Hardware**, expand Memory and select **High** from the **Shares** drop-down menu.
- 5 Click **OK**.
- 6 Right-click the marketing virtual machine (**VM-Marketing**) and select **Edit Settings**.
- 7 Under **Virtual Hardware**, expand CPU and change the **Reservation** value to the desired number.
- 8 Click **OK**.

If you select the cluster's **Resource Reservation** tab and click **CPU**, you should see that shares for **VM-QA** are twice that of the other virtual machine. Also, because the virtual machines have not been powered on, the **Reservation Used** fields have not changed.

## Admission Control

When you power on a virtual machine, the system checks the amount of CPU and memory resources that have not yet been reserved. Based on the available unreserved resources, the system determines whether it can guarantee the reservation for which the virtual machine is configured (if any). This process is called admission control.

If enough unreserved CPU and memory are available, or if there is no reservation, the virtual machine is powered on. Otherwise, an **Insufficient Resources** warning appears.

---

**NOTE** In addition to the user-specified memory reservation, for each virtual machine there is also an amount of overhead memory. This extra memory commitment is included in the admission control calculation.

---

When the vSphere DPM feature is enabled, hosts might be placed in standby mode (that is, powered off) to reduce power consumption. The unreserved resources provided by these hosts are considered available for admission control. If a virtual machine cannot be powered on without these resources, a recommendation to power on sufficient standby hosts is made.





## CPU Virtualization Basics

---

CPU virtualization emphasizes performance and runs directly on the processor whenever possible. The underlying physical resources are used whenever possible and the virtualization layer runs instructions only as needed to make virtual machines operate as if they were running directly on a physical machine.

CPU virtualization is not the same thing as emulation. ESXi does not use emulation to run virtual CPUs. With emulation, all operations are run in software by an emulator. A software emulator allows programs to run on a computer system other than the one for which they were originally written. The emulator does this by emulating, or reproducing, the original computer's behavior by accepting the same data or inputs and achieving the same results. Emulation provides portability and runs software designed for one platform across several platforms.

When CPU resources are overcommitted, the ESXi host time-slices the physical processors across all virtual machines so each virtual machine runs as if it has its specified number of virtual processors. When an ESXi host runs multiple virtual machines, it allocates to each virtual machine a share of the physical resources. With the default resource allocation settings, all virtual machines associated with the same host receive an equal share of CPU per virtual CPU. This means that a single-processor virtual machines is assigned only half of the resources of a dual-processor virtual machine.

This chapter includes the following topics:

- [“Software-Based CPU Virtualization,”](#) on page 17
- [“Hardware-Assisted CPU Virtualization,”](#) on page 18
- [“Virtualization and Processor-Specific Behavior,”](#) on page 18
- [“Performance Implications of CPU Virtualization,”](#) on page 18

### Software-Based CPU Virtualization

With software-based CPU virtualization, the guest application code runs directly on the processor, while the guest privileged code is translated and the translated code executes on the processor.

The translated code is slightly larger and usually executes more slowly than the native version. As a result, guest programs, which have a small privileged code component, run with speeds very close to native. Programs with a significant privileged code component, such as system calls, traps, or page table updates can run slower in the virtualized environment.

## Hardware-Assisted CPU Virtualization

Certain processors provide hardware assistance for CPU virtualization.

When using this assistance, the guest can use a separate mode of execution called guest mode. The guest code, whether application code or privileged code, runs in the guest mode. On certain events, the processor exits out of guest mode and enters root mode. The hypervisor executes in the root mode, determines the reason for the exit, takes any required actions, and restarts the guest in guest mode.

When you use hardware assistance for virtualization, there is no need to translate the code. As a result, system calls or trap-intensive workloads run very close to native speed. Some workloads, such as those involving updates to page tables, lead to a large number of exits from guest mode to root mode. Depending on the number of such exits and total time spent in exits, hardware-assisted CPU virtualization can speed up execution significantly.

## Virtualization and Processor-Specific Behavior

Although VMware software virtualizes the CPU, the virtual machine detects the specific model of the processor on which it is running.

Processor models might differ in the CPU features they offer, and applications running in the virtual machine can make use of these features. Therefore, it is not possible to use vMotion<sup>®</sup> to migrate virtual machines between systems running on processors with different feature sets. You can avoid this restriction, in some cases, by using Enhanced vMotion Compatibility (EVC) with processors that support this feature. See the *vCenter Server and Host Management* documentation for more information.

## Performance Implications of CPU Virtualization

CPU virtualization adds varying amounts of overhead depending on the workload and the type of virtualization used.

An application is CPU-bound if it spends most of its time executing instructions rather than waiting for external events such as user interaction, device input, or data retrieval. For such applications, the CPU virtualization overhead includes the additional instructions that must be executed. This overhead takes CPU processing time that the application itself can use. CPU virtualization overhead usually translates into a reduction in overall performance.

For applications that are not CPU-bound, CPU virtualization likely translates into an increase in CPU use. If spare CPU capacity is available to absorb the overhead, it can still deliver comparable performance in terms of overall throughput.

ESXi supports up to 128 virtual processors (CPUs) for each virtual machine.

---

**NOTE** Deploy single-threaded applications on uniprocessor virtual machines, instead of on SMP virtual machines that have multiple CPUs, for the best performance and resource use.

Single-threaded applications can take advantage only of a single CPU. Deploying such applications in dual-processor virtual machines does not speed up the application. Instead, it causes the second virtual CPU to use physical resources that other virtual machines could otherwise use.

---

# Administering CPU Resources

---

You can configure virtual machines with one or more virtual processors, each with its own set of registers and control structures.

When a virtual machine is scheduled, its virtual processors are scheduled to run on physical processors. The VMkernel Resource Manager schedules the virtual CPUs on physical CPUs, thereby managing the virtual machine's access to physical CPU resources. ESXi supports virtual machines with up to 128 virtual CPUs.

This chapter includes the following topics:

- [“View Processor Information,”](#) on page 19
- [“Specifying CPU Configuration,”](#) on page 19
- [“Multicore Processors,”](#) on page 20
- [“Hyperthreading,”](#) on page 20
- [“Using CPU Affinity,”](#) on page 22
- [“Host Power Management Policies,”](#) on page 23

## View Processor Information

You can access information about current CPU configuration in the vSphere Web Client.

### Procedure

- 1 Browse to the host in the vSphere Web Client navigator.
- 2 Click the **Manage** tab and click **Settings**.
- 3 Select **Processors** to view the information about the number and type of physical processors and the number of logical processors.

---

**NOTE** In hyperthreaded systems, each hardware thread is a logical processor. For example, a dual-core processor with hyperthreading enabled has two cores and four logical processors.

---

## Specifying CPU Configuration

You can specify CPU configuration to improve resource management. However, if you do not customize CPU configuration, the ESXi host uses defaults that work well in most situations.

You can specify CPU configuration in the following ways:

- Use the attributes and special features available through the vSphere Web Client. The vSphere Web Client allows you to connect to the ESXi host or a vCenter Server system.

- Use advanced settings under certain circumstances.
- Use the vSphere SDK for scripted CPU allocation.
- Use hyperthreading.

## Multicore Processors

Multicore processors provide many advantages for a host performing multitasking of virtual machines.

Intel and AMD have each developed processors which combine two or more processor cores into a single integrated circuit (often called a package or socket). VMware uses the term socket to describe a single package which can have one or more processor cores with one or more logical processors in each core.

A dual-core processor, for example, can provide almost double the performance of a single-core processor, by allowing two virtual CPUs to execute at the same time. Cores within the same processor are typically configured with a shared last-level cache used by all cores, potentially reducing the need to access slower main memory. A shared memory bus that connects a physical processor to main memory can limit performance of its logical processors if the virtual machines running on them are running memory-intensive workloads which compete for the same memory bus resources.

Each logical processor of each processor core can be used independently by the ESXi CPU scheduler to execute virtual machines, providing capabilities similar to SMP systems. For example, a two-way virtual machine can have its virtual processors running on logical processors that belong to the same core, or on logical processors on different physical cores.

The ESXi CPU scheduler can detect the processor topology and the relationships between processor cores and the logical processors on them. It uses this information to schedule virtual machines and optimize performance.

The ESXi CPU scheduler can interpret processor topology, including the relationship between sockets, cores, and logical processors. The scheduler uses topology information to optimize the placement of virtual CPUs onto different sockets to maximize overall cache utilization, and to improve cache affinity by minimizing virtual CPU migrations.

In some cases, such as when an SMP virtual machine exhibits significant data sharing between its virtual CPUs, this default behavior might be sub-optimal. For such workloads, it can be beneficial to schedule all of the virtual CPUs on the same socket, with a shared last-level cache, even when the ESXi host is undercommitted. In such scenarios, you can override the default behavior of spreading virtual CPUs across packages by including the following configuration option in the virtual machine's `.vmx` configuration file: `sched.cpu.vsmptConsolidate="TRUE"`.

## Hyperthreading

Hyperthreading technology allows a single physical processor core to behave like two logical processors. The processor can run two independent applications at the same time. To avoid confusion between logical and physical processors, Intel refers to a physical processor as a socket, and the discussion in this chapter uses that terminology as well.

Intel Corporation developed hyperthreading technology to enhance the performance of its Pentium IV and Xeon processor lines. Hyperthreading technology allows a single processor core to execute two independent threads simultaneously.

While hyperthreading does not double the performance of a system, it can increase performance by better utilizing idle resources leading to greater throughput for certain important workload types. An application running on one logical processor of a busy core can expect slightly more than half of the throughput that it obtains while running alone on a non-hyperthreaded processor. Hyperthreading performance improvements are highly application-dependent, and some applications might see performance degradation with hyperthreading because many processor resources (such as the cache) are shared between logical processors.

---

**NOTE** On processors with Intel Hyper-Threading technology, each core can have two logical processors which share most of the core's resources, such as memory caches and functional units. Such logical processors are usually called threads.

---

Many processors do not support hyperthreading and as a result have only one thread per core. For such processors, the number of cores also matches the number of logical processors. The following processors support hyperthreading and have two threads per core.

- Processors based on the Intel Xeon 5500 processor microarchitecture.
- Intel Pentium 4 (HT-enabled)
- Intel Pentium EE 840 (HT-enabled)

## Hyperthreading and ESXi Hosts

A host that is enabled for hyperthreading should behave similarly to a host without hyperthreading. You might need to consider certain factors if you enable hyperthreading, however.

ESXi hosts manage processor time intelligently to guarantee that load is spread smoothly across processor cores in the system. Logical processors on the same core have consecutive CPU numbers, so that CPUs 0 and 1 are on the first core together, CPUs 2 and 3 are on the second core, and so on. Virtual machines are preferentially scheduled on two different cores rather than on two logical processors on the same core.

If there is no work for a logical processor, it is put into a halted state, which frees its execution resources and allows the virtual machine running on the other logical processor on the same core to use the full execution resources of the core. The VMware scheduler properly accounts for this halt time, and charges a virtual machine running with the full resources of a core more than a virtual machine running on a half core. This approach to processor management ensures that the server does not violate any of the standard ESXi resource allocation rules.

Consider your resource management needs before you enable CPU affinity on hosts using hyperthreading. For example, if you bind a high priority virtual machine to CPU 0 and another high priority virtual machine to CPU 1, the two virtual machines have to share the same physical core. In this case, it can be impossible to meet the resource demands of these virtual machines. Ensure that any custom affinity settings make sense for a hyperthreaded system.

## Enable Hyperthreading

To enable hyperthreading, you must first enable it in your system's BIOS settings and then turn it on in the vSphere Web Client. Hyperthreading is enabled by default.

Consult your system documentation to determine whether your CPU supports hyperthreading.

### Procedure

- 1 Ensure that your system supports hyperthreading technology.
- 2 Enable hyperthreading in the system BIOS.

Some manufacturers label this option **Logical Processor**, while others call it **Enable Hyperthreading**.

- 3 Ensure that hyperthreading is enabled for the ESXi host.
  - a Browse to the host in the vSphere Web Client navigator.
  - b Click the **Manage** tab and click **Settings**.
  - c Under **System**, click **Advanced System Settings** and select **VMkernel.Boot.hyperthreading**.  
Hyperthreading is enabled if the value is **true**.
- 4 Under **Hardware**, click **Processors** to view the number of Logical processors.  
Hyperthreading is enabled.

## Using CPU Affinity

By specifying a CPU affinity setting for each virtual machine, you can restrict the assignment of virtual machines to a subset of the available processors in multiprocessor systems. By using this feature, you can assign each virtual machine to processors in the specified affinity set.

CPU affinity specifies virtual machine-to-processor placement constraints and is different from the relationship created by a VM-VM or VM-Host affinity rule, which specifies virtual machine-to-virtual machine host placement constraints.

In this context, the term CPU refers to a logical processor on a hyperthreaded system and refers to a core on a non-hyperthreaded system.

The CPU affinity setting for a virtual machine applies to all of the virtual CPUs associated with the virtual machine and to all other threads (also known as worlds) associated with the virtual machine. Such virtual machine threads perform processing required for emulating mouse, keyboard, screen, CD-ROM, and miscellaneous legacy devices.

In some cases, such as display-intensive workloads, significant communication might occur between the virtual CPUs and these other virtual machine threads. Performance might degrade if the virtual machine's affinity setting prevents these additional threads from being scheduled concurrently with the virtual machine's virtual CPUs. Examples of this include a uniprocessor virtual machine with affinity to a single CPU or a two-way SMP virtual machine with affinity to only two CPUs.

For the best performance, when you use manual affinity settings, VMware recommends that you include at least one additional physical CPU in the affinity setting to allow at least one of the virtual machine's threads to be scheduled at the same time as its virtual CPUs. Examples of this include a uniprocessor virtual machine with affinity to at least two CPUs or a two-way SMP virtual machine with affinity to at least three CPUs.

## Assign a Virtual Machine to a Specific Processor

Using CPU affinity, you can assign a virtual machine to a specific processor. This allows you to restrict the assignment of virtual machines to a specific available processor in multiprocessor systems.

### Procedure

- 1 Find the virtual machine in the vSphere Web Client inventory.
  - a To find a virtual machine, select a data center, folder, cluster, resource pool, or host.
  - b Click the **Related Objects** tab and click **Virtual Machines**.
- 2 Right-click the virtual machine and click **Edit Settings**.
- 3 Under Virtual Hardware, expand **CPU**.

- 4 Under Scheduling Affinity, select physical processor affinity for the virtual machine.  
Use '-' for ranges and ',' to separate values.  
For example, "0, 2, 4-7" would indicate processors 0, 2, 4, 5, 6 and 7.
- 5 Select the processors where you want the virtual machine to run and click **OK**.

## Potential Issues with CPU Affinity

Before you use CPU affinity, you might need to consider certain issues.

Potential issues with CPU affinity include:

- For multiprocessor systems, ESXi systems perform automatic load balancing. Avoid manual specification of virtual machine affinity to improve the scheduler's ability to balance load across processors.
- Affinity can interfere with the ESXi host's ability to meet the reservation and shares specified for a virtual machine.
- Because CPU admission control does not consider affinity, a virtual machine with manual affinity settings might not always receive its full reservation.

Virtual machines that do not have manual affinity settings are not adversely affected by virtual machines with manual affinity settings.

- When you move a virtual machine from one host to another, affinity might no longer apply because the new host might have a different number of processors.
- The NUMA scheduler might not be able to manage a virtual machine that is already assigned to certain processors using affinity.
- Affinity can affect the host's ability to schedule virtual machines on multicore or hyperthreaded processors to take full advantage of resources shared on such processors.

## Host Power Management Policies

ESXi can take advantage of several power management features that the host hardware provides to adjust the trade-off between performance and power use. You can control how ESXi uses these features by selecting a power management policy.

In general, selecting a high-performance policy provides more absolute performance, but at lower efficiency (performance per watt). Lower-power policies provide less absolute performance, but at higher efficiency.

ESXi provides five power management policies. If the host does not support power management, or if the BIOS settings specify that the host operating system is not allowed to manage power, only the Not Supported policy is available.

You select a policy for a host using the vSphere Web Client. If you do not select a policy, ESXi uses Balanced by default.

**Table 4-1.** CPU Power Management Policies

Power Management Policy	Description
Not supported	The host does not support any power management features or power management is not enabled in the BIOS.
High Performance	The VMkernel detects certain power management features, but will not use them unless the BIOS requests them for power capping or thermal events.
Balanced (Default)	The VMkernel uses the available power management features conservatively to reduce host energy consumption with minimal compromise to performance.

**Table 4-1.** CPU Power Management Policies (Continued)

Power Management Policy	Description
Low Power	The VMkernel aggressively uses available power management features to reduce host energy consumption at the risk of lower performance.
Custom	The VMkernel bases its power management policy on the values of several advanced configuration parameters. You can set these parameters in the vSphere Web Client Advanced Settings dialog box.

When a CPU runs at lower frequency, it can also run at lower voltage, which saves power. This type of power management is typically called Dynamic Voltage and Frequency Scaling (DVFS). ESXi attempts to adjust CPU frequencies so that virtual machine performance is not affected.

When a CPU is idle, ESXi can take advantage of deep halt states (known as C-states). The deeper the C-state, the less power the CPU uses, but the longer it takes for the CPU to resume running. When a CPU becomes idle, ESXi applies an algorithm to predict how long it will be in an idle state and chooses an appropriate C-state to enter. In power management policies that do not use deep C-states, ESXi uses only the shallowest halt state (C1) for idle CPUs.

## Select a CPU Power Management Policy

You set the CPU power management policy for a host using the vSphere Web Client.

### Prerequisites

Verify that the BIOS settings on the host system allow the operating system to control power management (for example, **OS Controlled**).

---

**NOTE** Some systems have Processor Clocking Control (PCC) technology, which allows ESXi to manage power on the host system even if the host BIOS settings do not specify OS Controlled mode. With this technology, ESXi does not manage P-states directly. Instead, the host cooperates with the BIOS to determine the processor clock rate. HP systems that support this technology have a BIOS setting called Cooperative Power Management that is enabled by default.

---

If the host hardware does not allow the operating system to manage power, only the Not Supported policy is available. (On some systems, only the High Performance policy is available.)

### Procedure

- 1 Browse to the host in the vSphere Web Client navigator.
- 2 Click the **Manage** tab and click **Settings**.
- 3 Under Hardware, select **Power Management** and click the **Edit** button.
- 4 Select a power management policy for the host and click **OK**.

The policy selection is saved in the host configuration and can be used again at boot time. You can change it at any time, and it does not require a server reboot.

## Configure Custom Policy Parameters for Host Power Management

When you use the Custom policy for host power management, ESXi bases its power management policy on the values of several advanced configuration parameters.

### Prerequisites

Select **Custom** for the power management policy, as described in [“Select a CPU Power Management Policy,”](#) on page 24.



**Procedure**

- 1 Browse to the host in the vSphere Web Client navigator.
- 2 Click the **Manage** tab and click **Settings**.
- 3 Under System, select **Advanced System Settings**.
- 4 In the right pane, you can edit the power management parameters that affect the Custom policy.

Power management parameters that affect the Custom policy have descriptions that begin with **In Custom policy**. All other power parameters affect all power management policies.

- 5 Select the parameter and click the **Edit** button.

---

**NOTE** The default values of power management parameters match the Balanced policy.

---

Parameter	Description
<b>Power.UsePStates</b>	Use ACPI P-states to save power when the processor is busy.
<b>Power.MaxCpuLoad</b>	Use P-states to save power on a CPU only when the CPU is busy for less than the given percentage of real time.
<b>Power.MinFreqPct</b>	Do not use any P-states slower than the given percentage of full CPU speed.
<b>Power.UseStallCtr</b>	Use a deeper P-state when the processor is frequently stalled waiting for events such as cache misses.
<b>Power.TimerHz</b>	Controls how many times per second ESXi reevaluates which P-state each CPU should be in.
<b>Power.UseCStates</b>	Use deep ACPI C-states (C2 or below) when the processor is idle.
<b>Power.CStateMaxLatency</b>	Do not use C-states whose latency is greater than this value.
<b>Power.CStateResidencyCoef</b>	When a CPU becomes idle, choose the deepest C-state whose latency multiplied by this value is less than the host's prediction of how long the CPU will remain idle. Larger values make ESXi more conservative about using deep C-states, while smaller values are more aggressive.
<b>Power.CStatePredictionCoef</b>	A parameter in the ESXi algorithm for predicting how long a CPU that becomes idle will remain idle. Changing this value is not recommended.
<b>Power.PerfBias</b>	Performance Energy Bias Hint (Intel-only). Sets an MSR on Intel processors to an Intel-recommended value. Intel recommends 0 for high performance, 6 for balanced, and 15 for low power. Other values are undefined.

- 6 Click **OK**.



# Memory Virtualization Basics

---

Before you manage memory resources, you should understand how they are being virtualized and used by ESXi.

The VMkernel manages all physical RAM on the host. The VMkernel dedicates part of this managed physical RAM for its own use. The rest is available for use by virtual machines.

The virtual and physical memory space is divided into blocks called pages. When physical memory is full, the data for virtual pages that are not present in physical memory are stored on disk. Depending on processor architecture, pages are typically 4 KB or 2 MB. See [“Advanced Memory Attributes,”](#) on page 114.

This chapter includes the following topics:

- [“Virtual Machine Memory,”](#) on page 27
- [“Memory Overcommitment,”](#) on page 28
- [“Memory Sharing,”](#) on page 28
- [“Types of Memory Virtualization,”](#) on page 29

## Virtual Machine Memory

Each virtual machine consumes memory based on its configured size, plus additional overhead memory for virtualization.

The configured size is the amount of memory that is presented to the guest operating system. This is different from the amount of physical RAM that is allocated to the virtual machine. The latter depends on the resource settings (shares, reservation, limit) and the level of memory pressure on the host.

For example, consider a virtual machine with a configured size of 1GB. When the guest operating system boots, it detects that it is running on a dedicated machine with 1GB of physical memory. In some cases, the virtual machine might be allocated the full 1GB. In other cases, it might receive a smaller allocation. Regardless of the actual allocation, the guest operating system continues to behave as though it is running on a dedicated machine with 1GB of physical memory.

<b>Shares</b>	Specify the relative priority for a virtual machine if more than the reservation is available.
<b>Reservation</b>	Is a guaranteed lower bound on the amount of physical RAM that the host reserves for the virtual machine, even when memory is overcommitted. Set the reservation to a level that ensures the virtual machine has sufficient memory to run efficiently, without excessive paging.

After a virtual machine consumes all of the memory within its reservation, it is allowed to retain that amount of memory and this memory is not reclaimed, even if the virtual machine becomes idle. Some guest operating systems (for example, Linux) might not access all of the configured memory immediately after booting. Until the virtual machines consumes all of the memory within its reservation, VMkernel can allocate any unused portion of its reservation to other virtual machines. However, after the guest's workload increases and the virtual machine consumes its full reservation, it is allowed to keep this memory.

### Limit

Is an upper bound on the amount of physical RAM that the host can allocate to the virtual machine. The virtual machine's memory allocation is also implicitly limited by its configured size.

## Memory Overcommitment

For each running virtual machine, the system reserves physical RAM for the virtual machine's reservation (if any) and for its virtualization overhead.

The total configured memory sizes of all virtual machines may exceed the amount of available physical memory on the host. However, it doesn't necessarily mean memory is overcommitted. Memory is overcommitted when the combined working memory footprint of all virtual machines exceed that of the host memory sizes.

Because of the memory management techniques the ESXi host uses, your virtual machines can use more virtual RAM than there is physical RAM available on the host. For example, you can have a host with 2GB memory and run four virtual machines with 1GB memory each. In that case, the memory is overcommitted. For instance, if all four virtual machines are idle, the combined consumed memory may be well below 2GB. However, if all 4GB virtual machines are actively consuming memory, then their memory footprint may exceed 2GB and the ESXi host will become overcommitted.

Overcommitment makes sense because, typically, some virtual machines are lightly loaded while others are more heavily loaded, and relative activity levels vary over time.

To improve memory utilization, the ESXi host transfers memory from idle virtual machines to virtual machines that need more memory. Use the Reservation or Shares parameter to preferentially allocate memory to important virtual machines. This memory remains available to other virtual machines if it is not in use. ESXi implements various mechanisms such as ballooning, memory sharing, memory compression and swapping to provide reasonable performance even if the host is not heavily memory overcommitted.

An ESXi host can run out of memory if virtual machines consume all reservable memory in a memory overcommitted environment. Although the powered on virtual machines are not affected, a new virtual machine might fail to power on due to lack of memory.

---

**NOTE** All virtual machine memory overhead is also considered reserved.

---

In addition, memory compression is enabled by default on ESXi hosts to improve virtual machine performance when memory is overcommitted as described in "[Memory Compression](#)," on page 41.

## Memory Sharing

Memory sharing is a proprietary ESXi technique that can help achieve greater memory density on a host.

Memory sharing relies on the observation that several virtual machines might be running instances of the same guest operating system, have the same applications or components loaded, or contain common data. In such cases, a host uses a proprietary Transparent Page Sharing (TPS) technique to eliminate redundant copies of memory pages. With memory sharing, a workload running in virtual machines often consumes

less memory than it would when running on physical machines. As a result, higher levels of overcommitment can be supported efficiently. The amount of memory saved by memory sharing depends on whether the workload consists of nearly identical machines which might free up more memory, while a more diverse workload might result in a significantly lower percentage of memory savings.

**NOTE** Due to security concerns, inter-virtual machine transparent page sharing is disabled by default and page sharing is being restricted to intra-virtual machine memory sharing. This means page sharing does not occur across virtual machines and only occurs inside of a virtual machine. See [“Sharing Memory Across Virtual Machines,”](#) on page 40 for more information.

## Types of Memory Virtualization

There are two types of memory virtualization: Software-based and hardware-assisted memory virtualization.

Because of the extra level of memory mapping introduced by virtualization, ESXi can effectively manage memory across all virtual machines. Some of the physical memory of a virtual machine might be mapped to shared pages or to pages that are unmapped, or swapped out.

A host performs virtual memory management without the knowledge of the guest operating system and without interfering with the guest operating system’s own memory management subsystem.


The VMM for each virtual machine maintains a mapping from the guest operating system’s physical memory pages to the physical memory pages on the underlying machine. (VMware refers to the underlying host physical pages as “machine” pages and the guest operating system’s physical pages as “physical” pages.)

Each virtual machine sees a contiguous, zero-based, addressable physical memory space. The underlying machine memory on the server used by each virtual machine is not necessarily contiguous.

For both software-based and hardware-assisted memory virtualization, the guest virtual to guest physical addresses are managed by the guest operating system. The hypervisor is only responsible for translating the guest physical addresses to machine addresses. Software-based memory virtualization combines the guest’s virtual to machine addresses in software and saves them in the shadow page tables managed by the hypervisor. Hardware-assisted memory virtualization utilizes the hardware facility to generate the combined mappings with the guest’s page tables and the nested page tables maintained by the hypervisor.

The diagram illustrates the ESXi implementation of memory virtualization.

**Figure 5-1.** ESXi Memory Mapping



- The boxes represent pages, and the arrows show the different memory mappings.
- The arrows from guest virtual memory to guest physical memory show the mapping maintained by the page tables in the guest operating system. (The mapping from virtual memory to linear memory for x86-architecture processors is not shown.)
- The arrows from guest physical memory to machine memory show the mapping maintained by the VMM.

- The dashed arrows show the mapping from guest virtual memory to machine memory in the shadow page tables also maintained by the VMM. The underlying processor running the virtual machine uses the shadow page table mappings.

## Software-Based Memory Virtualization

ESXi virtualizes guest physical memory by adding an extra level of address translation.

- The VMM maintains the combined virtual-to-machine page mappings in the shadow page tables. The shadow page tables are kept up to date with the guest operating system's virtual-to-physical mappings and physical-to-machine mappings maintained by the VMM.
- The VMM intercepts virtual machine instructions that manipulate guest operating system memory management structures so that the actual memory management unit (MMU) on the processor is not updated directly by the virtual machine.
- The shadow page tables are used directly by the processor's paging hardware.
- There is non-trivial computation overhead for maintaining the coherency of the shadow page tables. The overhead is more pronounced when the number of virtual CPUs increases.

This approach to address translation allows normal memory accesses in the virtual machine to execute without adding address translation overhead, after the shadow page tables are set up. Because the translation look-aside buffer (TLB) on the processor caches direct virtual-to-machine mappings read from the shadow page tables, no additional overhead is added by the VMM to access the memory. Note that software MMU has a higher overhead memory requirement than hardware MMU. Hence, in order to support software MMU, the maximum overhead supported for virtual machines in the VMkernel needs to be increased. In some cases, software memory virtualization may have some performance benefit over hardware-assisted approach if the workload induces a huge amount of TLB misses.

## Performance Considerations

The use of two sets of page tables has these performance implications.

- No overhead is incurred for regular guest memory accesses.
- Additional time is required to map memory within a virtual machine, which happens when:
  - The virtual machine operating system is setting up or updating virtual address to physical address mappings.
  - The virtual machine operating system is switching from one address space to another (context switch).
- Like CPU virtualization, memory virtualization overhead depends on workload.

## Hardware-Assisted Memory Virtualization

Some CPUs, such as AMD SVM-V and the Intel Xeon 5500 series, provide hardware support for memory virtualization by using two layers of page tables.

The first layer of page tables stores guest virtual-to-physical translations, while the second layer of page tables stores guest physical-to-machine translation. The TLB (translation look-aside buffer) is a cache of translations maintained by the processor's memory management unit (MMU) hardware. A TLB miss is a miss in this cache and the hardware needs to go to memory (possibly many times) to find the required translation. For a TLB miss to a certain guest virtual address, the hardware looks at both page tables to translate guest virtual address to machine address. The first layer of page tables is maintained by the guest operating system. The VMM only maintains the second layer of page tables.

## Performance Considerations

When you use hardware assistance, you eliminate the overhead for software memory virtualization. In particular, hardware assistance eliminates the overhead required to keep shadow page tables in synchronization with guest page tables. However, the TLB miss latency when using hardware assistance is significantly higher. By default the hypervisor uses large pages in hardware assisted modes to reduce the cost of TLB misses. As a result, whether or not a workload benefits by using hardware assistance primarily depends on the overhead the memory virtualization causes when using software memory virtualization. If a workload involves a small amount of page table activity (such as process creation, mapping the memory, or context switches), software virtualization does not cause significant overhead. Conversely, workloads with a large amount of page table activity are likely to benefit from hardware assistance.

The performance of hardware MMU has improved since it was first introduced with extensive caching implemented in hardware. Using software memory virtualization techniques, the frequency of context switches in a typical guest may happen from 100 to 1000 times per second. Each context switch will trap the VMM in software MMU. Hardware MMU approaches avoid this issue.

By default the hypervisor uses large pages in hardware assisted modes to reduce the cost of TLB misses. The best performance is achieved by using large pages in both guest virtual to guest physical and guest physical to machine address translations.

The option `LPage.LPageAlwaysTryForNPT` can change the policy for using large pages in guest physical to machine address translations. For more information, see [“Advanced Memory Attributes,”](#) on page 114.

---

**NOTE** Binary translation only works with software-based memory virtualization.

---





# Administering Memory Resources

---

Using the vSphere Web Client you can view information about and make changes to memory allocation settings. To administer your memory resources effectively, you must also be familiar with memory overhead, idle memory tax, and how ESXi hosts reclaim memory.

When administering memory resources, you can specify memory allocation. If you do not customize memory allocation, the ESXi host uses defaults that work well in most situations.

You can specify memory allocation in several ways.

- Use the attributes and special features available through the vSphere Web Client. The vSphere Web Client allows you to connect to the ESXi host or vCenter Server system.
- Use advanced settings.
- Use the vSphere SDK for scripted memory allocation.

This chapter includes the following topics:

- [“Understanding Memory Overhead,”](#) on page 33
- [“How ESXi Hosts Allocate Memory,”](#) on page 34
- [“Memory Reclamation,”](#) on page 35
- [“Using Swap Files,”](#) on page 36
- [“Sharing Memory Across Virtual Machines,”](#) on page 40
- [“Memory Compression,”](#) on page 41
- [“Measuring and Differentiating Types of Memory Usage,”](#) on page 42
- [“Memory Reliability,”](#) on page 43
- [“About System Swap,”](#) on page 43

## Understanding Memory Overhead

Virtualization of memory resources has some associated overhead.

ESXi virtual machines can incur two kinds of memory overhead.

- The additional time to access memory within a virtual machine.
- The extra space needed by the ESXi host for its own code and data structures, beyond the memory allocated to each virtual machine.

ESXi memory virtualization adds little time overhead to memory accesses. Because the processor's paging hardware uses page tables (shadow page tables for software-based approach or two level page tables for hardware-assisted approach) directly, most memory accesses in the virtual machine can execute without address translation overhead.

The memory space overhead has two components.

- A fixed, system-wide overhead for the VMkernel.
- Additional overhead for each virtual machine.

Overhead memory includes space reserved for the virtual machine frame buffer and various virtualization data structures, such as shadow page tables. Overhead memory depends on the number of virtual CPUs and the configured memory for the guest operating system.

## Overhead Memory on Virtual Machines

Virtual machines require a certain amount of available overhead memory to power on. You should be aware of the amount of this overhead.

The amount of overhead memory needed for a virtual machine depends on a large number of factors, including the number of vCPUs and memory size, the number and types of devices, the execution mode that the monitor is using and the hardware version of the virtual machine. The version of vSphere you are using can also effect the amount of memory needed. VMX automatically calculates the amount of overhead memory needed for a virtual machine.

In order to find out how much overhead memory is needed for your specific configuration, first power on the virtual machine in question. Look in the vmware.log file. When the virtual machine powers on, the amount of overhead memory it needs is printed to the log. Search within the log for VMEM to see the initial and precise amount of overhead memory reserved for the virtual machine.

## How ESXi Hosts Allocate Memory

A host allocates the memory specified by the `Limit` parameter to each virtual machine, unless memory is overcommitted. ESXi never allocates more memory to a virtual machine than its specified physical memory size.

For example, a 1GB virtual machine might have the default limit (unlimited) or a user-specified limit (for example 2GB). In both cases, the ESXi host never allocates more than 1GB, the physical memory size that was specified for it.

When memory is overcommitted, each virtual machine is allocated an amount of memory somewhere between what is specified by **Reservation** and what is specified by **Limit**. The amount of memory granted to a virtual machine above its reservation usually varies with the current memory load.

A host determines allocations for each virtual machine based on the number of shares allocated to it and an estimate of its recent working set size.

- Shares — ESXi hosts use a modified proportional-share memory allocation policy. Memory shares entitle a virtual machine to a fraction of available physical memory.
- Working set size — ESXi hosts estimate the working set for a virtual machine by monitoring memory activity over successive periods of virtual machine execution time. Estimates are smoothed over several time periods using techniques that respond rapidly to increases in working set size and more slowly to decreases in working set size.

This approach ensures that a virtual machine from which idle memory is reclaimed can ramp up quickly to its full share-based allocation when it starts using its memory more actively.

Memory activity is monitored to estimate the working set sizes for a default period of 60 seconds. To modify this default, adjust the `Mem.SamplePeriod` advanced setting. See [“Set Advanced Host Attributes,”](#) on page 113.

## Memory Tax for Idle Virtual Machines

If a virtual machine is not actively using all of its currently allocated memory, ESXi charges more for idle memory than for memory that is in use. This is done to help prevent virtual machines from hoarding idle memory.

The idle memory tax is applied in a progressive fashion. The effective tax rate increases as the ratio of idle memory to active memory for the virtual machine rises. (In earlier versions of ESXi that did not support hierarchical resource pools, all idle memory for a virtual machine was taxed equally.)

You can modify the idle memory tax rate with the `Mem.IdleTax` option. Use this option, together with the `Mem.SamplePeriod` advanced attribute, to control how the system determines target memory allocations for virtual machines. See [“Set Advanced Host Attributes,”](#) on page 113.

---

**NOTE** In most cases, changes to `Mem.IdleTax` are not necessary nor appropriate.

---

## VMX Swap Files

Virtual machine executable (VMX) swap files allow the host to greatly reduce the amount of overhead memory reserved for the VMX process.

---

**NOTE** VMX swap files are not related to the swap to host swap cache feature or to regular host-level swap files.

---

ESXi reserves memory per virtual machine for a variety of purposes. Memory for the needs of certain components, such as the virtual machine monitor (VMM) and virtual devices, is fully reserved when a virtual machine is powered on. However, some of the overhead memory that is reserved for the VMX process can be swapped. The VMX swap feature reduces the VMX memory reservation significantly (for example, from about 50MB or more per virtual machine to about 10MB per virtual machine). This allows the remaining memory to be swapped out when host memory is overcommitted, reducing overhead memory reservation for each virtual machine.

The host creates VMX swap files automatically, provided there is sufficient free disk space at the time a virtual machine is powered on.

## Memory Reclamation

ESXi hosts can reclaim memory from virtual machines.

A host allocates the amount of memory specified by a reservation directly to a virtual machine. Anything beyond the reservation is allocated using the host's physical resources or, when physical resources are not available, handled using special techniques such as ballooning or swapping. Hosts can use two techniques for dynamically expanding or contracting the amount of memory allocated to virtual machines.


- ESXi systems use a memory balloon driver (`vmmemctl`), loaded into the guest operating system running in a virtual machine. See [“Memory Balloon Driver,”](#) on page 36.
- ESXi system swaps out a page from a virtual machine to a server swap file without any involvement by the guest operating system. Each virtual machine has its own swap file.

## Memory Balloon Driver

The memory balloon driver (`vmmemctl`) collaborates with the server to reclaim pages that are considered least valuable by the guest operating system.

The driver uses a proprietary ballooning technique that provides predictable performance that closely matches the behavior of a native system under similar memory constraints. This technique increases or decreases memory pressure on the guest operating system, causing the guest to use its own native memory management algorithms. When memory is tight, the guest operating system determines which pages to reclaim and, if necessary, swaps them to its own virtual disk.

**Figure 6-1.** Memory Ballooning in the Guest Operating System




---

**NOTE** You must configure the guest operating system with sufficient swap space. Some guest operating systems have additional limitations.

---

If necessary, you can limit the amount of memory `vmmemctl` reclaims by setting the `sched.mem.maxmemctl` parameter for a specific virtual machine. This option specifies the maximum amount of memory that can be reclaimed from a virtual machine in megabytes (MB). See [“Set Advanced Virtual Machine Attributes,”](#) on page 116.

## Using Swap Files

You can specify the location of your guest swap file, reserve swap space when memory is overcommitted, and delete a swap file.

ESXi hosts use swapping to forcibly reclaim memory from a virtual machine when the `vmmemctl` driver is not available or is not responsive.

- It was never installed.
- It is explicitly disabled.
- It is not running (for example, while the guest operating system is booting).
- It is temporarily unable to reclaim memory quickly enough to satisfy current system demands.

- It is functioning properly, but maximum balloon size is reached.

Standard demand-paging techniques swap pages back in when the virtual machine needs them.

## Swap File Location

By default, the swap file is created in the same location as the virtual machine's configuration file, which could either be on a VMFS datastore, a vSAN datastore or a VVol datastore. On a vSAN datastore or a VVol datastore, the swap file is created as a separate vSAN or VVol object.

A swap file is created by the ESXi host when a virtual machine is powered on. If this file cannot be created, the virtual machine cannot power on. Instead of accepting the default, you can also:

- Use per-virtual machine configuration options to change the datastore to another shared storage location.
- Use host-local swap, which allows you to specify a datastore stored locally on the host. This allows you to swap at a per-host level, saving space on the SAN. However, it can lead to a slight degradation in performance for vSphere vMotion because pages swapped to a local swap file on the source host must be transferred across the network to the destination host. Please note that currently vSAN and VVol datastores cannot be specified for host-local swap.

## Enable Host-Local Swap for a DRS Cluster

Host-local swap allows you to specify a datastore stored locally on the host as the swap file location. You can enable host-local swap for a DRS cluster.

### Procedure

- 1 Browse to the cluster in the vSphere Web Client navigator.
- 2 Click the **Manage** tab and click **Settings**.
- 3 Under Configuration, click **General** to view the swap file location and click **Edit** to change it.
- 4 Select the **Datastore specified by host** option and click **OK**.
- 5 Browse to one of the hosts in the cluster in the vSphere Web Client navigator.
- 6 Click the **Manage** tab and click **Settings**.
- 7 Under Virtual Machines, select **Virtual Machine Swapfile Location**.
- 8 Click **Edit** and select the local datastore to use and click **OK**.
- 9 Repeat [Step 5](#) through [Step 8](#) for each host in the cluster.

Host-local swap is now enabled for the DRS cluster.

## Enable Host-Local Swap for a Standalone Host

Host-local swap allows you to specify a datastore stored locally on the host as the swap file location. You can enable host-local swap for a standalone host.

### Procedure

- 1 Browse to the host in the vSphere Web Client navigator.
- 2 Click the **Manage** tab and click **Settings**.
- 3 Under Virtual Machines, select **Virtual Machine Swapfile Location**.
- 4 Click **Edit** and select **Selected Datastore**.
- 5 Select a local datastore from the list and click **OK**.

Host-local swap is now enabled for the standalone host.

## Swap Space and Memory Overcommitment

You must reserve swap space for any unreserved virtual machine memory (the difference between the reservation and the configured memory size) on per-virtual machine swap files.

This swap reservation is required to ensure that the ESXi host is able to preserve virtual machine memory under any circumstances. In practice, only a small fraction of the host-level swap space might be used.

If you are overcommitting memory with ESXi, to support the intra-guest swapping induced by ballooning, ensure that your guest operating systems also have sufficient swap space. This guest-level swap space must be greater than or equal to the difference between the virtual machine's configured memory size and its Reservation.



**CAUTION** If memory is overcommitted, and the guest operating system is configured with insufficient swap space, the guest operating system in the virtual machine can fail.

---

To prevent virtual machine failure, increase the size of the swap space in your virtual machines.

- Windows guest operating systems— Windows operating systems refer to their swap space as paging files. Some Windows operating systems try to increase the size of paging files automatically, if there is sufficient free disk space.

See your Microsoft Windows documentation or search the Windows help files for “paging files.” Follow the instructions for changing the size of the virtual memory paging file.

- Linux guest operating system — Linux operating systems refer to their swap space as swap files. For information on increasing swap files, see the following Linux man pages:
  - `mkswap` — Sets up a Linux swap area.
  - `swapon` — Enables devices and files for paging and swapping.

Guest operating systems with a lot of memory and small virtual disks (for example, a virtual machine with 8GB RAM and a 2GB virtual disk) are more susceptible to having insufficient swap space.

---

**NOTE** Do not store swap files on thin-provisioned LUNs. Running a virtual machine with a swap file that is stored on a thin-provisioned LUN can cause swap file growth failure, which can lead to termination of the virtual machine.

---

When you create a large swap file (for example, larger than 100GB), the amount of time it takes for the virtual machine to power on can increase significantly. To avoid this, set a high reservation for large virtual machines.

You can also place swap files on less costly storage using host-local swap files.

## Configure Virtual Machine Swapfile Properties for the Host

Configure a swapfile location for the host to determine the default location for virtual machine swapfiles in the vSphere Web Client.

By default, swapfiles for a virtual machine are located on a datastore in the folder that contains the other virtual machine files. However, you can configure your host to place virtual machine swapfiles on an alternative datastore.

You can use this option to place virtual machine swapfiles on lower-cost or higher-performance storage. You can also override this host-level setting for individual virtual machines.

Setting an alternative swapfile location might cause migrations with vMotion to complete more slowly. For best vMotion performance, store the virtual machine on a local datastore rather than in the same directory as the virtual machine swapfiles. If the virtual machine is stored on a local datastore, storing the swapfile with the other virtual machine files will not improve vMotion.

### Prerequisites

Required privilege: **Host machine.Configuration.Storage partition configuration**

### Procedure

- 1 Browse to the host in the vSphere Web Client navigator.
- 2 Select the **Manage** tab and click **Settings**.
- 3 Under Virtual Machines, click **Swap file location**.

The selected swapfile location is displayed. If configuration of the swapfile location is not supported on the selected host, the tab indicates that the feature is not supported.

If the host is part of a cluster, and the cluster settings specify that swapfiles are to be stored in the same directory as the virtual machine, you cannot edit the swapfile location from the host **Manage** tab. To change the swapfile location for such a host, edit the cluster settings.

- 4 Click **Edit**.
- 5 Select where to store the swapfile.

Option	Description
<b>Virtual machine directory</b>	Stores the swapfile in the same directory as the virtual machine configuration file.
<b>Use a specific datastore</b>	Stores the swapfile in the location you specify. If the swapfile cannot be stored on the datastore that the host specifies, the swapfile is stored in the same folder as the virtual machine.

- 6 (Optional) If you select **Use a specific datastore**, select a datastore from the list.
- 7 Click **OK**.

The virtual machine swapfile is stored in the location you selected.

## Configure a Virtual Machine Swapfile Location for a Cluster

By default, swapfiles for a virtual machine are located on a datastore in the folder that contains the other virtual machine files. However, you can instead configure the hosts in your cluster to place virtual machine swapfiles on an alternative datastore of your choice.

You can configure an alternative swapfile location to place virtual machine swapfiles on either lower-cost or higher-performance storage, depending on your needs.

### Prerequisites

Before you configure a virtual machine swapfile location for a cluster, you must configure the virtual machine swapfile locations for the hosts in the cluster as described in [“Configure Virtual Machine Swapfile Properties for the Host,”](#) on page 38.

### Procedure

- 1 Browse to the cluster in the vSphere Web Client.
- 2 Click the **Manage** tab and click **Settings**.
- 3 Select **Configuration > General**.

- 4 Next to Swap file location, click **Edit**.
- 5 Select where to store the swapfile.

Option	Description
<b>Virtual machine directory</b>	Stores the swapfile in the same directory as the virtual machine configuration file.
<b>Datastore specified by host</b>	Stores the swapfile in the location specified in the host configuration. If the swapfile cannot be stored on the datastore that the host specifies, the swapfile is stored in the same folder as the virtual machine.

- 6 Click **OK**.

## Delete Swap Files

If a host fails, and that host had running virtual machines that were using swap files, those swap files continue to exist and consume many gigabytes of disk space. You can delete the swap files to eliminate this problem.

### Procedure

- 1 Restart the virtual machine that was on the host that failed.
- 2 Stop the virtual machine.

The swap file for the virtual machine is deleted.

## Sharing Memory Across Virtual Machines

Many ESXi workloads present opportunities for sharing memory across virtual machines (as well as within a single virtual machine).

ESXi memory sharing runs as a background activity that scans for sharing opportunities over time. The amount of memory saved varies over time. For a fairly constant workload, the amount generally increases slowly until all sharing opportunities are exploited.

To determine the effectiveness of memory sharing for a given workload, try running the workload, and use `resxtop` or `esxtop` to observe the actual savings. Find the information in the `PSHARE` field of the interactive mode in the Memory page.

Use the `Mem.ShareScanTime` and `Mem.ShareScanGHZ` advanced settings to control the rate at which the system scans memory to identify opportunities for sharing memory.

You can also configure sharing for individual virtual machines by setting the `sched.mem.pshare.enable` option.

Due to security concerns, inter-virtual machine transparent page sharing is disabled by default and page sharing is being restricted to intra-virtual machine memory sharing. This means page sharing does not occur across virtual machines and only occurs inside of a virtual machine. The concept of salting has been introduced to help address concerns system administrators may have over the security implications of transparent page sharing. Salting can be used to allow more granular management of the virtual machines participating in transparent page sharing than was previously possible. With the new salting settings, virtual machines can share pages only if the salt value and contents of the pages are identical. A new host config option `Mem.ShareForceSalting` can be configured to enable or disable salting.

See [Chapter 15, "Advanced Attributes,"](#) on page 113 for information on how to set advanced options.



## Memory Compression

ESXi provides a memory compression cache to improve virtual machine performance when you use memory overcommitment. Memory compression is enabled by default. When a host's memory becomes overcommitted, ESXi compresses virtual pages and stores them in memory.

Because accessing compressed memory is faster than accessing memory that is swapped to disk, memory compression in ESXi allows you to overcommit memory without significantly hindering performance. When a virtual page needs to be swapped, ESXi first attempts to compress the page. Pages that can be compressed to 2 KB or smaller are stored in the virtual machine's compression cache, increasing the capacity of the host.

You can set the maximum size for the compression cache and disable memory compression using the Advanced Settings dialog box in the vSphere Web Client.

### Enable or Disable the Memory Compression Cache

Memory compression is enabled by default. You can use Advanced System Settings in the vSphere Web Client to enable or disable memory compression for a host.

#### Procedure

- 1 Browse to the host in the vSphere Web Client navigator.
- 2 Click the **Manage** tab and click **Settings**.
- 3 Under System, select **Advanced System Settings**.
- 4 Locate Mem.MemZipEnable and click the **Edit** button.
- 5 Enter 1 to enable or enter 0 to disable the memory compression cache.
- 6 Click **OK**.

### Set the Maximum Size of the Memory Compression Cache

You can set the maximum size of the memory compression cache for the host's virtual machines.

You set the size of the compression cache as a percentage of the memory size of the virtual machine. For example, if you enter 20 and a virtual machine's memory size is 1000 MB, ESXi can use up to 200MB of host memory to store the compressed pages of the virtual machine.

If you do not set the size of the compression cache, ESXi uses the default value of 10 percent.

#### Procedure

- 1 Browse to the host in the vSphere Web Client navigator.
- 2 Click the **Manage** tab and click **Settings**.
- 3 Under System, select **Advanced System Settings**.
- 4 Locate Mem.MemZipMaxPct and click the **Edit** button.  
The value of this attribute determines the maximum size of the compression cache for the virtual machine.
- 5 Enter the maximum size for the compression cache.  
The value is a percentage of the size of the virtual machine and must be between 5 and 100 percent.
- 6 Click **OK**.

## Measuring and Differentiating Types of Memory Usage


The **Performance** tab of the vSphere Web Client displays a number of metrics that can be used to analyze memory usage.

Some of these memory metrics measure guest physical memory while other metrics measure machine memory. For instance, two types of memory usage that you can examine using performance metrics are guest physical memory and machine memory. You measure guest physical memory using the Memory Granted metric (for a virtual machine) or Memory Shared (for a host). To measure machine memory, however, use Memory Consumed (for a virtual machine) or Memory Shared Common (for a host). Understanding the conceptual difference between these types of memory usage is important for knowing what these metrics are measuring and how to interpret them.

The VMkernel maps guest physical memory to machine memory, but they are not always mapped one-to-one. Multiple regions of guest physical memory might be mapped to the same region of machine memory (in the case of memory sharing) or specific regions of guest physical memory might not be mapped to machine memory (when the VMkernel swaps out or balloons guest physical memory). In these situations, calculations of guest physical memory usage and machine memory usage for an individual virtual machine or a host differ.

Consider the example in the following figure, which shows two virtual machines running on a host. Each block represents 4 KB of memory and each color/letter represents a different set of data on a block.

**Figure 6-2.** Memory Usage Example



The performance metrics for the virtual machines can be determined as follows:

- To determine Memory Granted (the amount of guest physical memory that is mapped to machine memory) for virtual machine 1, count the number of blocks in virtual machine 1's guest physical memory that have arrows to machine memory and multiply by 4 KB. Since there are five blocks with arrows, Memory Granted would be 20 KB.
- Memory Consumed is the amount of machine memory allocated to the virtual machine, accounting for savings from shared memory. First, count the number of blocks in machine memory that have arrows from virtual machine 1's guest physical memory. There are three such blocks, but one block is shared with virtual machine 2. So count two full blocks plus half of the third and multiply by 4 KB for a total of 10 KB Memory Consumed.

The important difference between these two metrics is that Memory Granted counts the number of blocks with arrows at the guest physical memory level and Memory Consumed counts the number of blocks with arrows at the machine memory level. The number of blocks differs between the two levels due to memory sharing and so Memory Granted and Memory Consumed differ. This is not problematic and shows that memory is being saved through sharing or other reclamation techniques.

A similar result is obtained when determining Memory Shared and Memory Shared Common for the host.

- Memory Shared for the host is the sum of each virtual machine's Memory Shared. Calculate this by looking at each virtual machine's guest physical memory and counting the number of blocks that have arrows to machine memory blocks that themselves have more than one arrow pointing at them. There are six such blocks in the example, so Memory Shared for the host is 24 KB.
- Memory Shared Common is the amount of machine memory that is shared by virtual machines. To determine this, look at the machine memory and count the number of blocks that have more than one arrow pointing at them. There are three such blocks, so Memory Shared Common is 12 KB.

Memory Shared is concerned with guest physical memory and looks at the origin of the arrows. Memory Shared Common, however, deals with machine memory and looks at the destination of the arrows.

The memory metrics that measure guest physical memory and machine memory might appear contradictory. In fact, they are measuring different aspects of a virtual machine's memory usage. By understanding the differences between these metrics, you can better utilize them to diagnose performance issues.

## Memory Reliability

Memory reliability, also known as error isolation, allows ESXi to stop using parts of memory when it determines that a failure might occur, as well as when a failure did occur.

When enough corrected errors are reported at a particular address, ESXi stops using this address to prevent the corrected error from becoming an uncorrected error.

Memory reliability provides a better VMkernel reliability despite corrected and uncorrected errors in RAM. It also enables the system to avoid using memory pages that might contain errors.

### Correcting an Error Isolation Notification

With memory reliability, VMkernel stops using pages that receive an error isolation notification.

The user receives an event in the vSphere Web Client when VMkernel recovers from an uncorrectable memory error, when VMkernel retires a significant percentage of system memory due to a large number of correctable errors, or if there is a large number of pages that are unable to retire.

#### Procedure

- 1 Vacate the host.
- 2 Migrate the virtual machines.
- 3 Run memory related hardware tests.

## About System Swap

System swap is a memory reclamation process that can take advantage of unused memory resources across an entire system.

System swap allows the system to reclaim memory from memory consumers that are not virtual machines. When system swap is enabled you have a tradeoff between the impact of reclaiming the memory from another process and the ability to assign the memory to a virtual machine that can use it. The amount of space required for the system swap is 1GB.

Memory is reclaimed by taking data out of memory and writing it to background storage. Accessing the data from background storage is slower than accessing data from memory, so it is important to carefully select where to store the swapped data.

ESXi determines automatically where the system swap should be stored, this is the **Preferred swap file location**. This decision can be aided by selecting a certain set of options. The system selects the best possible enabled option. If none of the options are feasible then system swap is not activated.

The available options are:

- **Datastore** - Allow the use of the datastore specified. Please note that a vSAN datastore or a VVol datastore cannot be specified for system swap files.
- **Host Swap Cache** - Allow the use of part of the host swap cache.
- **Preferred swap file location** - Allow the use of the preferred swap file location configured for the host.

## Configure System Swap

You can customize the options that determine the system swap location.

### Prerequisites

Select the **Enabled** check box in the Edit System Swap Settings dialog box.

### Procedure

- 1 Browse to the host in the vSphere Web Client navigator.
- 2 Click the **Manage** tab.
- 3 Click **Settings**, and click **System Swap**.
- 4 Click **Edit**.
- 5 Select the check boxes for each option that you want to enable.
- 6 If you select the **datastore** option, select a datastore from the drop-down menu.
- 7 Click **OK**.

## View Graphics Information

---

You can access information about host graphics hardware capability for multiple virtual machines. You can view information about the graphics card and view the virtual machines that use the graphics card.

Virtual machines are listed only if they are turned on and if the graphics card is of the shared type.

### Prerequisites

Verify that the virtual machines are turned on.

### Procedure

- 1 In the vSphere Web Client, navigate to the host.
- 2 Click the **Manage** tab and click **Settings**.
- 3 Under Hardware, select **Graphics**, and select the graphics card to view the virtual machines that use it.  
You can view the name of the graphics card, the vendor, and the total memory allocated to the card.



# Managing Storage I/O Resources

---

vSphere Storage I/O Control allows cluster-wide storage I/O prioritization, which allows better workload consolidation and helps reduce extra costs associated with over provisioning.

Storage I/O Control extends the constructs of shares and limits to handle storage I/O resources. You can control the amount of storage I/O that is allocated to virtual machines during periods of I/O congestion, which ensures that more important virtual machines get preference over less important virtual machines for I/O resource allocation.

When you enable Storage I/O Control on a datastore, ESXi begins to monitor the device latency that hosts observe when communicating with that datastore. When device latency exceeds a threshold, the datastore is considered to be congested and each virtual machine that accesses that datastore is allocated I/O resources in proportion to their shares. You set shares per virtual machine. You can adjust the number for each based on need.

Configuring Storage I/O Control is a two-step process:

- 1 Enable Storage I/O Control for the datastore.
- 2 Set the number of storage I/O shares and upper limit of I/O operations per second (IOPS) allowed for each virtual machine.

By default, all virtual machine shares are set to Normal (1000) with unlimited IOPS.

---

**NOTE** Storage I/O Control is enabled by default on Storage DRS-enabled datastore clusters.

---

This chapter includes the following topics:

- [“Storage I/O Control Requirements,”](#) on page 47
- [“Storage I/O Control Resource Shares and Limits,”](#) on page 48
- [“Set Storage I/O Control Resource Shares and Limits,”](#) on page 49
- [“Enable Storage I/O Control,”](#) on page 49
- [“Set Storage I/O Control Threshold Value,”](#) on page 50
- [“Storage DRS Integration with Storage Profiles,”](#) on page 51

## Storage I/O Control Requirements

Storage I/O Control has several requirements and limitations.

- Datastores that are Storage I/O Control-enabled must be managed by a single vCenter Server system.
- Storage I/O Control is supported on Fibre Channel-connected, iSCSI-connected, and NFS-connected storage. Raw Device Mapping (RDM) is not supported.

- Storage I/O Control does not support datastores with multiple extents.
- Before using Storage I/O Control on datastores that are backed by arrays with automated storage tiering capabilities, check the *VMware Storage/SAN Compatibility Guide* to verify whether your automated tiered storage array has been certified to be compatible with Storage I/O Control.

Automated storage tiering is the ability of an array (or group of arrays) to migrate LUNs/volumes or parts of LUNs/volumes to different types of storage media (SSD, FC, SAS, SATA) based on user-set policies and current I/O patterns. No special certification is required for arrays that do not have these automatic migration/tiering features, including those that provide the ability to manually migrate data between different types of storage media.

## Storage I/O Control Resource Shares and Limits

You allocate the number of storage I/O shares and upper limit of I/O operations per second (IOPS) allowed for each virtual machine. When storage I/O congestion is detected for a datastore, the I/O workloads of the virtual machines accessing that datastore are adjusted according to the proportion of virtual machine shares each virtual machine has.

Storage I/O shares are similar to those used for memory and CPU resource allocation, which are described in [“Resource Allocation Shares,”](#) on page 11. These shares represent the relative importance of a virtual machine with regard to the distribution of storage I/O resources. Under resource contention, virtual machines with higher share values have greater access to the storage array. When you allocate storage I/O resources, you can limit the IOPS that are allowed for a virtual machine. By default, these are unlimited.

The benefits and drawbacks of setting resource limits are described in [“Resource Allocation Limit,”](#) on page 12. If the limit you want to set for a virtual machine is in terms of MB per second instead of IOPS, you can convert MB per second into IOPS based on the typical I/O size for that virtual machine. For example, to restrict a backup application with 64KB I/Os to 10MB per second, set a limit of 160 IOPS.

### View Storage I/O Control Shares and Limits

You can view the shares and limits for all virtual machines running on a datastore. Viewing this information allows you to compare the settings of all virtual machines that are accessing the datastore, regardless of the cluster in which they are running.

#### Procedure

- 1 Browse to the datastore in the vSphere Web Client navigator.
- 2 Under the Related Objects tab, click the **Virtual Machines** tab.

The tab displays each virtual machine running on the datastore and the associated shares value, and percentage of datastore shares.

### Monitor Storage I/O Control Shares

Use the datastore **Performance** tab to monitor how Storage I/O Control handles the I/O workloads of the virtual machines accessing a datastore based on their shares.

Datastore performance charts allow you to monitor the following information:

- Average latency and aggregated IOPS on the datastore
- Latency among hosts
- Queue depth among hosts
- Read/write IOPS among hosts
- Read/write latency among virtual machine disks
- Read/write IOPS among virtual machine disks



**Procedure**

- 1 Browse to the datastore in the vSphere Web Client navigator.
- 2 Under the Monitor tab, click the **Performance** tab
- 3 From the **View** drop-down menu, select **Performance**.

For more information, see the *vSphere Monitoring and Performance* documentation.

## Set Storage I/O Control Resource Shares and Limits

Allocate storage I/O resources to virtual machines based on importance by assigning a relative amount of shares to the virtual machine.

Unless virtual machine workloads are very similar, shares do not necessarily dictate allocation in terms of I/O operations or megabytes per second. Higher shares allow a virtual machine to keep more concurrent I/O operations pending at the storage device or datastore compared to a virtual machine with lower shares. Two virtual machines might experience different throughput based on their workloads.

**Procedure**

- 1 Find the virtual machine in the vSphere Web Client inventory.
  - a To find a virtual machine, select a data center, folder, cluster, resource pool, or host.
  - b Click the **Related Objects** tab and click **Virtual Machines**.
- 2 Right-click the virtual machine and click **Edit Settings**.
- 3 Click the **Virtual Hardware** tab and expand **Hard disk**.
- 4 Select a virtual hard disk from the list.
- 5 Under **Shares**, click the drop-down menu and select the relative amount of shares to allocate to the virtual machine (Low, Normal, or High).

You can select **Custom** to enter a user-defined shares value.

- 6 Under **Limit - IOPS**, click the drop-down menu and enter the upper limit of storage resources to allocate to the virtual machine.

IOPS are the number of I/O operations per second. By default, IOPS are unlimited. You select Low (500), Normal (1000), or High (2000), or you can select Custom to enter a user-defined number of shares.

- 7 Click **OK**.

Shares and limits are reflected under Related Objects on the **Virtual Machines** tab for the host and cluster.

## Enable Storage I/O Control

When you enable Storage I/O Control, ESXi monitors datastore latency and throttles the I/O load if the datastore average latency exceeds the threshold.

**Procedure**

- 1 Browse to the datastore in the vSphere Web Client navigator.
- 2 Click the **Manage** tab.
- 3 Click **Settings** and click **General**.
- 4 Click **Edit** for **Datastore Capabilities**.
- 5 Select the **Enable Storage I/O Control** check box.
- 6 Click **OK**.

Under **Datastore Capabilities**, Storage I/O Control is enabled for the datastore.

## Set Storage I/O Control Threshold Value

The congestion threshold value for a datastore is the upper limit of latency that is allowed for a datastore before Storage I/O Control begins to assign importance to the virtual machine workloads according to their shares.

You do not need to adjust the threshold setting in most environments.



**CAUTION** Storage I/O Control will not function correctly if you share the same spindles on two different datastores.

If you change the congestion threshold setting, set the value based on the following considerations.

- A higher value typically results in higher aggregate throughput and weaker isolation. Throttling will not occur unless the overall average latency is higher than the threshold.
- If throughput is more critical than latency, do not set the value too low. For example, for Fibre Channel disks, a value below 20ms could lower peak disk throughput. A very high value (above 50ms) might allow very high latency without any significant gain in overall throughput.
- A lower value will result in lower device latency and stronger virtual machine I/O performance isolation. Stronger isolation means that the shares controls are enforced more often. Lower device latency translates into lower I/O latency for the virtual machines with the highest shares, at the cost of higher I/O latency experienced by the virtual machines with fewer shares.
- A very low value (lower than 20ms) will result in lower device latency and isolation among I/Os at the potential cost of a decrease in aggregate datastore throughput.
- Setting the value extremely high or extremely lowly results in poor isolation.

### Prerequisites

Verify that Storage I/O Control is enabled.

### Procedure

- 1 Browse to the datastore in the vSphere Web Client navigator.
- 2 Click the **Manage** tab and click **Settings**.
- 3 Click **General**.
- 4 Click **Edit** for **Datastore Capabilities**.
- 5 Select the **Enable Storage I/O Control** check box.

Storage I/O Control automatically sets the latency threshold that corresponds to the estimated latency when the datastore is operating at 90% of its peak throughput.

- 6 (Optional) Adjust the **Congestion Threshold**.
  - ◆ Select a value from the **Percentage of peak throughput** drop-down menu.

The percentage of peak throughput value indicates the estimated latency threshold when the datastore is using that percentage of its estimated peak throughput.

- ◆ Select a value from the **Manual** drop-down menu.

The value must be between 5ms and 100ms. Setting improper congestion threshold values can be detrimental to the performance of the virtual machines on the datastore.

- 7 (Optional) Click **Reset to defaults** to restore the congestion threshold setting to the default value (30ms).

## Storage DRS Integration with Storage Profiles

Storage Policy Based Management (SPBM) allows you to specify the policy for a virtual machine which is enforced by Storage DRS. A datastore cluster can have set of datastores with different capability profiles. If the virtual machines have storage profiles associated with them, Storage DRS can enforce placement based on underlying datastore capabilities.

As part of Storage DRS integration with storage profiles, the Storage DRS cluster level advanced option `EnforceStorageProfiles` is introduced. Advanced option `EnforceStorageProfiles` takes one of these integer values: 0,1 or 2. The default value is 0. When the option is set to 0, it indicates that there is no storage profile or policy enforcement on the Storage DRS cluster. When the option is set to 1, it indicates that there is a storage profile or policy soft enforcement on the Storage DRS cluster. This is analogous with DRS soft rules. Storage DRS will comply with storage profile or policy in the optimum level. Storage DRS will violate the storage profile compliant if it is required to do so. Storage DRS affinity rules will have higher precedence over storage profiles only when storage profile enforcement is set to 1. When the option is set to 2, it indicates that there is a storage profile or policy hard enforcement on the Storage DRS cluster. This is analogous with DRS hard rules. Storage DRS will not violate the storage profile or policy compliant. Storage profiles will have higher precedence over affinity rules. Storage DRS will generate fault: could not fix anti-affinity rule violation

### Prerequisites

By default, Storage DRS will not enforce storage policies associated with a virtual machine. Please configure `EnforceStorageProfiles` option according to your requirements. The options are Default (0), Soft (1) or Hard (2).

### Procedure

- 1 Log in to the vSphere Web client as an Administrator.
- 2 In the vSphere Web Client, click on the Storage DRS cluster, then select **Manage > Settings > Storage DRS**.
- 3 Click **Edit > Advanced Options > Configuration parameters** and select **Add**.
- 4 Click in the area under the Option heading and type `EnforceStorageProfiles`
- 5 Click in the area under the Value heading to the right of the previously entered advanced option name and type the value of either 0, 1 or 2.
- 6 Click **OK**.



## Managing Resource Pools

A resource pool is a logical abstraction for flexible management of resources. Resource pools can be grouped into hierarchies and used to hierarchically partition available CPU and memory resources.


Each standalone host and each DRS cluster has an (invisible) root resource pool that groups the resources of that host or cluster. The root resource pool does not appear because the resources of the host (or cluster) and the root resource pool are always the same.

Users can create child resource pools of the root resource pool or of any user-created child resource pool. Each child resource pool owns some of the parent's resources and can, in turn, have a hierarchy of child resource pools to represent successively smaller units of computational capability.

A resource pool can contain child resource pools, virtual machines, or both. You can create a hierarchy of shared resources. The resource pools at a higher level are called parent resource pools. Resource pools and virtual machines that are at the same level are called siblings. The cluster itself represents the root resource pool. If you do not create child resource pools, only the root resource pools exist.

In the following example, RP-QA is the parent resource pool for RP-QA-UI. RP-Marketing and RP-QA are siblings. The three virtual machines immediately below RP-Marketing are also siblings.

**Figure 9-1.** Parents, Children, and Siblings in Resource Pool Hierarchy



For each resource pool, you specify reservation, limit, shares, and whether the reservation should be expandable. The resource pool resources are then available to child resource pools and virtual machines.

This chapter includes the following topics:

- [“Why Use Resource Pools?”](#) on page 54
- [“Create a Resource Pool,”](#) on page 55
- [“Edit a Resource Pool,”](#) on page 56
- [“Add a Virtual Machine to a Resource Pool,”](#) on page 56
- [“Remove a Virtual Machine from a Resource Pool,”](#) on page 57
- [“Remove a Resource Pool,”](#) on page 58
- [“Resource Pool Admission Control,”](#) on page 58

## Why Use Resource Pools?

Resource pools allow you to delegate control over resources of a host (or a cluster), but the benefits are evident when you use resource pools to compartmentalize all resources in a cluster. Create multiple resource pools as direct children of the host or cluster and configure them. You can then delegate control over the resource pools to other individuals or organizations.

Using resource pools can result in the following benefits.

- Flexible hierarchical organization—Add, remove, or reorganize resource pools or change resource allocations as needed.
- Isolation between pools, sharing within pools—Top-level administrators can make a pool of resources available to a department-level administrator. Allocation changes that are internal to one departmental resource pool do not unfairly affect other unrelated resource pools.
- Access control and delegation—When a top-level administrator makes a resource pool available to a department-level administrator, that administrator can then perform all virtual machine creation and management within the boundaries of the resources to which the resource pool is entitled by the current shares, reservation, and limit settings. Delegation is usually done in conjunction with permissions settings.
- Separation of resources from hardware—If you are using clusters enabled for DRS, the resources of all hosts are always assigned to the cluster. That means administrators can perform resource management independently of the actual hosts that contribute to the resources. If you replace three 2GB hosts with two 3GB hosts, you do not need to make changes to your resource allocations.

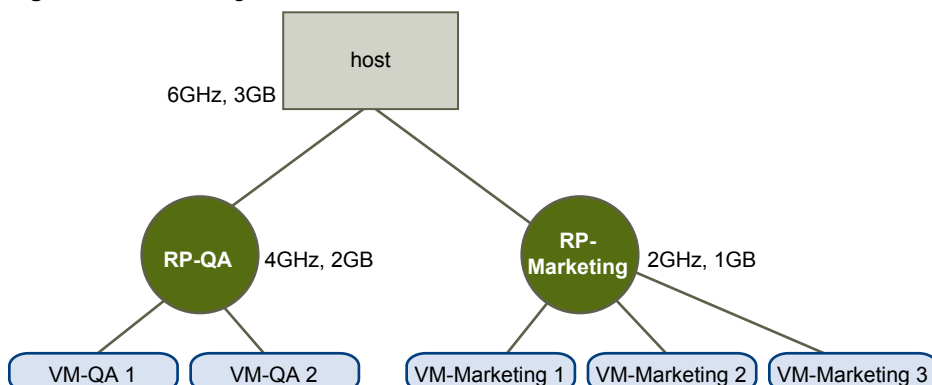
This separation allows administrators to think more about aggregate computing capacity and less about individual hosts.

- Management of sets of virtual machines running a multitier service—Group virtual machines for a multitier service in a resource pool. You do not need to set resources on each virtual machine. Instead, you can control the aggregate allocation of resources to the set of virtual machines by changing settings on their enclosing resource pool.

For example, assume a host has a number of virtual machines. The marketing department uses three of the virtual machines and the QA department uses two virtual machines. Because the QA department needs larger amounts of CPU and memory, the administrator creates one resource pool for each group. The administrator sets **CPU Shares to High** for the QA department pool and to **Normal** for the Marketing department pool so that the QA department users can run automated tests. The second resource pool with fewer CPU and memory resources is sufficient for the lighter load of the marketing staff. Whenever the QA department is not fully using its allocation, the marketing department can use the available resources.

The numbers in the following figure show the effective allocations to the resource pools.

**Figure 9-2.** Allocating Resources to Resource Pools



## Create a Resource Pool

You can create a child resource pool of any ESXi host, resource pool, or DRS cluster.

---

**Note** If a host has been added to a cluster, you cannot create child resource pools of that host. If the cluster is enabled for DRS, you can create child resource pools of the cluster.

---

When you create a child resource pool, you are prompted for resource pool attribute information. The system uses admission control to make sure you cannot allocate resources that are not available.

### Prerequisites

The vSphere Web Client is connected to the vCenter Server system.

### Procedure

- 1 In the vSphere Web Client navigator, select a parent object for the resource pool (a host, another resource pool, or a DRS cluster).
- 2 Right-click the object and select **New Resource Pool**.
- 3 Type a name to identify the resource pool.
- 4 Specify how to allocate CPU and memory resources.

The CPU resources for your resource pool are the guaranteed physical resources the host reserves for a resource pool. Normally, you accept the default and let the host handle resource allocation.

Option	Description
<b>Shares</b>	Specify shares for this resource pool with respect to the parent's total resources. Sibling resource pools share resources according to their relative share values bounded by the reservation and limit. <ul style="list-style-type: none"> <li>■ Select <b>Low</b>, <b>Normal</b>, or <b>High</b> to specify share values respectively in a 1:2:4 ratio.</li> <li>■ Select <b>Custom</b> to give each virtual machine a specific number of shares, which expresses a proportional weight.</li> </ul>
<b>Reservation</b>	Specify a guaranteed CPU or memory allocation for this resource pool. Defaults to 0. A nonzero reservation is subtracted from the unreserved resources of the parent (host or resource pool). The resources are considered reserved, regardless of whether virtual machines are associated with the resource pool.
<b>Expandable Reservation</b>	When the check box is selected (default), expandable reservations are considered during admission control. If you power on a virtual machine in this resource pool, and the combined reservations of the virtual machines are larger than the reservation of the resource pool, the resource pool can use resources from its parent or ancestors.
<b>Limit</b>	Specify the upper limit for this resource pool's CPU or memory allocation. You can usually accept the default ( <b>Unlimited</b> ). To specify a limit, deselect the <b>Unlimited</b> check box.

- 5 Click **OK**.

After you create a resource pool, you can add virtual machines to it. A virtual machine's shares are relative to other virtual machines (or resource pools) with the same parent resource pool.

## Example: Creating Resource Pools

Assume that you have a host that provides 6GHz of CPU and 3GB of memory that must be shared between your marketing and QA departments. You also want to share the resources unevenly, giving one department (QA) a higher priority. This can be accomplished by creating a resource pool for each department and using the **Shares** attribute to prioritize the allocation of resources.

The example shows how to create a resource pool with the ESXi host as the parent resource.

- 1 In the Create Resource Pool dialog box, type a name for the QA department's resource pool (for example, RP-QA).
- 2 Specify **Shares** of **High** for the CPU and memory resources of RP-QA.
- 3 Create a second resource pool, RP-Marketing.  
Leave Shares at **Normal** for CPU and memory.
- 4 Click **OK**.

If there is resource contention, RP-QA receives 4GHz and 2GB of memory, and RP-Marketing 2GHz and 1GB. Otherwise, they can receive more than this allotment. Those resources are then available to the virtual machines in the respective resource pools.

## Edit a Resource Pool

After you create the resource pool, you can edit its CPU and memory resource settings.

### Procedure

- 1 Browse to the resource pool in the vSphere Web Client navigator.
- 2 Click the **Manage** tab and click **Settings**.
- 3 (Optional) You can change all attributes of the selected resource pool as described in [“Create a Resource Pool,”](#) on page 55.
  - ◆ Under CPU Resources, click **Edit** to change CPU resource settings.
  - ◆ Under Memory Resources, click **Edit** to change memory resource settings.
- 4 Click **OK** to save your changes.

## Add a Virtual Machine to a Resource Pool

When you create a virtual machine, you can specify a resource pool location as part of the creation process. You can also add an existing virtual machine to a resource pool.

When you move a virtual machine to a new resource pool:

- The virtual machine's reservation and limit do not change.
- If the virtual machine's shares are high, medium, or low, %Shares adjusts to reflect the total number of shares in use in the new resource pool.
- If the virtual machine has custom shares assigned, the share value is maintained.

---

**NOTE** Because share allocations are relative to a resource pool, you might have to manually change a virtual machine's shares when you move it into a resource pool so that the virtual machine's shares are consistent with the relative values in the new resource pool. A warning appears if a virtual machine would receive a very large (or very small) percentage of total shares.

---



- Under **Monitor**, the information displayed in the **Resource Reservation** tab about the resource pool's reserved and unreserved CPU and memory resources changes to reflect the reservations associated with the virtual machine (if any).

---

**NOTE** If a virtual machine has been powered off or suspended, it can be moved but overall available resources (such as reserved and unreserved CPU and memory) for the resource pool are not affected.

---

#### Procedure

- 1 Find the virtual machine in the vSphere Web Client inventory.
  - a To find a virtual machine, select a data center, folder, cluster, resource pool, or host.
  - b Click the **Related Objects** tab and click **Virtual Machines**.
- 2 Right-click the virtual machine and click **Migrate**.
  - You can move the virtual machine to another host.
  - You can move the virtual machine's storage to another datastore.
  - You can move the virtual machine to another host and move its storage to another datastore.
- 3 Select a resource pool in which to run the virtual machine.
- 4 Review your selections and click **Finish**.

If a virtual machine is powered on, and the destination resource pool does not have enough CPU or memory to guarantee the virtual machine's reservation, the move fails because admission control does not allow it. An error dialog box displays available and requested resources, so you can consider whether an adjustment might resolve the issue.

## Remove a Virtual Machine from a Resource Pool

You can remove a virtual machine from a resource pool either by moving the virtual machine to another resource pool or deleting it.

When you remove a virtual machine from a resource pool, the total number of shares associated with the resource pool decreases, so that each remaining share represents more resources. For example, assume you have a pool that is entitled to 6GHz, containing three virtual machines with shares set to **Normal**. Assuming the virtual machines are CPU-bound, each gets an equal allocation of 2GHz. If one of the virtual machines is moved to a different resource pool, the two remaining virtual machines each receive an equal allocation of 3GHz.

#### Procedure

- 1 Browse to the resource pool in the vSphere Web Client navigator.
- 2 Choose one of the following methods to remove the virtual machine from a resource pool.
  - Right-click the virtual machine and select **Migrate** to move the virtual machine to another resource pool.  
You do not need to power off the virtual machine before you move it.
  - Right-click the virtual machine and select **Delete**.  
You must power off the virtual machine before you can completely remove it.

## Remove a Resource Pool

You can remove a resource pool from the inventory.

### Procedure

- 1 In the vSphere Web Client, right-click the resource pool and Select **Delete**.  
A confirmation dialog box appears.
- 2 Click **Yes** to remove the resource pool.

## Resource Pool Admission Control

When you power on a virtual machine in a resource pool, or try to create a child resource pool, the system performs additional admission control to ensure the resource pool's restrictions are not violated.

Before you power on a virtual machine or create a resource pool, ensure that sufficient resources are available using the **Resource Reservation** tab in the vSphere Web Client. The **Available Reservation** value for CPU and memory displays resources that are unreserved.

How available CPU and memory resources are computed and whether actions are performed depends on the **Reservation Type**.

**Table 9-1.** Reservation Types

Reservation Type	Description
<b>Fixed</b>	The system checks whether the selected resource pool has sufficient unreserved resources. If it does, the action can be performed. If it does not, a message appears and the action cannot be performed.
<b>Expandable</b> (default)	The system considers the resources available in the selected resource pool and its direct parent resource pool. If the parent resource pool also has the <b>Expandable Reservation</b> option selected, it can borrow resources from its parent resource pool. Borrowing resources occurs recursively from the ancestors of the current resource pool as long as the <b>Expandable Reservation</b> option is selected. Leaving this option selected offers more flexibility, but, at the same time provides less protection. A child resource pool owner might reserve more resources than you anticipate.

The system does not allow you to violate preconfigured **Reservation** or **Limit** settings. Each time you reconfigure a resource pool or power on a virtual machine, the system validates all parameters so all service-level guarantees can still be met.

### Expandable Reservations Example 1

This example shows you how a resource pool with expandable reservations works.

Assume an administrator manages pool P, and defines two child resource pools, S1 and S2, for two different users (or groups).

The administrator knows that users want to power on virtual machines with reservations, but does not know how much each user will need to reserve. Making the reservations for S1 and S2 expandable allows the administrator to more flexibly share and inherit the common reservation for pool P.

Without expandable reservations, the administrator needs to explicitly allocate S1 and S2 a specific amount. Such specific allocations can be inflexible, especially in deep resource pool hierarchies and can complicate setting reservations in the resource pool hierarchy.

Expandable reservations cause a loss of strict isolation. S1 can start using all of P's reservation, so that no memory or CPU is directly available to S2.


## Expandable Reservations Example 2

This example shows how a resource pool with expandable reservations works.

Assume the following scenario, as shown in the figure.

- Parent pool RP-MOM has a reservation of 6GHz and one running virtual machine VM-M1 that reserves 1GHz.
- You create a child resource pool RP-KID with a reservation of 2GHz and with **Expandable Reservation** selected.
- You add two virtual machines, VM-K1 and VM-K2, with reservations of 2GHz each to the child resource pool and try to power them on.
- VM-K1 can reserve the resources directly from RP-KID (which has 2GHz).
- No local resources are available for VM-K2, so it borrows resources from the parent resource pool, RP-MOM. RP-MOM has 6GHz minus 1GHz (reserved by the virtual machine) minus 2GHz (reserved by RP-KID), which leaves 3GHz unreserved. With 3GHz available, you can power on the 2GHz virtual machine.


**Figure 9-3.** Admission Control with Expandable Resource Pools: Successful Power-On



Now, consider another scenario with VM-M1 and VM-M2.

- Power on two virtual machines in RP-MOM with a total reservation of 3GHz.
- You can still power on VM-K1 in RP-KID because 2GHz are available locally.
- When you try to power on VM-K2, RP-KID has no unreserved CPU capacity so it checks its parent. RP-MOM has only 1GHz of unreserved capacity available (5GHz of RP-MOM are already in use—3GHz reserved by the local virtual machines and 2GHz reserved by RP-KID). As a result, you cannot power on VM-K2, which requires a 2GHz reservation.

**Figure 9-4.** Admission Control with Expandable Resource Pools: Power-On Prevented





## Creating a DRS Cluster

---

A cluster is a collection of ESXi hosts and associated virtual machines with shared resources and a shared management interface. Before you can obtain the benefits of cluster-level resource management you must create a cluster and enable DRS.

Depending on whether or not Enhanced vMotion Compatibility (EVC) is enabled, DRS behaves differently when you use vSphere Fault Tolerance (vSphere FT) virtual machines in your cluster.

**Table 10-1.** DRS Behavior with vSphere FT Virtual Machines and EVC

<b>EVC</b>	<b>DRS (Load Balancing)</b>	<b>DRS (Initial Placement)</b>
Enabled	Enabled (Primary and Secondary VMs)	Enabled (Primary and Secondary VMs)
Disabled	Disabled (Primary and Secondary VMs)	Disabled (Primary VMs) Fully Automated (Secondary VMs)

This chapter includes the following topics:

- [“Admission Control and Initial Placement,”](#) on page 62
- [“Virtual Machine Migration,”](#) on page 63
- [“DRS Cluster Requirements,”](#) on page 65
- [“Configuring DRS with Virtual Flash,”](#) on page 66
- [“Create a Cluster,”](#) on page 66
- [“Edit a Cluster,”](#) on page 67
- [“Create a DRS Cluster,”](#) on page 68
- [“Set a Custom Automation Level for a Virtual Machine,”](#) on page 69
- [“Disable DRS,”](#) on page 70
- [“Restore a Resource Pool Tree,”](#) on page 70

## Admission Control and Initial Placement

When you attempt to power on a single virtual machine or a group of virtual machines in a DRS-enabled cluster, vCenter Server performs admission control. It checks that there are enough resources in the cluster to support the virtual machine(s).

If the cluster does not have sufficient resources to power on a single virtual machine, or any of the virtual machines in a group power-on attempt, a message appears. Otherwise, for each virtual machine, DRS generates a recommendation of a host on which to run the virtual machine and takes one of the following actions

- Automatically executes the placement recommendation.
- Displays the placement recommendation, which the user can then choose to accept or override.

---

**NOTE** No initial placement recommendations are given for virtual machines on standalone hosts or in non-DRS clusters. When powered on, they are placed on the host where they currently reside.

---

### Single Virtual Machine Power On

In a DRS cluster, you can power on a single virtual machine and receive initial placement recommendations.

When you power on a single virtual machine, you have two types of initial placement recommendations:

- A single virtual machine is being powered on and no prerequisite steps are needed.

The user is presented with a list of mutually exclusive initial placement recommendations for the virtual machine. You can select only one.

- A single virtual machine is being powered on, but prerequisite actions are required.

These actions include powering on a host in standby mode or the migration of other virtual machines from one host to another. In this case, the recommendations provided have multiple lines, showing each of the prerequisite actions. The user can either accept this entire recommendation or cancel powering on the virtual machine.

### Group Power On

You can attempt to power on multiple virtual machines at the same time (group power on).

Virtual machines selected for a group power-on attempt do not have to be in the same DRS cluster. They can be selected across clusters but must be within the same data center. It is also possible to include virtual machines located in non-DRS clusters or on standalone hosts. These are powered on automatically and not included in any initial placement recommendation.

The initial placement recommendations for group power-on attempts are provided on a per-cluster basis. If all of the placement-related actions for a group power-on attempt are in automatic mode, the virtual machines are powered on with no initial placement recommendation given. If placement-related actions for any of the virtual machines are in manual mode, the powering on of all of the virtual machines (including those that are in automatic mode) is manual and is included in an initial placement recommendation.

For each DRS cluster that the virtual machines being powered on belong to, there is a single recommendation, which contains all of the prerequisites (or no recommendation). All such cluster-specific recommendations are presented together under the Power On Recommendations tab.

When a nonautomatic group power-on attempt is made, and virtual machines not subject to an initial placement recommendation (that is, those on standalone hosts or in non-DRS clusters) are included, vCenter Server attempts to power them on automatically. If these power ons are successful, they are listed under the Started Power-Ons tab. Any virtual machines that fail to power on are listed under the Failed Power-Ons tab.

## Example: Group Power On

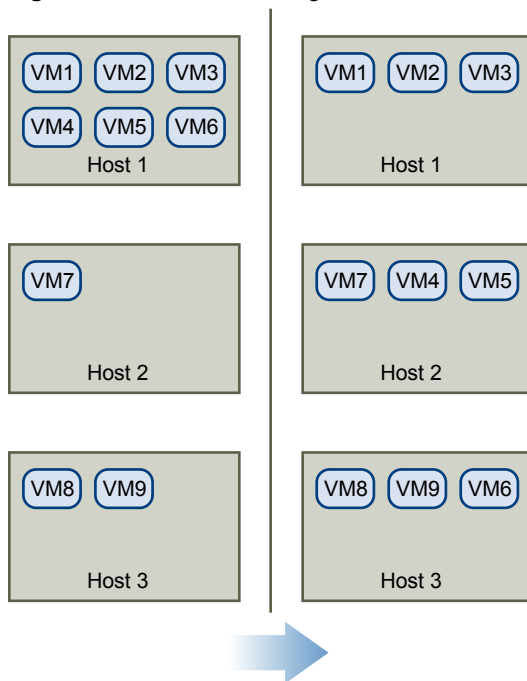
The user selects three virtual machines in the same data center for a group power-on attempt. The first two virtual machines (VM1 and VM2) are in the same DRS cluster (Cluster1), while the third virtual machine (VM3) is on a standalone host. VM1 is in automatic mode and VM2 is in manual mode. For this scenario, the user is presented with an initial placement recommendation for Cluster1 (under the Power On Recommendations tab) which consists of actions for powering on VM1 and VM2. An attempt is made to power on VM3 automatically and, if successful, it is listed under the Started Power-Ons tab. If this attempt fails, it is listed under the Failed Power-Ons tab.

## Virtual Machine Migration

Although DRS performs initial placements so that load is balanced across the cluster, changes in virtual machine load and resource availability can cause the cluster to become unbalanced. To correct such imbalances, DRS generates migration recommendations.

If DRS is enabled on the cluster, load can be distributed more uniformly to reduce the degree of this imbalance. For example, the three hosts on the left side of the following figure are unbalanced. Assume that Host 1, Host 2, and Host 3 have identical capacity, and all virtual machines have the same configuration and load (which includes reservation, if set). However, because Host 1 has six virtual machines, its resources might be overused while ample resources are available on Host 2 and Host 3. DRS migrates (or recommends the migration of) virtual machines from Host 1 to Host 2 and Host 3. On the right side of the diagram, the properly load balanced configuration of the hosts that results appears.

**Figure 10-1.** Load Balancing



When a cluster becomes unbalanced, DRS makes recommendations or migrates virtual machines, depending on the default automation level:

- If the cluster or any of the virtual machines involved are manual or partially automated, vCenter Server does not take automatic actions to balance resources. Instead, the Summary page indicates that migration recommendations are available and the DRS Recommendations page displays recommendations for changes that make the most efficient use of resources across the cluster.

- If the cluster and virtual machines involved are all fully automated, vCenter Server migrates running virtual machines between hosts as needed to ensure efficient use of cluster resources.

---

**NOTE** Even in an automatic migration setup, users can explicitly migrate individual virtual machines, but vCenter Server might move those virtual machines to other hosts to optimize cluster resources.

---

By default, automation level is specified for the whole cluster. You can also specify a custom automation level for individual virtual machines.

## DRS Migration Threshold

The DRS migration threshold allows you to specify which recommendations are generated and then applied (when the virtual machines involved in the recommendation are in fully automated mode) or shown (if in manual mode). This threshold is also a measure of how much cluster imbalance across host (CPU and memory) loads is acceptable.

You can move the threshold slider to use one of five settings, ranging from Conservative to Aggressive. The five migration settings generate recommendations based on their assigned priority level. Each setting you move the slider to the right allows the inclusion of one more lower level of priority. The Conservative setting generates only priority-one recommendations (mandatory recommendations), the next level to the right generates priority-two recommendations and higher, and so on, down to the Aggressive level which generates priority-five recommendations and higher (that is, all recommendations.)

A priority level for each migration recommendation is computed using the load imbalance metric of the cluster. This metric is displayed as Current host load standard deviation in the cluster's Summary tab in the vSphere Web Client. A higher load imbalance leads to higher-priority migration recommendations. For more information about this metric and how a recommendation priority level is calculated, see the VMware Knowledge Base article "Calculating the priority level of a VMware DRS migration recommendation."

After a recommendation receives a priority level, this level is compared to the migration threshold you set. If the priority level is less than or equal to the threshold setting, the recommendation is either applied (if the relevant virtual machines are in fully automated mode) or displayed to the user for confirmation (if in manual or partially automated mode.)

## Migration Recommendations

If you create a cluster with a default manual or partially automated mode, vCenter Server displays migration recommendations on the DRS Recommendations page.

The system supplies as many recommendations as necessary to enforce rules and balance the resources of the cluster. Each recommendation includes the virtual machine to be moved, current (source) host and destination host, and a reason for the recommendation. The reason can be one of the following:

- Balance average CPU loads or reservations.
- Balance average memory loads or reservations.
- Satisfy resource pool reservations.
- Satisfy an affinity rule.
- Host is entering maintenance mode or standby mode.

---

**NOTE** If you are using the vSphere Distributed Power Management (DPM) feature, in addition to migration recommendations, DRS provides host power state recommendations.

---



## DRS Cluster Requirements

Hosts that are added to a DRS cluster must meet certain requirements to use cluster features successfully.

### Shared Storage Requirements

A DRS cluster has certain shared storage requirements.

Ensure that the managed hosts use shared storage. Shared storage is typically on a SAN, but can also be implemented using NAS shared storage.

See the *vSphere Storage* documentation for information about other shared storage.

### Shared VMFS Volume Requirements

A DRS cluster has certain shared VMFS volume requirements.

Configure all managed hosts to use shared VMFS volumes.

- Place the disks of all virtual machines on VMFS volumes that are accessible by source and destination hosts.
- Ensure the VMFS volume is sufficiently large to store all virtual disks for your virtual machines.
- Ensure all VMFS volumes on source and destination hosts use volume names, and all virtual machines use those volume names for specifying the virtual disks.

---

**NOTE** Virtual machine swap files also need to be on a VMFS accessible to source and destination hosts (just like .vmdk virtual disk files). This requirement does not apply if all source and destination hosts are ESX Server 3.5 or higher and using host-local swap. In that case, vMotion with swap files on unshared storage is supported. Swap files are placed on a VMFS by default, but administrators might override the file location using advanced virtual machine configuration options.

---

### Processor Compatibility Requirements

A DRS cluster has certain processor compatibility requirements.

To avoid limiting the capabilities of DRS, you should maximize the processor compatibility of source and destination hosts in the cluster.

vMotion transfers the running architectural state of a virtual machine between underlying ESXi hosts. vMotion compatibility means that the processors of the destination host must be able to resume execution using the equivalent instructions where the processors of the source host were suspended. Processor clock speeds and cache sizes might vary, but processors must come from the same vendor class (Intel versus AMD) and the same processor family to be compatible for migration with vMotion.

Processor families are defined by the processor vendors. You can distinguish different processor versions within the same family by comparing the processors' model, stepping level, and extended features.

Sometimes, processor vendors have introduced significant architectural changes within the same processor family (such as 64-bit extensions and SSE3). VMware identifies these exceptions if it cannot guarantee successful migration with vMotion.

vCenter Server provides features that help ensure that virtual machines migrated with vMotion meet processor compatibility requirements. These features include:

- Enhanced vMotion Compatibility (EVC) – You can use EVC to help ensure vMotion compatibility for the hosts in a cluster. EVC ensures that all hosts in a cluster present the same CPU feature set to virtual machines, even if the actual CPUs on the hosts differ. This prevents migrations with vMotion from failing due to incompatible CPUs.

Configure EVC from the Cluster Settings dialog box. The hosts in a cluster must meet certain requirements for the cluster to use EVC. For information about EVC and EVC requirements, see the *vCenter Server and Host Management* documentation.

- CPU compatibility masks – vCenter Server compares the CPU features available to a virtual machine with the CPU features of the destination host to determine whether to allow or disallow migrations with vMotion. By applying CPU compatibility masks to individual virtual machines, you can hide certain CPU features from the virtual machine and potentially prevent migrations with vMotion from failing due to incompatible CPUs.

## vMotion Requirements for DRS Clusters

A DRS cluster has certain vMotion requirements.

To enable the use of DRS migration recommendations, the hosts in your cluster must be part of a vMotion network. If the hosts are not in the vMotion network, DRS can still make initial placement recommendations.

To be configured for vMotion, each host in the cluster must meet the following requirements:

- vMotion does not support raw disks or migration of applications clustered using Microsoft Cluster Service (MSCS).
- vMotion requires a private Gigabit Ethernet migration network between all of the vMotion enabled managed hosts. When vMotion is enabled on a managed host, configure a unique network identity object for the managed host and connect it to the private migration network.

## Configuring DRS with Virtual Flash

DRS can manage virtual machines that have virtual flash reservations.

Virtual flash capacity appears as a statistic that is regularly reported from the host to the vSphere Web Client. Each time DRS runs, it uses the most recent capacity value reported.

You can configure one virtual flash resource per host. This means that during virtual machine power-on time, DRS does not need to select between different virtual flash resources on a given host.

DRS selects a host that has sufficient available virtual flash capacity to start the virtual machine. If DRS cannot satisfy the virtual flash reservation of a virtual machine, it cannot be powered-on. DRS treats a powered-on virtual machine with a virtual flash reservation as having a soft affinity with its current host. DRS will not recommend such a virtual machine for vMotion except for mandatory reasons, such as putting a host in maintenance mode, or to reduce the load on an over utilized host.

## Create a Cluster

A cluster is a group of hosts. When a host is added to a cluster, the host's resources become part of the cluster's resources. The cluster manages the resources of all hosts within it. Clusters enable the vSphere High Availability (HA) and vSphere Distributed Resource Scheduler (DRS) solutions. You can also enable Virtual SAN on a cluster.

### Prerequisites

- Verify that you have sufficient permissions to create a cluster object.
- Verify that a data center exists in the inventory.

### Procedure

- 1 Browse to a data center in the vSphere Web Client navigator.
- 2 Right-click the data center and select **New Cluster**.

- 3 Enter a name for the cluster.
- 4 Select DRS and vSphere HA cluster features.

Option	Description
<b>To use DRS with this cluster</b>	<ol style="list-style-type: none"> <li>a Select the DRS <b>Turn ON</b> check box.</li> <li>b Select an automation level and a migration threshold.</li> </ol>
<b>To use HA with this cluster</b>	<ol style="list-style-type: none"> <li>a Select the vSphere HA <b>Turn ON</b> check box.</li> <li>b Select whether to enable host monitoring and admission control.</li> <li>c If admission control is enabled, specify a policy.</li> <li>d Select a VM Monitoring option.</li> <li>e Specify the virtual machine monitoring sensitivity.</li> </ol>

- 5 Select an Enhanced vMotion Compatibility (EVC) setting.  
 EVC ensures that all hosts in a cluster present the same CPU feature set to virtual machines, even if the actual CPUs on the hosts differ. This prevents migrations with vMotion from failing due to incompatible CPUs.
- 6 Enable Virtual SAN.
  - a Select the Virtual SAN **Turn ON** check box.
  - b Specify whether to add disks automatically or manually to the Virtual SAN cluster.
- 7 Click **OK**.

The cluster is added to the inventory.

#### What to do next

Add hosts and resource pools to the cluster.

For information about Virtual SAN and how to use Virtual SAN clusters, see the *vSphere Storage* publication.

## Edit a Cluster

You can enable cluster features and edit settings for existing clusters.

#### Prerequisites

Required privileges: **Host.Inventory.Modify Cluster** or **Host.Inventory.Remove Cluster** or **Host.Inventory.Rename Cluster** or **Host.Inventory.Move Cluster**

#### Procedure

- 1 Browse to the cluster in the vSphere Web Client navigator.
- 2 Select the **Manage** tab and click **Settings**.
- 3 Under Services select **vSphere DRS** or **vSphere HA** and click **Edit** to edit cluster services.
- 4 Click **OK**.

## Create a DRS Cluster

When you add a host to a DRS cluster, the host's resources become part of the cluster's resources. In addition to this aggregation of resources, with a DRS cluster you can support cluster-wide resource pools and enforce cluster-level resource allocation policies.

The following cluster-level resource management capabilities are also available.

### Load Balancing

The distribution and usage of CPU and memory resources for all hosts and virtual machines in the cluster are continuously monitored. DRS compares these metrics to an ideal resource utilization given the attributes of the cluster's resource pools and virtual machines, the current demand, and the imbalance target. It then performs (or recommends) virtual machine migrations accordingly. See ["Virtual Machine Migration,"](#) on page 63. When you first power on a virtual machine in the cluster, DRS attempts to maintain proper load balancing by either placing the virtual machine on an appropriate host or making a recommendation. See ["Admission Control and Initial Placement,"](#) on page 62.

### Power management

When the vSphere Distributed Power Management (DPM) feature is enabled, DRS compares cluster- and host-level capacity to the demands of the cluster's virtual machines, including recent historical demand. It places (or recommends placing) hosts in standby power mode if sufficient excess capacity is found or powering on hosts if capacity is needed. Depending on the resulting host power state recommendations, virtual machines might need to be migrated to and from the hosts as well. See ["Managing Power Resources,"](#) on page 80.

### Affinity Rules

You can control the placement of virtual machines on hosts within a cluster, by assigning affinity rules. See ["Using DRS Affinity Rules,"](#) on page 84.

### Prerequisites

You can create a cluster without a special license, but you must have a license to enable a cluster for vSphere DRS (or vSphere HA).

### Procedure

- 1 Right-click a data center in the vSphere Web Client and select **New Cluster**.
- 2 Name the cluster in the **Name** text box.  
This name appears in the vSphere Web Client navigator.
- 3 Select the DRS **Turn ON** check box.
- 4 Select a default automation level for DRS.

Automation Level	Action
<b>Manual</b>	<ul style="list-style-type: none"> <li>■ Initial placement: Recommended host(s) is displayed.</li> <li>■ Migration: Recommendation is displayed.</li> </ul>
<b>Partially Automated</b>	<ul style="list-style-type: none"> <li>■ Initial placement: Automatic.</li> <li>■ Migration: Recommendation is displayed.</li> </ul>
<b>Fully Automated</b>	<ul style="list-style-type: none"> <li>■ Initial placement: Automatic.</li> <li>■ Migration: Recommendation is executed automatically.</li> </ul>

- 5 Set the migration threshold for DRS.

- 6 (Optional) Select the vSphere HA **Turn ON** check box to enable vSphere HA.  
vSphere HA allows you to:
  - Enable host monitoring.
  - Enable admission control.
  - Specify the type of policy that admission control should enforce.
  - Adjust the monitoring sensitivity of virtual machine monitoring.
- 7 If appropriate, enable Enhanced vMotion Compatibility (EVC) and select the mode it should operate in.
- 8 Click **OK** to complete cluster creation.

A new cluster does not include any hosts or virtual machines.

To add hosts and virtual machines to the cluster see [“Adding Hosts to a Cluster,”](#) on page 71.

## Set a Custom Automation Level for a Virtual Machine

After you create a DRS cluster, you can customize the automation level for individual virtual machines to override the cluster’s default automation level.

For example, you can select **Manual** for specific virtual machines in a cluster with full automation, or **Partially Automated** for specific virtual machines in a manual cluster.

If a virtual machine is set to **Disabled**, vCenter Server does not migrate that virtual machine or provide migration recommendations for it. This is known as pinning the virtual machine to its registered host.

---

**NOTE** If you have not enabled Enhanced vMotion Compatibility (EVC) for the cluster, fault tolerant virtual machines are set to DRS disabled. They appear on this screen, but you cannot assign an automation mode to them.

---

### Procedure

- 1 Browse to the cluster in the vSphere Web Client navigator.
- 2 Click the **Manage** tab and click **Settings**.
- 3 Under Services, select **vSphere DRS** and click **Edit**. Expand DRS Automation.
- 4 Select the **Enable individual virtual machine automation levels** check box.
- 5 To temporarily disable any individual virtual machine overrides, deselect the **Enable individual virtual machine automation levels** check box.

Virtual machine settings are restored when the check box is selected again.

- 6 To temporarily suspend all vMotion activity in a cluster, put the cluster in manual mode and deselect the **Enable individual virtual machine automation levels** check box.
- 7 Select one or more virtual machines.
- 8 Click the **Automation Level** column and select an automation level from the drop-down menu.

Option	Description
<b>Manual</b>	Placement and migration recommendations are displayed, but do not run until you manually apply the recommendation.
<b>Fully Automated</b>	Placement and migration recommendations run automatically.
<b>Partially Automated</b>	Initial placement is performed automatically. Migration recommendations are displayed, but do not run.
<b>Disabled</b>	vCenter Server does not migrate the virtual machine or provide migration recommendations for it.

- 9 Click **OK**.

---

**Note** Other VMware products or features, such as vSphere vApp and vSphere Fault Tolerance, might override the automation levels of virtual machines in a DRS cluster. Refer to the product-specific documentation for details.

---

## Disable DRS

You can turn off DRS for a cluster.

When DRS is disabled, the cluster's resource pool hierarchy and affinity rules are not reestablished when DRS is turned back on. If you disable DRS, the resource pools are removed from the cluster. To avoid losing the resource pools, save a snapshot of the resource pool tree on your local machine. You can use the snapshot to restore the resource pool when you enable DRS.

### Procedure

- 1 Browse to the cluster in the vSphere Web Client navigator.
- 2 Click the **Manage** tab and click **Settings**.
- 3 Under **vSphere DRS**, click **Edit**.
- 4 Deselect the **Turn On vSphere DRS** check box.
- 5 Click **OK** to turn off DRS.
- 6 (Optional) Choose an option to save the resource pool.
  - Click **Yes** to save a resource pool tree snapshot on a local machine.
  - Click **No** to turn off DRS without saving a resource pool tree snapshot.

## Restore a Resource Pool Tree

You can restore a previously saved resource pool tree snapshot.

### Prerequisites

- vSphere DRS must be turned ON.
- You can restore a snapshot only on the same cluster that it was taken.
- No other resource pools are present in the cluster.

### Procedure

- 1 Browse to the cluster in the vSphere Web Client navigator.
- 2 Right-click on the cluster and select **Restore Resource Pool Tree**.
- 3 Click **Browse**, and locate the snapshot file on your local machine.
- 4 Click **Open**.
- 5 Click **OK** to restore the resource pool tree.

# Using DRS Clusters to Manage Resources

# 11

After you create a DRS cluster, you can customize it and use it to manage resources.

To customize your DRS cluster and the resources it contains you can configure affinity rules and you can add and remove hosts and virtual machines. When a cluster's settings and resources have been defined, you should ensure that it is and remains a valid cluster. You can also use a valid DRS cluster to manage power resources and interoperate with vSphere HA.

This chapter includes the following topics:

- [“Adding Hosts to a Cluster,”](#) on page 71
- [“Adding Virtual Machines to a Cluster,”](#) on page 73
- [“Removing Virtual Machines from a Cluster,”](#) on page 73
- [“Removing a Host from a Cluster,”](#) on page 74
- [“DRS Cluster Validity,”](#) on page 75
- [“Managing Power Resources,”](#) on page 80
- [“Using DRS Affinity Rules,”](#) on page 84

## Adding Hosts to a Cluster

The procedure for adding hosts to a cluster is different for hosts managed by the same vCenter Server (managed hosts) than for hosts not managed by that server.

After a host has been added, the virtual machines deployed to the host become part of the cluster and DRS can recommend migration of some virtual machines to other hosts in the cluster.

### Add a Managed Host to a Cluster

When you add a standalone host already being managed by vCenter Server to a DRS cluster, the host's resources become associated with the cluster.

You can decide whether you want to associate existing virtual machines and resource pools with the cluster's root resource pool or graft the resource pool hierarchy.

---

**NOTE** If a host has no child resource pools or virtual machines, the host's resources are added to the cluster but no resource pool hierarchy with a top-level resource pool is created.

---

#### Procedure

- 1 Browse to the host in the vSphere Web Client navigator.
- 2 Right-click the host and select **Move To**.

- 3 Select a cluster.
- 4 Click **OK** to apply the changes.
- 5 Select what to do with the host's virtual machines and resource pools.
  - **Put this host's virtual machines in the cluster's root resource pool**  
vCenter Server removes all existing resource pools of the host and the virtual machines in the host's hierarchy are all attached to the root. Because share allocations are relative to a resource pool, you might have to manually change a virtual machine's shares after selecting this option, which destroys the resource pool hierarchy.
  - **Create a resource pool for this host's virtual machines and resource pools**  
vCenter Server creates a top-level resource pool that becomes a direct child of the cluster and adds all children of the host to that new resource pool. You can supply a name for that new top-level resource pool. The default is **Grafted from <host\_name>**.

The host is added to the cluster.

## Add an Unmanaged Host to a Cluster

You can add an unmanaged host to a cluster. Such a host is not currently managed by the same vCenter Server system as the cluster and it is not visible in the vSphere Web Client.

### Procedure

- 1 Browse to the cluster in the vSphere Web Client navigator.
- 2 Right-click the cluster and select **Add Host**.
- 3 Enter the host name, user name, and password, and click **Next**.
- 4 View the summary information and click **Next**.
- 5 Assign an existing or a new license key and click **Next**.
- 6 (Optional) You can enable lockdown mode to prevent remote users from logging directly into the host. If you do not enable lockdown mode, you can configure this option later by editing Security Profile in host settings.
- 7 Select what to do with the host's virtual machines and resource pools.
  - **Put this host's virtual machines in the cluster's root resource pool**  
vCenter Server removes all existing resource pools of the host and the virtual machines in the host's hierarchy are all attached to the root. Because share allocations are relative to a resource pool, you might have to manually change a virtual machine's shares after selecting this option, which destroys the resource pool hierarchy.
  - **Create a resource pool for this host's virtual machines and resource pools**  
vCenter Server creates a top-level resource pool that becomes a direct child of the cluster and adds all children of the host to that new resource pool. You can supply a name for that new top-level resource pool. The default is **Grafted from <host\_name>**.
- 8 Review settings and click **Finish**.

The host is added to the cluster.



## Adding Virtual Machines to a Cluster

You can add a virtual machine to a cluster in a number of ways.

- When you add a host to a cluster, all virtual machines on that host are added to the cluster.
- When a virtual machine is created, the New Virtual Machine wizard prompts you for the location to place the virtual machine. You can select a standalone host or a cluster and you can select any resource pool inside the host or cluster.
- You can migrate a virtual machine from a standalone host to a cluster or from a cluster to another cluster using the Migrate Virtual Machine wizard. To start this wizard, right-click the virtual machine name and select **Migrate**.

## Move a Virtual Machine to a Cluster

You can move a virtual machine to a cluster.

### Procedure

- 1 Find the virtual machine in the vSphere Web Client inventory.
  - a To find a virtual machine, select a data center, folder, cluster, resource pool, or host.
  - b Click the **Related Objects** tab and click **Virtual Machines**.
- 2 Right-click on the virtual machine and select **Move To**.
- 3 Select a cluster.
- 4 Click OK.

## Removing Virtual Machines from a Cluster

You can remove virtual machines from a cluster.

You can remove a virtual machine from a cluster in two ways.

- When you remove a host from a cluster, all of the powered-off virtual machines that you do not migrate to other hosts are removed as well. You can remove a host only if it is in maintenance mode or disconnected. If you remove a host from a DRS cluster, the cluster can become yellow because it is overcommitted.
- You can migrate a virtual machine from a cluster to a standalone host or from a cluster to another cluster using the Migrate Virtual Machine wizard. To start this wizard right-click the virtual machine name and select **Migrate**.

## Move a Virtual Machine Out of a Cluster

You can move a virtual machine out of a cluster.

### Procedure

- 1 Find the virtual machine in the vSphere Web Client inventory.
  - a To find a virtual machine, select a data center, folder, cluster, resource pool, or host.
  - b Click the **Related Objects** tab and click **Virtual Machines**.
- 2 Right-click the virtual machine and select **Migrate**.
- 3 Select **Change datastore** and click **Next**.

- 4 Select a datastore and click **Next**.
- 5 Click **Finish**.

If the virtual machine is a member of a DRS cluster rules group, vCenter Server displays a warning before it allows the migration to proceed. The warning indicates that dependent virtual machines are not migrated automatically. You have to acknowledge the warning before migration can proceed.

## Removing a Host from a Cluster

When you remove a host from a DRS cluster, you affect resource pool hierarchies, virtual machines, and you might create invalid clusters. Consider the affected objects before you remove the host.

- **Resource Pool Hierarchies** – When you remove a host from a cluster, the host retains only the root resource pool, even if you used a DRS cluster and decided to graft the host resource pool when you added the host to the cluster. In that case, the hierarchy remains with the cluster. You can create a host-specific resource pool hierarchy.

---

**NOTE** Ensure that you remove the host from the cluster by first placing it in maintenance mode. If you instead disconnect the host before removing it from the cluster, the host retains the resource pool that reflects the cluster hierarchy.

---

- **Virtual Machines** – A host must be in maintenance mode before you can remove it from the cluster and for a host to enter maintenance mode all powered-on virtual machines must be migrated off that host. When you request that a host enter maintenance mode, you are also asked whether you want to migrate all the powered-off virtual machines on that host to other hosts in the cluster.
- **Invalid Clusters** – When you remove a host from a cluster, the resources available for the cluster decrease. If the cluster has enough resources to satisfy the reservations of all virtual machines and resource pools in the cluster, the cluster adjusts resource allocation to reflect the reduced amount of resources. If the cluster does not have enough resources to satisfy the reservations of all resource pools, but there are enough resources to satisfy the reservations for all virtual machines, an alarm is issued and the cluster is marked yellow. DRS continues to run.

## Place a Host in Maintenance Mode

You place a host in maintenance mode when you need to service it, for example, to install more memory. A host enters or leaves maintenance mode only as the result of a user request.

Virtual machines that are running on a host entering maintenance mode need to be migrated to another host (either manually or automatically by DRS) or shut down. The host is in a state of **Entering Maintenance Mode** until all running virtual machines are powered down or migrated to different hosts. You cannot power on virtual machines or migrate virtual machines to a host entering maintenance mode.

When no more running virtual machines are on the host, the host's icon changes to include **under maintenance** and the host's Summary panel indicates the new state. While in maintenance mode, the host does not allow you to deploy or power on a virtual machine.

---

**NOTE** DRS does not recommend (or perform, in fully automated mode) any virtual machine migrations off of a host entering maintenance or standby mode if the vSphere HA failover level would be violated after the host enters the requested mode.

---

### Procedure

- 1 Browse to the host in the vSphere Web Client navigator.
- 2 Right-click the host and select **Enter Maintenance Mode**.
  - If the host is part of a partially automated or manual DRS cluster, a list of migration recommendations for virtual machines running on the host appears.

- If the host is part of an automated DRS cluster, virtual machines are migrated to different hosts when the host enters maintenance mode.
- 3 If applicable, click **Yes**.

The host is in maintenance mode until you select **Exit Maintenance Mode**.

## Remove a Host from a Cluster

You can remove hosts from a cluster.

### Procedure

- 1 Browse to the host in the vSphere Web Client navigator.
- 2 Right-click the host and select **Enter Maintenance Mode**.

When the host is in maintenance mode, move it to a different inventory location, either the top-level data center or to a different cluster.

- 3 Right-click the host and select **Move To**.
- 4 Select a new location for the host and click **OK**.

When you move the host, its resources are removed from the cluster. If you grafted the host's resource pool hierarchy onto the cluster, that hierarchy remains with the cluster.

### What to do next

After you remove a host from a cluster, you can perform the following tasks.

- Remove the host from vCenter Server.
- Run the host as a standalone host under vCenter Server.
- Move the host into another cluster.

## Using Standby Mode

When a host machine is placed in standby mode, it is powered off.

Normally, hosts are placed in standby mode by the vSphere DPM feature to optimize power usage. You can also place a host in standby mode manually. However, DRS might undo (or recommend undoing) your change the next time it runs. To force a host to remain off, place it in maintenance mode and power it off.

## DRS Cluster Validity

The vSphere Web Client indicates whether a DRS cluster is valid, overcommitted (yellow), or invalid (red).

DRS clusters become overcommitted or invalid for several reasons.

- A cluster might become overcommitted if a host fails.
- A cluster becomes invalid if vCenter Server is unavailable and you power on virtual machines using the vSphere Web Client.
- A cluster becomes invalid if the user reduces the reservation on a parent resource pool while a virtual machine is in the process of failing over.
- If changes are made to hosts or virtual machines using the vSphere Web Client while vCenter Server is unavailable, those changes take effect. When vCenter Server becomes available again, you might find that clusters have turned red or yellow because cluster requirements are no longer met.

When considering cluster validity scenarios, you should understand these terms.

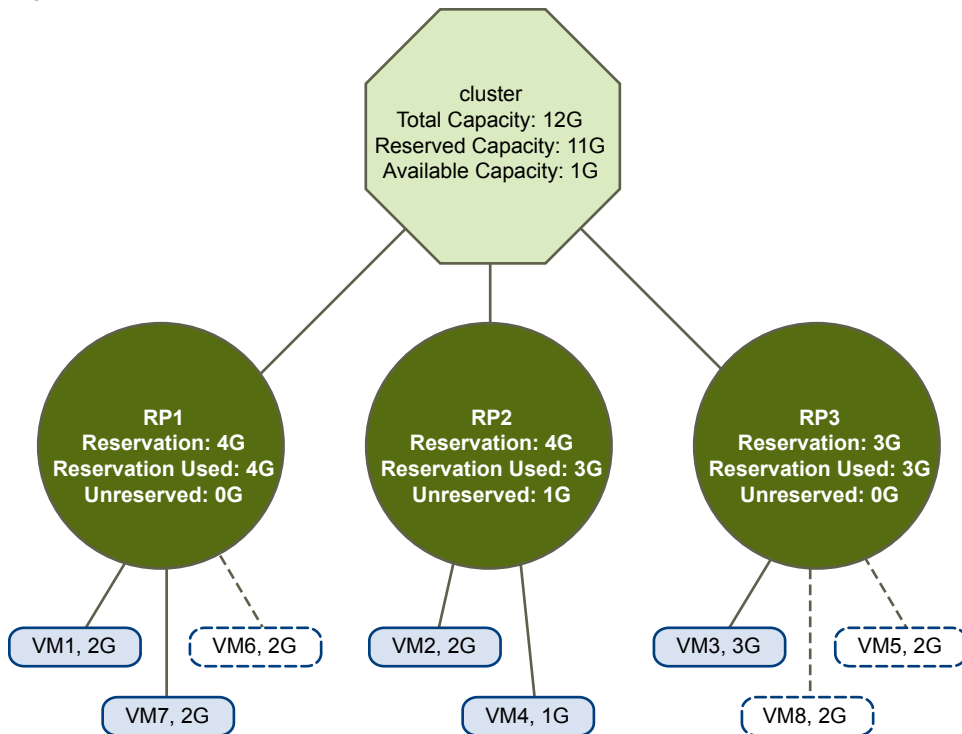
<b>Reservation</b>	A fixed, guaranteed allocation for the resource pool input by the user.
<b>Reservation Used</b>	The sum of the reservation or reservation used (whichever is larger) for each child resource pool, added recursively.
<b>Unreserved</b>	This nonnegative number differs according to resource pool type. <ul style="list-style-type: none"> <li>■ Nonexpandable resource pools: Reservation minus reservation used.</li> <li>■ Expandable resource pools: (Reservation minus reservation used) plus any unreserved resources that can be borrowed from its ancestor resource pools.</li> </ul>

## Valid DRS Clusters

A valid cluster has enough resources to meet all reservations and to support all running virtual machines.

The following figure shows an example of a valid cluster with fixed resource pools and how its CPU and memory resources are computed.

**Figure 11-1.** Valid Cluster with Fixed Resource Pools




The cluster has the following characteristics:

- A cluster with total resources of 12GHz.
- Three resource pools, each of type **Fixed** (**Expandable Reservation** is not selected).
- The total reservation of the three resource pools combined is 11GHz (4+4+3 GHz). The total is shown in the **Reserved Capacity** field for the cluster.
- RP1 was created with a reservation of 4GHz. Two virtual machines. (VM1 and VM7) of 2GHz each are powered on (**Reservation Used**: 4GHz). No resources are left for powering on additional virtual machines. VM6 is shown as not powered on. It consumes none of the reservation.

- RP2 was created with a reservation of 4GHz. Two virtual machines of 1GHz and 2GHz are powered on (**Reservation Used**: 3GHz). 1GHz remains unreserved.
- RP3 was created with a reservation of 3GHz. One virtual machine with 3GHz is powered on. No resources for powering on additional virtual machines are available.

The following figure shows an example of a valid cluster with some resource pools (RP1 and RP3) using reservation type **Expandable**.

**Figure 11-2.** Valid Cluster with Expandable Resource Pools



A valid cluster can be configured as follows:

- A cluster with total resources of 16GHz.
- RP1 and RP3 are of type **Expandable**, RP2 is of type **Fixed**.
- The total reservation used of the three resource pools combined is 16GHz (6GHz for RP1, 5GHz for RP2, and 5GHz for RP3). 16GHz shows up as the **Reserved Capacity** for the cluster at top level.
- RP1 was created with a reservation of 4GHz. Three virtual machines of 2GHz each are powered on. Two of those virtual machines (for example, VM1 and VM7) can use RP1's reservations, the third virtual machine (VM6) can use reservations from the cluster's resource pool. (If the type of this resource pool were **Fixed**, you could not power on the additional virtual machine.)
- RP2 was created with a reservation of 5GHz. Two virtual machines of 1GHz and 2GHz are powered on (**Reservation Used**: 3GHz). 2GHz remains unreserved.


RP3 was created with a reservation of 5GHz. Two virtual machines of 3GHz and 2GHz are powered on. Even though this resource pool is of type **Expandable**, no additional 2GHz virtual machine can be powered on because the parent's extra resources are already used by RP1.

## Overcommitted DRS Clusters

A cluster becomes overcommitted (yellow) when the tree of resource pools and virtual machines is internally consistent but the cluster does not have the capacity to support all resources reserved by the child resource pools.

There will always be enough resources to support all running virtual machines because, when a host becomes unavailable, all its virtual machines become unavailable. A cluster typically turns yellow when cluster capacity is suddenly reduced, for example, when a host in the cluster becomes unavailable. VMware recommends that you leave adequate additional cluster resources to avoid your cluster turning yellow.

**Figure 11-3.** Yellow Cluster



In this example:

- A cluster with total resources of 12GHz coming from three hosts of 4GHz each.
- Three resource pools reserving a total of 12GHz.
- The total reservation used by the three resource pools combined is 12GHz (4+5+3 GHz). That shows up as the **Reserved Capacity** in the cluster.
- One of the 4GHz hosts becomes unavailable, so total resources reduce to 8GHz.
- At the same time, VM4 (1GHz) and VM3 (3GHz), which were running on the host that failed, are no longer running.
- The cluster is now running virtual machines that require a total of 6GHz. The cluster still has 8GHz available, which is sufficient to meet virtual machine requirements.

The resource pool reservations of 12GHz can no longer be met, so the cluster is marked as yellow.

## Invalid DRS Clusters

A cluster enabled for DRS becomes invalid (red) when the tree is no longer internally consistent, that is, resource constraints are not observed.

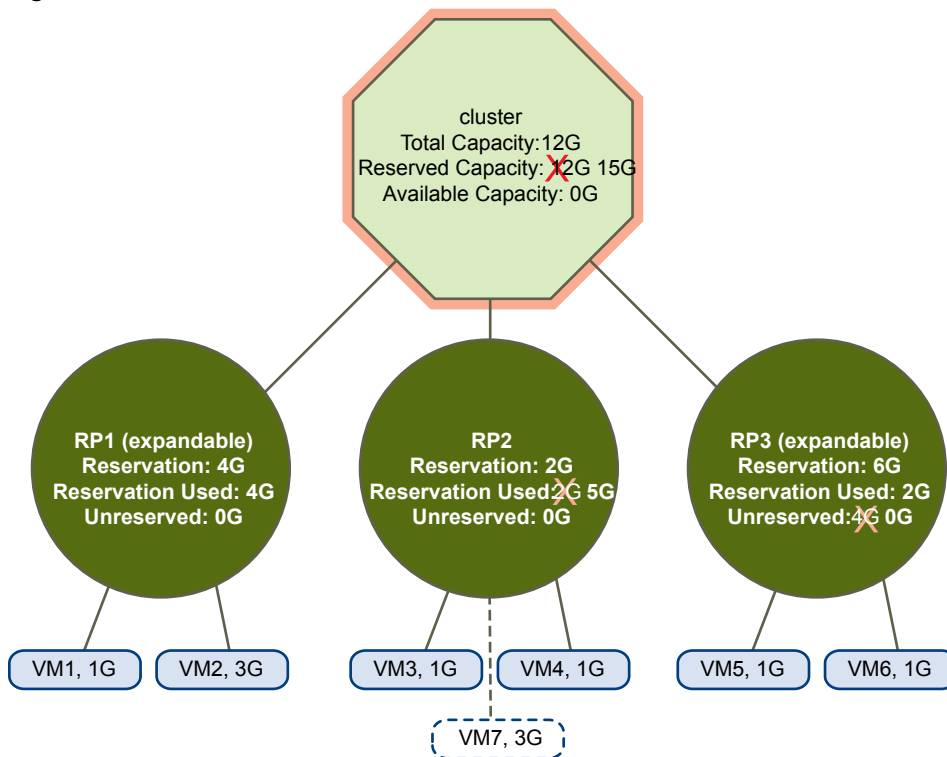
The total amount of resources in the cluster does not affect whether the cluster is red. A cluster can be red, even if enough resources exist at the root level, if there is an inconsistency at a child level.

You can resolve a red DRS cluster problem either by powering off one or more virtual machines, moving virtual machines to parts of the tree that have sufficient resources, or editing the resource pool settings in the red part. Adding resources typically helps only when you are in the yellow state.

A cluster can also turn red if you reconfigure a resource pool while a virtual machine is failing over. A virtual machine that is failing over is disconnected and does not count toward the reservation used by the parent resource pool. You might reduce the reservation of the parent resource pool before the failover completes. After the failover is complete, the virtual machine resources are again charged to the parent resource pool. If the pool's usage becomes larger than the new reservation, the cluster turns red.

If a user is able to start a virtual machine (in an unsupported way) with a reservation of 3GHz under resource pool 2, the cluster would become red, as shown in the following figure.

**Figure 11-4.** Red Cluster



## Managing Power Resources

The vSphere Distributed Power Management (DPM) feature allows a DRS cluster to reduce its power consumption by powering hosts on and off based on cluster resource utilization.

vSphere DPM monitors the cumulative demand of all virtual machines in the cluster for memory and CPU resources and compares this to the total available resource capacity of all hosts in the cluster. If sufficient excess capacity is found, vSphere DPM places one or more hosts in standby mode and powers them off after migrating their virtual machines to other hosts. Conversely, when capacity is deemed to be inadequate, DRS brings hosts out of standby mode (powers them on) and uses vMotion to migrate virtual machines to them. When making these calculations, vSphere DPM considers not only current demand, but it also honors any user-specified virtual machine resource reservations.

---

**NOTE** ESXi hosts cannot automatically be brought out of standby mode unless they are running in a cluster managed by vCenter Server.

---

vSphere DPM can use one of three power management protocols to bring a host out of standby mode: Intelligent Platform Management Interface (IPMI), Hewlett-Packard Integrated Lights-Out (iLO), or Wake-On-LAN (WOL). Each protocol requires its own hardware support and configuration. If a host does not support any of these protocols it cannot be put into standby mode by vSphere DPM. If a host supports multiple protocols, they are used in the following order: IPMI, iLO, WOL.

---

**NOTE** Do not disconnect a host in standby mode or move it out of the DRS cluster without first powering it on, otherwise vCenter Server is not able to power the host back on.

---

### Configure IPMI or iLO Settings for vSphere DPM

IPMI is a hardware-level specification and Hewlett-Packard iLO is an embedded server management technology. Each of them describes and provides an interface for remotely monitoring and controlling computers.

You must perform the following procedure on each host.

#### Prerequisites

Both IPMI and iLO require a hardware Baseboard Management Controller (BMC) to provide a gateway for accessing hardware control functions, and allow the interface to be accessed from a remote system using serial or LAN connections. The BMC is powered-on even when the host itself is powered-off. If properly enabled, the BMC can respond to remote power-on commands.

If you plan to use IPMI or iLO as a wake protocol, you must configure the BMC. BMC configuration steps vary according to model. See your vendor's documentation for more information. With IPMI, you must also ensure that the BMC LAN channel is configured to be always available and to allow operator-privileged commands. On some IPMI systems, when you enable "IPMI over LAN" you must configure this in the BIOS and specify a particular IPMI account.

vSphere DPM using only IPMI supports MD5- and plaintext-based authentication, but MD2-based authentication is not supported. vCenter Server uses MD5 if a host's BMC reports that it is supported and enabled for the Operator role. Otherwise, plaintext-based authentication is used if the BMC reports it is supported and enabled. If neither MD5 nor plaintext authentication is enabled, IPMI cannot be used with the host and vCenter Server attempts to use Wake-on-LAN.

#### Procedure

- 1 Browse to the host in the vSphere Web Client navigator.
- 2 Click the **Manage** tab and click **Settings**.
- 3 Under System, click **Power Management**.



- 4 Click **Edit**.
- 5 Enter the following information.
  - User name and password for a BMC account. (The user name must have the ability to remotely power the host on.)
  - IP address of the NIC associated with the BMC, as distinct from the IP address of the host. The IP address should be static or a DHCP address with infinite lease.
  - MAC address of the NIC associated with the BMC.
- 6 Click **OK**.

## Test Wake-on-LAN for vSphere DPM

The use of Wake-on-LAN (WOL) for the vSphere DPM feature is fully supported, if you configure and successfully test it according to the VMware guidelines. You must perform these steps before enabling vSphere DPM for a cluster for the first time or on any host that is being added to a cluster that is using vSphere DPM.

### Prerequisites

Before testing WOL, ensure that your cluster meets the prerequisites.

- Your cluster must contain at least two hosts that are version ESX 3.5 (or ESX 3i version 3.5) or later.
- Each host's vMotion networking link must be working correctly. The vMotion network should also be a single IP subnet, not multiple subnets separated by routers.
- The vMotion NIC on each host must support WOL. To check for WOL support, first determine the name of the physical network adapter corresponding to the VMkernel port by selecting the host in the inventory panel of the vSphere Web Client, selecting the **Configuration** tab, and clicking **Networking**. After you have this information, click on **Network Adapters** and find the entry corresponding to the network adapter. The **Wake On LAN Supported** column for the relevant adapter should show Yes.
- To display the WOL-compatibility status for each NIC on a host, select the host in the inventory panel of the vSphere Web Client, select the **Configuration** tab, and click **Network Adapters**. The NIC must show Yes in the **Wake On LAN Supported** column.
- The switch port that each WOL-supporting vMotion NIC is plugged into should be set to auto negotiate the link speed, and not set to a fixed speed (for example, 1000 Mb/s). Many NICs support WOL only if they can switch to 100 Mb/s or less when the host is powered off.

After you verify these prerequisites, test each ESXi host that is going to use WOL to support vSphere DPM. When you test these hosts, ensure that the vSphere DPM feature is disabled for the cluster.



**CAUTION** Ensure that any host being added to a vSphere DPM cluster that uses WOL as a wake protocol is tested and disabled from using power management if it fails the testing. If this is not done, vSphere DPM might power off hosts that it subsequently cannot power back up.

### Procedure

- 1 Browse to the host in the vSphere Web Client navigator.
- 2 Right-click the host and select **Enter Standby Mode**  
This action powers down the host.
- 3 Right-click the host and select **Power On** to attempt to bring it out of standby mode.
- 4 Observe whether or not the host successfully powers back on.

- 5 For any host that fails to exit standby mode successfully, perform the following steps.
  - a Select the host in the vSphere Web Client navigator and select the **Manage** tab.
  - b Under Power Management, click **Edit** to adjust the power management policy.
 After you do this, vSphere DPM does not consider that host a candidate for being powered off.

## Enabling vSphere DPM for a DRS Cluster

After you have performed configuration or testing steps required by the wake protocol you are using on each host, you can enable vSphere DPM.

Configure the power management automation level, threshold, and host-level overrides. These settings are configured under **Power Management** in the cluster's Settings dialog box.

You can also create scheduled tasks to enable and disable DPM for a cluster using the Schedule Task: Change Cluster Power Settings wizard.

---

**NOTE** If a host in your DRS cluster has USB devices connected, disable DPM for that host. Otherwise, DPM might turn off the host and sever the connection between the device and the virtual machine that was using it.

---

### Automation Level

Whether the host power state and migration recommendations generated by vSphere DPM are executed automatically or not depends upon the power management automation level selected for the feature.

The automation level is configured under **Power Management** in the cluster's Settings dialog box.

---

**NOTE** The power management automation level is not the same as the DRS automation level.

---

**Table 11-1.** Power Management Automation Level

Option	Description
Off	The feature is disabled and no recommendations will be made.
Manual	Host power operation and related virtual machine migration recommendations are made, but not automatically executed. These recommendations appear on the cluster's <b>DRS</b> tab in the vSphere Web Client.
Automatic	Host power operations are automatically executed if related virtual machine migrations can all be executed automatically.

### vSphere DPM Threshold

The power state (host power on or off) recommendations generated by the vSphere DPM feature are assigned priorities that range from priority-one recommendations to priority-five recommendations.

These priority ratings are based on the amount of over- or under-utilization found in the DRS cluster and the improvement that is expected from the intended host power state change. A priority-one recommendation is mandatory, while a priority-five recommendation brings only slight improvement.

The threshold is configured under **Power Management** in the cluster's Settings dialog box. Each level you move the vSphere DPM Threshold slider to the right allows the inclusion of one more lower level of priority in the set of recommendations that are executed automatically or appear as recommendations to be manually executed. At the Conservative setting, vSphere DPM only generates priority-one recommendations, the next level to the right only priority-two and higher, and so on, down to the Aggressive level which generates priority-five recommendations and higher (that is, all recommendations.)

---

**NOTE** The DRS threshold and the vSphere DPM threshold are essentially independent. You can differentiate the aggressiveness of the migration and host-power-state recommendations they respectively provide.

---

## Host-Level Overrides

When you enable vSphere DPM in a DRS cluster, by default all hosts in the cluster inherit its vSphere DPM automation level.

You can override this default for an individual host by selecting the Host Options page of the cluster's Settings dialog box and clicking its **Power Management** setting. You can change this setting to the following options:

- Disabled
- Manual
- Automatic

---

**NOTE** Do not change a host's Power Management setting if it has been set to Disabled due to failed exit standby mode testing.

---

After enabling and running vSphere DPM, you can verify that it is functioning properly by viewing each host's **Last Time Exited Standby** information displayed on the Host Options page in the cluster Settings dialog box and on the **Hosts** tab for each cluster. This field shows a timestamp and whether vCenter Server Succeeded or Failed the last time it attempted to bring the host out of standby mode. If no such attempt has been made, the field displays Never.

---

**NOTE** Times for the **Last Time Exited Standby** text box are derived from the vCenter Server event log. If this log is cleared, the times are reset to Never.

---

## Monitoring vSphere DPM

You can use event-based alarms in vCenter Server to monitor vSphere DPM.

The most serious potential error you face when using vSphere DPM is the failure of a host to exit standby mode when its capacity is needed by the DRS cluster. You can monitor for instances when this error occurs by using the preconfigured **Exit Standby Error** alarm in vCenter Server. If vSphere DPM cannot bring a host out of standby mode (vCenter Server event `DrsExitStandbyModeFailedEvent`), you can configure this alarm to send an alert email to the administrator or to send notification using an SNMP trap. By default, this alarm is cleared after vCenter Server is able to successfully connect to that host.

To monitor vSphere DPM activity, you can also create alarms for the following vCenter Server events.

**Table 11-2.** vCenter Server Events

Event Type	Event Name
Entering Standby mode (about to power off host)	DrsEnteringStandbyModeEvent
Successfully entered Standby mode (host power off succeeded)	DrsEnteredStandbyModeEvent

**Table 11-2.** vCenter Server Events (Continued)

Event Type	Event Name
Exiting Standby mode (about to power on the host)	DrsExitingStandbyModeEvent
Successfully exited Standby mode (power on succeeded)	DrsExitedStandbyModeEvent

For more information about creating and editing alarms, see the *vSphere Monitoring and Performance* documentation.

If you use monitoring software other than vCenter Server, and that software triggers alarms when physical hosts are powered off unexpectedly, you might have a situation where false alarms are generated when vSphere DPM places a host into standby mode. If you do not want to receive such alarms, work with your vendor to deploy a version of the monitoring software that is integrated with vCenter Server. You could also use vCenter Server itself as your monitoring solution, because starting with vSphere 4.x, it is inherently aware of vSphere DPM and does not trigger these false alarms.

## Using DRS Affinity Rules

You can control the placement of virtual machines on hosts within a cluster by using affinity rules.

You can create two types of rules.

- Used to specify affinity or anti-affinity between a group of virtual machines and a group of hosts. An affinity rule specifies that the members of a selected virtual machine DRS group can or must run on the members of a specific host DRS group. An anti-affinity rule specifies that the members of a selected virtual machine DRS group cannot run on the members of a specific host DRS group.

See “[VM-Host Affinity Rules](#),” on page 86 for information about creating and using this type of rule.

- Used to specify affinity or anti-affinity between individual virtual machines. A rule specifying affinity causes DRS to try to keep the specified virtual machines together on the same host, for example, for performance reasons. With an anti-affinity rule, DRS tries to keep the specified virtual machines apart, for example, so that when a problem occurs with one host, you do not lose both virtual machines.

See “[VM-VM Affinity Rules](#),” on page 85 for information about creating and using this type of rule.

When you add or edit an affinity rule, and the cluster's current state is in violation of the rule, the system continues to operate and tries to correct the violation. For manual and partially automated DRS clusters, migration recommendations based on rule fulfillment and load balancing are presented for approval. You are not required to fulfill the rules, but the corresponding recommendations remain until the rules are fulfilled.

To check whether any enabled affinity rules are being violated and cannot be corrected by DRS, select the cluster's **DRS** tab and click **Faults**. Any rule currently being violated has a corresponding fault on this page. Read the fault to determine why DRS is not able to satisfy the particular rule. Rules violations also produce a log event.

---

**NOTE** VM-VM and VM-Host affinity rules are different from an individual host's CPU affinity rules.

---

## Create a Host DRS Group

A VM-Host affinity rule establishes an affinity (or anti-affinity) relationship between a virtual machine DRS group with a host DRS group. You must create both of these groups before you can create a rule that links them.

### Procedure

- 1 Browse to the cluster in the vSphere Web Client navigator.
- 2 Click the **Manage** tab.

- 3 Click **Settings**, and click **DRS Groups**.
- 4 In the DRS Groups section, click **Add**.
- 5 In the Create DRS Group dialog box, type a name for the group.
- 6 Select Host DRS Group from the **Type** drop down box and click **Add**.
- 7 Click the check box next to a host to add it. Continue this process until all desired hosts have been added.
- 8 Click **OK**.

#### What to do next

Using this host DRS group, you can create a VM-Host affinity rule that establishes an affinity (or anti-affinity) relationship with an appropriate virtual machine DRS group.

[“Create a Virtual Machine DRS Group,”](#) on page 85

[“Create a VM-Host Affinity Rule,”](#) on page 87

## Create a Virtual Machine DRS Group

Affinity rules establish an affinity (or anti-affinity) relationship between DRS groups. You must create DRS groups before you can create a rule that links them.

#### Procedure

- 1 Browse to the cluster in the vSphere Web Client navigator.
- 2 Click the **Manage** tab.
- 3 Click **Settings**, and click **DRS Groups**.
- 4 In the DRS Groups section, click **Add**.
- 5 In the Create DRS Group dialog box, type a name for the group.
- 6 Select VM DRS Group from the **Type** drop down box and click **Add**.
- 7 Click the check box next to a virtual machine to add it. Continue this process until all desired virtual machines have been added.
- 8 Click **OK**.

#### What to do next

[“Create a Host DRS Group,”](#) on page 84

[“Create a VM-Host Affinity Rule,”](#) on page 87

[“Create a VM-VM Affinity Rule,”](#) on page 86

## VM-VM Affinity Rules

A VM-VM affinity rule specifies whether selected individual virtual machines should run on the same host or be kept on separate hosts. This type of rule is used to create affinity or anti-affinity between individual virtual machines that you select.

When an affinity rule is created, DRS tries to keep the specified virtual machines together on the same host. You might want to do this, for example, for performance reasons.

With an anti-affinity rule, DRS tries to keep the specified virtual machines apart. You could use such a rule if you want to guarantee that certain virtual machines are always on different physical hosts. In that case, if a problem occurs with one host, not all virtual machines would be placed at risk.

## Create a VM-VM Affinity Rule

You can create VM-VM affinity rules to specify whether selected individual virtual machines should run on the same host or be kept on separate hosts.

---

**NOTE** If you use the vSphere HA Specify Failover Hosts admission control policy and designate multiple failover hosts, VM-VM affinity rules are not supported.

---

### Procedure

- 1 Browse to the cluster in the vSphere Web Client navigator.
- 2 Click the **Manage** tab.
- 3 Click **Settings** and click **DRS Rules**.
- 4 Click **Add**.
- 5 In the Create DRS Rule dialog box, type a name for the rule.
- 6 From the **Type** drop-down menu, select either **Keep Virtual Machines Together** or **Separate Virtual Machines**.
- 7 Click **Add**.
- 8 Select at least two virtual machines to which the rule will apply and click **OK**.
- 9 Click **OK**.

## VM-VM Affinity Rule Conflicts

You can create and use multiple VM-VM affinity rules, however, this might lead to situations where the rules conflict with one another.

If two VM-VM affinity rules are in conflict, you cannot enable both. For example, if one rule keeps two virtual machines together and another rule keeps the same two virtual machines apart, you cannot enable both rules. Select one of the rules to apply and disable or remove the conflicting rule.

When two VM-VM affinity rules conflict, the older one takes precedence and the newer rule is disabled. DRS only tries to satisfy enabled rules and disabled rules are ignored. DRS gives higher precedence to preventing violations of anti-affinity rules than violations of affinity rules.

## VM-Host Affinity Rules

A VM-Host affinity rule specifies whether or not the members of a selected virtual machine DRS group can run on the members of a specific host DRS group.

Unlike a VM-VM affinity rule, which specifies affinity (or anti-affinity) between individual virtual machines, a VM-Host affinity rule specifies an affinity relationship between a group of virtual machines and a group of hosts. There are 'required' rules (designated by "must") and 'preferential' rules (designated by "should").

A VM-Host affinity rule includes the following components.

- One virtual machine DRS group.
- One host DRS group.
- A designation of whether the rule is a requirement ("must") or a preference ("should") and whether it is affinity ("run on") or anti-affinity ("not run on").

Because VM-Host affinity rules are cluster-based, the virtual machines and hosts that are included in a rule must all reside in the same cluster. If a virtual machine is removed from the cluster, it loses its DRS group affiliation, even if it is later returned to the cluster.

## Create a VM-Host Affinity Rule

You can create VM-Host affinity rules to specify whether or not the members of a selected virtual machine DRS group can run on the members of a specific host DRS group.

### Prerequisites

Create the virtual machine and host DRS groups to which the VM-Host affinity rule applies.

### Procedure

- 1 Browse to the cluster in the vSphere Web Client navigator.
- 2 Click the **Manage** tab.
- 3 Click **Settings** and click **DRS Rules**.
- 4 Click **Add**.
- 5 In the Create DRS Rule dialog box, type a name for the rule.
- 6 From the **Type** drop down menu, select **Virtual Machines to Hosts**.
- 7 Select the virtual machine DRS group and the host DRS group to which the rule applies.
- 8 Select a specification for the rule.
  - **Must run on hosts in group.** Virtual machines in VM Group 1 must run on hosts in Host Group A.
  - **Should run on hosts in group.** Virtual machines in VM Group 1 should, but are not required, to run on hosts in Host Group A.
  - **Must not run on hosts in group.** Virtual machines in VM Group 1 must never run on host in Host Group A.
  - **Should not run on hosts in group.** Virtual machines in VM Group 1 should not, but might, run on hosts in Host Group A.
- 9 Click **OK**.

## Using VM-Host Affinity Rules

You use a VM-Host affinity rule to specify an affinity relationship between a group of virtual machines and a group of hosts. When using VM-Host affinity rules, you should be aware of when they could be most useful, how conflicts between rules are resolved, and the importance of caution when setting required affinity rules.

One use case where VM-Host affinity rules are helpful is when the software you are running in your virtual machines has licensing restrictions. You can place such virtual machines into a DRS group and then create a rule that requires them to run on a host DRS group that contains only host machines that have the required licenses.

---

**NOTE** When you create a VM-Host affinity rule that is based on the licensing or hardware requirements of the software running in your virtual machines, you are responsible for ensuring that the groups are properly set up. The rule does not monitor the software running in the virtual machines nor does it know what non-VMware licenses are in place on which ESXi hosts.

---

If you create more than one VM-Host affinity rule, the rules are not ranked, but are applied equally. Be aware that this has implications for how the rules interact. For example, a virtual machine that belongs to two DRS groups, each of which belongs to a different required rule, can run only on hosts that belong to both of the host DRS groups represented in the rules.

When you create a VM-Host affinity rule, its ability to function in relation to other rules is not checked. So it is possible for you to create a rule that conflicts with the other rules you are using. When two VM-Host affinity rules conflict, the older one takes precedence and the newer rule is disabled. DRS only tries to satisfy enabled rules and disabled rules are ignored.

DRS, vSphere HA, and vSphere DPM never take any action that results in the violation of required affinity rules (those where the virtual machine DRS group 'must run on' or 'must not run on' the host DRS group). Accordingly, you should exercise caution when using this type of rule because of its potential to adversely affect the functioning of the cluster. If improperly used, required VM-Host affinity rules can fragment the cluster and inhibit the proper functioning of DRS, vSphere HA, and vSphere DPM.

A number of cluster functions are not performed if doing so would violate a required affinity rule.

- DRS does not evacuate virtual machines to place a host in maintenance mode.
- DRS does not place virtual machines for power-on or load balance virtual machines.
- vSphere HA does not perform failovers.
- vSphere DPM does not optimize power management by placing hosts into standby mode.

To avoid these situations, exercise caution when creating more than one required affinity rule or consider using VM-Host affinity rules that are preferential only (those where the virtual machine DRS group 'should run on' or 'should not run on' the host DRS group). Ensure that the number of hosts in the cluster with which each virtual machine is affined is large enough that losing a host does not result in a lack of hosts on which the virtual machine can run. Preferential rules can be violated to allow the proper functioning of DRS, vSphere HA, and vSphere DPM.

---

**NOTE** You can create an event-based alarm that is triggered when a virtual machine violates a VM-Host affinity rule. In the vSphere Web Client, add a new alarm for the virtual machine and select **VM is violating VM-Host Affinity Rule** as the event trigger. For more information about creating and editing alarms, see the vSphere Monitoring and Performance documentation.

---



## Creating a Datastore Cluster

---

A datastore cluster is a collection of datastores with shared resources and a shared management interface. Datastore clusters are to datastores what clusters are to hosts. When you create a datastore cluster, you can use vSphere Storage DRS to manage storage resources.

---

**NOTE** Datastore clusters are referred to as storage pods in the vSphere API.

---

When you add a datastore to a datastore cluster, the datastore's resources become part of the datastore cluster's resources. As with clusters of hosts, you use datastore clusters to aggregate storage resources, which enables you to support resource allocation policies at the datastore cluster level. The following resource management capabilities are also available per datastore cluster.

<b>Space utilization load balancing</b>	You can set a threshold for space use. When space use on a datastore exceeds the threshold, Storage DRS generates recommendations or performs Storage vMotion migrations to balance space use across the datastore cluster.
<b>I/O latency load balancing</b>	You can set an I/O latency threshold for bottleneck avoidance. When I/O latency on a datastore exceeds the threshold, Storage DRS generates recommendations or performs Storage vMotion migrations to help alleviate high I/O load.
<b>Anti-affinity rules</b>	You can create anti-affinity rules for virtual machine disks. For example, the virtual disks of a certain virtual machine must be kept on different datastores. By default, all virtual disks for a virtual machine are placed on the same datastore.

This chapter includes the following topics:

- [“Initial Placement and Ongoing Balancing,”](#) on page 90
- [“Storage Migration Recommendations,”](#) on page 90
- [“Create a Datastore Cluster,”](#) on page 90
- [“Enable and Disable Storage DRS,”](#) on page 91
- [“Set the Automation Level for Datastore Clusters,”](#) on page 91
- [“Setting the Aggressiveness Level for Storage DRS,”](#) on page 92
- [“Datastore Cluster Requirements,”](#) on page 93
- [“Adding and Removing Datastores from a Datastore Cluster,”](#) on page 94

## Initial Placement and Ongoing Balancing

Storage DRS provides initial placement and ongoing balancing recommendations to datastores in a Storage DRS-enabled datastore cluster.

Initial placement occurs when Storage DRS selects a datastore within a datastore cluster on which to place a virtual machine disk. This happens when the virtual machine is being created or cloned, when a virtual machine disk is being migrated to another datastore cluster, or when you add a disk to an existing virtual machine.

Initial placement recommendations are made in accordance with space constraints and with respect to the goals of space and I/O load balancing. These goals aim to minimize the risk of over-provisioning one datastore, storage I/O bottlenecks, and performance impact on virtual machines.

Storage DRS is invoked at the configured frequency (by default, every eight hours) or when one or more datastores in a datastore cluster exceeds the user-configurable space utilization thresholds. When Storage DRS is invoked, it checks each datastore's space utilization and I/O latency values against the threshold. For I/O latency, Storage DRS uses the 90th percentile I/O latency measured over the course of a day to compare against the threshold.

## Storage Migration Recommendations

vCenter Server displays migration recommendations on the Storage DRS Recommendations page for datastore clusters that have manual automation mode.

The system provides as many recommendations as necessary to enforce Storage DRS rules and to balance the space and I/O resources of the datastore cluster. Each recommendation includes the virtual machine name, the virtual disk name, the name of the datastore cluster, the source datastore, the destination datastore, and a reason for the recommendation.

- Balance datastore space use
- Balance datastore I/O load

Storage DRS makes mandatory recommendations for migration in the following situations:

- The datastore is out of space.
- Anti-affinity or affinity rules are being violated.
- The datastore is entering maintenance mode and must be evacuated.

In addition, optional recommendations are made when a datastore is close to running out of space or when adjustments should be made for space and I/O load balancing.

Storage DRS considers moving virtual machines that are powered off or powered on for space balancing. Storage DRS includes powered-off virtual machines with snapshots in these considerations.

## Create a Datastore Cluster

You can manage datastore cluster resources using Storage DRS.

### Procedure

- 1 Browse to Datacenters in the vSphere Web Client navigator.
- 2 Right-click the datacenter object and select **New Datastore Cluster**.
- 3 Follow the prompts to complete the New Datastore Cluster wizard.
- 4 Click **Finish**.

## Enable and Disable Storage DRS

Storage DRS allows you to manage the aggregated resources of a datastore cluster. When Storage DRS is enabled, it provides recommendations for virtual machine disk placement and migration to balance space and I/O resources across the datastores in the datastore cluster.

When you enable Storage DRS, you enable the following functions.

- Space load balancing among datastores within a datastore cluster.
- I/O load balancing among datastores within a datastore cluster.
- Initial placement for virtual disks based on space and I/O workload.

The Enable Storage DRS check box in the Datastore Cluster Settings dialog box enables or disables all of these components at once. If necessary, you can disable I/O-related functions of Storage DRS independently of space balancing functions.

When you disable Storage DRS on a datastore cluster, Storage DRS settings are preserved. When you enable Storage DRS, the settings for the datastore cluster are restored to the point where Storage DRS was disabled.

### Procedure

- 1 Browse to the datastore cluster in the vSphere Web Client navigator.
- 2 Click the **Manage** tab and click **Settings**.
- 3 Under Services, select **Storage DRS** and click **Edit**.
- 4 Select **Turn ON vSphere DRS** and click **OK**.
- 5 (Optional) To disable only I/O-related functions of Storage DRS, leaving space-related controls enabled, perform the following steps.
  - a Under **Storage DRS** select **Edit**.
  - b Deselect the **Enable I/O metric for Storage DRS** check box and click **OK**.

## Set the Automation Level for Datastore Clusters

The automation level for a datastore cluster specifies whether or not placement and migration recommendations from Storage DRS are applied automatically.

### Procedure

- 1 Browse to the datastore cluster in the vSphere Web Client navigator.
- 2 Click the **Manage** tab and click **Settings**.
- 3 Under Services, select **DRS** and click **Edit**.
- 4 Expand DRS Automation and select an automation level.

Manual is the default automation level.

Option	Description
<b>No Automation (Manual Mode)</b>	Placement and migration recommendations are displayed, but do not run until you manually apply the recommendation.
<b>Partially Automated</b>	Placement recommendations run automatically and migration recommendations are displayed, but do not run until you manually apply the recommendation.
<b>Fully Automated</b>	Placement and migration recommendations run automatically.

- 5 Click **OK**.

## Setting the Aggressiveness Level for Storage DRS

The aggressiveness of Storage DRS is determined by specifying thresholds for space used and I/O latency.

Storage DRS collects resource usage information for the datastores in a datastore cluster. vCenter Server uses this information to generate recommendations for placement of virtual disks on datastores.

When you set a low aggressiveness level for a datastore cluster, Storage DRS recommends Storage vMotion migrations only when absolutely necessary, for example, if when I/O load, space utilization, or their imbalance is high. When you set a high aggressiveness level for a datastore cluster, Storage DRS recommends migrations whenever the datastore cluster can benefit from space or I/O load balancing.

In the vSphere Web Client, you can use the following thresholds to set the aggressiveness level for Storage DRS:

<b>Space Utilization</b>	Storage DRS generates recommendations or performs migrations when the percentage of space utilization on the datastore is greater than the threshold you set in the vSphere Web Client.
<b>I/O Latency</b>	Storage DRS generates recommendations or performs migrations when the 90th percentile I/O latency measured over a day for the datastore is greater than the threshold.

You can also set advanced options to further configure the aggressiveness level of Storage DRS.

<b>Space utilization difference</b>	This threshold ensures that there is some minimum difference between the space utilization of the source and the destination. For example, if the space used on datastore A is 82% and datastore B is 79%, the difference is 3. If the threshold is 5, Storage DRS will not make migration recommendations from datastore A to datastore B.
<b>I/O load balancing invocation interval</b>	After this interval, Storage DRS runs to balance I/O load.
<b>I/O imbalance threshold</b>	Lowering this value makes I/O load balancing less aggressive. Storage DRS computes an I/O fairness metric between 0 and 1, which 1 being the fairest distribution. I/O load balancing runs only if the computed metric is less than $1 - (\text{I/O imbalance threshold} / 100)$ .

## Set Storage DRS Runtime Rules

Set Storage DRS triggers and configure advanced options for the datastore cluster.

### Procedure

- 1 (Optional) Select or deselect the **Enable I/O metric for SDRS recommendations** check box to enable or disable I/O metric inclusion.

When you disable this option, vCenter Server does not consider I/O metrics when making Storage DRS recommendations. When you disable this option, you disable the following elements of Storage DRS:

- I/O load balancing among datastores within a datastore cluster.
- Initial placement for virtual disks based on I/O workload. Initial placement is based on space only.

## 2 (Optional) Set Storage DRS thresholds.

You set the aggressiveness level of Storage DRS by specifying thresholds for used space and I/O latency.

- Use the Utilized Space slider to indicate the maximum percentage of consumed space allowed before Storage DRS is triggered. Storage DRS makes recommendations and performs migrations when space use on the datastores is higher than the threshold.
- Use the I/O Latency slider to indicate the maximum I/O latency allowed before Storage DRS is triggered. Storage DRS makes recommendations and performs migrations when latency is higher than the threshold.

---

**NOTE** The Storage DRS I/O Latency threshold for the datastore cluster should be lower than or equal to the Storage I/O Control congestion threshold.

---

## 3 (Optional) Configure advanced options.

- No recommendations until utilization difference between source and destination is: Use the slider to specify the space utilization difference threshold. Utilization is usage \* 100/capacity.

This threshold ensures that there is some minimum difference between the space utilization of the source and the destination. For example, if the space used on datastore A is 82% and datastore B is 79%, the difference is 3. If the threshold is 5, Storage DRS will not make migration recommendations from datastore A to datastore B.

- Check imbalances every: Specify how often Storage DRS should assess space and I/O load balancing.
- I/O imbalance threshold: Use the slider to indicate the aggressiveness of I/O load balancing. Lowering this value makes I/O load balancing less aggressive. Storage DRS computes an I/O fairness metric between 0 and 1, which 1 being the fairest distribution. I/O load balancing runs only if the computed metric is less than  $1 - (I/O \text{ imbalance threshold} / 100)$ .

4 Click **OK**.

## Datastore Cluster Requirements

Datastores and hosts that are associated with a datastore cluster must meet certain requirements to use datastore cluster features successfully.

Follow these guidelines when you create a datastore cluster.

- Datastore clusters must contain similar or interchangeable datastores.

A datastore cluster can contain a mix of datastores with different sizes and I/O capacities, and can be from different arrays and vendors. However, the following types of datastores cannot coexist in a datastore cluster.

- NFS and VMFS datastores cannot be combined in the same datastore cluster.
- Replicated datastores cannot be combined with non-replicated datastores in the same Storage-DRS-enabled datastore cluster.
- All hosts attached to the datastores in a datastore cluster must be ESXi 5.0 and later. If datastores in the datastore cluster are connected to ESX/ESXi 4.x and earlier hosts, Storage DRS does not run.
- Datastores shared across multiple data centers cannot be included in a datastore cluster.
- As a best practice, do not include datastores that have hardware acceleration enabled in the same datastore cluster as datastores that do not have hardware acceleration enabled. Datastores in a datastore cluster must be homogeneous to guarantee hardware acceleration-supported behavior.

## Adding and Removing Datastores from a Datastore Cluster

You add and remove datastores to and from an existing datastore cluster.

You can add to a datastore cluster any datastore that is mounted on a host in the vSphere Web Client inventory, with the following exceptions:

- All hosts attached to the datastore must be ESXi 5.0 and later.
- The datastore cannot be in more than one data center in the same instance of the vSphere Web Client.

When you remove a datastore from a datastore cluster, the datastore remains in the vSphere Web Client inventory and is not unmounted from the host.

# Using Datastore Clusters to Manage Storage Resources

# 13

After you create a datastore cluster, you can customize it and use it to manage storage I/O and space utilization resources.

This chapter includes the following topics:

- [“Using Storage DRS Maintenance Mode,”](#) on page 95
- [“Applying Storage DRS Recommendations,”](#) on page 97
- [“Change Storage DRS Automation Level for a Virtual Machine,”](#) on page 98
- [“Set Up Off-Hours Scheduling for Storage DRS,”](#) on page 98
- [“Storage DRS Anti-Affinity Rules,”](#) on page 99
- [“Clear Storage DRS Statistics,”](#) on page 102
- [“Storage vMotion Compatibility with Datastore Clusters,”](#) on page 103

## Using Storage DRS Maintenance Mode

You place a datastore in maintenance mode when you need to take it out of use to service it. A datastore enters or leaves maintenance mode only as the result of a user request.

Maintenance mode is available to datastores within a Storage DRS-enabled datastore cluster. Standalone datastores cannot be placed in maintenance mode.

Virtual disks that are located on a datastore that is entering maintenance mode must be migrated to another datastore, either manually or using Storage DRS. When you attempt to put a datastore in maintenance mode, the **Placement Recommendations** tab displays a list of migration recommendations, datastores within the same datastore cluster where virtual disks can be migrated. On the **Faults** tab, vCenter Server displays a list of the disks that cannot be migrated and the reasons why. If Storage DRS affinity or anti-affinity rules prevent disks from being migrated, you can choose to enable the Ignore Affinity Rules for Maintenance option.

The datastore is in a state of Entering Maintenance Mode until all virtual disks have been migrated.

### Place a Datastore in Maintenance Mode

If you need to take a datastore out of service, you can place the datastore in Storage DRS maintenance mode.

#### Prerequisites

Storage DRS is enabled on the datastore cluster that contains the datastore that is entering maintenance mode.

No CD-ROM image files are stored on the datastore.

There are at least two datastores in the datastore cluster.

### Procedure

- 1 Browse to the datastore in the vSphere Web Client navigator.
- 2 Right-click the datastore and select **Enter Storage DRS Maintenance Mode**.  
A list of recommendations appears for datastore maintenance mode migration.
- 3 (Optional) On the Placement Recommendations tab, deselect any recommendations you do not want to apply.

---

**NOTE** The datastore cannot enter maintenance mode without evacuating all disks. If you deselect recommendations, you must manually move the affected virtual machines.

---

- 4 If necessary, click **Apply Recommendations**.  
vCenter Server uses Storage vMotion to migrate the virtual disks from the source datastore to the destination datastore and the datastore enters maintenance mode.

The datastore icon might not be immediately updated to reflect the datastore's current state. To update the icon immediately, click **Refresh**.

## Ignore Storage DRS Affinity Rules for Maintenance Mode

Storage DRS affinity or anti-affinity rules might prevent a datastore from entering maintenance mode. You can ignore these rules when you put a datastore in maintenance mode.

When you enable the Ignore Affinity Rules for Maintenance option for a datastore cluster, vCenter Server ignores Storage DRS affinity and anti-affinity rules that prevent a datastore from entering maintenance mode.

Storage DRS rules are ignored only for evacuation recommendations. vCenter Server does not violate the rules when making space and load balancing recommendations or initial placement recommendations.

### Procedure

- 1 Browse to the datastore cluster in the vSphere Web Client navigator.
- 2 Click the **Manage** tab and click **Settings**.
- 3 Under Services, select **DRS** and click **Edit**.
- 4 Expand **Advanced Options** and click **Add**.
- 5 In the Option column, type **IgnoreAffinityRulesForMaintenance**.
- 6 In the Value column, type **1** to enable the option.  
Type **0** to disable the option.
- 7 Click **OK**.

The Ignore Affinity Rules for Maintenance Mode option is applied to the datastore cluster.



## Applying Storage DRS Recommendations

Storage DRS collects resource usage information for all datastores in a datastore cluster. Storage DRS uses the information to generate recommendations for virtual machine disk placement on datastores in a datastore cluster.

Storage DRS recommendations appear on the **Storage DRS** tab in the vSphere Web Client datastore view. Recommendations also appear when you attempt to put a datastore into Storage DRS maintenance mode. When you apply Storage DRS recommendations, vCenter Server uses Storage vMotion to migrate virtual machine disks to other datastores in the datastore cluster to balance the resources.

You can apply a subset of the recommendations by selecting the Override Suggested DRS Recommendations check box and selecting each recommendation to apply.

**Table 13-1.** Storage DRS Recommendations

Label	Description
Priority	Priority level (1-5) of the recommendation. (Hidden by default.)
Recommendation	Action being recommended by Storage DRS.
Reason	Why the action is needed.
Space Utilization % Before (source) and (destination)	Percentage of space used on the source and destination datastores before migration.
Space Utilization % After (source) and (destination)	Percentage of space used on the source and destination datastores after migration.
I/O Latency Before (source)	Value of I/O latency on the source datastore before migration.
I/O Latency Before (destination)	Value of I/O latency on the destination datastore before migration.

## Refresh Storage DRS Recommendations

Storage DRS migration recommendations appear on the **Storage DRS** tab in the vSphere Web Client. You can refresh these recommendations by running Storage DRS.

### Prerequisites

At least one datastore cluster must exist in the vSphere Web Client inventory.

Enable Storage DRS for the datastore cluster. The **Storage DRS** tab appears only if Storage DRS is enabled.

### Procedure

- 1 In the vSphere Web Client datastore view, select the datastore cluster and click the **Storage DRS** tab.
- 2 Select the **Recommendations** view and click the **Run Storage DRS** link in the upper right corner.

The recommendations are updated. The Last Updated timestamp displays the time when Storage DRS recommendations were refreshed.

## Change Storage DRS Automation Level for a Virtual Machine

You can override the datastore cluster-wide automation level for individual virtual machines. You can also override default virtual disk affinity rules.

### Procedure

- 1 Browse to the datastore cluster in the vSphere Web Client navigator.
- 2 Click the **Manage** tab and click **Settings**.
- 3 Under **VM Overrides**, select **Add**.
- 4 Select a virtual machine.
- 5 Click the Automation level drop-down menu, and select an automation level for the virtual machine.

Option	Description
<b>Default (Manual)</b>	Placement and migration recommendations are displayed, but do not run until you manually apply the recommendation.
<b>Fully Automated</b>	Placement and migration recommendations run automatically.
<b>Disabled</b>	vCenter Server does not migrate the virtual machine or provide migration recommendations for it.

- 6 Click the **Keep VMDKs together**, drop-down menu to override default VMDK affinity.  
See [“Override VMDK Affinity Rules,”](#) on page 101.
- 7 Click **OK**.

## Set Up Off-Hours Scheduling for Storage DRS

You can create a scheduled task to change Storage DRS settings for a datastore cluster so that migrations for fully automated datastore clusters are more likely to occur during off-peak hours.

You can create a scheduled task to change the automation level and aggressiveness level for a datastore cluster. For example, you might configure Storage DRS to run less aggressively during peak hours, when performance is a priority, to minimize the occurrence of storage migrations. During non-peak hours, Storage DRS can run in a more aggressive mode and be invoked more frequently.

### Prerequisites

Enable Storage DRS.

### Procedure

- 1 Browse to the datastore cluster in the vSphere Web Client navigator.
- 2 Click the **Manage** tab and click **Settings**.
- 3 Under **DRS** click the **Schedule DRS** button.
- 4 In the Edit Datastore Cluster dialog box, click **SDRS Scheduling**.

- 5 Expand **DRS Automation**.
  - a Select an automation level.
  - b Set the Migration threshold.  
Use the Migration slider to select the priority level of vCenter Server recommendations that adjust the cluster's load balance.
  - c Select whether to enable Virtual Machine Automation.  
Override for individual virtual machines can be set from the VM Overrides page.
- 6 Expand **Power Managment**.
  - a Select an automation level.
  - b Set the DPM threshold.  
Use the DPM slider to select the power recommendations that vCenter Server will apply.
- 7 Type a Task name.
- 8 Type a description of the Task you have created.
- 9 Under Configured Scheduler, click **Change** and select the time for the task to run and click **OK**.
- 10 Type an email address to send a notification email to when the task is complete.
- 11 Click **OK**.

The scheduled task runs at the specified time.

## Storage DRS Anti-Affinity Rules

You can create Storage DRS anti-affinity rules to control which virtual disks should not be placed on the same datastore within a datastore cluster. By default, a virtual machine's virtual disks are kept together on the same datastore.

When you create an anti-affinity rule, it applies to the relevant virtual disks in the datastore cluster. Anti-affinity rules are enforced during initial placement and Storage DRS-recommendation migrations, but are not enforced when a migration is initiated by a user.

---

**NOTE** Anti-affinity rules do not apply to CD-ROM ISO image files that are stored on a datastore in a datastore cluster, nor do they apply to swapfiles that are stored in user-defined locations.

---

### **Inter-VM Anti-Affinity Rules**

Specify which virtual machines should never be kept on the same datastore. See [“Create Inter-VM Anti-Affinity Rules,”](#) on page 100.

### **Intra-VM Anti-Affinity Rules**

Specify which virtual disks associated with a particular virtual machine must be kept on different datastores. See [“Create Intra-VM Anti-Affinity Rules,”](#) on page 101.

If you move a virtual disk out of the datastore cluster, the affinity or anti-affinity rule no longer applies to that disk.

When you move virtual disk files into a datastore cluster that has existing affinity and anti-affinity rules, the following behavior applies:

- Datastore Cluster B has an intra-VM affinity rule. When you move a virtual disk out of Datastore Cluster A and into Datastore Cluster B, any rule that applied to the virtual disk for a given virtual machine in Datastore Cluster A no longer applies. The virtual disk is now subject to the intra-VM affinity rule in Datastore Cluster B.

- Datastore Cluster B has an inter-VM anti-affinity rule. When you move a virtual disk out of Datastore Cluster A and into Datastore Cluster B, any rule that applied to the virtual disk for a given virtual machine in Datastore Cluster A no longer applies. The virtual disk is now subject to the inter-VM anti-affinity rule in Datastore Cluster B.
- Datastore Cluster B has an intra-VM anti-affinity rule. When you move a virtual disk out of Datastore Cluster A and into Datastore Cluster B, the intra-VM anti-affinity rule does not apply to the virtual disk for a given virtual machine because the rule is limited to only specified virtual disks in Datastore Cluster B.

---

**Note** Storage DRS rules might prevent a datastore from entering maintenance mode. You can choose to ignore Storage DRS rules for maintenance mode by enabling the Ignore Affinity Rules for Maintenance option.

---

## Create Inter-VM Anti-Affinity Rules

You can create an anti-affinity rule to indicate that all virtual disks of certain virtual machines must be kept on different datastores. The rule applies to individual datastore clusters.

Virtual machines that participate in an inter-VM anti-affinity rule in a datastore cluster must be associated with an intra-VM affinity rule in the datastore cluster. The virtual machines must also comply with the intra-VM affinity rule.

If a virtual machine is subject to an inter-VM anti-affinity rule, the following behavior applies:

- Storage DRS places the virtual machine's virtual disks according to the rule.
- Storage DRS migrates the virtual disks using vMotion according to the rule, even if the migration is for a mandatory reason such as putting a datastore in maintenance mode.
- If the virtual machine's virtual disk violates the rule, Storage DRS makes migration recommendations to correct the error or reports the violation as a fault if it cannot make a recommendation that will correct the error.

No inter-VM anti-affinity rules are defined by default.

### Procedure

- 1 Browse to the datastore cluster in the vSphere Web Client navigator.
- 2 Click the **Manage** tab and click **Settings**.
- 3 Under Configuration, select **Rules**.
- 4 Click **Add**.
- 5 Type a name for the rule.
- 6 From the Type menu, select **VM anti-affinity**.
- 7 Click **Add**.
- 8 Click **Select Virtual Machine**.
- 9 Select at least two virtual machines and click **OK**.
- 10 Click **OK** to save the rule.

## Create Intra-VM Anti-Affinity Rules

You can create a VMDK anti-affinity rule for a virtual machine that indicates which of its virtual disks must be kept on different datastores.

VMDK anti-affinity rules apply to the virtual machine for which the rule is defined, not to all virtual machines. The rule is expressed as a list of virtual disks that are to be separated from one another.

If you attempt to set an intra-VM anti-affinity rule and an intra-VM affinity rule for a virtual machine, vCenter Server rejects the most recently defined rule.

If a virtual machine is subject to a VMDK anti-affinity rule, the following behavior applies:

- Storage DRS places the virtual machine's virtual disks according to the rule.
- Storage DRS migrates the virtual disks using vMotion according to the rule, even if the migration is for a mandatory reason such as putting a datastore in maintenance mode.
- If the virtual machine's virtual disk violates the rule, Storage DRS makes migration recommendations to correct the error or reports the violation as a fault if it cannot make a recommendation that will correct the error.

No intra-VM anti-affinity rules are defined by default.

### Procedure

- 1 Browse to the datastore cluster in the vSphere Web Client navigator.
- 2 Click the **Manage** tab and click **Settings**.
- 3 Under Configuration, select **Rules**.
- 4 Click **Add**.
- 5 Type a name for the rule.
- 6 From the Type menu, select **VMDK anti-affinity**.
- 7 Click **Add**.
- 8 Click **Select Virtual Machine**.
- 9 Select a virtual machine and click **OK**.
- 10 Select at least two virtual disks to which the rule applies and click **OK**.
- 11 Click **OK** to save the rule.

## Override VMDK Affinity Rules

VMDK affinity rules indicate that all virtual disks in a datastore cluster that are associated with a particular virtual machine are located on the same datastore in the datastore cluster. The rules apply to individual datastore clusters.

VMDK affinity rules are enabled by default for all virtual machines that are in a datastore cluster. You can override the default setting for the datastore cluster or for individual virtual machines.

Virtual machines that are subject to VMDK affinity rules have the following behavior:

- Storage DRS places the virtual machine's virtual disks according to the rule.
- Storage DRS migrates the virtual disks using vMotion according to the rule, even if the migration is for a mandatory reason such as putting a datastore in maintenance mode.

- If the virtual machine's virtual disk violates the rule, Storage DRS makes migration recommendations to correct the error or reports the violation as a fault if it cannot make a recommendation that will correct the error.

When you add a datastore to a datastore cluster that is enabled for Storage DRS, the VMDK affinity rule is disabled for any virtual machine that has virtual disks on that datastore if it also has virtual disks on other datastores.

#### Procedure

- 1 Browse to the datastore cluster in the vSphere Web Client navigator.
- 2 Click the **Manage** tab and click **Settings**.
- 3 Under Configuration, select **VM Overrides**.
- 4 Click **Add**.
- 5 Use the + button to select virtual machines.
- 6 Click the **Keep VMDKs together** drop-down menu and select **No**.
- 7 Click **OK**.

## Clear Storage DRS Statistics

To diagnose problems with Storage DRS, you can clear Storage DRS statistics before you manually run Storage DRS.

---

**IMPORTANT** When you enable the option to clear Storage DRS statistics, statistics are cleared every time Storage DRS runs until you disable the option. Always disable the option after you diagnose the Storage DRS problem.

---

#### Prerequisites

Enable Storage DRS for the datastore cluster.

#### Procedure

- 1 Enable the **ClearIoStatsOnSdrsRun** option.
  - a Browse to the datastore cluster in the vSphere Web Client navigator.
  - b Click the **Manage** tab and click **Settings**.
  - c Under Services, select **DRS** and click **Edit**.
  - d Expand **Advanced Options** and click **Add**.
  - e In the Option column, type **ClearIoStatsOnSdrsRun**.
  - f In the corresponding Value text box, type **1**.
  - g Click **OK**.

- 2 Run Storage DRS on the datastore cluster.

The current Storage DRS statistics for all datastores and virtual disks in all datastore clusters in the vSphere Web Client inventory are cleared, but no new statistics are collected.

- 3 Change the **ClearIoStatsOnSdrsRun** flag value to **0** to disable it.
- 4 Run Storage DRS again.

Storage DRS runs normally. Allow several hours for the new setting to take effect.

## Storage vMotion Compatibility with Datastore Clusters

A datastore cluster has certain vSphere Storage vMotion<sup>®</sup> requirements.

- The host must be running a version of ESXi that supports Storage vMotion.
- The host must have write access to both the source datastore and the destination datastore.
- The host must have enough free memory resources to accommodate Storage vMotion.
- The destination datastore must have sufficient disk space.
- The destination datastore must not be in maintenance mode or entering maintenance mode.





## Using NUMA Systems with ESXi

---

ESXi supports memory access optimization for Intel and AMD Opteron processors in server architectures that support NUMA (non-uniform memory access).

After you understand how ESXi NUMA scheduling is performed and how the VMware NUMA algorithms work, you can specify NUMA controls to optimize the performance of your virtual machines.

This chapter includes the following topics:

- [“What is NUMA?”](#) on page 105
- [“How ESXi NUMA Scheduling Works,”](#) on page 106
- [“VMware NUMA Optimization Algorithms and Settings,”](#) on page 107
- [“Resource Management in NUMA Architectures,”](#) on page 108
- [“Using Virtual NUMA,”](#) on page 108
- [“Specifying NUMA Controls,”](#) on page 110

### What is NUMA?

NUMA systems are advanced server platforms with more than one system bus. They can harness large numbers of processors in a single system image with superior price to performance ratios.

For the past decade, processor clock speed has increased dramatically. A multi-gigahertz CPU, however, needs to be supplied with a large amount of memory bandwidth to use its processing power effectively. Even a single CPU running a memory-intensive workload, such as a scientific computing application, can be constrained by memory bandwidth.

This problem is amplified on symmetric multiprocessing (SMP) systems, where many processors must compete for bandwidth on the same system bus. Some high-end systems often try to solve this problem by building a high-speed data bus. However, such a solution is expensive and limited in scalability.

NUMA is an alternative approach that links several small, cost-effective nodes using a high-performance connection. Each node contains processors and memory, much like a small SMP system. However, an advanced memory controller allows a node to use memory on all other nodes, creating a single system image. When a processor accesses memory that does not lie within its own node (remote memory), the data must be transferred over the NUMA connection, which is slower than accessing local memory. Memory access times are not uniform and depend on the location of the memory and the node from which it is accessed, as the technology’s name implies.

### Challenges for Operating Systems

Because a NUMA architecture provides a single system image, it can often run an operating system with no special optimizations.

The high latency of remote memory accesses can leave the processors under-utilized, constantly waiting for data to be transferred to the local node, and the NUMA connection can become a bottleneck for applications with high-memory bandwidth demands.

Furthermore, performance on such a system can be highly variable. It varies, for example, if an application has memory located locally on one benchmarking run, but a subsequent run happens to place all of that memory on a remote node. This phenomenon can make capacity planning difficult.

Some high-end UNIX systems provide support for NUMA optimizations in their compilers and programming libraries. This support requires software developers to tune and recompile their programs for optimal performance. Optimizations for one system are not guaranteed to work well on the next generation of the same system. Other systems have allowed an administrator to explicitly decide on the node on which an application should run. While this might be acceptable for certain applications that demand 100 percent of their memory to be local, it creates an administrative burden and can lead to imbalance between nodes when workloads change.

Ideally, the system software provides transparent NUMA support, so that applications can benefit immediately without modifications. The system should maximize the use of local memory and schedule programs intelligently without requiring constant administrator intervention. Finally, it must respond well to changing conditions without compromising fairness or performance.

## How ESXi NUMA Scheduling Works

ESXi uses a sophisticated NUMA scheduler to dynamically balance processor load and memory locality or processor load balance.

- 1 Each virtual machine managed by the NUMA scheduler is assigned a home node. A home node is one of the system's NUMA nodes containing processors and local memory, as indicated by the System Resource Allocation Table (SRAT).
- 2 When memory is allocated to a virtual machine, the ESXi host preferentially allocates it from the home node. The virtual CPUs of the virtual machine are constrained to run on the home node to maximize memory locality.
- 3 The NUMA scheduler can dynamically change a virtual machine's home node to respond to changes in system load. The scheduler might migrate a virtual machine to a new home node to reduce processor load imbalance. Because this might cause more of its memory to be remote, the scheduler might migrate the virtual machine's memory dynamically to its new home node to improve memory locality. The NUMA scheduler might also swap virtual machines between nodes when this improves overall memory locality.

Some virtual machines are not managed by the ESXi NUMA scheduler. For example, if you manually set the processor or memory affinity for a virtual machine, the NUMA scheduler might not be able to manage this virtual machine. Virtual machines that are not managed by the NUMA scheduler still run correctly. However, they don't benefit from ESXi NUMA optimizations.

The NUMA scheduling and memory placement policies in ESXi can manage all virtual machines transparently, so that administrators do not need to address the complexity of balancing virtual machines between nodes explicitly.

The optimizations work seamlessly regardless of the type of guest operating system. ESXi provides NUMA support even to virtual machines that do not support NUMA hardware, such as Windows NT 4.0. As a result, you can take advantage of new hardware even with legacy operating systems.

A virtual machine that has more virtual processors than the number of physical processor cores available on a single hardware node can be managed automatically. The NUMA scheduler accommodates such a virtual machine by having it span NUMA nodes. That is, it is split up as multiple NUMA clients, each of which is assigned to a node and then managed by the scheduler as a normal, non-spanning client. This can improve the performance of certain memory-intensive workloads with high locality. For information on configuring the behavior of this feature, see [“Advanced Virtual Machine Attributes,”](#) on page 116.

ESXi 5.0 and later includes support for exposing virtual NUMA topology to guest operating systems. For more information about virtual NUMA control, see [“Using Virtual NUMA,”](#) on page 108.

## VMware NUMA Optimization Algorithms and Settings

This section describes the algorithms and settings used by ESXi to maximize application performance while still maintaining resource guarantees.

### Home Nodes and Initial Placement

When a virtual machine is powered on, ESXi assigns it a home node. A virtual machine runs only on processors within its home node, and its newly allocated memory comes from the home node as well.

Unless a virtual machine’s home node changes, it uses only local memory, avoiding the performance penalties associated with remote memory accesses to other NUMA nodes.

When a virtual machine is powered on, it is assigned an initial home node so that the overall CPU and memory load among NUMA nodes remains balanced. Because internode latencies in a large NUMA system can vary greatly, ESXi determines these internode latencies at boot time and uses this information when initially placing virtual machines that are wider than a single NUMA node. These wide virtual machines are placed on NUMA nodes that are close to each other for lowest memory access latencies.

Initial placement-only approaches are usually sufficient for systems that run only a single workload, such as a benchmarking configuration that remains unchanged as long as the system is running. However, this approach is unable to guarantee good performance and fairness for a datacenter-class system that supports changing workloads. Therefore, in addition to initial placement, ESXi 5.0 does dynamic migration of virtual CPUs and memory between NUMA nodes for improving CPU balance and increasing memory locality.

### Dynamic Load Balancing and Page Migration

ESXi combines the traditional initial placement approach with a dynamic rebalancing algorithm. Periodically (every two seconds by default), the system examines the loads of the various nodes and determines if it should rebalance the load by moving a virtual machine from one node to another.

This calculation takes into account the resource settings for virtual machines and resource pools to improve performance without violating fairness or resource entitlements.

The rebalancer selects an appropriate virtual machine and changes its home node to the least loaded node. When it can, the rebalancer moves a virtual machine that already has some memory located on the destination node. From that point on (unless it is moved again), the virtual machine allocates memory on its new home node and it runs only on processors within the new home node.

Rebalancing is an effective solution to maintain fairness and ensure that all nodes are fully used. The rebalancer might need to move a virtual machine to a node on which it has allocated little or no memory. In this case, the virtual machine incurs a performance penalty associated with a large number of remote memory accesses. ESXi can eliminate this penalty by transparently migrating memory from the virtual machine’s original node to its new home node:

- 1 The system selects a page (4KB of contiguous memory) on the original node and copies its data to a page in the destination node.
- 2 The system uses the virtual machine monitor layer and the processor’s memory management hardware to seamlessly remap the virtual machine’s view of memory, so that it uses the page on the destination node for all further references, eliminating the penalty of remote memory access.

When a virtual machine moves to a new node, the ESXi host immediately begins to migrate its memory in this fashion. It manages the rate to avoid overtaxing the system, particularly when the virtual machine has little remote memory remaining or when the destination node has little free memory available. The memory migration algorithm also ensures that the ESXi host does not move memory needlessly if a virtual machine is moved to a new node for only a short period.

When initial placement, dynamic rebalancing, and intelligent memory migration work in conjunction, they ensure good memory performance on NUMA systems, even in the presence of changing workloads. When a major workload change occurs, for instance when new virtual machines are started, the system takes time to readjust, migrating virtual machines and memory to new locations. After a short period, typically seconds or minutes, the system completes its readjustments and reaches a steady state.

## Transparent Page Sharing Optimized for NUMA

Many ESXi workloads present opportunities for sharing memory across virtual machines.

For example, several virtual machines might be running instances of the same guest operating system, have the same applications or components loaded, or contain common data. In such cases, ESXi systems use a proprietary transparent page-sharing technique to securely eliminate redundant copies of memory pages. With memory sharing, a workload running in virtual machines often consumes less memory than it would when running on physical machines. As a result, higher levels of overcommitment can be supported efficiently.

Transparent page sharing for ESXi systems has also been optimized for use on NUMA systems. On NUMA systems, pages are shared per-node, so each NUMA node has its own local copy of heavily shared pages. When virtual machines use shared pages, they don't need to access remote memory.

---

**NOTE** This default behavior is the same in all previous versions of ESX and ESXi.

---

## Resource Management in NUMA Architectures

You can perform resource management with different types of NUMA architecture.

With the proliferation of highly multicore systems, NUMA architectures are becoming more popular as these architectures allow better performance scaling of memory intensive workloads. All modern Intel and AMD systems have NUMA support built into the processors. Additionally, there are traditional NUMA systems like the IBM Enterprise X-Architecture that extend Intel and AMD processors with NUMA behavior with specialized chipset support.

Typically, you can use BIOS settings to enable and disable NUMA behavior. For example, in AMD Opteron-based HP Proliant servers, NUMA can be disabled by enabling node interleaving in the BIOS. If NUMA is enabled, the BIOS builds a system resource allocation table (SRAT) which ESXi uses to generate the NUMA information used in optimizations. For scheduling fairness, NUMA optimizations are not enabled for systems with too few cores per NUMA node or too few cores overall. You can modify the `numa.rebalancecorestotal` and `numa.rebalancecoresnode` options to change this behavior.

## Using Virtual NUMA

vSphere 5.0 and later includes support for exposing virtual NUMA topology to guest operating systems, which can improve performance by facilitating guest operating system and application NUMA optimizations.

Virtual NUMA topology is available to hardware version 8 virtual machines and is enabled by default when the number of virtual CPUs is greater than eight. You can also manually influence virtual NUMA topology using advanced configuration options.

You can affect the virtual NUMA topology with two settings in the vSphere Web Client: number of virtual sockets and number of cores per socket for a virtual machine. If the number of cores per socket (`cpuid.coresPerSocket`) is greater than one, and the number of virtual cores in the virtual machine is greater than 8, the virtual NUMA node size matches the virtual socket size. If the number of cores per socket is less than or equal to one, virtual NUMA nodes are created to match the topology of the first physical host where the virtual machine is powered on.

When the number of virtual CPUs and the amount of memory used grow proportionately, you can use the default values. For virtual machines that consume a disproportionately large amount of memory, you can override the default values in one of the following ways:

- Increase the number of virtual CPUs, even if this number of virtual CPUs is not used. See [“Change the Number of Virtual CPUs,”](#) on page 109.
- Use advanced options to control virtual NUMA topology and its mapping over physical NUMA topology. See [“Virtual NUMA Controls,”](#) on page 109.

## Change the Number of Virtual CPUs

You can configure a virtual machine that runs on an ESXi host to have up to 128 virtual CPUs.

---

**NOTE** When you configure your virtual machine for multicore virtual CPU settings, you must ensure that your configuration complies with the requirements of the guest operating system EULA.

---

### Procedure

- 1 Find the virtual machine in the vSphere Web Client inventory.
  - a To find a virtual machine, select a data center, folder, cluster, resource pool, or host.
  - b Click the **Related Objects** tab and click **Virtual Machines**.
- 2 Right-click the virtual machine and click **Edit Settings**.
- 3 On the **Virtual Hardware** tab, expand the **CPU** section.
- 4 Select a value from the **CPU** drop-down menu.
- 5 Select a value from the **Cores per Socket** drop-down menu.
 

The resulting total number of cores is a number equal to or less than the number of logical CPUs on the host.
- 6 Click **OK**.

## Virtual NUMA Controls

For virtual machines with disproportionately large memory consumption, you can use advanced options to manually override the default virtual CPU settings.

You can add these advanced options to the virtual machine configuration file.

**Table 14-1.** Advanced Options for Virtual NUMA Controls

Option	Description	Default Value
<code>cpuid.coresPerSocket</code>	Determines the number of virtual cores per virtual CPU socket. If the value is greater than 1, also determines the size of virtual NUMA nodes if a virtual machine has a virtual NUMA topology. You can set this option if you know the exact virtual NUMA topology for each physical host.	1
<code>numa.vcpu.maxPerVirtualNode</code>	If <code>cpuid.coresPerSocket</code> is too restrictive as a power of two, you can set <code>numa.vcpu.maxPerVirtualNode</code> directly. In this case, do not set <code>cpuid.coresPerSocket</code> .	8
<code>numa.autosize</code>	When you set this option, the virtual NUMA topology has the same number of virtual CPUs per virtual node as there are cores on each physical node.	FALSE

**Table 14-1.** Advanced Options for Virtual NUMA Controls (Continued)

Option	Description	Default Value
<code>numa.autosize.once</code>	When you create a virtual machine template with these settings, the settings are guaranteed to remain the same every time you subsequently power on the virtual machine. The virtual NUMA topology will be reevaluated if the configured number of virtual CPUs on the virtual machine is modified.	TRUE
<code>numa.vcpu.min</code>	Minimum number of virtual CPUs in a virtual machine that are required in order to generate a virtual NUMA topology.	9

**NOTE** When you set `numa.autosize` to TRUE, and if the configuration is set up manually or with a script, some guests might not be able to handle dynamic changes.

For example, a Linux application configured with the `numactl` system utility is set up and tested on one physical host with four cores per node. The host requires two NUMA nodes for a virtual machine with eight virtual CPUs. If the same virtual machine is run on a system with eight cores per node, which might occur during a vMotion operation, and `numa.autosize` is set to TRUE, only one virtual NUMA node will be created (rather than two virtual NUMA nodes). When `numactl` references the second virtual NUMA node, the operation will fail.

To avoid this, scripts should be intelligent enough to first query `numactl --hardware`. Otherwise, you must set the NUMA topology explicitly or allow the default `numa.autosize.once` setting to take effect.

## Specifying NUMA Controls

If you have applications that use a lot of memory or have a small number of virtual machines, you might want to optimize performance by specifying virtual machine CPU and memory placement explicitly.

Specifying controls is useful if a virtual machine runs a memory-intensive workload, such as an in-memory database or a scientific computing application with a large data set. You might also want to optimize NUMA placements manually if the system workload is known to be simple and unchanging. For example, an eight-processor system running eight virtual machines with similar workloads is easy to optimize explicitly.

**NOTE** In most situations, the ESXi host's automatic NUMA optimizations result in good performance.

ESXi provides three sets of controls for NUMA placement, so that administrators can control memory and processor placement of a virtual machine.

The vSphere Web Client lets you specify the following options.

<b>NUMA Node Affinity</b>	When you set this option, NUMA can schedule a virtual machine only on the nodes specified in the affinity.
<b>CPU Affinity</b>	When you set this option, a virtual machine uses only the processors specified in the affinity.
<b>Memory Affinity</b>	When you set this option, the server allocates memory only on the specified nodes.

A virtual machine is still managed by NUMA when you specify NUMA node affinity, but its virtual CPUs can be scheduled only on the nodes specified in the NUMA node affinity. Likewise, memory can be obtained only from the nodes specified in the NUMA node affinity. When you specify CPU or memory affinities, a virtual machine ceases to be managed by NUMA. NUMA management of these virtual machines is effective when you remove the CPU and memory affinity constraints.

Manual NUMA placement might interfere with ESXi resource management algorithms, which distribute processor resources fairly across a system. For example, if you manually place 10 virtual machines with processor-intensive workloads on one node, and manually place only 2 virtual machines on another node, it is impossible for the system to give all 12 virtual machines equal shares of systems resources.

## Associate Virtual Machines with Specific Processors

You might be able to improve the performance of the applications on a virtual machine by pinning its virtual CPUs to fixed processors. This allows you to prevent the virtual CPUs from migrating across NUMA nodes.

### Procedure

- 1 Find the virtual machine in the vSphere Web Client inventory.
  - a To find a virtual machine, select a data center, folder, cluster, resource pool, or host.
  - b Click the **Related Objects** tab and click **Virtual Machines**.
- 2 Right-click the virtual machine and click **Edit Settings**.
- 3 Select the **Virtual Hardware** tab, and expand **CPU**.
- 4 Under Scheduling Affinity, set the CPU affinity to the preferred processors.

---

**NOTE** You must manually select all processors in the NUMA node. CPU affinity is specified on a per-processor, not on a per-node, basis.

---

## Associate Memory Allocations with Specific NUMA Nodes Using Memory Affinity

You can specify that all future memory allocations on a virtual machine use pages associated with specific NUMA nodes (also known as manual memory affinity).

---

**NOTE** Specify nodes to be used for future memory allocations only if you have also specified CPU affinity. If you make manual changes only to the memory affinity settings, automatic NUMA rebalancing does not work properly.

---

### Procedure

- 1 Browse to the virtual machine in the vSphere Web Client navigator.
- 2 Click the **Manage** tab.
- 3 Click **Settings**, and click **VM Hardware**.
- 4 Click **Edit**.
- 5 Select the **Virtual Hardware** tab, and expand **Memory**.
- 6 Under NUMA Memory Affinity, set memory affinity.

### Example: Binding a Virtual Machine to a Single NUMA Node

The following example illustrates manually binding the last four physical CPUs to a single NUMA node for a two-way virtual machine on an eight-way server.

The CPUs (for example, 4, 5, 6, and 7) are the physical CPU numbers.

- 1 In the vSphere Web Client navigator, right-click the virtual machine and select **Edit Settings**.
- 2 Select **Options** and click **Advanced**.
- 3 Click the **Configuration Parameters** button.

- 4 In the vSphere Web Client, turn on CPU affinity for processors 4, 5, 6, and 7.

Then, you want this virtual machine to run only on node 1.

- 1 In the vSphere Web Client inventory panel, select the virtual machine and select **Edit Settings**.
- 2 Select **Options** and click **Advanced**.
- 3 Click the **Configuration Parameters** button.
- 4 In the vSphere Web Client, set memory affinity for the NUMA node to 1.

Completing these two tasks ensures that the virtual machine runs only on NUMA node 1 and, when possible, allocates memory from the same node.

## Associate Virtual Machines with Specified NUMA Nodes

When you associate a NUMA node with a virtual machine to specify NUMA node affinity, you constrain the set of NUMA nodes on which NUMA can schedule a virtual machine's virtual CPU and memory.

---

**NOTE** When you constrain NUMA node affinities, you might interfere with the ability of the ESXi NUMA scheduler to rebalance virtual machines across NUMA nodes for fairness. Specify NUMA node affinity only after you consider the rebalancing issues.

---

### Procedure

- 1 Browse to the cluster in the vSphere Web Client object navigator.
- 2 Click the **Manage** tab and click **Settings**.
- 3 Under VM Options, click the **Edit** button.
- 4 Select the VM Options tab and expand **Advanced**.
- 5 Under Configuration Parameters, click the **Edit Configuration** button.
- 6 Click **Add Row** to add a new option.
- 7 In the Name column, enter **numa.nodeAffinity**.
- 8 In the Value column, enter the NUMA nodes where the virtual machine can be scheduled.

Use a comma-separated list for multiple nodes. For example, enter **0,1** to constrain the virtual machine resource scheduling to NUMA nodes 0 and 1.

- 9 Click **OK**.
- 10 Click **OK** to close the Edit VM dialog box.



# Advanced Attributes

---

You can set advanced attributes for hosts or individual virtual machines to help you customize resource management.

In most cases, adjusting the basic resource allocation settings (reservation, limit, shares) or accepting default settings results in appropriate resource allocation. However, you can use advanced attributes to customize resource management for a host or a specific virtual machine.

This chapter includes the following topics:

- [“Set Advanced Host Attributes,”](#) on page 113
- [“Set Advanced Virtual Machine Attributes,”](#) on page 116
- [“Latency Sensitivity,”](#) on page 118
- [“About Reliable Memory,”](#) on page 118

## Set Advanced Host Attributes

You can set advanced attributes for a host.



**CAUTION** Changing advanced options is considered unsupported unless VMware technical support or a KB article instruct you to do so. In all other cases, changing these options is considered unsupported. In most cases, the default settings produce the optimum result.

---

### Procedure

- 1 Browse to the host in the vSphere Web Client navigator.
- 2 Click the **Manage** tab and click **Settings**.
- 3 Under **System**, click **Advanced System Settings**.
- 4 In Advanced Settings, select the appropriate item.
- 5 Click the **Edit** button to edit the value.
- 6 Click **OK**.

## Advanced Memory Attributes

You can use the advanced memory attributes to customize memory resource usage.

**Table 15-1.** Advanced Memory Attributes

Attribute	Description	Default
Mem.ShareForceSalting	<p>Mem.ShareForceSalting 0: Inter-virtual machine Transparent Page Sharing (TPS) behavior is still retained. The value of VMX option <b>sched.mem.pshare.salt</b> is ignored even if present.</p> <p>Mem.ShareForceSalting 1: By default the salt value is taken from <b>sched.mem.pshare.salt</b>. If not specified, it falls back to old TPS (inter-VM) behavior by considering salt values for the virtual machine as 0.</p> <p>Mem.ShareForceSalting 2: By default the salt value is taken from <b>sched.mem.pshare.salt</b> if present, or <b>vc.uuid</b>. If it does not exist, then the page sharing algorithm generates random and unique value for salting per virtual machine, which is not configurable by users.</p>	2
Mem.SamplePeriod	Specifies the periodic time interval, measured in seconds of the virtual machine's execution time, over which memory activity is monitored to estimate working set sizes.	60
Mem.BalancePeriod	Specifies the periodic time interval, in seconds, for automatic memory reallocations. Significant changes in the amount of free memory also trigger reallocations.	15
Mem.IdleTax	Specifies the idle memory tax rate, as a percentage. This tax effectively charges virtual machines more for idle memory than for memory they are actively using. A tax rate of 0 percent defines an allocation policy that ignores working sets and allocates memory strictly based on shares. A high tax rate results in an allocation policy that allows idle memory to be reallocated away from virtual machines that are unproductively hoarding it.	75
Mem.ShareScanGHz	Specifies the maximum amount of memory pages to scan (per second) for page sharing opportunities for each GHz of available host CPU resource. For example, defaults to 4 MB/sec per 1 GHz.	4
Mem.ShareScanTime	Specifies the time, in minutes, within which an entire virtual machine is scanned for page sharing opportunities. Defaults to 60 minutes.	60
Mem.CtlMaxPercent	Limits the maximum amount of memory reclaimed from any virtual machine using the memory balloon driver ( <b>vmmemctl</b> ), based on a percentage of its configured memory size. Specify 0 to disable reclamation for all virtual machines.	65
Mem.AllocGuestLargePage	Enables backing of guest large pages with host large pages. Reduces TLB misses and improves performance in server workloads that use guest large pages. 0=disable.	1
Mem.AllocUsePSharePool and Mem.AllocUseGuestPool	Reduces memory fragmentation by improving the probability of backing guest large pages with host large pages. If host memory is fragmented, the availability of host large pages is reduced. 0 = disable.	15
Mem.MemZipEnable	Enables memory compression for the host. 0 = disable.	1
Mem.MemZipMaxPct	Specifies the maximum size of the compression cache in terms of the maximum percentage of each virtual machine's memory that can be stored as compressed memory.	10
LPage.LPageDefragEnable	Enables large page defragmentation. 0 = disable.	1
LPage.LPageDefragRateVM	Maximum number of large page defragmentation attempts per second per virtual machine. Accepted values range from 1 to 1024.	32

**Table 15-1.** Advanced Memory Attributes (Continued)

Attribute	Description	Default
LPage.LPageDefragRateTotal	Maximum number of large page defragmentation attempts per second. Accepted values range from 1 to 10240.	256
LPage.LPageAlwaysTryForNPT	Try to allocate large pages for nested page tables (called 'RVI' by AMD or 'EPT' by Intel). If you enable this option, all guest memory is backed with large pages in machines that use nested page tables (for example, AMD Barcelona). If NPT is not available, only some portion of guest memory is backed with large pages. 0= disable.	1

## Advanced NUMA Attributes

You can use the advanced NUMA attributes to customize NUMA usage.

**Table 15-2.** Advanced NUMA Attributes

Attribute	Description	Default
Numa.RebalancePeriod	Controls the frequency of rebalance periods, specified in milliseconds. More frequent rebalancing can increase CPU overheads, particularly on machines with a large number of running virtual machines. More frequent rebalancing can also improve fairness.	2000
Numa.MigImbalanceThreshold	The NUMA rebalancer computes the CPU imbalance between nodes, accounting for the difference between each virtual machine's CPU time entitlement and its actual consumption. This option controls the minimum load imbalance between nodes needed to trigger a virtual machine migration, in percent.	10
Numa.RebalanceEnable	Enable NUMA rebalancing and scheduling. Set this option to 0 to disable all NUMA rebalancing and initial placement of virtual machines, effectively disabling the NUMA scheduling system.	1
Numa.RebalanceCoresTotal	Specifies the minimum number of total processor cores on the host required to enable the NUMA rebalancer.	4
Numa.RebalanceCoresNode	Specifies the minimum number of processor cores per node required to enable the NUMA rebalancer. This option and Numa.RebalanceCoresTotal are useful when disabling NUMA rebalancing on small NUMA configurations (for example, two-way Opteron hosts), where the small number of total or per-node processors can compromise scheduling fairness when you enable NUMA rebalancing.	2
Numa.AutoMemAffinity	Automatically set memory affinity for virtual machines that have CPU affinity set.	1
Numa.PageMigEnable	Automatically migrate pages between NUMA nodes to improve memory locality. Page migration rates set manually are still in effect.	1

## Set Advanced Virtual Machine Attributes

You can set advanced attributes for a virtual machine.

### Procedure

- 1 Find the virtual machine in the vSphere Web Client inventory.
  - a To find a virtual machine, select a data center, folder, cluster, resource pool, or host.
  - b Click the **Related Objects** tab and click **Virtual Machines**.
- 2 Right-click the virtual machine and select **Edit Settings**.
- 3 Click **VM Options**.
- 4 Expand **Advanced**.
- 5 Under Configuration Parameters, click the **Edit Configuration** button.
- 6 In the dialog box that appears, click **Add Row** to enter a new parameter and its value.
- 7 Click **OK**.

## Advanced Virtual Machine Attributes

You can use the advanced virtual machine attributes to customize virtual machine configuration.

**Table 15-3.** Advanced Virtual Machine Attributes

Attribute	Description	Default
sched.mem.maxmemctl	Maximum amount of memory reclaimed from the selected virtual machine by ballooning, in megabytes (MB). If the ESXi host needs to reclaim additional memory, it is forced to swap. Swapping is less desirable than ballooning.	-1 (Unlimited)
sched.mem.pshare.enable	Enables memory sharing for a selected virtual machine.  This boolean value defaults to True. If you set it to False for a virtual machine, this turns off memory sharing.	True
sched.mem.pshare.salt	A salt value is a configurable VMX option for each virtual machine. If this option is not present in the virtual machine's VMX file, then the value of <b>vc.uuid vmx</b> option is taken as the default value. Since the <b>vc.uuid</b> is unique to each virtual machine, by default transparent page sharing happens only among the pages belonging to a particular virtual machine (intra-VM). If a group of virtual machines are considered trustworthy, it is possible to share pages among them by setting a common salt value for all of those virtual machines (inter-VM).	user configurable

**Table 15-3.** Advanced Virtual Machine Attributes (Continued)

Attribute	Description	Default
sched.swap.persist	Specifies whether the virtual machine's swap files should persist or be deleted when the virtual machine is powered off. By default, the system creates the swap file for a virtual machine when the virtual machine is powered on, and deletes the swap file when the virtual machine is powered off.	False
sched.swap.dir	Directory location of the virtual machine's swap file. Defaults to the virtual machine's working directory, that is, the directory that contains its configuration file. This directory must remain on a host that is accessible to the virtual machine. If you move the virtual machine (or any clones created from it), you might need to reset this attribute.	Equals <code>workingDir</code>

## Advanced Virtual NUMA Attributes

You can use the advanced virtual NUMA attributes to customize virtual NUMA usage.

**Table 15-4.** Advanced NUMA Attributes

Attribute	Description	Default
cpuid.coresPerSocket	Determines the number of virtual cores per virtual CPU socket. If the value is greater than 1, also determines the size of virtual NUMA nodes if a virtual machine has a virtual NUMA topology. You can set this option if you know the exact virtual NUMA topology for each physical host.	1
numa.autosize	When you set this option, the virtual NUMA topology has the same number of virtual CPUs per virtual node as there are cores on each physical node.	FALSE
numa.autosize.once	When you create a virtual machine template with these settings, the settings are guaranteed to remain the same every time you subsequently power on the virtual machine. The virtual NUMA topology will be reevaluated if the configured number of virtual CPUs on the virtual machine is modified.	TRUE
numa.vcpu.maxPerVirtualNode	If <code>cpuid.coresPerSocket</code> is too restrictive as a power of two, you can set <code>numa.vcpu.maxPerVirtualNode</code> directly. In this case, do not set <code>cpuid.coresPerSocket</code> .	8
numa.vcpu.min	Minimum number of virtual CPUs in a virtual machine that are required in order to generate a virtual NUMA topology.	9
numa.vcpu.maxPerMachineNode	Maximum number of virtual CPUs that belong to the same virtual machine that can be scheduled on a NUMA node at the same time. Use this attribute to ensure maximum bandwidth, by forcing different NUMA clients on different NUMA nodes.	Number of cores per node on the physical host where a virtual machine is running.
numa.vcpu.maxPerClient	Maximum number of virtual CPUs in a NUMA client. A client is a group of virtual CPUs that are NUMA-managed as a single entity. By default, each virtual NUMA node is a NUMA client, but if a virtual NUMA node is larger than a physical NUMA node, a single virtual NUMA node can be backed by multiple NUMA clients.	Equals <code>numa.vcpu.maxPerMachineNode</code>

**Table 15-4.** Advanced NUMA Attributes (Continued)

Attribute	Description	Default
numa.nodeAffinity	Constrains the set of NUMA nodes on which a virtual machine's virtual CPU and memory can be scheduled. <b>NOTE</b> When you constrain NUMA node affinities, you might interfere with the ability of the NUMA scheduler to rebalance virtual machines across NUMA nodes for fairness. Specify NUMA node affinity only after you consider the rebalancing issues.	
numa.mem.interleave	Specifies whether the memory allocated to a virtual machine is statically interleaved across all the NUMA nodes on which its constituent NUMA clients are running and there is no virtual NUMA topology exposed.	True

## Latency Sensitivity

You can adjust the latency sensitivity of a virtual machine to optimize the scheduling delay for latency sensitive applications.

ESXi is optimized to deliver high throughput. You can optimize your virtual machine to meet the low latency requirement of latency sensitive applications. Examples of latency sensitive applications are VOIP or media player applications, or applications that require frequent access to the mouse or keyboard devices.

### Adjust Latency Sensitivity

You can adjust the latency sensitivity of a virtual machine.

#### Procedure

- 1 Find the virtual machine in the vSphere Web Client inventory.
  - a To find a virtual machine, select a data center, folder, cluster, resource pool, or host.
  - b Click the **Related Objects** tab and click **Virtual Machines**.
- 2 Right-click the virtual machine and click **Edit Settings**.
- 3 Click **VM Options** and click **Advanced**.
- 4 Select a setting from the **Latency Sensitivity** drop-down menu.
- 5 Click **OK**.

## About Reliable Memory

ESXi supports reliable memory.

Some systems have reliable memory, which is a part of memory that is less likely to have hardware memory errors than other parts of the memory in the system. If the hardware exposes information about the different levels of reliability, ESXi might be able to achieve higher system reliability.

### View Reliable Memory

View Reliable Memory in the vSphere Web Client.

You can view whether or not the license permits reliable memory.

#### Procedure

- 1 Browse to the host in the vSphere Web Client navigator.

- 2 Click the **Manage** tab and click **Settings**.
- 3 Under **System** select **Licensing**.
- 4 Under **Features** verify Reliable Memory is displayed.

**What to do next**

You can look up how much memory is considered reliable by using the ESXCLI `hardware memory get` command.





## Fault Definitions

---

DRS faults indicate the reasons that prevent the generation of DRS actions (or the recommendation of those actions in manual mode).

The DRS faults are defined within this section.

This chapter includes the following topics:

- [“Virtual Machine is Pinned,”](#) on page 122
- [“Virtual Machine not Compatible with any Host,”](#) on page 122
- [“VM/VM DRS Rule Violated when Moving to another Host,”](#) on page 122
- [“Host Incompatible with Virtual Machine,”](#) on page 122
- [“Host has Virtual Machine that Violates VM/VM DRS Rules,”](#) on page 122
- [“Host has Insufficient Capacity for Virtual Machine,”](#) on page 122
- [“Host in Incorrect State,”](#) on page 122
- [“Host has Insufficient Number of Physical CPUs for Virtual Machine,”](#) on page 123
- [“Host has Insufficient Capacity for Each Virtual Machine CPU,”](#) on page 123
- [“The Virtual Machine is in vMotion,”](#) on page 123
- [“No Active Host in Cluster,”](#) on page 123
- [“Insufficient Resources,”](#) on page 123
- [“Insufficient Resources to Satisfy Configured Failover Level for HA,”](#) on page 123
- [“No Compatible Hard Affinity Host,”](#) on page 123
- [“No Compatible Soft Affinity Host,”](#) on page 123
- [“Soft Rule Violation Correction Disallowed,”](#) on page 123
- [“Soft Rule Violation Correction Impact,”](#) on page 124

## Virtual Machine is Pinned

This fault occurs when DRS cannot move a virtual machine because DRS is disabled on it. That is, the virtual machine is "pinned" on its registered host.

## Virtual Machine not Compatible with any Host

This fault occurs when DRS cannot find a host that can run the virtual machine.

This might occur, for example, if no host can satisfy the virtual machine's CPU or memory resource needs or if no host currently has network or storage access needed by the virtual machine.

To address this problem, provide a host that can meet the virtual machine's requirements.

## VM/VM DRS Rule Violated when Moving to another Host

This fault occurs when more than one virtual machines running on the same host and share affinity rules with each other cannot be moved to another host.

This might occur because not all the virtual machines can vMotion off the current host. For example, one of the virtual machines in the group is DRS-disabled.

To prevent this, check for reasons why some virtual machines in the group cannot vMotion.

## Host Incompatible with Virtual Machine

This fault occurs when DRS considers migrating a virtual machine to a host, but finds that the host is incompatible with the given virtual machine.

This might occur because the target host does not have access to the network or storage connection needed by the virtual machine. Another reason this fault occurs is if the target host has a CPU that differs sufficiently from the current host so that using vMotion amongst the hosts is not supported.

To avoid this, create clusters such that all hosts are configured consistently and vMotion is compatible amongst the hosts.

Another reason the host is incompatible with the virtual machine is that there is a required VM/Host DRS rule in place that instructs DRS to never place this virtual machine on this host.

## Host has Virtual Machine that Violates VM/VM DRS Rules

This fault occurs when the virtual machine, when powered on or moved by starting vMotion, would violate a VM/VM DRS rule.

The virtual machine can still be manually powered on or moved with vMotion, but vCenter Server cannot automatically do so.

## Host has Insufficient Capacity for Virtual Machine

This fault occurs when the host does not have enough CPU or memory capacity for running the virtual machine.

## Host in Incorrect State

This fault occurs when the host is entering maintenance or standby state when needed for DRS action to occur.

To address this fault, cancel the request for the host to enter standby or maintenance mode.

## Host has Insufficient Number of Physical CPUs for Virtual Machine

This fault occurs when the host hardware does not enough physical CPUs (hyperthreads) to support the number of virtual CPUs in the virtual machine.

## Host has Insufficient Capacity for Each Virtual Machine CPU

This fault occurs when the host does not have enough CPU capacity for running the virtual machine.

## The Virtual Machine is in vMotion

This fault occurs when DRS cannot move a virtual machine because it is in vMotion.

## No Active Host in Cluster

This fault occurs when the cluster in which the virtual machine is being moved does not contain any hosts that are connected and in a non-maintenance state.

This can occur, for example, if all the hosts are disconnected or in maintenance mode.

## Insufficient Resources

This fault occurs when an attempted operation conflicts with a resource configuration policy.

This fault may occur, for example, if a power-on operation reserves more memory than is allocated to a resource pool.

Retry the operation after adjusting the resources to allow more memory.

## Insufficient Resources to Satisfy Configured Failover Level for HA

This fault occurs when the HA configuration of CPU or memory resources reserved for failover is violated or cannot be met by the DRS operation under consideration.

This fault is reported when:

- The host is requested to enter maintenance or standby mode.
- The Virtual machine violates failover when it attempts to power on.

## No Compatible Hard Affinity Host

No host is available for the virtual machine that satisfies its mandatory VM/Host DRS affinity or anti-affinity rules.

## No Compatible Soft Affinity Host

No host is available for the virtual machine that satisfied its preferred VM/Host DRS affinity or anti-affinity rules.

## Soft Rule Violation Correction Disallowed

DRS migration threshold is set at mandatory-only.

This does not allow the generation of DRS actions to correct non-mandatory VM/Host DRS affinity rules.

## **Soft Rule Violation Correction Impact**

Correcting the non-mandatory VM/Host DRS affinity rule does not occur because it impacts performance.

# DRS Troubleshooting Information

---

This information describes vSphere® Distributed Resource Scheduler (DRS) problems for particular categories: cluster, host, and virtual machine problems.

This chapter includes the following topics:

- [“Cluster Problems,”](#) on page 125
- [“Host Problems,”](#) on page 128
- [“Virtual Machine Problems,”](#) on page 131

## Cluster Problems

Cluster problems can prevent DRS from performing optimally or from reporting faults.

### Load Imbalance on Cluster

A cluster has a load imbalance of resources.

#### Problem

A cluster might become unbalanced because of uneven resource demands from virtual machines and unequal capacities of hosts.

#### Cause

The following are possible reasons why the cluster has a load imbalance:

- The migration threshold is too high.  
A higher threshold makes the cluster a more likely candidate for load imbalance.
- VM/VM or VM/Host DRS rules prevent virtual machines from being moved.
- DRS is disabled for one or more virtual machines.
- A device is mounted to one or more virtual machines preventing DRS from moving the virtual machine in order to balance the load.
- Virtual machines are not compatible with the hosts to which DRS would move them. That is, at least one of the hosts in the cluster is incompatible for the virtual machines that would be migrated. For example, if host A's CPU is not vMotion-compatible with host B's CPU, then host A becomes incompatible for powered-on virtual machines running on host B.
- It would be more detrimental for the virtual machine's performance to move it than for it to run where it is currently located. This may occur when loads are unstable or the migration cost is high compared to the benefit gained from moving the virtual machine.

- vMotion is not enabled or set up for the hosts in the cluster.

**Solution**

Address the problem that is causing the load imbalance.

## Cluster is Yellow

The cluster is yellow due to a shortage of resources.

**Problem**

If the cluster does not have enough resources to satisfy the reservations of all resource pools and virtual machines, but does have enough resources to satisfy the reservations of all running virtual machines, DRS continues to run and the cluster is yellow.

**Cause**

A cluster can become yellow if the host resources are removed from the cluster (for example, if a host fails).

**Solution**

Add host resources to the cluster or reduce the resource pool reservations.

## Cluster is Red Because of Inconsistent Resource Pool

A DRS cluster becomes red when it is invalid. It may become red because the resource pool tree is not internally consistent.

**Problem**

If the cluster resource pool tree is not internally consistent (for example, the sum of the children's reservations is greater than the parent pool's nonexpandable reservation), the cluster does not have enough resources to satisfy the reservations of all running virtual machines making the cluster red.

**Cause**

This can occur if vCenter Server is unavailable or if resource pool settings are changed while a virtual machine is in a failover state.

**Solution**

Revert the associated changes or otherwise revise the resource pool settings.

## Cluster is Red Because Failover Capacity is Violated

A DRS cluster becomes red when it is invalid. It may become red because failover capacity is violated.

**Problem**

The cluster attempts to failover virtual machines in case of host failure, but is not guaranteed to have enough resources available to failover all virtual machines covered by the failover requirements.

**Cause**

If a cluster enabled for HA loses so many resources that it can no longer fulfill its failover requirements, a message appears and the cluster's status changes to red.

**Solution**

Review the list of configuration issues in the yellow box at the top of the cluster Summary page and address the issue that is causing the problem.

## No Hosts are Powered Off When Total Cluster Load is Low

Hosts are not powered off when the total cluster load is low.

### Problem

Hosts are not powered off when the total cluster load is low because extra capacity is needed for HA failover reservations.

### Cause

Hosts might not be powered off for the following reasons:

- The `MinPoweredOn{Cpu|Memory}Capacity` advanced options settings need to be met.
- Virtual machines cannot be consolidated onto fewer hosts due to their resource reservations, VM/Host DRS rules, VM/VM DRS rules, not being DRS-enabled, or not being compatible with the hosts having available capacity.
- Loads are unstable.
- DRS migration threshold is at the highest setting and only allows mandatory moves.
- vMotion is unable to run because it is not configured.
- DPM is disabled on the hosts that might be powered off.
- Hosts are not compatible for virtual machines to be moved to another host.
- Host does not have Wake On LAN, IPMI, or iLO technology. Either one is required for DPM to enter a host in standby.

### Solution

Address the issue that prevents hosts from being powered off when the total cluster load is low.

## Hosts are Powered Off When Total Cluster Load is High

Hosts are powered off when total cluster load is high.

### Problem

DRS determined that virtual machines could be run on a fewer number of hosts without degrading the host or virtual machine performance. DRS is also constrained from moving the virtual machines running on the highly-utilized hosts to the hosts scheduled for power-off.

### Cause

This occurs when the total cluster load is too high.

### Solution

Reduce the cluster load.

## DRS Seldom or Never Performs vMotion Migrations

DRS seldom or never performs vMotion migrations.

### Problem

DRS does not perform vMotion migrations.

**Cause**

DRS never performs vMotion migrations when one or more of the following issues is present on the cluster.

- DRS is disabled on the cluster.
- The hosts do not have shared storage.
- The hosts in the cluster do not contain a vMotion network.
- DRS is manual and no one has approved the migration.

DRS seldom performs vMotion when one or more of the following issues is present on the cluster:

- Loads are unstable, or vMotion takes a long time, or both. A move is not appropriate.
- DRS seldom or never migrates virtual machines.
- DRS migration threshold is set too high.

DRS moves virtual machines for the following reasons:

- Evacuation of host that a user requested enter maintenance or standby mode.
- VM/Host DRS rules or VM/VM DRS rules.
- Reservation violations.
- Load imbalance.
- Power management.

**Solution**

Address the issues that are causing DRS to avoid performing vMotion migrations.

## Host Problems

Host problems might cause DRS to not perform as expected.

### DRS Recommends Host be Powered On to Increase Capacity When Total Cluster Load Is Low

The host should be powered on to help provide more capacity for the cluster or help hosts that are overloaded.

**Problem**

DRS recommends that the host be powered on to increase capacity when the total cluster load is low.

**Cause**

The recommendation might be made because:

- The cluster is a DRS-HA cluster. Additional powered-on hosts are needed in order to provide more failover capability.
- Some hosts are overloaded and virtual machines on currently powered-on hosts can be moved to hosts in standby mode to balance the load.
- The capacity is needed to meet the `MinPoweredOn{Cpu|Memory}Capacity` advanced options.

**Solution**

Power on the host.



## Total Cluster Load Is High

The total cluster load is high.

### Problem

When the total cluster load is high, DRS does not power on the host.

### Cause

The following are possible reasons why DRS does not power on the host:

- VM/VM DRS rules or VM/Host DRS rules prevent the virtual machine from being moved to this host.
- Virtual machines are pinned to their current hosts, hence DRS cannot move these virtual machines to hosts in standby mode to balance the load.
- DRS or DPM is in manual mode and the recommendations were not applied.
- No virtual machines on highly utilized hosts will be moved to that host.
- DPM is disabled on the host because of a user setting or host previously failing to successfully exit standby.

### Solution

Address that issue that prevents DRS from powering on the host.

## Total Cluster Load Is Low

The total cluster load is low.

### Problem

When the total cluster load is low, DRS does not power off the host.

### Cause

The following are possible reasons why DRS does not power off the host:

- Distributed Power Management (DPM) detected better candidates to power off.
- vSphere HA needs extra capacity for failover.
- The load is not low enough to trigger the host to power off.
- DPM projects that the load will increase.
- DPM is not enabled for the host.
- DPM threshold is set too high.
- While DPM is enabled for the host, no suitable power-on mechanism is present for the host.
- DRS cannot evacuate the host.
- The DRS migration threshold is at the highest setting and only performs mandatory moves.

### Solution

Address the issue that is preventing DRS from powering off the host.

## DRS Does Not Evacuate a Host Requested to Enter Maintenance or Standby Mode

DRS does not evacuate a host requested to enter maintenance mode or standby mode.

### Problem

When you attempt to put a host into maintenance or standby mode, DRS does not evacuate the host as expected.

### Cause

vSphere HA is enabled and evacuating this host might violate HA failover capacity.

### Solution

There is no solution. If appropriate, disable vSphere HA before you attempt to put the host into maintenance mode or standby mode.

## DRS Does Not Move Any Virtual Machines onto a Host

DRS does not move any virtual machines onto a host.

### Problem

DRS does not recommend migration of virtual machine to a host that has been added to a DRS-enabled cluster.

### Cause

After a host has been added to a DRS-enabled cluster, the virtual machines deployed to the host become part of the cluster. DRS can recommend migration of some virtual machines to this host just added to the cluster. If that does not occur, there may be problems with vMotion, host compatibility, or affinity rules. The following are possible reasons:

- vMotion is not configured or enabled on this host.
- Virtual machines on other hosts are not compatible with this host.
- The host does not have sufficient resources for any virtual machine.
- Moving any virtual machines to this host would violate a VM/VM DRS rule or VM/Host DRS rule.
- This host is reserved for HA failover capacity.
- A device is mounted to the virtual machine.
- The vMotion threshold is too high.
- DRS is disabled for the virtual machines, hence the virtual machine could not be moved onto the destination host.

### Solution

Address the issue that prevents DRS from moving virtual machines onto a host.

## DRS Does Not Move Any Virtual Machines from a Host

DRS does not move any virtual machines from a host.

### Problem

Virtual machines are not moved from this host.

**Cause**

This may be because of problems with vMotion, DRS, or host compatibility. The following are the possible reasons:

- vMotion is not configured or enabled on this host.
- DRS is disabled for the virtual machines on this host.
- Virtual machines on this host are not compatible with any other hosts.
- No other hosts have sufficient resources for any virtual machines on this host.
- Moving any virtual machines from this host would violate a VM/VM DRS rule or VM/Host DRS rule.
- DRS is disabled for one or more virtual machines on the host.
- A device is mounted to the virtual machine.

**Solution**

Address the issues that are preventing DRS from moving virtual machines from the host.

## Virtual Machine Problems

Virtual machine problems might cause DRS to not perform as expected.

### Insufficient CPU or Memory Resources

The virtual machine does not receive enough CPU or memory resources.

**Problem**

In some cases, the virtual machine's demand is greater than its resource entitlement. When this occurs, the virtual machine doesn't receive enough CPU or memory resources.

**Cause**

The following sections describe the factors that influence the entitlement for a virtual machine.

- |  |  |
|--|--|
| <b>Cluster is Yellow or Red</b>          | If the cluster is yellow or red, the capacity is insufficient to meet the resource reservations configured for all virtual machines and resource pools in the cluster. The particular virtual machine might be one that is not receiving its reservation. Check the status of the cluster (red or yellow) and resolve the situation. |
| <b>Resource Limit is Too Restrictive</b> | The virtual machine, its parent resource pool, or its resource pool ancestors might have a configured resource limit that is too restrictive. Check whether demand is equal to or greater than any configured limits.  |

**Cluster is Overloaded**

The cluster on which the virtual machine is running might have insufficient resources. Also, the virtual machine's share value is such that other virtual machines are granted proportionally more of the resources. To determine the demand is larger than the capacity, check the cluster statistics.

**Host is Overloaded**

To determine if the host's resources are oversubscribed, check the host statistics. If they are oversubscribed, consider why DRS is not moving any of the virtual machines now running on the host to other hosts. This condition might exist for the following reasons:

- The VM/VM DRS rules and VM/Host DRS rules require the current virtual machine-to-host mapping. If such rules are configured in the cluster, consider disabling one or more of them. Then run DRS and check whether the situation is corrected.
- DRS cannot move this virtual machine or enough of the other virtual machines to other hosts to free up capacity. DRS will not move a virtual machine for any of the following reasons:
  - DRS is disabled for the virtual machine.
  - A host device is mounted to the virtual machine.
  - Either of its resource reservations is so large that the virtual machine cannot run on any other host in the cluster.
  - The virtual machine is not compatible with any other host in the cluster.

Check whether any of these conditions exist for the virtual machine. If none exist, the conditions might exist for other virtual machines in the cluster. If this is the case, DRS cannot balance the cluster to address the virtual machine's demand.

- Decrease the DRS migration threshold setting and check whether the situation is resolved.
- Increase the virtual machine's reservation.

**Solution**

Address the problem that is causing the virtual machine to not receive enough CPU or memory resources.

**VM/VM DRS Rule or VM/Host DRS Rule Violated**

DRS rules specify which host a virtual machine must or must not reside on, or which virtual machines must be or must not be on the same host.

**Problem**

A VM/VM DRS rule or a VM/Host DRS rule is violated.

**Cause**

VM/VM DRS rules specify that selected virtual machines should be placed on the same host (affinity) or that virtual machines be placed on different hosts (anti-affinity). VM/Host DRS rules specify that selected virtual machines should be placed on specified hosts (affinity) or that selected virtual machines should not be placed on specified hosts (anti-affinity).

When a VM/VM DRS rule or VM/Host DRS rule is violated, it might be because DRS cannot move some or all of the virtual machines in the rule. The reservation of the virtual machine or other virtual machines in the affinity rule, or their parent resource pools, might prevent DRS from locating all virtual machines on the same host.

#### **Solution**

- Check the DRS faults panel for faults associated with affinity rules.
- Compute the sum of the reservations of all the virtual machines in the affinity rule. If that value is greater than the available capacity on any host, the rule cannot be satisfied.
- Compute the sum of the reservations of their parent resource pools. If that value is greater than the available capacity of any host, the rule cannot be satisfied if the resources are obtained from a single host.

## **Virtual Machine Power On Operation Fails**

An error message appears stating that the virtual machine fails to power on.

#### **Problem**

The virtual machine fails to power on.

#### **Cause**

The virtual machine might fail to power on because of insufficient resources or because there are no compatible hosts for the virtual machine.

#### **Solution**

If the cluster does not have sufficient resources to power on a single virtual machine or any of the virtual machines in a group power-on attempt, check the resources required by the virtual machine against those available in the cluster or its parent resource pool. If necessary, reduce the reservations of the virtual machine to be powered-on, reduce the reservations of its sibling virtual machines, or increase the resources available in the cluster or its parent resource pool.

## **DRS Does Not Move the Virtual Machine**

DRS does not move the virtual machine when it is initially powered on despite insufficient resources on the host.

#### **Problem**

When you power on a virtual machine, DRS does not migrate it as expected when there are not enough resources on the host where the virtual machine is registered.

#### **Cause**

The following are possible reasons why DRS does not move the virtual machine.

- DRS is disabled on the virtual machine.
- The virtual machine has a device mounted.
- The virtual machine is not compatible with any other hosts.
- No other hosts have a sufficient number of physical CPUs or capacity for each CPU for the virtual machine.
- No other hosts have sufficient CPU or memory resources to satisfy the reservations and required memory of this virtual machine.
- Moving the virtual machine will violate an affinity or anti-affinity rule.

- The DRS automation level of the virtual machine is manual and the user does not approve the migration recommendation.
- DRS will not move fault tolerance-enabled virtual machines.

**Solution**

Address the issue that prevents DRS from moving the virtual machine.

# Index

## Numerics

3D Graphics card **45**

## A

admission control

  CPU **23**

  resource pools **58**

  with expandable resource pools **59**

advanced attributes

  hosts **113**

  memory **114**

  NUMA **115**

  Storage I/O Control **50**

  virtual machines **116, 118**

  virtual NUMA **117**

affinity rules

  creating **86, 87**

  intra-VM **101**

  Storage DRS **99**

alarms **83**

AMD Opteron-based systems **105, 115**

anti-affinity rules **122**

applications

  CPU-bound **18**

  single-threaded **18**

automation level

  datastore clusters **91**

  Storage DRS **98**

  virtual machines **69**

## B

ballooning, memory **36**

Baseboard Management Controller (BMC) **80**

## C

cluster, host inactive **123**

cluster issues

  cluster load high **127**

  DRS does not vMotion **127**

  host not powered off **127**

  hosts power off **127**

  load imbalance **125**

  low cluster load **127**

  red cluster because failover capacity  
    violated **126**

  red cluster because inconsistent resource  
    pool **126**

  yellow cluster **126**

cluster problems **125**

cluster settings

  affinity rules **86, 87**

  virtual machine swapfile location **39**

clustering, datastores **89, 94**

clusters

  creating **66**

  editing **67**

congestion threshold **50**

correcting error isolation **43**

CPU

  admission control **23**

  managing allocation **17, 19**

  overcommitment **17**

  virtual **109**

CPU resources, editing **13**

CPU affinity

  hyperthreading **21**

  NUMA nodes **111**

  potential issues **23**

CPU capacity insufficient **123**

CPU configuration **19**

CPU power efficiency **23**

CPU virtualization

  hardware-assisted **18**

  software-based **17**

CPU-bound applications **18**

  creating, clusters **66**

  custom automation mode, DRS **69**

## D

datastore clusters

  about **89**

  adding datastores **94**

  as resource providers **9**

  automation level **91**

  creating **90–92**

  maintenance mode **95**

  removing datastores **94**

datastores

  compatibility **93**

  maintenance mode **95**

Distributed Power Management, See DPM

**DPM**

- and admission control **15**
- automation level **82**
- enabling **82**
- Last Time Exited Standby **83**
- monitoring **83**
- overrides **83**
- threshold **82**

**DRS**

- creating rules **86**
- disable **70**
- enabling **67**
- fully automated **68**
- group power on **62**
- initial placement **61, 62**
- load balancing **61**
- manual **68**
- migration **61**
- migration recommendations **64**
- migration thresholds **64**
- partially automated **68**
- single virtual machine power on **62**
- virtual machine migration **63**
- vMotion network **65**

**DRS groups**

- host **84**
- virtual machine **85**

**DRS clusters**

- adding managed hosts **71**
- adding unmanaged hosts **72**
- as resource providers **9**
- creating **68**
- managing resources with **71**
- prerequisites **65**
- processor compatibility **65**
- shared storage **65**
- shared VMFS volume **65**
- validity **75**

dual-processor virtual machine **17**

dynamic load balancing, NUMA **107**

Dynamic Voltage and Frequency Scaling (DVFS) **23**

**E**

Enhanced vMotion Compatibility (EVC) **18, 65, 68**

entering maintenance mode, host **74**

Exit Standby Error **83**

expandable reservations, example **58**

**F**

fault definitions **121**

**faults**

- affinity rules **122**
- anti-affinity rules **122**
- CPU capacity **123**
- host inactive **123**
- host incorrect state **122**
- incompatible host **122**
- insufficient capacity **122**
- insufficient resources **123**
- no compatible hard affinity host **123**
- no compatible soft affinity host **123**
- non-compatible virtual machine **122**
- pcpus insufficient **123**
- soft rule violation correction disallowed **123**
- soft rule violation correction impact **124**
- virtual machine pinned **122**
- vMotion **123**

fully automated DRS **68**

**G**

grafted, resource pool **71, 72**

**graphics**

- 3D **45**
- GPU **45**

group power on **62**

**groups**

- DRS host **84**
- DRS virtual machine **85**

**H**

home nodes, NUMA **107**

**host**

- entering maintenance mode **74**
- memory allocation **34**

host DRS groups **84**

host incorrect state **122**

**host issues**

- high cluster load **129**
- host not powered off **129**
- host not powered on **129**
- low cluster load **128, 129**
- maintenance mode **130**
- standby mode **130**

virtual machines not moved by DRS **130**

host power management, custom policy **24**

host power management policies, setting **24**

host problems **128**

**host-local swap**

- DRS cluster **37**
- standalone host **37**



- hosts
  - adding to DRS clusters **71, 72**
  - advanced attributes **113**
  - as resource providers **9**
  - removing from a DRS cluster **75**
  - virtual machine swapfile location **38**
- hyperthreading
  - and hosts **21**
  - CPU affinity **21**
  - disabling **19**
  - enabling **21**
  - performance implications **20**
- I**
- idle memory tax **35**
- iLO, configuring **80**
- incompatible host **122**
- initial placement
  - NUMA **107**
  - Storage DRS **90**
- insufficient capacity **122**
- insufficient resources **123**
- Intelligent Platform Management Interface (IPMI), configuring **80**
- inter-VM anti-affinity rules, creating **100**
- intra-VM anti-affinity rules **101**
- introduction **125**
- invalid DRS clusters **79**
- L**
- Last Time Exited Standby **83**
- latency sensitivity **118**
- Latency Sensitivity setting **118**
- limit **12**
- load balancing
  - datastores **89**
  - Storage DRS **92**
  - virtual machines **63**
- logical processors **19, 20**
- LPage.LPageAlwaysTryForNPT **114**
- LPage.LPageDefragEnable **114**
- LPage.LPageDefragRateTotal **114**
- LPage.LPageDefragRateVM **114**
- M**
- maintenance mode
  - datastores **95, 96**
  - hosts **74**
  - hosts entering **74**
  - ignore affinity rules **96**
- manual DRS **68**
- Mem.AllocGuestLargePage **114**
- Mem.AllocUseGuestPool **114**
- Mem.AllocUsePSharePool **114**
- Mem.BalancePeriod **114**
- Mem.CtlMaxPercent **114**
- Mem.IdleTax **35, 114**
- Mem.MemZipEnable **114**
- Mem.MemZipMaxPct **114**
- Mem.SamplePeriod **34, 114**
- Mem.ShareScanGHz **40, 114**
- Mem.ShareScanTime **40, 114**
- memory
  - advanced attributes **114**
  - balloon driver **36**
  - managing allocation **33**
  - overcommitment **28, 38**
  - overhead **27**
  - overhead, understanding **33**
  - reclaiming unused **35**
  - sharing **28**
  - sharing across virtual machines **40**
  - virtual machines **35**
- memory allocation **43**
- memory allocationmemory allocation **44**
- memory resources, editing **13**
- memory virtualization
  - hardware-assisted **30**
  - software-based **30**
- memory affinity, NUMA nodes **111**
- memory compression **41**
- memory compression cache
  - disabling **41**
  - enabling **41**
  - set size **41**
- memory idle tax **35**
- memory reclamation **35**
- memory reliability **43**
- memory usage **42**
- migration recommendations **64**
- migration thresholds, DRS **64**
- monitoring, Storage I/O Control **48**
- monitoring software **83**
- multicore processors **20**
- N**
- NUMA
  - advanced attributes **115**
  - AMD Opteron-based systems **108**
  - challenges for operating systems **105**
  - CPU affinity **111, 112**
  - description **105**
  - dynamic load balancing **107**
  - home nodes **107**
  - home nodes and initial placement **107**

- IBM Enterprise X-Architecture **108**
  - manual controls **110**
  - memory affinity **112**
  - memory page sharing **108**
  - optimization algorithms **107**
  - page migration **107**
  - scheduling **106**
  - supported architectures **108**
  - transparent page sharing **108**
  - virtual **108, 109**
  - Numa.AutoMemAffinity **115**
  - Numa.MigImbalanceThreshold **115**
  - Numa.PageMigEnable **115**
  - Numa.RebalanceCoresNode **115**
  - Numa.RebalanceCoresTotal **115**
  - Numa.RebalanceEnable **115**
  - Numa.RebalancePeriod **115**
- O**
- overcommitted DRS clusters **78**
  - overhead memory **27, 43**
- P**
- page migration, NUMA **107**
  - parent resource pool **53**
  - partially automated DRS **68**
  - pcpus insufficient **123**
  - performance, CPU-bound applications **18**
  - physical processors **19**
  - policies
    - CPU power management **23**
    - host power management **24**
  - power management policies, CPU **23**
  - power on, single virtual machine **62**
  - preface **7**
  - processor-specific behavior **18**
  - processors
    - assigning **22**
    - compatibility requirements **65**
- R**
- recommendations, Storage DRS **97**
  - red DRS clusters **79**
  - Reliable Memory **118**
  - requirements, datastore clusters **93**
  - reservation **12**
  - resource allocation settings
    - changing **14**
    - limit **12**
    - reservation **12**
    - shares **11**
    - suggestions **13**
  - resource consumers **10**
  - resource management
    - customizing **113**
    - goals **10**
  - resource pool tree snapshot **70**
  - resource pools
    - adding virtual machines **56**
    - admission control **58**
    - advantages **54**
    - creating **55**
    - deleting **58**
    - editing attributes of **56**
    - grafted **71, 72**
    - parent **53**
    - removing **58**
    - removing virtual machines **57**
    - root resource pool **53**
    - siblings **53**
  - resource providers **9**
  - resource types **9**
  - restore a resource pool tree snapshot **70**
  - root resource pool **53**
- S**
- sched.mem.maxmemctl **36, 116**
  - sched.mem.pshare.enable **116**
  - sched.swap.dir **116**
  - sched.swap.persist **116**
  - shares, Storage I/O Control **47**
  - shares and limits, Storage I/O Control **48**
  - sharing memory **28**
  - siblings **53**
  - single virtual machine power on **62**
  - single-processor virtual machine **17**
  - single-threaded applications **18**
  - SMP virtual machines **18**
  - standby mode, Last Time Exited Standby **83**
  - Storage DRS
    - about **91**
    - affinity rules **99**
    - aggressiveness **92**
    - anti-affinity rules **100**
    - automation level **98**
    - disabling **91**
    - enabling **91**
    - I/O load balancing **92**
    - initial placement **90**
    - maintenance mode **95**
    - recommendations **90, 97**
    - scheduled task **98**
    - space load balancing **92**

- statistics **102**
    - thresholds **92**
  - Storage I/O Control
    - enabling **49**
    - limitations **47**
    - monitoring **48**
    - requirements **47**
    - shares and limits **47–49**
    - threshold **50**
  - storage migration recommendations **90**
  - storage requirements **65**
  - Storage vMotion
    - datastore cluster compatibility **103**
    - recommendations **90**
  - swap file
    - deleting **40**
    - location **37**
    - using **36**
  - swap files, VMX **35**
  - swap space **38**
  - swap to VMX **35**
  - swapfile location **39**
  - swapfiles **38**
  - system resource allocation, editing **13**
  - System Resource Allocation Table (SRAT) **106**
  - system swap **43**
  - system swap data **44**
- T**
- threshold, Storage I/O Control **50**
- V**
- valid DRS clusters **76**
  - vCenter Server events **83**
  - virtual flash **66**
  - virtual CPUs, changing number **109**
  - virtual machine affinity **84–87**
  - virtual machine anti-affinity **84–87**
  - virtual machine DRS groups **85**
  - Virtual Machine File System (VMFS) **65, 116**
  - virtual machine issues
    - affinity rules **132**
    - anti-affinity rules **132**
    - CPU **131**
    - memory resources **131**
    - power on failure **133**
    - virtual machine not moved by DRS **133**
  - virtual machine monitor (VMM) **27**
  - virtual machine pinned **122**
  - virtual machine problems **131**
  - virtual machine swapfile location **39**
  - virtual machines
    - adding to resource pools **56**
    - advanced attributes **116**
    - as resource consumers **10**
    - assigning to a specific processor **22**
    - configuration file **65**
    - DRS automation level **69**
    - memory **27, 35**
    - migration **63**
    - monitor **30**
    - number of virtual processors **18**
    - overhead memory **34**
    - removing from a DRS cluster **73**
    - removing from resource pools **57**
    - sharing memory across **40**
    - swapfile location **38, 39**
  - virtual NUMA
    - advanced attributes **117**
    - advanced options **109**
    - controlling **109**
  - VM-Host affinity rule
    - conflicts **87**
    - using **87**
  - VM-VM affinity rule, conflicts **86**
  - VMFS (Virtual Machine File System) **65, 116**
  - VMFS volume requirements **65**
  - VMM **27, 30**
  - vmmemctl, Mem.CtlMaxPercent **114**
  - vMotion **123**
  - VMX swap files **35**
  - vSphere HA, enabling **67**
  - vSphere DRS, creating rules **87**
- W**
- wake protocols **80**
  - Wake-on-LAN (WOL), testing **81**
  - working set size **34**
- Y**
- yellow DRS clusters **78**

