



Implementing and Configuring Dynamic Memory

Updated: October 2010

© 2010 Microsoft Corporation. All rights reserved. This document is developed prior to the product's release to manufacturing, and as such, we cannot guarantee that all details included herein will be exactly as what is found in the shipping product. The information contained in this document represents the current view of Microsoft Corporation on the issues discussed as of the date of publication. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information presented after the date of publication. The information represents the product at the time this document was printed and should be used for planning purposes only. Information subject to change at any time without prior notice. This whitepaper is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS SUMMARY.

Microsoft, Hyper-V, Windows, the Windows logo, Windows Vista, Windows 7 and Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

All other trademarks are property of their respective owners.

Executive Summary

Dynamic Memory is a new feature of Hyper-V™ introduced in Service Pack 1 for Windows Server® 2008 R2 that enables Hyper-V hosts to dynamically adjust the amount of memory available to virtual machines in response to changing workloads. The benefits of Dynamic Memory include higher virtual machine consolidation ratios and increased flexibility for managing virtualized workloads. This whitepaper explains how Dynamic Memory works, requirements for implementing Dynamic Memory, how to enable and configure Dynamic Memory, and some considerations and best practices for implementing Dynamic Memory.

Disclaimer

The information contained in this document relates to pre-release software product, which may be substantially modified before its first commercial release. Accordingly, the information may not accurately describe or reflect the software product when first commercially released. This document is provided for informational purposes only, and Microsoft makes no warranties, express or implied, with respect to this document or the information contained in it.

Understanding Dynamic Memory

Hyper-V in Windows Server 2008 and Windows Server 2008 R2 employed a static memory model where the administrator of a Hyper-V host would assign the amount of physical memory for each virtual machine on the host. Then when a virtual machine is started on the host, the assigned memory is allocated to the virtual machine and never changes unless the virtual machine is shut down and a different amount of memory is assigned to it. For Microsoft customers who want to achieve higher virtual machine consolidation ratios for server consolidation or VDI scenarios, the static memory model of Hyper-V on these earlier platforms

becomes a limiting factor. To meet the needs of such customers, Microsoft is introducing a new feature called Dynamic Memory in Service Pack 1 for Windows Server 2008 R2.

Dynamic Memory allows the amount of physical memory allocated to a virtual machine to be dynamically adjusted in response to the changing workload needs of the virtual machine. Instead of assigning a specific amount of memory to a virtual machine, the administrator instead configures a range of memory, memory priority and other settings that Hyper-V then uses to determine how much memory to allocate to the virtual machine in real time. For more information on Dynamic Memory configuration settings, see the section titled "Enabling and Configuring Dynamic Memory" later in this whitepaper.

The amount of physical memory on a Hyper-V host is the key factor in determining how many virtual machines can concurrently be running on the host. Dynamic Memory enables customers to better utilize the memory resources of Hyper-V hosts by balancing how memory is distributed between running virtual machines. Memory can be dynamically reallocated between different virtual machines in response to the changing workloads of these machines. Dynamic Memory thus enables more efficient use of memory while maintaining consistent workload performance and scalability.

Implementing Dynamic Memory means that higher levels of server consolidation can be achieved with minimal impact on performance. Dynamic Memory also means larger numbers of virtual desktops per Hyper-V host for VDI scenarios. The net result for both scenarios is more efficient use of expensive server hardware resources, which can translate into easier management and lower costs.

How Dynamic Memory Works

At a high level, Dynamic Memory works by pooling together all physical memory on the host, minus a certain amount of memory that's reserved for exclusive use by the parent partition. Memory from this pool is then securely distributed

among virtual machines running on the host, while the amount of memory allocated to a particular virtual machine dynamically adjusts in real time based on the needs of the virtual machine.

The Hyper-V host and the enlightened guest operating system running in the virtual machine continually communicate with each other to determine the current memory needs of the virtual machine. If the workload of the virtual machine increases and more memory is needed in order to maintain consistent performance, memory is dynamically added to the virtual machine. If the workload of the virtual machine decreases or other virtual machines on the host having higher memory weight request additional memory, memory is dynamically removed from the virtual machine.

This addition or removal of memory is implemented using the driver enlightened architecture (VSP/VSC/VMBus) of Hyper-V. On the host side, the Virtual Service Provider (VSP) arbitrates the allocation of physical memory resources between the virtual machines running on the host. On the virtual machine side, the Virtual Service Consumer (VSC) collects the information to determine virtual machine's memory needs and executes necessary operations to add or remove memory.

In order to be able to dynamically add memory to a virtual machine, Dynamic Memory requires that the virtual machine's guest operating system include a kernel enlightenment that supports Dynamic Memory. For a list of guest operating systems supported by Dynamic Memory, see the section titled "Dynamic Memory Requirements" later in this whitepaper.

In order to be able to dynamically remove memory from a virtual machine, Dynamic Memory employs a process called ballooning. Ballooning works by having the virtual machine and the host agree that certain memory pages will no longer be accessible by the guest. This effectively reduces the amount of memory available to the guest, although system tools such as Task Manager running in the guest operating system will not display any decrease in memory available to the guest. Should the virtual machine later request additional memory from the host, memory pages that were previously ballooned will be unballooned as

needed until all ballooned pages have been returned to the guest. Should the guest continue to request additional memory after this, Dynamic Memory allocates more memory to the virtual machine. Note that memory is not instantaneously removed from a virtual machine unless it is instantaneously required for use by another virtual machine.

Understanding Memory Pressure and Memory Buffer

Dynamic Memory determines the amount of memory needed by a virtual machine by calculating something called memory pressure. To perform this calculation, Hyper-V looks at the total committed memory of the guest operating system running in the virtual machine and then calculates pressure as the ratio of how much memory the virtual machine wants to how much it has. The amount of memory that Hyper-V then assigns to the virtual machine equals total committed memory plus some additional memory to be used as a buffer. This additional amount of memory is configurable on a per-virtual machine basis as a percentage amount. For example, configuring a buffer value of 50% means that an additional memory of up to 50% of committed memory can be allocated to the virtual machine. The guest operating system of the virtual machine typically uses this additional memory for its system file cache to boost the performance of the operating system and applications.

As an example, assume that Hyper-V has determined that the total committed memory a virtual machine needs is 4 GB. If the Dynamic Memory buffer value has been configured as 50% then Hyper-V can make up to 2 GB of additional memory available to the virtual machine for use by the file system cache. The total physical memory being used by the virtual machine may therefore be as high as 6 GB. However, Dynamic Memory does not guarantee that the total committed memory amount is always assigned to the virtual machine. Neither does Dynamic Memory guarantee that the additional memory amount configured as a buffer value is always assigned to the virtual memory. This is

because the actual amount of memory assigned to a virtual machine depends upon the memory pressure being exerted upon the host by the memory needs of other virtual machines running on the host.

For information on how to configure Dynamic Memory buffering, see the section titled "Enabling and Configuring Dynamic Memory" later in this whitepaper.

Understanding Dynamic Memory Weight

Dynamic Memory also allows you to assign a memory weight value to each virtual machine. In a production environment, some virtual machines may have different performance requirements than others because of the different workloads they perform. Virtual machines that have higher performance requirements should have a higher weight for requesting additional memory when they need it than less important virtual machines running on the same host. Because Dynamic Memory now allows the physical memory of a Hyper-V host to become a dynamically schedulable resource similar to the CPU resources of the host, prioritizing memory allocation on a per-virtual machine basis is now possible.

Memory weight comes into play only when all available physical memory has been allocated to virtual machines on the host. If there is still lots of unused memory on the host, changing the memory weight values for different virtual machines will have no effect. If unused memory is scarce however, increasing the memory weight value of one virtual machine will cause the memory of other virtual machines to shrink as the host reallocates memory from them to the one having increased weight. Another way of saying this is that the higher the memory weight of a virtual machine, the lower the pressure and the smaller the amount of memory removed from the machine when the physical memory of the host is being fully utilized.

For information on how to configure memory weight, see the section titled "Enabling and Configuring Dynamic Memory" later in this whitepaper.

Dynamic Memory Requirements

Dynamic Memory has requirements on both the host side and the guest operating system side. The following sections summarize these requirements for the beta release of Service Pack 1.

Host Requirements for Dynamic Memory

In order to be able to use the Dynamic Memory feature on a Hyper-V host, Service Pack 1 must be applied to one of the following virtualization platforms:

- Windows Server 2008 R2 with the Hyper-V server role installed
- Microsoft Hyper-V Server 2008 R2

Guest Requirements for Dynamic Memory

The following Windows server operating systems support Dynamic Memory when installed as the guest operating system on a virtual machine:

- Windows Server 2008 R2 Standard Edition SP1*
- Windows Server 2008 R2 Enterprise Edition SP1
- Windows Server 2008 R2 Datacenter Edition SP1
- Windows Server 2008 R2 Web Edition SP1*
- Windows Server 2008 Standard Edition SP2*
- Windows Server 2008 Enterprise Edition SP2
- Windows Server 2008 Datacenter Edition SP2
- Windows Server 2008 Web Edition SP2*
- Windows Server 2003 R2 Standard Edition SP2 or higher*

- Windows Server 2003 R2 Enterprise Edition SP2 or higher
- Windows Server 2003 R2 Datacenter Edition SP2 or higher
- Windows Server 2003 R2 Web Edition SP2 or higher*
- Windows Server 2003 Standard Edition SP2 or higher*
- Windows Server 2003 Enterprise Edition SP2 or higher
- Windows Server 2003 Datacenter Edition SP2 or higher
- Windows Server 2003 Web Edition SP2 or higher*

Note: The Beta release of Service Pack 1 does not support Dynamic Memory for the operating systems marked with an asterisk (*) above. However, support for Dynamic Memory for these operating systems will be added in a future release of SP1.

Note: Dynamic Memory is supported for both the x86 and x64 architectures of Windows Server 2003, Windows Server 2003 R2 and Windows Server 2008.

The following Windows client operating systems for both x86 and x64 architecture also support Dynamic Memory when installed as the guest operating system on a virtual machine:

- Windows® 7 Enterprise Edition
- Windows 7 Ultimate Edition
- Windows Vista® Enterprise Edition SP2
- Windows Vista Ultimate Edition SP2

Enabling and Configuring Dynamic Memory

Enabling and configuring Dynamic Memory on a virtual machine involves the following steps:

- Upgrading Integration Components
- Enabling Dynamic Memory
- Configuring Startup RAM
- Configuring Maximum RAM
- Configuring Memory Buffer
- Configuring Memory Weight

The sections that follow provide more information concerning each of these steps.

Upgrading Integration Components

Before you can use Dynamic Memory on a virtual machine, the Integration Components installed in the guest operating system must first be upgraded to the latest version (the version available on a Hyper-V R2 SP1 host).

If the guest operating system is running a version of Windows 7 or Windows Server 2008 R2 that supports Dynamic Memory, there are two ways you can choose from to upgrade the Integration Components in the guest operating system:

- By applying Service Pack 1 for Windows 7 and Windows Server 2008 R2 to the guest operating system of the running virtual machine.

- By selecting the Action menu option in the Virtual Machine Connection tool and selecting Insert Integration Services Setup Disk in the running virtual machine.

For other supported guest operating systems, only the second method above can be used to upgrade the Integration Components in the guest. Note that both of these methods require rebooting the guest operating system which may result in service interruption for applications or users who access the virtual machine.

Note that if you fail to upgrade the Integration Components on a guest, the settings for Dynamic Memory will still be configurable for the virtual machine. However, when you change the Memory Management setting from Static to Enabled and configure Startup RAM and other Dynamic Memory settings for the virtual machine, all of the settings you configure are ignored except for Startup RAM which then determines the fixed amount of memory assigned to the machine. In other words, the Dynamic Memory settings for the virtual machine can be configured but they don't do anything—a virtual machine that doesn't have the latest Integration Components can only have a fixed amount of memory assigned to it.

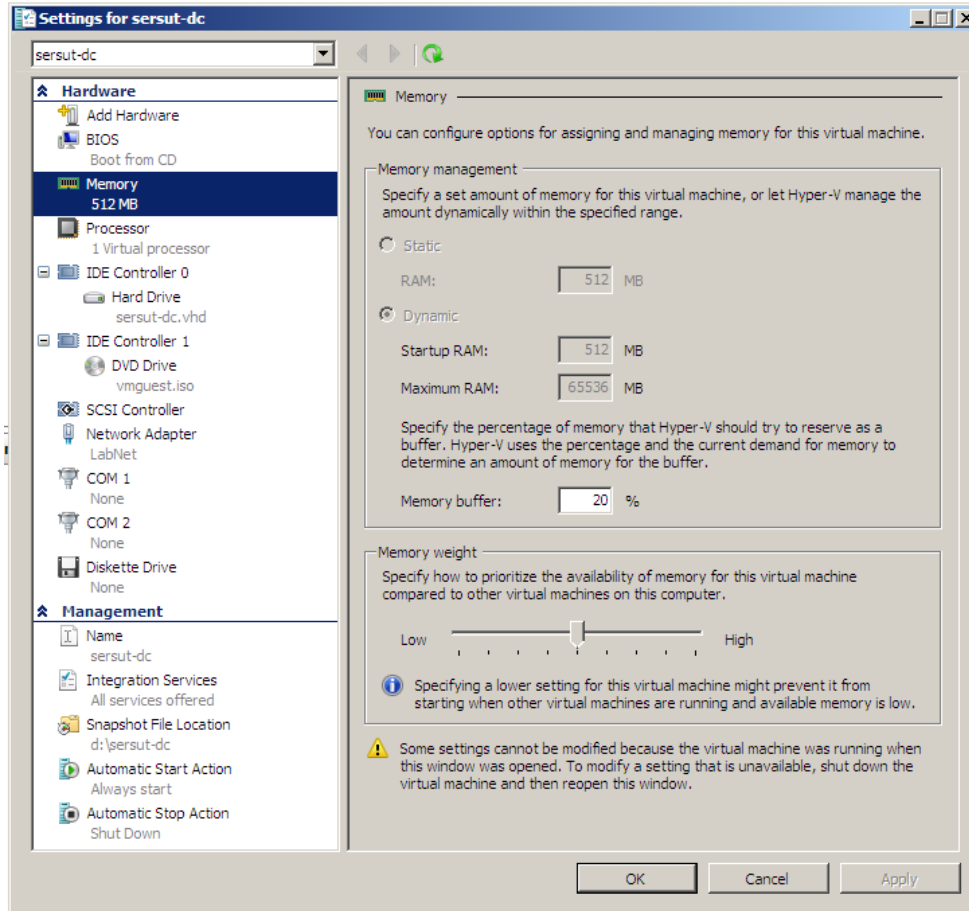
Note also that the settings for configuring Dynamic Memory are also exposed for virtual machines whose guest operating systems do not support Dynamic Memory, for example Windows XP. However, if you try and enable Dynamic Memory for a virtual machine running Windows XP, the guest operating system will only ever be able to access the amount of memory specified by the value of the Startup RAM setting.

Enabling Dynamic Memory

To enable Dynamic Memory on a virtual machine, you must first shut down the virtual machine. Note that shutting down a virtual machine may result in service interruption for applications or users who access the virtual machine. Once the virtual machine is turned off, you can enable and configure Dynamic Memory by

opening the Settings for the virtual machine and selecting Memory under Hardware as shown in Figure 1:

Figure 1: Enabling and configuring Dynamic Memory on a virtual machine.



To enable Dynamic Memory on the virtual machine, change the Memory Management setting from Static to Enabled. If desired you can later disable Dynamic Memory on the machine by changing the setting back to Static to specify that a fixed amount of memory be used by the machine. Note that you can have a mix of different virtual machines on the same Hyper-V R2 SP1 with some virtual machines having Dynamic Memory enabled on them and other virtual machines having fixed amounts of memory assigned to them.

Once you have enabled Dynamic Memory on a virtual machine, you need to configure the Startup RAM, Maximum RAM, Memory Buffer and Memory Weight for the virtual machine. These settings are discussed next in the sections below. To configure Startup RAM or Maximum RAM, you must first shut down the virtual machine which may result in service interruption for applications or users who access the virtual machine. Memory Buffer and Memory Weight however can be configured even when the virtual machine is running.

Configuring Startup RAM

The Startup RAM setting specifies the lower bound for how much physical memory can be allocated to the virtual machine. The memory allocated to the virtual machine can never fall below this value. The default value for Startup RAM is 512 MB. The maximum and minimum values for Start RAM are 65,536 MB and 8 MB respectively. Table 1 shows the recommended value for Startup RAM for different guest operating systems.

Table 1: Recommended value for Startup RAM for different guest operating systems.

Guest operating systems	Recommended value for Startup RAM
Windows 7	512 MB
Windows Vista	
Windows Server 2008	
Windows Server 2008 R2	
Windows Server 2003	128 MB

Windows Server 2003 R2

An important consideration when configuring Startup RAM is that you specify a value that will be sufficient to allow the guest operating system in the virtual machine to boot. You should review carefully the above recommended memory requirements for the guest operating system, taking into consideration as well as any additional memory requirements that may result from applications or services that need to start up on the operating system. Once you have determined the minimum amount of memory the virtual machine needs in order for the guest operating system to boot and services to start up properly, configure the Startup RAM setting for the virtual machine as may be required. Other considerations when configuring Startup RAM are described later in the sections titled "Considerations When Implementing Dynamic Memory" and "Tuning Dynamic Memory" and "Best Practices for Implementing Dynamic Memory" later in this whitepaper.

Configuring Maximum RAM

The Maximum RAM setting specifies upper bound for how much physical memory can be allocated to the virtual machine. The memory allocated to the virtual machine can never grow above this value. The default value for this setting for all guest operating systems is 65,536 MB or 64 GB. This is also the maximum value for this setting, while the minimum amount for this setting is the value specified for Startup RAM. For additional considerations when configuring Maximum RAM, see the section titled "Best Practices for Implementing Dynamic Memory" later in this whitepaper.

Configuring Memory Buffer

The Memory Buffer setting specifies the percentage of memory, based on the workload of the virtual machine, that Hyper-V should try to reserve as a buffer.

The possible values for this setting range from 5% to 2000% with the default value being 20%. For example, if the committed memory in the guest is 1000 MB and the Memory Buffer setting is 20%, Hyper-V will attempt to allocate an additional $20\% * 1000 \text{ MB} = 200 \text{ MB}$ making a total of 1200 MB of physical memory allocated to the virtual machine. If however the Memory Buffer setting in this example is increased to 80%, Hyper-V will attempt to allocate an additional $80\% * 1000 \text{ MB} = 800 \text{ MB}$ making a total of 1800 MB of physical memory allocated to the virtual machine. Whether the host is actually able to reserve such memory in practice depends upon the amount of unused physical memory available and whether memory will need to be removed from other virtual machines running on the host in order to satisfy the workload needs of the virtual machine under consideration.

Additional information on how Dynamic Memory buffering works are explained in the section titled "Understanding Memory Pressure and Memory Buffer" earlier in this whitepaper. For more information on configuring this setting, see the sections titled "Tuning Dynamic Memory" and "Best Practices for Implementing Dynamic Memory" later in this whitepaper.

Configuring Memory Weight

The Memory Weight setting specifies the relative priority the virtual machine has compared to other virtual machines running on the host with respect to the distribution of pooled memory resources. The possible values for this setting range from 1 being the lowest priority to 10,000 being the highest priority. The default value for this setting is 5,000 and corresponds to the halfway mark of the slider.

In general, this setting will have little or no effect as long as unused physical memory is available on the host. Should memory becomes a scarce resource however, Hyper-V will use Memory Weight to determine which virtual machines should receive more memory and which should have memory removed. In other words, increasing the memory weight for a virtual machine will ensure that

available physical memory is assigned to the virtual machine before being assigned to other virtual machines when physical memory becomes scarce on the host. For more information on configuring this setting, see the section titled "Best Practices for Implementing Dynamic Memory" later in this whitepaper.

Monitoring Dynamic Memory

Once Dynamic Memory has been enabled and configured on a virtual machine, you should monitor how memory is being allocated to the virtual machine in response to varying workload. You can monitor how memory is being allocated using two approaches:

- Using the three new columns available in the Virtual Machines pane of Hyper-V Manager.
- Using the new performance counters included in Service Pack 1 for Windows Server 2008 R2.

By using these two techniques for monitoring how Hyper-V allocates memory to a virtual machine, you can determine whether your current Dynamic Memory settings for the machine are optimal or need further adjustment to ensure consistent and predictable performance for the guest operating system and applications installed in the machine. For more information on tuning Dynamic Memory to optimize workload performance, see the section titled "Tuning Dynamic Memory" later in this chapter.

Monitoring Dynamic Memory Using Hyper-V Manager

When Service Pack 1 is applied to a Hyper-V R2 host, the Memory column in the Virtual Machines pane of Hyper-V Manager is replaced with two new columns as shown in Figure 2 below:

- **Assigned Memory** – This displays the total memory currently allocated to each running virtual machine out of the pooled physical memory on the host.
- **Memory Demand** – This displays how much memory the virtual machine needs at this time to meet the requirements of the active processes running in the virtual machine. This amount is equal to the the total committed memory within the VM.
- **Memory Status** - shows how much of the buffer amount specified for the virtual machine is available to the virtual machine at this time. Status is reported as one of the following:
 - **OK:** This indicates that there is enough physical memory available to give the virtual machine the full amount of memory buffer.
 - **Low:** This indicates that the current amount of memory assigned to the virtual machine as a buffer is lower than amount that Hyper-V determines should be available as a buffer.
 - **Warning:** This indicates that there is not enough physical memory available to assign any memory buffer to the virtual machine.

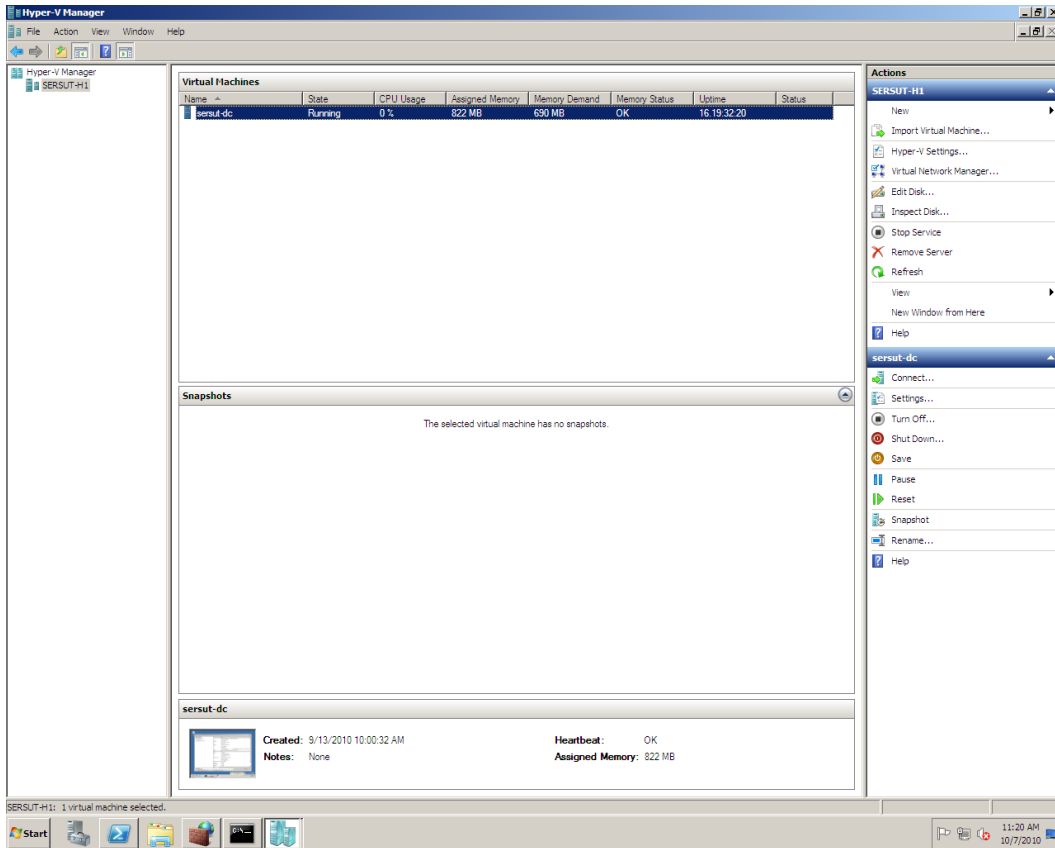


Figure 2: The Virtual Machines pane of Hyper-V Manager for a Hyper-V R2 SP1 host.

Monitoring Dynamic Memory Using Performance Counters

Dynamic Memory adds new performance counters for monitoring how memory is allocated to virtual machines. The two new groups of performance counters for monitoring Dynamic Memory are:

- Hyper-V Dynamic Memory Balancer
- Hyper-V Dynamic Memory VM

Tables 2 and 3 list the individual performance counters included in these groups along with their description.

Table 2: Performance counters included in the Hyper-V Dynamic Memory Balancer counter group.

Performance Counter	Description
Added Memory	The cumulative amount of memory added to VMs.
Available Memory	The amount of memory left on the node.
Average Pressure	The average pressure on the balancer node.
Memory Add Operations	The total number of add operations.
Memory Remove Operations	The total number of remove operations.
Removed Memory	The cumulative amount of memory removed from VMs.

Table 3: Performance counters included in the Hyper-V Dynamic Memory VM counter group.

Performance Counter	Description
Added Memory	The cumulative amount of memory added to the VM.
Average Pressure	The average pressure in the VM.

Current Pressure	The current pressure in the VM.
Guest Visible Physical Memory	The amount of memory visible in the VM.
Maximum Pressure	The maximum pressure band in the VM.
Memory Add Operations	The total number of add operations for the VM.
Memory Remove Operations	The total number of remove operations for the VM.
Minimum Pressure	The minimum pressure band in the VM.
Physical Memory	The current amount of memory in the VM.
Removed Memory	The cumulative amount of memory removed from the VM.

Tuning Dynamic Memory

Dynamic Memory is designed to enable Hyper-V to allocate physical memory to virtual machines as efficiently and dynamically as possible with minimal performance impact. If you are using Dynamic Memory however and are not satisfied with the performance of your virtual machines, you can take various steps to further tune how memory is allocated to the machines. For example, if a virtual machine is performing poorly because too much memory is being removed from it, try increasing the value of the Startup RAM setting for the virtual machine. Doing this will limit the amount of memory that can be removed from the virtual machine when other virtual machines with higher memory priorities are requesting more memory from the host.

Another example of how to tune Dynamic Memory might be when you have a virtual machine that consistently performs below your expectation. In this case, try increasing the value of the Memory Buffer setting as this enables Hyper-V to allocate additional memory to the virtual machine. In general, the more memory allocated to the system file cache, the better the performance for IO intensive applications and services. On the other hand, configuring Memory Buffer settings too high for your virtual machines can result in inefficient distribution of pooled physical memory among virtual machines.

If you have tried adjusting various Dynamic Memory settings in order to improve the performance of a virtual machine but are still not satisfied, try increasing the page file size in the machine's guest operating system. Increasing the page file size can allow larger amounts of memory to be added to the virtual machine if this is needed because of sudden increase in workload on the machine. Increasing the page file size can also enable the virtual machine to run better when available memory becomes scarce on the machine. Note that while Windows generally provides better page file management when configured to manage the page file by itself, if the virtual machine has a workload where Windows' algorithm doesn't provide optimal performance you can try configuring a custom page file size yourself. To do this, change the virtual memory settings in the guest from System Managed Size to Custom Size and specify an Initial Size and Maximum Size for the page file as desired. Be sure to baseline the performance of your virtual machine before and after changing the page file management method to determine whether this form of tuning has been effective or not.

Considerations When Implementing Dynamic Memory

The following are some further considerations you should be aware of when you plan on implementing Dynamic Memory in your Hyper-V environment.

Service Interruption Considerations

As indicated previously in this whitepaper, enabling and configuring Dynamic Memory requires shutting down the virtual machine and this may result in service interruption for applications or users who access the virtual machine. Because of this, you should plan carefully how you will configure Dynamic Memory on each virtual machine so you can enable and configure the feature on all virtual machines during your regular server maintenance window and thereby minimize downtime.

Application Considerations

Certain applications running in a virtual machine's guest operating system may not perform as well after Dynamic Memory is enabled for the virtual machine. Types of applications that fall into this category may include the following:

- Applications that perform their own memory management by taking over certain aspects of memory management from the operating system. Such applications typically grab as much memory as they possibly can in order to ensure the application's best performance which can cause the amount of memory allocated to their virtual machine to grow until it reaches the amount specified by the Maximum RAM setting.
- Applications where memory allocation is a one shot operation that is performed either when the application starts for the first time or each time the application starts. Such applications cannot use any additional memory that might be allocated afterwards while the application is running. If you decide to enable Dynamic Memory for such applications, try increasing the Startup RAM setting to a value that enables the application to perform optimally under expected workloads. Be aware however that increasing this initial memory value reduces the overall flexibility and effectiveness of Dynamic Memory for virtual machines running on the host.

Downlevel Host Considerations

Once a virtual machine has been configured to use Dynamic Memory by installing the latest Integration Components on the guest operating system, the virtual machine will no longer work on pre-SP1 hosts and cannot be moved to such hosts.

Failover Clustering Considerations

Using Hyper-V together with Failover Clustering lets you make virtual machines highly available within your environment. By enabling Dynamic Memory in such a scenario you can more effectively utilize your server hardware. One risk in this scenario is that a virtual machine may fail to start after a hardware failure if other virtual machines do not relinquish memory when they are requested to do so. You can mitigate this risk however by calculating the amount of physical memory available in the cluster when a node has failed and ensure that (a) the sum value of all Startup RAM values for all virtual machines on the cluster does not exceed your calculated value, (b) the sum value of all Maximum RAM values for all virtual machines on the cluster does not exceed twice your calculated value, and (c) the performance of virtual machines when they are limited to this value is acceptable for your environment.

Best Practices for Implementing Dynamic Memory

The following are some best practices for how to implement Dynamic Memory in your Hyper-V environment.

Capacity Planning

One of the advantages provided by Dynamic Memory is that capacity planning no longer needs to be performed at the virtual machine level. Instead, begin by

using Table 1 earlier in this whitepaper to assign a value for Startup RAM to each virtual machine based on guest operating system. Then ensure sure that the total of Startup RAM for all virtual machines does not exceed the total physical memory of the host system by either adding more physical memory to the system or reducing the number of virtual machines on the host.

Other Recommendations

Try not to configure too high a value for the Startup RAM setting on a virtual machine. The reason is because Dynamic Memory as implemented in Hyper-V R2 SP1 can never decrease the memory allocated to the virtual machine below the value of the Startup RAM setting. The Dynamic Memory algorithm will always strive to make this minimum amount of memory available to the virtual machine.

If you determine that a virtual machine that has Dynamic Memory enabled is using too much memory on the host, try limiting the amount of memory used by the virtual machine by lowering the value of the Maximum RAM setting on the virtual machine. You can also lower this setting in environments where it is important to restrict the workloads of virtual machines. An example of this might be a hosting service provider that uses Hyper-V to host virtual machines for its customers.

If you are uncertain as to how large the workload for a virtual machine will be, or if the workload is expected to vary rapidly to a large degree, then consider increasing the value of the Maximum RAM setting and observing the result of this change on the performance of the virtual machine,

More Information

For more information on Dynamic Memory in Hyper-V R2 SP1, see the Windows Virtualization Team Blog at <http://blogs.technet.com/b/virtualization/>.