

Measuring Performance on Virtual Machines

Measuring Memory Performance

Use the following performance monitor counters to measure the impact of available memory on the performance of a guest operating system installed on a Hyper-V virtual machine:

- **Measure available memory on the Hyper-V host operating system** – The amount of physical memory available to the Hyper-V host operating system can be determined by monitoring the “\Memory\Available MBytes” performance monitor counter on the physical computer. This counter reports the amount of free physical memory available to the host operating system. Use the following rules of thumb when evaluating available physical memory available to the host operating system:
 - **\Memory\Available Mbytes** – Available MBytes measures the amount of physical memory available to processes running on the computer, as a percentage of physical memory installed on the computer. The following guidelines apply when measuring the value of this performance monitor counter:
 - 50% of free memory available or more = Healthy
 - 25% of free memory available = Monitor
 - 10% of free memory available = Warning
 - Less than 5% of free memory available = Critical, performance will be adversely affected
 - **\Memory\Pages/sec** – This performance monitor counter measures the rate at which pages are read from or written to disk to resolve hard page faults. To resolve hard page faults, the operating system must swap the contents of memory to disk, which negatively impacts performance. A high number of pages per second in correlation with low available physical memory may indicate a lack of physical memory. The following guidelines apply when measuring the value of this performance monitor counter:
 - Less than 500 = Healthy
 - 500 - 1000 = Monitor or Caution
 - Greater than 1000 = Critical, performance will be adversely affected

For more information about the impact of available physical memory on application server performance, see the Exchange Server 2003 Help topic “Ruling Out Memory-Bound Problems” at <http://go.microsoft.com/fwlink/?LinkId=121056>.

- **Measure available memory on the guest operating system** – Memory that is available to the guest operating systems can be measured with the same performance monitor counters used to measure memory available to the Hyper-V host operating system.

Measuring Network Performance

Hyper-V allows guest computers to share the same physical network adapter. While this helps to consolidate hardware, take care not to saturate the physical adapter. Use the following methods to ensure the health of the network used by the Hyper-V virtual machines:

- **Test network latency** – Ping each virtual machine to ensure adequate network latency. On local area networks, expect to receive less than 1ms response times.
- **Test for packet loss** – Use the pathping.exe utility to test packet loss between virtual machines. Pathping.exe measures packet loss on the network and is available with all versions of Windows Server since Windows Server 2000. Pathping.exe sends out a burst of 100 ping requests to each network node and calculates how many pings are returned. On local area networks there should be no loss of ping requests from the pathping.exe utility.
- **Test network file transfers** – Copy a 100MB file between virtual machines and measure the length of time required to complete the copy. On a healthy 100Mbit (megabit) network, a 100MB (megabyte) file should copy in 10 to 20 seconds. On a healthy 1Gbit network, a 100MB file should copy in about 3 to 5 seconds. Copy times outside of these parameters are indicative of a network problem. One common cause of poor network transfers occurs when the network adapter has “auto detected” a 10MB half-duplex network which prevents the network adapter from taking full advantage of available bandwidth.
- **Measure network utilization on the Hyper-V host operating system** – Use the following performance monitor counters to measure network utilization on the Hyper-V host operating system:
 - **\Network Interface(*)\Bytes Total/sec** – The percentage of network utilization is calculated by multiplying Bytes Total/sec by 8 to convert it to bits, multiply the result by 100, then divide by the network adapter’s current bandwidth. Use the following thresholds to evaluate network bandwidth utilization:
 - Less than 40% of the interface consumed = Healthy
 - 41%-64% of the interface consumed = Monitor or Caution
 - 65-100% of the interface consumed = Critical, performance will be adversely affected
 - **\Network Interface(*)\Output Queue Length** – The output queue length measures the number of threads waiting on the network adapter. If there are more

than 2 threads waiting on the network adapter, then the network may be a bottleneck. Common causes of this are poor network latency and/or high collision rates on the network. Use the following thresholds to evaluate output queue length:

- 0 = Healthy
- 1-2 = Monitor or Caution
- Greater than 2 = Critical, performance will be adversely affected.

Ensure that the network adapters for all computers (physical and virtual) in the solution are configured to use the same value for maximum transmission unit (MTU). For more information about configuring the MTU value see “Appendix A: TCP/IP Configuration Parameters” at <http://go.microsoft.com/fwlink/?LinkId=113716>.

If an output queue length of 2 or more is measured, consider adding one or more physical network adapters to the physical computer that hosts the virtual machines and bind the network adapters used by the guest operating systems to these physical network adapters.

- **Measure network utilization on the guest operating systems** – If a network adapter on the Hyper-V root partition is busy as indicated by the performance monitor counters mentioned above, then consider using the "\Hyper-V Virtual Network Adapter(*)\Bytes/sec" performance monitor counter to identify which virtual network adapters are consuming the most network utilization.

For more information about network performance analysis, see “Chapter 15 - Measuring .NET Application Performance” at <http://go.microsoft.com/fwlink/?LinkId=121073>.

Measuring Processor Performance

The following considerations apply when evaluating processor performance on a guest operating system installed on a Hyper-V virtual machine:

- **Guest operating system processors do not have a set affinity to physical processors/cores** – The hypervisor determines how physical resources are used. In the case of processor utilization, the hypervisor schedules the guest processor time to physical processor in the form of threads. This means the processor load of virtual machines will be spread across the processors of the physical computer. Furthermore, virtual machines cannot exceed the processor utilization of the configured number of logical processors, for example if a single virtual machine is configured to run with 2 logical processors on a physical computer with 8 processors/cores, then the virtual machine cannot exceed the processor capacity of the number of configured logical processors (in this case 2 processors).

- **Measure guest operating system processor utilization** – Traditionally, processor performance can be measured using the “\Processor(*)\% Processor Time” performance monitor counter. This is not an accurate counter for evaluating processor utilization of a guest operating system though because Hyper-V measures and reports this value relative to the number of processors allocated to the virtual machine. If more processors are allocated to running virtual machines than are actually present on the physical computer, the value returned by each guest operating system for the “\Processor(*)\% Processor Time” performance monitor counter will be low, even if in fact processor utilization is a bottleneck. This occurs because the virtual processors utilize the physical processors in a round-robin fashion. Each virtual processor will try and allocate itself a share of overall system resources, so in a 4 physical processor system, each virtual processor will by default try to utilize 25% of the system resources. If 8 virtual processors are created this means that collectively the virtual processors will attempt to utilize 200% of the server CPU capacity. In this case, each virtual processor will report a low utilization as measured by the “\Processor(*)\% Processor Time” performance monitor counter (relative to the level it expects) and the excessive context switching between the virtual processors will result in poor performance for each virtual machine. In this scenario, consider reducing the number of virtual processors allocated to Hyper-V virtual machines on the host operating system.

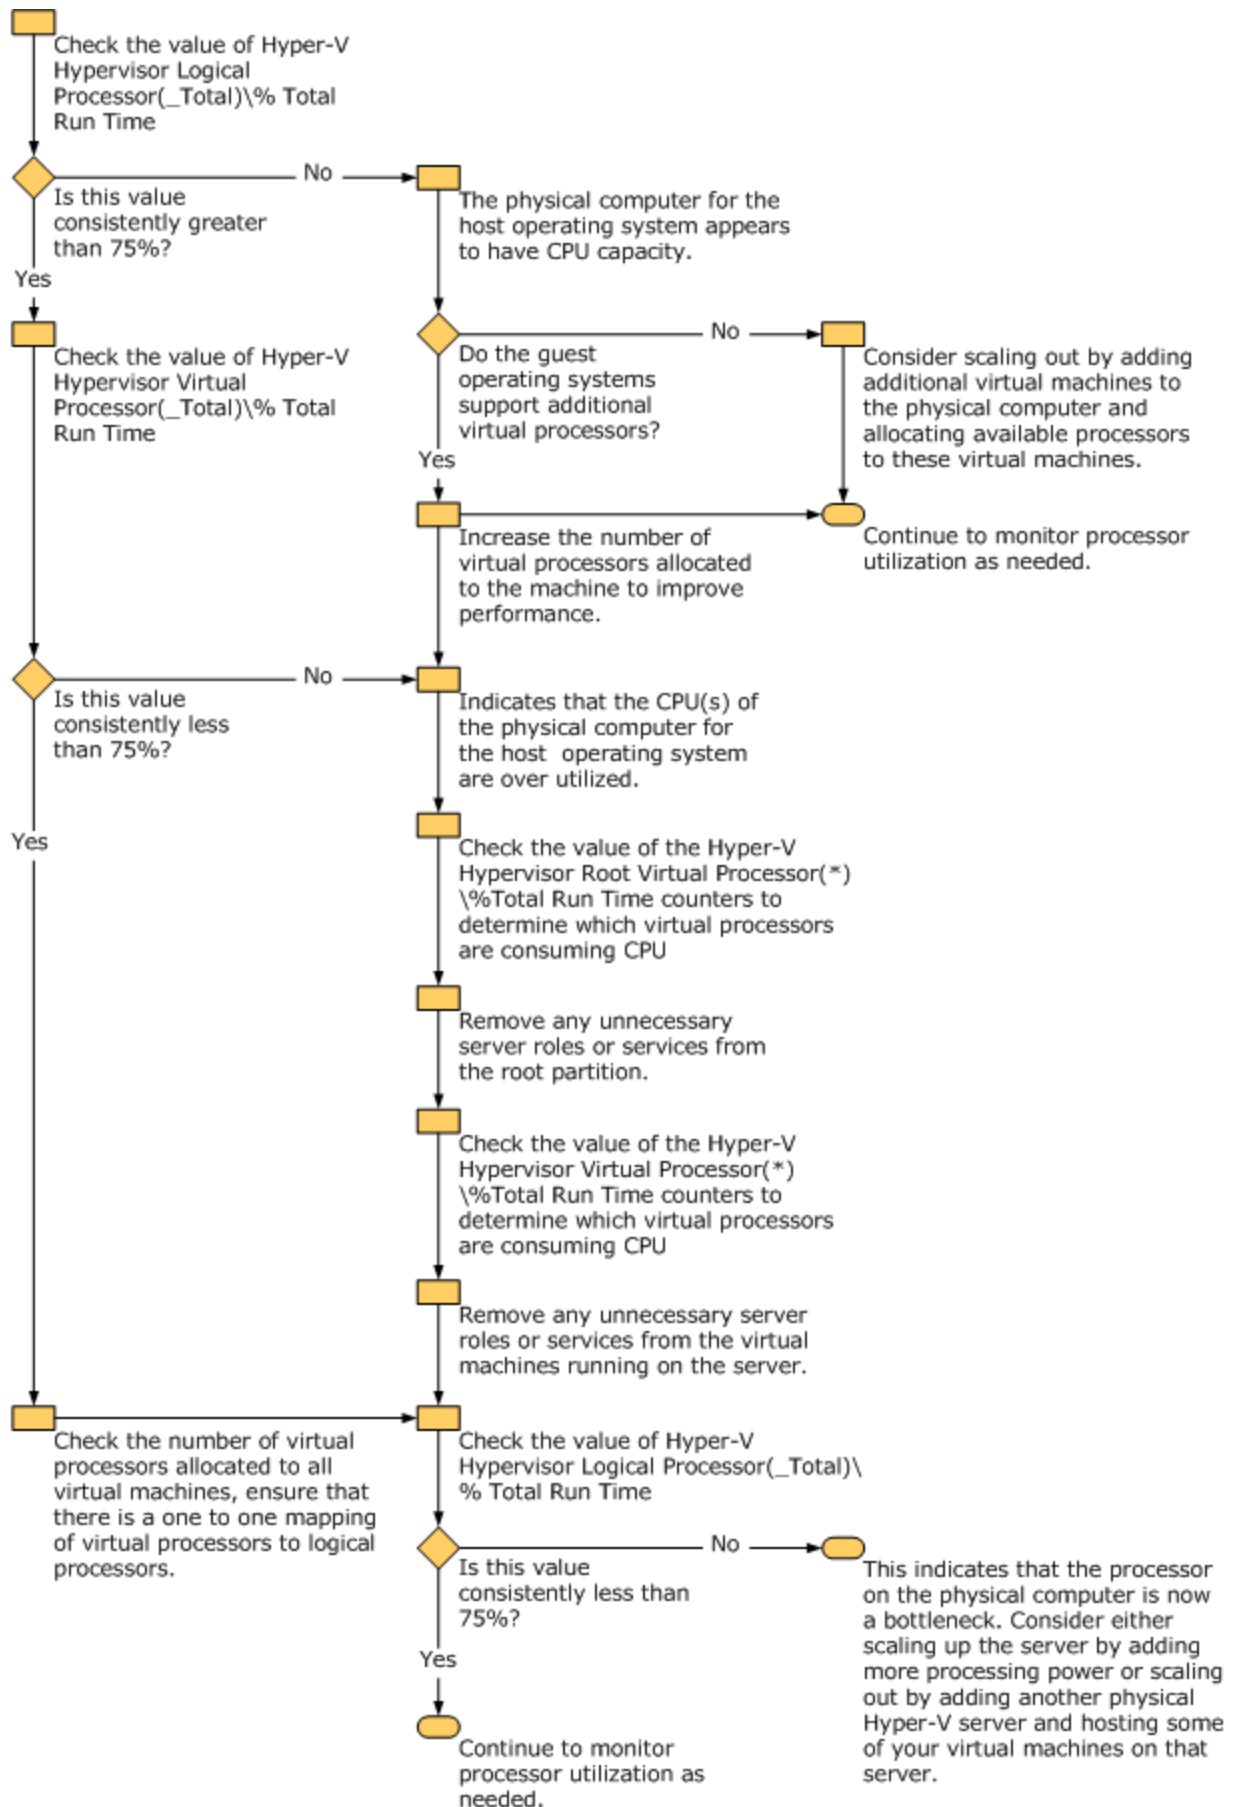
Hyper-V provides hypervisor performance objects to monitor the performance of both logical and virtual processors. A logical processor correlates directly to the number of processors or cores that are installed on the physical computer. For example, 2 quad core processors installed on the physical computer would correlate to 8 logical processors. Virtual processors are what the virtual machines actually use, and all execution in the root and child partitions occurs in virtual processors.

To accurately measure the processor utilization of a guest operating system, use the “\Hyper-V Hypervisor Logical Processor(_Total)\% Total Run Time” performance monitor counter on the Hyper-V host operating system. Use the following thresholds to evaluate guest operating system processor utilization using the “\Hyper-V Hypervisor Logical Processor(_Total)\% Total Run Time” performance monitor counter:

- Less than 60% consumed = Healthy
- 60% - 89% consumed = Monitor or Caution
- 90% - 100% consumed = Critical, performance will be adversely affected

To troubleshoot processor performance of guest operating systems on a Hyper-V environment, it is best to strive for a balance between the values reported by the host operating system for “\Hyper-V Hypervisor Logical Processor(_Total)\% Total Run Time” (LPTR) and “\Hyper-V Hypervisor Virtual Processor(_Total)\% Total Run Time” (VPTR). If LPTR is high and VPTR is low then verify that there are not more processors allocated to virtual machines than are physically available on the physical computer. Use the “\Hyper-V Hypervisor Virtual Processor(*)\% Guest Run Time” counters to determine which virtual processors are consuming CPU and de-allocate virtual processors from virtual machines as appropriate to configure a one to one mapping of virtual processors to logical processors. For more information about configuring a one to one mapping of

virtual processors to logical processors see the “Optimizing Processor Performance” section in [Optimizing Performance on Hyper-V](#). If VPTR is high and LPTR is low then consider allocating additional processors to virtual machines if there are available logical processors and if additional processors are supported by the guest operating system. In the case where VPTR is high, LPTR is low, there are available logical processors to allocate, but additional processors are not supported by the guest operating system, consider scaling out by adding additional virtual machines to the physical computer and allocating available processors to these virtual machines. In the case where both VPTR and LPTR are high, the configuration is pushing the limits of the physical computer and should consider scaling out by adding another physical computer and additional Hyper-V virtual machines to the environment. The flowchart below describes the process that should be used when troubleshooting processor performance in a Hyper-V environment.



- **Measure overall processor utilization of the Hyper-V environment using Hyper-V performance monitor counters** – For purposes of measuring processor utilization, the host operating system is logically viewed as just another guest operating system. Therefore, the “\Processor(*)\% Processor Time” monitor counter measures the processor utilization of the host operating system only. To measure total physical processor utilization of the host operating system and all guest operating systems, use the “\Hyper-V Hypervisor Logical Processor(_Total)\% Total Run Time” performance monitor counter. This counter measures the total percentage of time spent by the processor running the both the host operating system and all guest operating systems. Use the following thresholds to evaluate overall processor utilization of the Hyper-V environment using the “\Hyper-V Hypervisor Logical Processor(_Total)\% Total Run Time” performance monitor counter:
 - Less than 60% consumed = Healthy
 - 60% - 89% consumed = Monitor or Caution
 - 90% - 100% consumed = Critical, performance will be adversely affected