SofTrack®

Enterprise Software Audit and Control Platform

Administrator Supplement

VMWARE AND SOFTRACK SERVER AGENT

INTEGRITY Software

Revised October 10, 2018

VMware and SofTrack Server Agent

If your SofTrack Server Agent is hosted within a VMware (virtualized) environment there are certain configuration considerations.

Please review this section if your SofTrack Server Agent (SSA) has 100 or more SofTrack Local Workstation Agents (LWAs) in regular communication.

If your SSA installation is hosting fewer than 100 LWAs the following considerations may still apply.

Quick Summary of VMware recommended configuration

The following recommendations are based on our research. Further details are provided in the sections that follow.

✔ = new /revised items since last update of this document

- ✔ Use VMXNET3 (***do not use*** E1000 or VMXNET2)
  CRITICAL: Must use v1.8.3.1 or later:
  (VMware Tools 10.3.2 Release Notes) and
  (Download VMware Tools 10.3.2)
- ✔ Enable Multi-Queue Support for NICs on ESXi
- ✔ Enable Receive Side Scaling – requires compatible network adapter(s)
- ✔ Disable virtual interrupt coalescing for VMXNET3
- ✔ Configure all physical network devices to use Full Duplex (adapters and switches)
- ✔ Use auto-negotiation for all physical network devices for determining Ethernet speed
- ✔ Set VMXNET3 at Windows VM hosting SSA *Small RX Buffers* to use maximum (8192)
- ✔ Set VMXNET3 at Windows VM hosting SSA *RX Ring #1* to use maximum (4096)
- ✔ Set VMXNET3 at Windows VM hosting SSA *Receive Buffers* to use maximum (2048)
- ✔ Set Power Plan at Windows VM hosting SSA to *High Performance*
- ✔ Configure all Firewalls and Anti-virus to exclude monitoring of all TCP/IP Port 3884 traffic
- ✔ If Riverbed WAN Accelerators are used, configure each device with a pass-through rule for all inbound and outbound TCP/IP Port 3884 traffic (to prevent caching of Port 3884 traffic)
- ✔ The number of Logical CPUs (or Physical CPUs if not using Hyper-threading) is greater than the number of Virtual CPUs (if CPUs are oversubscribed there can be delays in processing network packets, even if CPU utilization remains low, network activity is time-critical)
- ✔ Do not oversubscribe physical RAM, network activity is time-critical and the SSA host must be fully functional at all times with no delay due to physical memory being oversubscribed.

SSA Memory requirements

The SofTrack Server Agent (SSA) host executable (stsrvc.exe) is a 32bit application and will generally make use of no more than 2GB of RAM. The kernel driver (stview64.sys) is 64bit, the other kernel driver (stview2k.sys) is 32bit – the version of Windows being used determines which kernel driver will be used. The kernel driver portion of the SSA generally uses less than 500MB of RAM – this memory usage is separate from the stsrvc.exe.

SofTrack utilizes "threads" to handle LWA communications, if you notice the number of threads regularly exceeding 1000 on your SSA host (i.e. open Windows Task Manager, add the *Threads* column and view usage of the STSRVC.EXE process – there will be two STSRVC.EXE if email notifications are enabled) this indicates a performance issue. The issue may be related to slow network performance, dropped packets, slow disk storage (high latency) or other application such as a firewall or anti-virus that is slowing inbound and outbound port 3884 packets.

Please note that adding RAM beyond 4GB to the guest VM will not be used by the SSA though it may be useful for other applications and virtual machines hosted on the same physical server.

Starting with SSA version v7.11m, if the SSA ever exhausts its ability to create threads to handle inbound LWA connections an entry similar to the following will be added to the ALERTS.LOG:

```
Jun 6, 2017 (12:15:00): SSA Host Thread Count = 1633
```

The ALERTS.LOG values can be viewed in the SofTrack Console, Settings tab, the values are shown in most recent to oldest.

Network communications requirements

The SofTrack Server Agent (SSA), when in regular communication with 100s or 1000s of SofTrack Local Workstation Agents (LWA) requires "server level" throughput performance regarding network capacity. Use of VMware for the server hosting the SSA requires specific configurations to ensure sufficient network capacity is enabled.

Poor network performance or high network latency on Windows virtual machines (2008925)

https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2008925

Symptoms

A virtual machine with one virtual CPU and a high CPU load, or a virtual machine with two or more virtual CPUs in general and a Windows guest operating system, may experience these symptoms:

- Poor network performance and/or high ping response times:
  - When receiving network traffic (regardless of the amount of data and type)
  - While under high CPU load, or sharing CPU resources with highly-utilized virtual machines

- Observed throughput may decrease to 512kb/s on Gigabit Ethernet. Timeouts and connectivity disruptions may also be observed.
- Ping replies may take up to 20 seconds.
- Sensitive services like database servers may perform poorly or time out.
- The number of virtual and physical network cards has no effect on this issue.
- This issue occurs with different virtual network adapter types (E1000, VMXNET2 and VMXNET3).
- Measured performance results (generated with tools like iperf) may worsen when adding more virtual CPUs to the virtual machine.

Cause

There are two possible causes for this issue:

- Power plan
  - Quick solution: Change the power plan from **Balanced** to **High Performance**

- High CPU ready time (also referred as "%RDY" or "%RDY time")

Checking CPU %RDY times

To determine if a virtual machine is impacted by high CPU %RDY times (value is consistently over 5%), use one of these methods:
Count all virtual CPUs on a particular host or cluster, and divide by the number of logical CPUs. A result of one or higher means that the host or cluster is overcommitted and should be investigated. Values of four or higher are considered overloaded and must be investigated **immediately**.

If the symptoms above are familiar please refer to the link at the top of this section for further details regarding Causes and Resolutions.

Further details regarding ratio of Virtual CPUs versus Logical/Physical CPUs:

Years ago Dell provided the following guidelines:

- 1:1 to 3:1 is no problem
- 3:1 to 5:1 may begin to cause performance degradation

• 6:1 or greater is often going to cause a problem

CPU Ready - From an overall host health standpoint with regard to CPU, this metric is, by far, the most important gauge. CPU Ready is a metric used to determine the length of time that a virtual machine is waiting for enough physical processors to become available in order to meet the demands of the virtual machine. If a virtual machine is allocated four Virtual CPUs (vCPUs), this metric will indicate the length of time that the virtual machine waited for four corresponding Logical/Physical CPUs (pCPUs) to become available at the same time. Further, if Physical CPUs are using Hyper-Threading VMware's label is *Logical CPU*. *Physical CPU* is used for a CPU not using Hyper-Threading.

Additional insights found on the following links:

vCPU and logical CPU sizing with Hyper-Threading explained
http://www.vmwarebits.com/content/vcpu-and-logical-cpu-sizing-hyper-threading-explained

Best Practices for Oversubscription of CPU, Memory and Storage in vSphere Virtual Environments: https://communities.vmware.com/servlet/JiveServlet/previewBody/21181-102-1-28328/vsphee-oversubscription-best-practices[1].pdf

In most engagements the SSA and LWAs are communicating over a physical network involving physical network adapters, switches and cables and/or wi-fi. If your environment matches this description it is required that the physical host(s) for the SSA are using server class network adapters that are on VMware's Hardware Compatibility List (please note: VMware typically will not install if an unapproved network adapter is used). And, if using 10Gb Ethernet it is presumed all your physical switches are able to provide that level of performance otherwise your network communications throughput will be reduced to the lowest performing device.

Configuring the speed and duplex of an ESXi or ESX host network adapter (1004089)
https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1004089 ESXi/ESX recommended settings for Gigabit-Ethernet speed and duplex while connecting to a physical switch port: Auto Negotiate <-> Auto Negotiate (For 1 Gbps) Auto Negotiate <-> Auto Negotiate (For 10 Gbps, supported only by ESXi/ESX 3.5 Update 1 and above) Note: Many drivers do not support forced 1000Mbps or 10000Mbps speeds, and require the auto-negotiation set for this to work correctly. Auto-negotiation is considered as the normal and official way that both Gigabit and 10-Gigabit networking is designed to function. For more information, see the IEEE 802.3ab, 802.3an, and 802.3ae standards. Many drivers do not allow forced 1000Mbps or 10000Mbps because it is not officially supported by the IEEE standards.
When working with 10 GB Fiber Channel over Ethernet (FCoE) configurations, Auto Negotiate may or may not be supported or recommended. For more information, consult your networking equipment vendor or administrator. 1000 MB / Full Duplex <-> 1000 MB / Full Duplex
VMware does not recommend mixing a hard-coded setting with Auto-negotiate Fast Ethernet – 100 MB /Full Duplex <-> 100 MB /Full Duplex For more information on Virtual Switch properties, see the Editing Virtual Switch properties section in the vSphere Networking Guide: http://pubs.vmware.com/vsphere-50/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-501-networking-guide.pdf

SofTrack's data transfers are typically small, under 20KB and most individual packets are under the size of single Ethernet Frame (1522 bytes). SofTrack uses TCP/IP streams for communication, which requires set up, and tear down sequences, a single transfer usually requires 9 packets and total transfer size is under 2500 bytes. There are larger transfers, for instance, quick inventory data and workstation audit records. The latter are transferred once per minute and depending on audit options (file activity and browser activity) the once-per-minute transfers can exceed 100KB. Logon / logoff activity is transferred on demand as the events occur and typically less than 2000 bytes total transfer size.

Further, most SSA reply buffers are small, however, when an LWA first connects, and if SofTrack's *Offline Mode* is configured, a larger transfer can occur from the SSA to the LWA, up to 100KB.

Predominately data is transferred "upstream" from the LWAs to the SSA, so the SSA host must be able to handle the potential load (samples below) the VMware configuration must enable the SSA host VM

to sufficiently participate in the bandwidth availability of the physical network adapter(s), and depending on your VMware configuration you may need to dedicate physical network adaptors to the SSA host to ensure it has the bandwidth it needs.

Sample SofTrack Packet dimensions
Once per minute the LWA will poll the SSA host: 1265 bytes, total once a minute poll {9 packets} If 800 workstations LWAs are communicating with the SSA host that would be about 1MB of traffic per minute, 7,200 packets
Each application launch monitored by the LWA is transferred to the SSA to request metering instructions. 2095 bytes (approx.) for each launch {9 packets} If 800 workstations, presuming a busy PC that can have 10 or more application launches per minute would be about 17MB of traffic per minute, 72,000 packets
Workstation Audit logging of file activity and/or browser activity could average 50KB per minute, if 800 workstations are engaged in workstation auditing, that would be approx. 40MB of traffic each minute, about 50,000 packets
Quick Inventory data is sent as configured, default is each time the LWA first starts and average data size transferred is 100KB, range typically is from 30KB to 500KB, if 800 workstations rebooted at the same time that would result in approx. 80MB of traffic, about 500,000 packets over a period of a few minutes

The amount of traffic indicated above for 800 workstations is not large, especially when considering a SSA host using dual 10Gb physical adapters. The number of small packets can be more significant and an insufficiently provided SSA host can become overwhelmed – further detailed in the Recommended section (next).

However, VMware has many abilities and depending on configuration, the physical bandwidth can be otherwise dedicated and consumed.

Because VMware allows sharing of network bandwidth, the virtual machine hosting the SSA must be provided with sufficient capacity to reduce the possibility of LWAs encountering communication failures to the SSA host.

Several VMware features can require significant amounts of network communications bandwidth including:

- vMotion
  - This traffic consists of live migration of active virtual machines from one physical server to another with zero downtime. For each 10Gb network adapter (link) up to 8 concurrent vMotion instances can be configured. For each 1Gb link up to 4 vMotion instances are allowed. It is important to ensure the physical devices (servers) used are of equal or superior capabilities.
- Management
  - This traffic flows through the vmknic and is typically not a high load.
- Fault Tolerant

    - This traffic flows through the vmknic and can have significant bandwidth requirements as it must replicate the I/O traffic and memory state information to a secondary virtual machine (which can be located on a separate physical server).
- Virtual Machine traffic
    - If using guest virtual machines (virtual networks through other virtual machines, i.e. using the vmnet adapter)
- IP based storage
    - ISCSI/NFS traffic flows through the vmknic and can have very significant bandwidth requirements; it is recommended to use end-to-end jumbo frames (i.e. Ethernet frame configuration) to improve IP storage performance.

Recommended VMware configuration

The number one issue we have found for VMware hosting the SSA is large delays initially accepting packet connections from the LWAs. The SSA will automatically wait up to 34 seconds from when an inbound connection request is received until the network stack allows the connection's acceptance. In a properly functioning SSA host there should be no delay between the arrival of a new connection and its acceptance. On an insufficiently configured VMware system we have noticed the transfer of a 150KB inventory file taking more than 20 minutes to occur due to continual dropped packets (average transfer time should be less than 20 seconds).

We recommend you review VMware knowledge base article 2039495:

Large packet loss at the guest operating system level on VMXNET3 vNIC in ESXi (2039495)
https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2039495

The pertinent details:

To reduce burst traffic or steady high flow traffic packet drops in Windows Server:

Establish the following Buffer Settings:

Start > **Control Panel** > **Device Manager**

Right-click **vmxnet3** and click **Properties**.

Click the **Advanced** tab.

Click **Small Rx Buffers** and increase the value. The default value is 512 and the maximum is 8192.

SSA Host recommendation is **8192** (non-paged kernel memory required: 12MB)

Click **Rx Ring #1** Size and increase the value. The default value is 1024 and the maximum is 4096.

SSA Host recommendation is **4096** (non-paged kernel memory required: 6MB)

Notes:
These changes will happen on the fly, so no reboot is required. However, any application sensitive to TCP session disruption can likely fail and have to be restarted. This applies to RDP, so it is better to make these changes in a console window and not remotely since your session may be dropped.

This issue is seen in the Windows guest operating system with a VMXNET3 vNIC. It can occur with versions besides 2008 R2 (and with any VM adapter, including **E1000** and **VMXNET2**).

The VMware knowledge base includes this comment:
It is important to increase the value of Small Rx Buffers and Rx Ring #1 gradually to avoid drastically increasing the memory overhead on the Windows server host and possibly causing performance issues if resources are close to capacity.
However, we disagree, we suggest setting values to the maximum allowed, memory *usage is insignificant.*

Microsoft provides this instruction:
Increasing Network Adapter Resources
For network adapters that allow manual configuration of resources, such as receive and send buffers, you should increase the allocated resources. Some network adapters set their receive buffers low to conserve allocated memory from the host. The low value results in dropped packets and decreased performance. Therefore, for ***receive-intensive scenarios (such as the SofTrack Server Agent)***, we recommend that you increase the *receive buffer* value to the maximum.

VMware knowledge base article continues:
If this issue occurs on only 2-3 virtual machines, set the value of Small Rx Buffers and Rx Ring #1 to the maximum value. Monitor virtual machine performance to see if this resolves the issue.

The Small Rx Buffers and Rx Ring #1 variables affect non-jumbo frame traffic only on the adapter.

The following write up is from a ServerFault user response and provides additional insight:

https://serverfault.com/questions/711693/vmxnet3-receive-buffer-sizing-and-memory-usage/853304

VMXNET3 receive buffer sizing and memory usage

What is the relationship between number of buffers and ring size?

They're related, but independent. The rx "ring" refers to a set of buffers in memory that are used as a queue to pass incoming network packets from the host (hypervisor) to the guest (Windows VM). The memory gets reserved in the guest by the network driver, and it gets mapped into host memory.

As new network packets come in on the host, they get put on the next available buffer in the ring. Then, the host triggers an IRQ in the guest, to which the guest driver responds by taking he packet off the ring, and dispatching it to the network stack of the guest OS, which presumably sends it to the guest

application intending to receive it. Assuming the packets are coming in slow enough, and the guest driver is processing them fast enough, there should always be a free slot in the ring. However, if packets are coming in too fast, or the guest is processing them too slowly, the ring can become full, and packets may be dropped.

Increasing the ring size can help mitigate this issue. If you increase it, more slots will be available in the ring at a time. This segues into the second setting, "Small Rx Buffers", which is the total amount of buffers available that can be used to fill the slots in the ring. There needs to be at least as many buffers as slots in the ring. Typically you want more. When the guest takes a buffer off the ring to give to the guest network stack, it may not always be immediately returned back to the driver. If that happens, having spare buffers to fill the ring means you can go longer without dropping packets.

The Rx Ring #1 / Small Rx Buffers are used for non-jumbo frames. If you have a default NIC configuration, that's the only ring that will be used.

How does one calculate the amount of memory used for given values of these settings?

Assuming you're talking about non-jumbo frames, each buffer needs to be big enough to store an entire network packet, roughly 1.5kb. So if you have 8192 buffers available, that would use 12MB. A larger ring will also use more memory, but the descriptors are small (bytes), so it's really the buffers you have to worry about.

Because these settings are on the NIC itself within the guest OS, I assume they are driver settings. This makes me think that the RAM used might be paged or non-paged pool.

Yes, it's a non-paged pool. If the ring buffers were paged, it would likely result in dropped packets while the buffers were being paged back in.

Are there concerns I'm not taking into account here?
I'm not sure this is relevant to your situation, but it might be worth noting that a larger ring will increase the cache footprint of the network rx path. In microbenchmarks, you will see that a larger ring usually hurts performance. That said, in real life applications, ***if a packet gets dropped***, that's usually a bigger deal than a small performance gain in speed bursts.

The above is text is from a write up from a ServerFault where a user responsed and provides additional insight, above text from this link:

https://serverfault.com/questions/711693/vmxnet3-receive-buffer-sizing-and-memory-usage/853304

Additional useful links:

https://www.vmguru.com/2015/12/vsphere-6-experiencing-high-packet-loss/
(includes thorough details of how to determine packet loss counts)

https://www.reddit.com/r/vmware/comments/52tdt1/vmxnet3_packet_loss_despite_rx_ring_tuning/

https://www.null-byte.org/vmware/random-packet-loss-with-vmware-esxi-5-1-virtual-machines-using-vcni/

https://helpdesk.flexradio.com/hc/en-us/articles/202118518-Optimizing-Ethernet-Adapter-Settings-for-Maximum-Performance

Microsoft recommends:

https://technet.microsoft.com/en-us/library/jj574151(v=ws.11).aspx

Increasing Network Adapter Resources

For network adapters that allow manual configuration of resources, such as receive and send buffers, you should increase the allocated resources. Some network adapters set their receive buffers low to conserve allocated memory from the host. The low value results in dropped packets and decreased performance. Therefore, for ***receive-intensive scenarios (such as the SofTrack Server Agent)***, we recommend that you increase the **receive buffer** value to the maximum.

To change:
Start > **Control Panel** > **Device Manager**

Right-click **vmxnet3** and click **Properties**.

Click the **Advanced** tab.

Locate the "**Receive Buffers**" value and increase it to the maximum. Tip: To quickly set the maximum value: manually enter the value 9999 and then click the "down arrow" and then click the "up arrow"

Windows Filtering Platform

If any applications such as Firewalls and Anti-virus are installed, each may be utilizing the Windows Filtering Platform. If any such applications are being used on the SSA host server, each can be a source of packet loss by impeding the flow of packets to the SSA due to packet inspection. If configurable, we suggest such applications be set to ignore all inbound and outbound port 3884 (decimal number) TCP/IP traffic.

VMware Solutions for Poor Network Performance

The following is an excerpt from this link:

https://pubs.vmware.com/vsphere-50/index.jsp?topic=%2Fcom.vmware.vsphere.monitoring.doc_50%2FGUID-C4EA86FD-95CB-4DE7-A9E3-63F6BFC1A268.html

Network performance is dependent on application workload and network configuration. Dropped network packets indicate a bottleneck in the network. Slow network performance can be a sign of load-balancing problems.

Problem

Network problems can manifest in many ways:

- Packets are being dropped.
- Network latency is high
- Data receive rate is low.

Cause

Network problems can have several causes:

- Virtual machine network resource shares are too few.
- Network packet size is too large, which results in high network latency. Use the VMware AppSpeed performance monitoring application or a third-party application to check network latency.
- Network packet size is too small, which increases the demand for the CPU resources needed for processing each packet. Host CPU, or possibly virtual machine CPU, resources are not enough to handle the load.

Solution

- Determine whether packets are being dropped by using **esxtop** or the advanced performance charts to examine the **droppedTx** and **droppedRx** network counter values. Verify that VMware Tools is installed on each virtual machine.
- Check the number of virtual machines assigned to each physical NIC. If necessary, perform load balancing by moving virtual machines to different vSwitches or by adding more NICs to the host. You can also move virtual machines to another host or increase the host CPU or virtual machine CPU.
- If possible, use **vmxnet3** NIC drivers, which are available with VMware Tools. They are optimized for high performance.
- If virtual machines running on the same host communicate with each other, connect them to the same vSwitch to avoid the cost of transferring packets over the physical network.
- Assign each physical NIC to a port group and a vSwitch.
- Use separate physical NICs to handle the different traffic streams, such as network packets generated by virtual machines, iSCSI protocols, VMotion tasks.
- Ensure that the physical NIC capacity is large enough to handle the network traffic on that vSwitch. If the capacity is not enough, consider using a high-bandwidth physical NIC (10Gbps) or moving some virtual machines to a vSwitch with a lighter load or to a new vSwitch.
- If packets are being dropped at the vSwitch port, increase the virtual network driver ring buffers where applicable.
- Verify that the reported speed and duplex settings for the physical NIC match the hardware expectations and that the hardware is configured to run at its maximum capability. For example, verify that NICs with 1Gbps are not reset to 100Mbps because they are connected to an older switch.
- Verify that all NICs are running in full duplex mode. Hardware connectivity issues might result in a NIC resetting itself to a lower speed or half duplex mode.
- Use vNICs that are TSO-capable, and verify that TSO-Jumbo Frames are enabled where possible.
- Ensure the Storage Device used by the SSA to store workstation (quick) inventory and

workstation audit records is not introducing its own latency

Enable Multi-Queue Support for NICs on ESXi

The following copied from this link. Additional relevant details on this link.

Multi-queue allows network performance to scale with the number of vCPUs and allows for parallel packet processing by creating multiple TX and RX queues.

Modify the .vmx file or access Advanced Settings to enable multi-queue.

Step 1: Enable multi-queue.

* Open the .vmx file.
* Add the following parameter:

```
ethernetX.pnicFeatures = "4"
```

Step 2: Enable receive-side scaling (RSS).

* Log in to the CLI on the ESXi host.
* Execute the following command:

```
$ vmkload_mod -u ixgbe
$ vmkload_mod ixgbe RSS="4,4,4,4,4,4"
```

Step 3: For the best performance, allocate additional CPU threads per ethernet/vSwitch device. This is limited by the amount of spare CPU resources available on the ESXi host.

* Open the .vmx file.
* Add the following parameter:

```
ethernetX.ctxPerDev = "1"
```

Enable Multi-Queue Receive Side Scaling in Windows

From Windows CMD prompt enter the following command:

```
netsh interface tcp set global rss=enabled
```

Or open **devmgmt.msc** and edit:

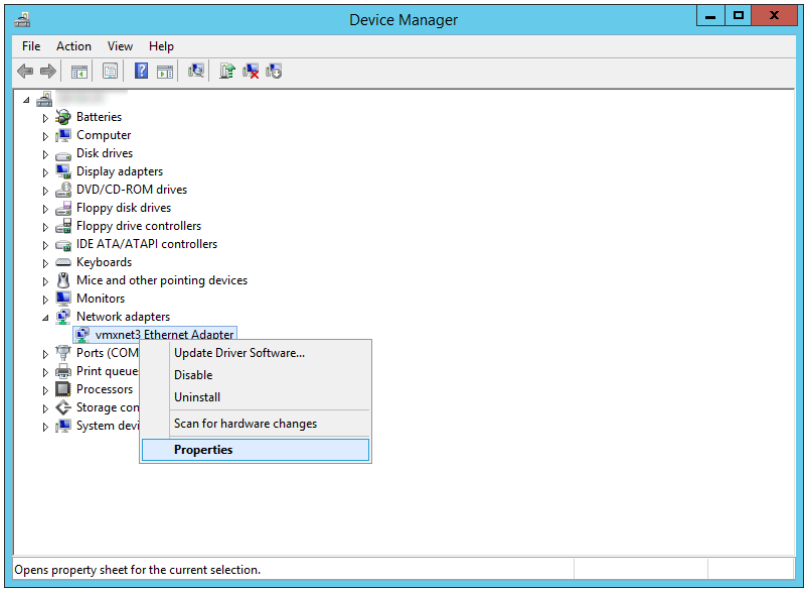

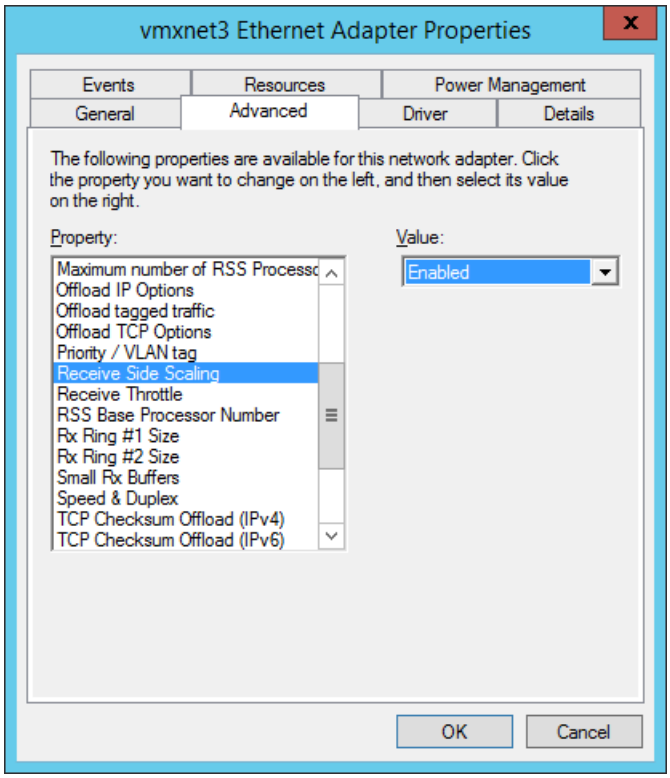

Click here for further details Windows Server 2008

Click here for further details Windows Server 2012 / 2012R2

Click here for further details Windows Server 2016

Click here for additional Windows Server 2016 details

Revised October 10, 2018

Disable virtual interrupt coalescing for VMXNET3

Via the vSphere Web Client, go to:

VM Settings => Options tab => Advanced General => Configuration Parameters

and add an entry for

`ethernetX.coalescingScheme`

with the value of: **`disabled`**.

An alternative way to disable virtual interrupt coalescing for all virtual NICs on the host which affects all VMs, not just the latency-sensitive ones, is by setting the advanced networking performance option:

Configuration => Advanced Settings => Net

`CoalesceDefaultOn to 0` (disabled).

Above from this link

Additional VMWARE documents regarding high performance configurations

https://docs.vmware.com/en/VMware-vCloud-NFV/2.0/vmware-tuning-vcloud-nfv-for-data-plane-intensive-workloads.pdf

https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/techpaper/vmware-tuning-telco-nfv-workloads-vsphere-white-paper.pdf

Poor network performance or high network latency on Windows virtual machines (2008925):
https://kb.vmware.com/s/article/2008925

Solutions for Poor Network Performance:
https://pubs.vmware.com/vsphere-50/index.jsp?topic=%2Fcom.vmware.vsphere.monitoring.doc_50%2FGUID-C4EA86FD-95CB-4DE7-A9E3-63F6BFC1A268.html

https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/techpaper/latency-sensitive-perf-vsphere55-white-paper.pdf