

# Monitoring Databases on VMware

Ensure Optimum Performance with the Correct Metrics

By Dean Richards, Manager, Sales Engineering

---

Confio Software  
4772 Walnut Street, Suite 100  
Boulder, CO 80301

[www.confio.com](http://www.confio.com)

## Introduction

Databases are one of the last frontiers for virtualizing into a VMware environment. They can be very resource intensive so many companies have been hesitant to move them into a virtual environment. Most of the hesitation has come from poor experiences during initial tests of VMware. Many of our customers are successfully running databases inside VMware and the key for them was monitoring the correct metrics that can help show where the bottlenecks are. This paper will review the VMware metrics that helped many companies have a far better virtualization experience.

## Database support on VMware

Many companies are also hesitant to move databases to VMware because of potential support issues. Typically most database vendors do not officially support or certify the database running on VMware. If there are issues, the vendor may ask for the problem to be reproduced on physical hardware, which can be problematic at times. However, there are no documented bugs in the vendor support sites related to their database running on VMware. Also, many companies are already running a database on VMware, even mission critical applications, which is another example that is seems acceptable for many environments.

## Database licensing on VMware

This paper will not discuss database licensing in detail, but when virtualizing databases in a VMware environment, understand that for many platforms, license fees can be charged for each CPU/Core the database could run on. If your VM running the database is operating within a cluster of 4 hosts that each have 8 cores, Oracle, Microsoft and other vendors may charge you for  $4 * 8$  cores or 32 cores. This is something to be aware of and talk to your vendor representatives for full details. For this licensing reason, many companies are configuring dedicated clusters and hosts for databases so licenses are not required for physical hosts not running the database.

## Terms and concepts

Before getting into the details of which metrics to use for monitoring, the following are some key VMware terms and concepts.

- ESX / ESXi – the hypervisor and foundation for virtualization. ESX is installed into a Linux O/S while ESXi takes the place of the O/S and is the only option for VMware 5.0 and higher.
- Physical Host – underlying hardware that ESX is installed.
- Virtual Machine (VM) – container inside the physical host that runs a guest O/S like Windows or Linux. This, for most intents, acts and looks like a normal physical machine. This is where the database software is installed and managed.
- vCenter Server – centralized management of VMware environment. Performance information is also collected and stored here.

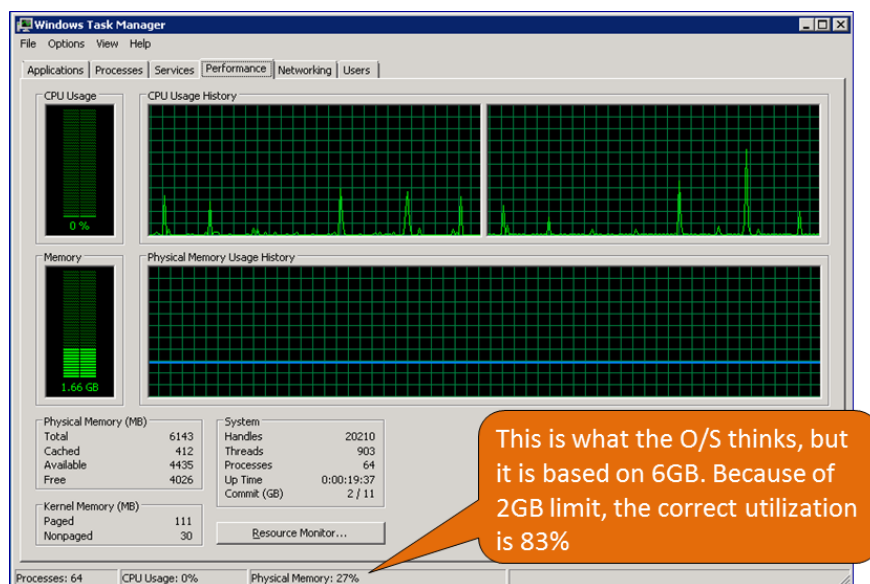
- vSphere Client – the GUI for connecting to vCenter and visualizing metrics via charts. The GUI is also used for managing the VMware environment.
- Cluster – made up of two or more physical hosts that share common storage.
- Distributed Resource Scheduler (DRS) – monitors each physical host’s resource consumption. Within a cluster, DRS tracks many resources and can move a VM (vMotion) to another physical host to keep resource consumption balanced.
- vMotion – live migration of a VM from one physical host to another. This can either be manual or automated so that resource consumption is balanced across all physical hosts in the cluster.

## vSphere GUI

For monitoring a database running within a VM, the vSphere GUI, or other tool like DPA with VM option from SolarWinds, will be the primary interface. There is both a Windows client (downloaded to your desktop) and a web client for accessing vSphere information. However, the web version is limited in functionality so the Windows GUI will most likely be the preferred option. To download the GUI, open a browser and point to <http://<vCenterServer>> (your VMware administrator can help with the name of this machine) and the download links will be provided.

## Guest O/S statistics

DBAs will often wonder or question why they need to use the vSphere GUI. The feeling is that O/S statistics being used in the non-virtualized world are still valid in the VMware world. However, this is not the case and is documented by VMware and O/S vendors alike. O/S metrics such as CPU Utilization can be incorrect due to the timing nature of VMware. Plus, if only the guest O/S metrics are used, the other layers (physical host, storage, network) will be missed which is crucial. Resource issues at the physical host or storage layer can severely limit the performance of a VM and this will not show up in the guest O/S metrics. Below is an example when guest Windows O/S statistics do not provide complete information (similar issues arise with Linux and other guest O/S as well):



In this example, the VM was configured with 6 GB RAM, but limited to 2 GB. The Windows Task Manager does not see the VMware limit put on the VM so it thinks it has 6 GB to operate with and hence shows a memory utilization of 27%. However, since the VM is limited to 2 GB, the real utilization of 83% which is a much different reality.

## Monitoring with vSphere

Now that some groundwork has been laid, the rest of this paper will discuss each of the four main resource areas of a VMware environment: Memory, CPU, Storage and Network. The vSphere GUI has the capabilities of showing 100's of metrics and provides no relative importance among them, so some DBAs have complained that it is difficult to find the root cause of an issue. Because of that common problem, this paper will provide the metrics the author feels are most important when monitoring databases running within a VM.

## VMware memory concepts

Databases can be very memory intensive applications and monitoring this resource effectively is required. Key concepts related to memory in VMware are discussed below:

- Configured Memory – amount of RAM given to the VM, similar to the amount of RAM configured on a desktop or other physical machines.
- Reservation – amount of RAM guaranteed to the VM. This amount of memory must be available on the physical host when a VM boots. The default is 0 meaning no physical RAM is given until needed, i.e. when the VM is booting the O/S, starting the database instance and when other applications launch. A good rule of thumb for a VM is to set the memory reservation greater than the size of the database memory required with room available for the guest O/S and any other applications.
- Limit – ceiling on the amount of RAM a VM can utilize. The default is unlimited up to the configured amount which is how most VMs running databases will be configured.
- Shares – the priority for a VM getting to the memory. If two VMs request memory at the same time, and memory is readily available on the physical host, those VMs would both get the memory as requested. However, if the physical host that contains those VMs is memory constrained at the time, the VM with a higher number of shares would get priority for the memory. Often, database VMs have a higher number of shares than other VMs.
- Ballooning – when a physical host is memory constrained, memory already given to a VM may be borrowed from unused memory on another VM. This borrowed memory is taken from unused or least recently used memory.
- Swapping – when a physical host is severely constrained and Ballooning is not sufficient, a portion of the VM's memory may be swapped out to give to another VM. This swapping process is not discriminate about the memory taken so memory currently being used might be swapped and written to disk.

- Shared Memory – VMware will recognize identical pages of memory across VMs. For example, if there are 10 VMs running the same version of the database vendor binaries, the memory used to store that information will most likely be identical for each and hence only stored once. All VMs will point to this shared physical memory.

## Monitoring memory

When monitoring memory (and other metrics discussed in this paper) within a VMware environment, the vSphere GUI can be used to see memory at both the VM and the physical host layers.

### VM memory metrics

Below is a view showing the resource allocation for both Memory and CPU within a VM.

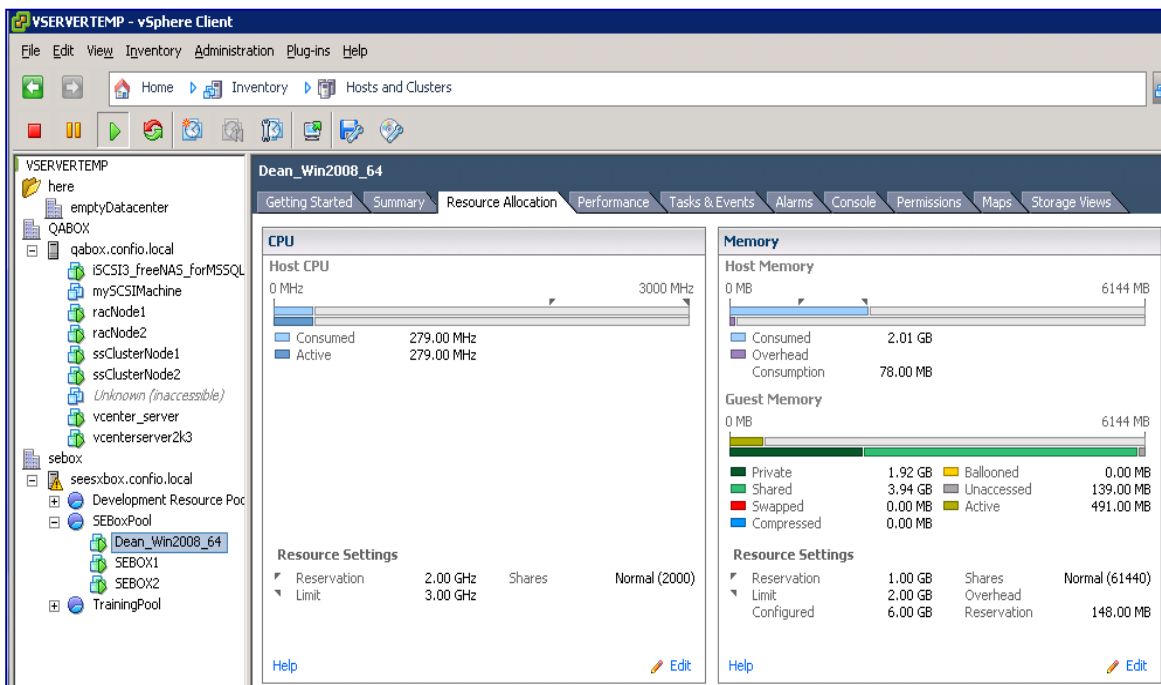


Figure 1. Resource Allocation View in vSphere

Shown on the Resource Allocation tab for a VM are the key concepts discussed in the above section. Values are shown for static settings like Configured Memory, Limit, Reservation as well as dynamic values for memory consumed, Ballooned, Swapped, etc. In the example above, this VM is configured with 6 GB RAM, has a reservation of 1 GB, a limit of 2 GB. It is currently consuming 2 GB of the RAM and actively using 491 MB (there is approximately 1.5 GB available for ballooning).

Another way to view memory metrics is to use charting within the vSphere GUI. Below is a multi-metric chart for a VM that shows the key metrics of Balloon, Swap Out, Usage and Active. Other metrics may be important but these are four of the more important VMware metrics for a VM running a database.

Note: use the Performance tab to create this chart by clicking on Advanced, then Chart Options and selecting the memory metrics of interest.



Figure 2. Key Memory Metrics for a VM

## Physical host memory metrics

Similar views are also available for the physical host that contains the VM. The primary metric for the physical host is total memory usage values.

## Monitoring memory rules

When reviewing the memory metrics shown above, here are a few simple rules to follow in order listed:

1. If swapping is occurring at the VM level, this indicates the physical host is under severe memory pressure. When a VM is swapping, applications appear to hang and other severe performance degradation is most likely occurring. Methods to fix this issue include:
  - a. Review the memory usage for the physical host. If it regularly exceeds 90%, more memory is required for the physical host to accommodate the current resource consumption.
  - b. Also review the reservations for all VMs to see if excessive memory has been given to a VM that does not need it. Lowering reservations may give back memory to the physical host.
  - c. Review workloads for each VM and tune to reduce memory consumption. For example, lower the database instance memory settings for less critical databases that may not need the amount of memory currently consumed.
  - d. vMotion VMs to another physical hosts to reduce memory constraints on the current physical host.

2. If ballooning is occurring, it also indicates the physical host is under memory pressure, but not quite to the level that swapping would indicate. Memory for the physical host is in short supply so perform similar steps above for swapping issues. The performance degradation may not be noticeable at this point, but severe issues may be on the horizon.
3. If VM Memory Utilization is high, this could indicate a pending issue. If the physical host has spare memory, consider increasing the configured amount or limits if being used. Also, review database specific memory settings to see if the consumption can be decreased.
4. If Host Memory Utilization is high, again this may indicate pending problems. Try to reduce the memory consumption of all VMs by lowering reservations and tuning the memory consumption of the applications on each VM. Also consider moving memory intensive VMs to another physical host.

## VMware CPU concepts

Databases tend not to be CPU intensive applications, however monitoring this resource is still very important. Some key concepts related to CPU consumption in VMware are discussed below:

- Configured CPU – the number of virtual CPUs (vCPU) allocated to the VM which is very similar to the number of CPU cores present in a physical machine.
- Reservation – guaranteed amount of CPU given to the VM.
- Limit – ceiling on the amount of CPU the VM has access to.
- Shares – the priority for a VM getting CPU resources. If two VMs request CPU at the same time, and CPU is readily available on the physical host, those VMs would both get the CPU as requested. However, if the physical host that contains those VMs is CPU constrained at the time, the VM with the higher number of shares would get priority.
- Clock Speed – consider thinking of CPU resources in terms of MHz or GHz, e.g. if the VM has been given 4 vCPU and the physical host CPUs are running at 3 GHz, translate that to a total of 12,000 MHz (3,000 MHz \* 4 cores). Doing this will help in understanding the settings for Reservation and Limit since they are specified in MHz.

## CPU sizing

Databases are not typically CPU bound so only configure a virtual machine with the number of vCPUs required. This is especially true when virtualizing a database from older hardware onto a newer hardware within the VMware environment. If the database instance is currently running in an environment with 8 cores, the wrong thing to do is request 8 cores or vCPU on the new VM if the clock speed is significantly faster. A simple conversion can be done with clock speeds, e.g. if the old hardware is running 8 cores of 1 GHz processors, the current environment is utilizing 8,000 MHz (8 cores \* 1,000 MHz) of total capacity. If the current CPU utilization peaks (use peaks and not averages) at 50%, the new environment will need to handle 4,000 MHz at the peaks. If the new hardware runs at 3 GHz, the new VM may only need 2 vCPUs to handle the load. Even if some headroom is required, 4 vCPU would be more than sufficient instead of 8.

## Monitoring CPU

### VM CPU metrics

Figure 1 above shows that the Resource Allocation tab in vSphere can also be used to view CPU configuration. Values for Configured CPU, Limit, Reservation as well as dynamic values for consumed CPU. In the example above, this VM is configured with 3,000 MHz, has a reservation of 2,000 MHz and is unlimited. It is currently consuming 496 MHz.

Just as with memory, another way to view memory metrics is to use charting within the vSphere GUI. Below is a multi-metric chart for a VM that specifically shows the key metrics of Usage and Ready Time. Other metrics may be important but these are four of the more important CPU metrics for a VM running a database.



Figure 3. Key CPU Metrics for a VM

### Physical host CPU metrics

On the physical host the CPU utilization metric is also an important metric to review.

### Monitoring CPU rules

When reviewing the CPU metrics shown above, there are a few simple rules to follow:

1. Ready time is a key measurement for a VM. It is very similar to a wait event or wait type in the database world, e.g. Ready Time signifies the VM is waiting for CPU resources. If ready time is above 10-20%, review the physical host CPU utilization. If that is high, the physical host is running out of CPU resources and more are needed. It may also be possible to reduce CPU resource consumption on all VMs by tuning or moving a heavy CPU user to another physical host.



2. If ready time is high but the CPU utilization on the physical host is not high, then the VM is being limited in some fashion. This may be caused by having too few vCPU so consider increasing the number available. There could also be a hard limit imposed reducing the available MHz available. For example, a VM may have 4 vCPU running at 3GHz for a total of 12,000 MHz, but there is a limit of 6,000 MHz limiting the CPU resources.
3. If VM CPU Utilization is high this may not be an issue now, but could be a pending problem.
4. If the Host CPU Utilization is high, this may not be an issue now, but could be a pending problem. If only one Host has high CPU utilization in a cluster, consider balancing the load across the other hosts. Running VMware's DRS in fully automated mode is a great way to automatically keep resources balanced across all hosts in a cluster.

## Monitoring storage

Storage for databases has been virtualized onto the SAN for quite some time and there is not much different in the VMware world. However, there are a few concepts and terms to be aware of:

- VMFS – Virtual Machine File System is a high performance cluster file system which has been designed and optimized for the virtualized environment.
- NAS – Network Attached Storage or NFS devices. All I/O is passed over the network directly to the storage device.
- RDM – Raw Device Mapping allows a special file in the VMFS file system to essentially point to a raw disk device. All I/O activity is passed directly onto the device. Many companies use Oracle ASM or SQL Server clustering on top of RDM devices in the VMware environment. If so, RDM device metrics are not available and monitoring must be done within ASM or at the raw device level.
- Datastore – a storage location for VM files that can be a VMFS volume, a directory on Network Attached Storage (NAS) or a local system path. Datastores scope is at the datacenter level and therefore does not change when VMs move between physical hosts. Datastores are configured per physical host, thus managing them is done at both levels. Each host is explicitly configured for the datastores it can access.
- Thin Provisioning – allows the VMware administrator to dedicate more storage than what is utilized. For example, if a 100 GB datastore is created with thin provisioning, 100 GB is not allocated up front. When adding or extending a database file on a device, space is allocated only as needed. This option can dramatically save space but might have performance ramifications in certain situations.
- Storage DRS and Storage vMotion – just like VMware's standard DRS that helps manage memory and CPU resources within a cluster, Storage DRS essentially allows the management of several storage volumes in a similar fashion. Ongoing load balancing is performed across the devices to ensure space and I/O bottlenecks are avoided.

## VM storage metrics

Below is a VM view showing the datastores a VM is currently using.

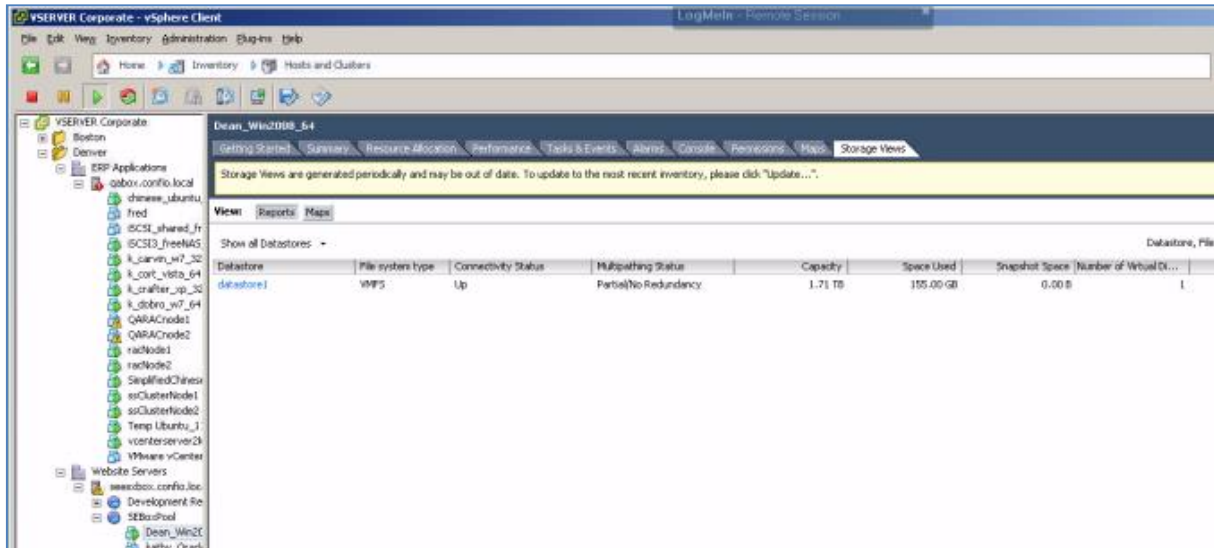


Figure 4. Storage View for a VM

In the screenshot above, the VM is configured with one VMFS datastore named “datastore1”. Also shown are capacity and space used information.



Figure 5. Datastore Metrics

## Monitoring datastore rules

When reviewing the datastore metrics shown above, there are a few simple rules to follow, in the order listed:

1. Latencies are typically the best way to monitor any type of storage and that follows in the VMware environment too. If read or write latencies are above 20-30 ms, contact storage / SAN administrators to understand why. For databases to perform well, latencies should be much lower than this. Consider asking for better and/or faster storage.
2. Read and Write rates may also drive the latency numbers above, i.e. if too many IOPS are being attempted, the storage may not be able to keep up. The rates shown above will also be key metrics to provide to the storage administrators so they know where to place the VM data.
3. Obviously datastore space must also be monitored and Figure 4 above shows where that data is available in the vSphere GUI.

## Monitoring network

Following are a few new concepts and terms to be aware of related to networks in VMware:

- vSwitch – a software version of standard network switch built into the VMware environment that is configured at the physical host level.
- Distributed vSwitch – a centralized configuration of a vSwitch that can be shared by many physical hosts and their VMs.

## VM network metrics

Below is a VM summary showing the network components a VM is currently using.

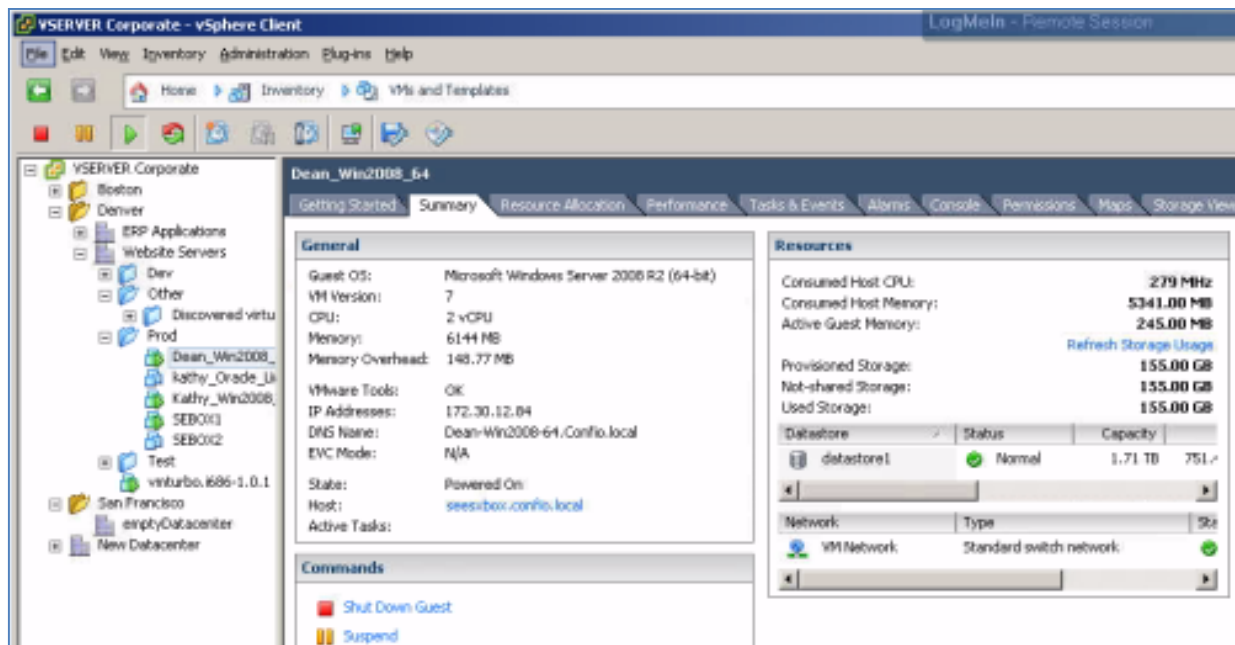


Figure 6. Network Configuration

In this case, the VM is using the standard vSwitch named VMNetwork. Below are the metrics for monitoring the network interface. Note that these metrics are found at the host level where the vSwitch is configured.

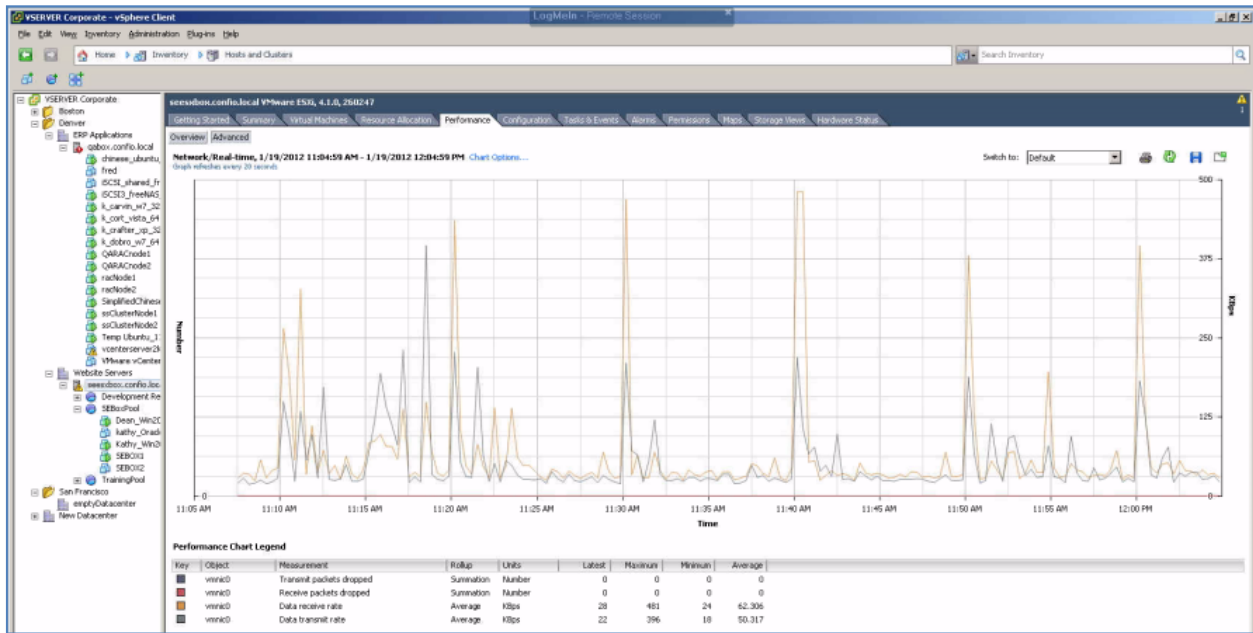


Figure 7. Host Network Metrics

### Monitoring network rules

When reviewing the network metrics shown above for a physical host, there are a few simple rules to follow in order listed:

1. If there are any packets being dropped (transmit or receive), this needs to be investigated and the probable causes are hardware problems or too much data being pushed through the network. When there are packets being dropped, review the transmit and receive rates compared to the physical limitations of the network hardware. If it is close to what the maximum that the network hardware can handle, tune the database to reduce the network load. For example, performing data loads across the network may put an undo load onto the hardware. Another option is to use network hardware that has more capacity.
2. If the receive and transmit rates are getting close to the network hardware limitations, this could be a pending issue. Watch this closely and specifically watch for dropped packets.

### Summary

Running databases in a VMware environment has been successful at many companies. If you are faced with a database virtualization project, and especially if past projects have failed, ensure the correct monitoring metrics are being used in troubleshooting scenarios. If only O/S metrics are being relied on, these may be wrong and cause the DBA team and others to make incorrect assumptions and decisions. Obtaining access to vSphere and/or installing DPA with VM option from SolarWinds is a key step to

getting a complete picture of performance at the VM, but also the physical host, storage and networks layers too. This paper also outlines the key metrics to review at each layer.

## **About Confio Software**

Confio Software, now a part of the SolarWinds family, builds award-winning database performance analysis tools for DBAs and developers. SolarWinds Database Performance Analyzer (formerly Confio Ignite) improves the productivity and efficiency of IT organizations. By resolving problems faster, speeding development cycles, and squeezing more performance out of expensive database systems, Database Performance Analyzer makes DBA and development teams more productive and valuable to the organization. Customers worldwide use our products to improve database performance on Oracle, SQL Server, Sybase and DB2 on physical and virtual machines.

For more information, please visit: [confio.com/performance/vm/ignite/](http://confio.com/performance/vm/ignite/)