

VMware Monitors

Overview

The VMware service monitors allow you to monitor and alert based on the performance and status your virtual resources. These monitors can watch for threshold violations with computing resources for VMs, ESX servers, and changes in power states.

Most of these service monitors use VMware vCenter-collected metrics that are made available to *up.time* through vSync. These VMware vCenter-monitored components, combined with more granular agent-based server monitoring, provide you with choice between breadth and depth.

vSphere Performance Monitors

The vSphere performance monitors allow you to monitor and alert on specific VMware vSphere components: datacenters, clusters, resource pools, and vApps.

The metrics collected through VMware vCenter servers can be used by *up.time* through vSync, and subsequently be used to trigger *up.time*'s own alerts and actions, allowing you to integrate both your vSphere-managed and non-virtual resources.

These performance monitors can answer questions such as the following:

- Is the CPU usage of VMs in a vApp, resource pool, cluster, or datacenter passing an acceptable level?
- Is the memory consumed by VMs in a vApp, resource pool, cluster, or datacenter passing an acceptable level?
- Is the number of ESX servers that are a part of the cluster or datacenter exceeding an acceptable number and threatening performance?

Datacenter and Cluster Performance

The Datacenter Performance and Cluster Performance monitors can trigger alerts on metrics collected through vSync.

Datacenter Performance and Cluster Performance Monitor Metrics

The following VMware vSphere metric types for datacenter or cluster performance can be used to configure thresholds in *up.time* :

Time Interval	A positive integer indicating the number of minutes' worth of performance data samples to average, then compare against threshold definitions (default: 30)
Number of Running VMs: warning threshold and critical threshold	Warning- and critical-level thresholds can be set, using positive integers, for the average number of VMs powered on during the time interval.
Number of Running Hosts: warning threshold and critical threshold	Warning- and critical-level thresholds can be set, using positive integers, for the average number of vSphere ESX servers powered on during the time interval.
CPU Consumed: warning threshold and critical threshold	Warning- and critical-level thresholds can be set, using positive integers, for the total percentage of CPU cycles consumed by VMs belonging to this datacenter or cluster.
Memory Consumed: warning threshold and critical threshold	Warning- and critical-level thresholds can be set, using positive integers, for the total percentage of memory consumed by VMs belonging to this datacenter or cluster.

Configuring Datacenter Performance or Cluster Performance Monitors

To configure a Datacenter Performance or Cluster Performance monitor, do the following:

1. Select the monitor from the *Add Service Monitor* window, in the *VMware Monitors* section.
2. Click *Continue* to begin configuring the service monitor.
3. Complete the monitor information fields.
See [Monitor Identification](#) for more information on configuring service monitor information fields.
4. In the *Cluster Performance Settings* section, configure the monitor's warning- and critical-level threshold values:
 - Time Interval
 - Number of Running VMs
 - Number of Running Hosts
 - CPU Consumed
 - Memory ConsumedFor more information on these metrics, see [Datacenter Performance and Cluster Performance Monitor Metrics](#).
For more information about setting thresholds and response time, see [Configuring Warning and Critical Thresholds](#).
5. Complete the following settings:
 - Timing Settings (see [Adding Monitor Timing Settings Information](#) for more information)
 - Alert Settings (see [Monitor Alert Settings](#) for more information)
 - Monitoring Period settings (see [Monitor Timing Settings](#) for more information)
 - Alert Profile settings (see [Alert Profiles](#) for more information)
 - Action Profile settings (see [Action Profiles](#) for more information)
6. Click *Finish*.

Resource Pool and vApp Performance

The Resource Pool Performance and vApp Performance monitors can trigger alerts on metrics collected through vSync.

Resource Pool Performance and vApp Performance Monitor Metrics

The following VMware vSphere metric types for resource pool and vApp performance can be used to configure thresholds in *up.time* :

Time Interval	A positive integer indicating the number of minutes' worth of performance data samples to average, then compare against threshold definitions (default: 30)
Number of Running VMs: warning threshold and critical threshold	Warning- and critical-level thresholds can be set, using positive integers, for the average number of VMs powered on during the time interval.
CPU Consumed: warning threshold and critical threshold	Warning- and critical-level thresholds can be set, using positive integers, for the total percentage of CPU cycles consumed by VMs belonging to this resource pool or vApp.
Memory Consumed: warning threshold and critical threshold	Warning- and critical-level thresholds can be set, using positive integers, for the total percentage of memory consumed by VMs belonging to this resource pool or vApp.

Configuring Resource Pool Performance or vApp Performance Monitors

To configure a Resource Pool Performance or vApp Performance monitor, do the following:

1. Select the monitor from the *Add Service Monitor* window, in the *VMware Monitors* section.
2. Click *Continue* to begin configuring the service monitor.
3. Complete the monitor information fields.
See [Monitor Identification](#) for more information on configuring service monitor information fields.
4. In the *Resource Pool Performance Settings* section, configure the monitor's warning- and critical-level threshold values:
 - Time Interval
 - Number of Running VMs
 - CPU Consumed
 - Memory ConsumedFor more information on these metrics, see [Resource Pool Performance and vApp Performance Monitor Metrics](#).
For more information about setting thresholds and response time, see [Configuring Warning and Critical Thresholds](#).
5. Complete the following settings:
 - Timing Settings (see [Adding Monitor Timing Settings Information](#) for more information)
 - Alert Settings (see [Monitor Alert Settings](#) for more information)
 - Monitoring Period settings (see [Monitor Timing Settings](#) for more information)
 - Alert Profile settings (see [Alert Profiles](#) for more information)
 - Action Profile settings (see [Action Profiles](#) for more information)
6. Click *Finish*.

ESX Server Monitors

ESX Server monitors focus on the ESX server host, as a physical computing resource, for monitoring and alerting.

There are currently three ESX-related monitors: ESX (Advanced Metrics), a monitor that uses an *up.time* agent on the ESX server; the vSphere ESX Server Performance monitor, which uses metrics transferred to *up.time* using vSync, and the legacy ESX Workload monitor, which is a legacy monitor that can no longer be added to *up.time* .

The metrics collected for these ESX server monitors can be used to trigger *up.time* alerts and actions. These performance monitors can answer questions such as the following:

- Are CPU or memory usage on the host too high?
- Are network and disk I/O usage or latency within acceptable limits?
- Are disk and network error rates too high?
- Are memory ballooning targets being exceeded?

vSphere ESX Server Performance

The vSphere ESX Server Performance monitor allows you to alert based on performance checks on ESX server Elements managed by VMware vSphere, but monitored in *up.time* via vSync.

vSphere ESX Server Performance Monitor Metrics

The following vSphere metric types for ESX server performance can be used to configure thresholds in *up.time* :

Time Interval	A positive integer indicating the number of minutes' worth of performance data samples to average, then compare against threshold definitions (default: 30)
---------------	---

CPU Check: value type, warning threshold, and critical threshold	Warning- and critical-level thresholds can be set, using positive integers, for average CPU usage as either percentage usage, or average MHz usage
Memory Check: value type, warning threshold, and critical threshold	Warning- and critical-level thresholds can be set, using positive integers, for one of the following value types: <ul style="list-style-type: none"> Usage (%) - percentage of total configured or available memory used Memory Consumed (MB) - amount of memory consumed by VMs on this host Memory Active (MB) - amount of memory actively used by VMs on this host Balloon Memory (MB) - amount of memory allocated by <i>vmmemctl</i> across all VMs on this host Zero Memory (KB) - memory that only contains 0s allocated to VMs
Swap Check: value type, warning threshold, and critical threshold	Warning- and critical-level thresholds can be set, using positive integers, for either swap space used (in MB), or swap rate (the combined swap-in rate and swap-out rate, in KBps, across all VMs on this host).
Disk Device I/O: coverage, value type, warning threshold, and critical threshold	Warning- and critical-level thresholds can be set, using positive integers, for one of the following value types: <ul style="list-style-type: none"> Usage (KBps) - aggregate disk I/O rate across all VMs on the host Physical Device Command Latency (ms) - average time to process a read and write from the physical device Queue Command Latency (ms) - average time spent in the VMkernel queue per SCSI command Command Latency (ms) - average time taken to process a SCSI command issued by the Guest OS to the VM <p>Checks are made against the average for all detected disk devices, or any individual device that is violating the threshold.</p>
Disk Device Errors Check: coverage, value type, warning threshold, and critical threshold	Warning- and critical-level thresholds can be set, using positive integers, for either the number of SCSI command aborts per minute, or the number of bus resets per minute. <p>Checks are made against the average for all detected disk devices, or any individual device that is violating the threshold.</p>
Network I/O: coverage, warning threshold, and critical threshold	Warning- and critical-level thresholds can be set, using positive integers, for the aggregate received and transmitted rate (in KBps). <p>Checks are made against the average for all detected network interfaces, or any individual network interface that is violating the threshold.</p>
Network Errors Check: coverage, value type, warning threshold, and critical threshold	Warning- and critical-level thresholds can be set, using positive integers, for the aggregate received and transmitted packets dropped per minute. <p>Checks are made against the average for all detected network interfaces, or any individual network interface that is violating the threshold.</p>

Configuring vSphere ESX Server Performance Monitors

To configure a vSphere ESX Server Performance monitor, do the following:

1. Select the monitor from the *Add Service Monitor* window, in the *VMware Monitors* section.
2. Click *Continue* to begin configuring the service monitor.
3. Complete the monitor information fields.

See [Monitor Identification](#) for more information on configuring service monitor information fields.



When selecting an Element associated with this service monitor, only ESX servers monitored in *up.time* via vSync will appear in the *Single System* list

4. In the *vSphere ESX Server Performance Settings* section, in the *Time Interval* sub-section, enter the number of minutes' worth of time samples that will be used to compare thresholds.
5. For the following metric categories, select the metric unit of measurement, then configure the monitor's warning- or critical-level threshold values:
 - CPU Usage
 - Memory
 - Swap
 - Disk Device I/O
 - Disk Errors
 - Network I/O
 - Network Errors

For more information on these metrics, see [vSphere ESX Server Performance Monitor Metrics](#).
For more information about setting thresholds, see [Configuring Warning and Critical Thresholds](#).
6. Complete the following settings:
 - Timing Settings (see [Adding Monitor Timing Settings Information](#) for more information)
 - Alert Settings (see [Monitor Alert Settings](#) for more information)
 - Monitoring Period settings (see [Monitor Timing Settings](#) for more information)
 - Alert Profile settings (see [Alert Profiles](#) for more information)
 - Action Profile settings (see [Action Profiles](#) for more information)
7. Click *Finish*.

ESX (Advanced Metrics)

The ESX (Advanced Metrics) monitor offers greater visibility into your ESX environment by expanding on the high level usage metrics for a virtual machine's CPU, memory, and disk activity.

ESX Advanced Metrics Monitor Metrics

The following ESX server metrics can be used to configure thresholds:

Percent Wait	Guest metric - The percentage of time that a virtual CPU was not running. A non-running CPU could be idle (halted) or waiting for an external event such as I/O.
Memory Balloon (Avg)	Guest metric - The average amount of memory, in KB, held by memory control for ballooning.
Memory Balloon Target	Guest metric - The total amount of memory, in KB, that can be used by memory control for ballooning.
Memory Overhead (Avg)	Guest metric - The average amount of additional host memory, in KB, allocated to the virtual machine.
Memory Swap In (Avg)	Guest metric - The average amount of memory, in KB, that was swapped in.
Memory Swap Out (Avg)	Guest metric - The average amount of memory, in KB, that was swapped out.
Memory Zero (Avg)	Guest metric - The average amount of memory, in KB, that was zeroed out.
Memory Swap Used (Avg)	Host metric - The average amount of memory, in KB, that was used by the swap file.
Memory Swap Target	Guest metric - The total amount of memory, in KB, that can be swapped.
Disk Total Latency	Host metric - The average time, in milliseconds, taken for disk commands by a guest OS. This is the sum of <i>kernelCommandLatency</i> and <i>physical deviceCommandLatency</i> .
Disk Kernel Latency	Host metric - The average time, in milliseconds, spent in the ESX Server <i>VMkernel</i> per command.
Disk Device Latency	Host metric - The average time, in milliseconds, taken to complete a command from the physical device.
Disk Queue Latency	Host metric - The average time, in milliseconds, spent in the ESX Server <i>VMkernel</i> queue per write.
Disk Commands Aborted	Host metric - The number of disk commands aborted during the defined interval.
Disk Commands Issued	Host metric - The number of disk commands issued during the defined interval.
Disk Bus Resets	Host metric - The number of bus resets during the defined interval.

Configuring ESX (Advanced Metrics) Monitors

To configure an ESX (Advanced Metrics) monitor, do the following:

1. Select the monitor from the *Add Service Monitor* window, in the *VMware Monitors* section.
2. Click *Continue* to begin configuring the service monitor.
3. Complete the monitor information fields.
See [Monitor Identification](#) for more information on configuring service monitor information fields.
4. In the *ESX (Advanced Metrics) Settings* section, configure the monitor's warning- and critical-level alerting thresholds by completing the following fields:
 - Percent Wait
 - Memory Balloon
 - Memory Balloon Target
 - Memory Overhead
 - Memory Swap In
 - Memory Swap Out
 - Memory Zero
 - Memory Swap Used
 - Memory Swap Target
 - Disk Total Latency
 - Disk Kernel Latency
 - Disk Device Latency
 - Disk Queue Latency
 - Disk Commands Aborted
 - Disk Commands Issued
 - Disk Bus Resets

- Response time
For more information on these metrics, see [ESX Advanced Metrics Monitor Metrics](#).
For more information about setting thresholds and response time, see [Configuring Warning and Critical Thresholds](#).
5. Complete the following settings:
 - Timing Settings (see [Adding Monitor Timing Settings Information](#) for more information)
 - Alert Settings (see [Monitor Alert Settings](#) for more information)
 - Monitoring Period settings (see [Monitor Timing Settings](#) for more information)
 - Alert Profile settings (see [Alert Profiles](#) for more information)
 - Action Profile settings (see [Action Profiles](#) for more information)
 6. Click *Finish*.

ESX Workload

The ESX Workload monitor collects a set of metrics from all of the instances that are running on an ESX v3 or v4 server over a specified time period. This monitor is a legacy monitor that cannot be added to your *up.time* configuration as a new service monitor. It exists in upgraded configurations that included it.

The monitor compares the highest values returned by the instances and then compares them to the thresholds that you set. If the values exceed the thresholds, *up.time* issues an alert. The monitor does not pinpoint the specific instance(s) that have exceeded the defined thresholds.

For example, you are monitoring an ESX server that is running three instances. You configured the ESX Workload monitor to collect data samples every 10 minutes, and to issue a warning when memory usage exceeds 300 MB. The three instances are using the following amounts of memory: 110 MB, 227 MB, and 315 MB. The ESX Workload monitor focuses on the value of 315 MB and, since it exceeds the warning threshold, issues an alert.

ESX Workload Monitor Metrics

The following metrics are used by the ESX Workload monitor:

Time Interval	The amount of time, in minutes, at which the monitor will collect data samples from the ESX server.
CPU Warning Threshold	The amount of CPU resources, measured in megahertz (MHz), that the instances on the ESX server must consume before <i>up.time</i> issues a warning.
CPU Critical Threshold	The amount of CPU resources, measured in megahertz MHz, that the instances on the ESX server must consume before <i>up.time</i> issues a critical alert.
Network Bandwidth Warning Threshold	The amount of network traffic in and out of the server, measured in megabits per second (Mbit/s), that must be exceeded before <i>up.time</i> issues a warning.
Network Bandwidth Critical Threshold	The amount of network traffic in and out of the server, measured in megabits per second (Mbit/s), that must be exceeded before <i>up.time</i> issues a critical alert.
Disk Usage Warning Threshold	The amount of data being written to the server's hard disk, measured in kilobytes per second (KB/s), that must be exceeded before <i>up.time</i> issues a warning.
Disk Usage Critical Threshold	The amount of data being written to the server's hard disk, measured in kilobytes per second (KB/s), that must be exceeded before <i>up.time</i> issues a critical alert.
Memory Usage Warning Threshold	The amount of overall system memory, measured in megabytes (MB), that must be exceeded before <i>up.time</i> issues a warning.
Memory Usage Critical Threshold	The amount of overall system memory, measured in megabytes (MB), that must be exceeded before <i>up.time</i> issues a critical alert.
Percent Ready Warning Threshold	The percentage of time that one or more instances running on an ESX server is ready to run, but cannot run because it cannot access the processor on the ESX server. If the valued returned from the server exceeds this threshold, then <i>up.time</i> issues a warning.
Percent Ready Critical Threshold	The percentage of time that one or more instances running on an ESX server is ready to run, but cannot run because it cannot access the processor on the ESX server. If the valued returned from the server exceeds this threshold, then <i>up.time</i> issues a critical alert.
Percent Used Warning Threshold	The percentage of CPU time that an instance running on an ESX server is using. If the valued returned from the server exceeds this threshold, then <i>up.time</i> issues a warning.
Percent Used Critical Threshold	The percentage of CPU time that an instance running on an ESX server is using. If the valued returned from the server exceeds this threshold, then <i>up.time</i> issues a critical alert.

Modifying an ESX Workload Monitor Configuration

To modify the configuration of a legacy ESX Workload monitor, do the following:

1. If required, change the monitor information fields.
See [Monitor Identification](#) for more information.

2. In the *ESX Workload Settings* section, modify any of the monitor's existing warning- or critical-level threshold values:
 - Time Interval
 - CPU Usage
 - Network Bandwidth Usage
 - Disk Usage
 - Memory Usage
 - Percent Ready
 - Percent Used

For more information on these metrics, see [ESX Workload Monitor Metrics](#).
For more information about setting thresholds, see [Configuring Warning and Critical Thresholds](#).
3. Complete the following settings:
 - Timing Settings (see [Adding Monitor Timing Settings Information](#) for more information)
 - Alert Settings (see [Monitor Alert Settings](#) for more information)
 - Monitoring Period settings (see [Monitor Timing Settings](#) for more information)
 - Alert Profile settings (see [Alert Profiles](#) for more information)
 - Action Profile settings (see [Action Profiles](#) for more information)
4. Click *Finish*.

Power State Monitors

The power state monitors help you manage both available computing resources within your clusters, resource pools, and other logical divisions in your vSphere-managed infrastructure, as well as power consumption in your physical datacenters. Power state changes to your hosts, and the VMs running on them, can be alerted and acted on.

The power state monitors can answer questions such as the following:

- Has a mission-critical VM been powered off?
- Did a routine maintenance procedure start and complete properly?
- Are enough expected VMs powering down during the weekend, indicating vSphere's Distributed Power Management is functioning correctly?

ESX Server Power State

The ESX Server Power State monitor watches for changes to the power states of an ESX server that is managed by VMware vSphere, and can run alert or action profiles based on the change.

ESX Server Power State Monitor Status Types

In *up.time*, vSphere hosts will be in one of the following states:

Powered On	The host is running.
Powered Off	The host was powered off by an administrator through the vSphere Client.
Put on Standby	The host was put in standby mode either explicitly by an administrator, or automatically by vSphere Distributed Power Management (DPM).
Put in Maintenance	The host state is determined to be "unknown" if it is disconnected or not responding, implying it is in maintenance.

Configuring ESX Server Power State Monitors

To configure an ESX Server Power State monitor, do the following:

1. Select the monitor from the *Add Service Monitor* window, in the *VMware Monitors* section.
2. Click *Continue* to begin configuring the service monitor.
3. Complete the monitor information fields.
See [Monitor Identification](#) for more information on configuring service monitor information fields.
For more information on these power states, see [ESX Server Power State Monitor Status Types](#).



When selecting an Element associated with this service monitor, only ESX servers monitored in *up.time* via vSync will appear in the *Single System* list.

4. In the main *ESX Server Power State Settings* section, in the Powered On sub-section, do the following:
 - In the *Set Status to* drop-down box, indicate what the monitor's *up.time* state will be when the ESX server's state is Powered On.
 - From the list, select which (if any) Alert Profiles are triggered when the host enters a powered-on state.
 - From the list select which (if any) Action Profiles will be triggered when the host enters a powered-on state.
5. In the *Powered Off* sub-section, do the following:
 - In the *Set Status to* drop-down box, indicate what the monitor's *up.time* state will be when the ESX server's state is Powered Off.
 - From the list, select which (if any) Alert Profiles are triggered when the host enters a powered-off state.
 - From the list select which (if any) Action Profiles will be triggered when the host enters a powered-off state.
6. In the *Put on Standby* sub-section, do the following:
 - In the *Set Status to* drop-down box, indicate what the monitor's *up.time* state will be when the ESX server's state is Standby.
 - From the list, select which (if any) Alert Profiles are triggered when the host enters a standby state.
 - From the list select which (if any) Action Profiles will be triggered when the host enters a standby state.
7. In the *Put in Maintenance* sub-section, do the following:
 - In the *Set Status to* drop-down box, indicate what the monitor's *up.time* state will be when the ESX server's state is Unknown.
 - From the list, select which (if any) Alert Profiles are triggered when the host enters an unknown state.
 - From the list select which (if any) Action Profiles will be triggered when the host enters an unknown state.

8. Complete the following settings:
 - Timing Settings (see [Adding Monitor Timing Settings Information](#) for more information)
 - Monitoring Period settings (see [Monitor Timing Settings](#) for more information)
9. Click *Finish*.

VM Instance Power State

The VM Instance Power State monitor watches for changes to the power states of a VM running on an ESX server that is managed by vSphere, and can run alert or action profiles based on the change.

See [Power State Monitors](#) for more information.

VM Instance Power State Monitor Status Types

A virtual machine's three basic power states are as follows:

Powered On	The virtual machine is running.
Powered Off	The virtual machine is not running.
Suspended	The virtual machine is not running, but a snapshot of its running applications and processes is retained.

Configuring VM Instance Power State Monitors

To configure a VM Instance Power State monitor, do the following:

1. Select the monitor from the *Add Service Monitor* window, in the *VMware Monitors* section.
2. Click *Continue* to begin configuring the service monitor.
3. Complete the monitor information fields.
See [Monitor Identification](#) for more information on configuring service monitor information fields.
For more information on these VM power states, see [VM Instance Power State Monitor Status Types](#).



When selecting a VM associated with this service monitor, only VMs monitored in *up.time* via vSync will appear in the *Single System* list.

4. In the main *VM Instance Power State Settings* section, in the Powered On sub-section, do the following:
 - In the *Set Status to* drop-down box, indicate what the monitor's *up.time* state will be when the VM's state is "powered on".
 - From the list, select which (if any) Alert Profiles are triggered when the host enters a powered-on state.
 - From the list select which (if any) Action Profiles will be triggered when the host enters a powered-on state.
5. In the *Powered Off* sub-section, do the following:
 - In the *Set Status to* drop-down box, indicate what the monitor's *up.time* state will be when the VM's state is "powered off".
 - From the list, select which (if any) Alert Profiles are triggered when the host enters a powered-off state.
 - From the list select which (if any) Action Profiles will be triggered when the host enters a powered-off state.
6. In the *Suspended* sub-section, do the following:
 - In the *Set Status to* drop-down box, indicate what the monitor's *up.time* state will be when the VM's state is "suspended".
 - From the list, select which (if any) Alert Profiles are triggered when the host enters a suspended state.
 - From the list select which (if any) Action Profiles will be triggered when the host enters a suspended state.
7. Complete the following settings:
 - Timing Settings (see [Adding Monitor Timing Settings Information](#) for more information)
 - Monitoring Period settings (see [Monitor Timing Settings](#) for more information)
8. Click *Finish*.