

## Machine Agent Settings for Server Visibility

### On this page:

- [Editing YML Files: Important Notes](#)
- [Configurable Server Visibility Settings](#)
- [Docker Visibility Settings](#)

You can modify the default configuration for Server Visibility. This page describes the configurable settings and their default values.

Configuration file to edit: `<machine-agent-home>/extensions/ServerMonitoring/conf/ServerMonitoring.yml`.

### Editing YML Files: Important Notes

Note the following:

- If you make any changes to `ServerMonitoring.yml`, or to any other YML file, make sure that the modified file meets standard YML syntax rules. Some important YML good practices include:
  - Do not include any tab characters. Use whitespace characters only to indent fields.
  - Use the exact same number of whitespace characters to indent entries that are on the same level.
  - Use a plain-text editor, not a WYSIWIG editor, to edit the file. Use a monowidth/monospaced font to view the contents.
  - Always save using UTF-8 encoding.
  - Test and verify the edited file using an online YML syntax tester.
- The safest way to edit a setting in this file is to
  - Copy the line you want to edit and paste the copy into a new line. Make sure that you select, copy, and paste the entire line.
  - Comment out the original line and edit the copy as desired.
- If you want to add or edit a regular expression in this file, it is good practice to test and verify the regular expression using an online regex tester.
- Before you make any changes to this file, read the option descriptions below and the comments in the file carefully for the setting you want to change.
- The agent updates dynamically in response to agent configuration property changes, so there is no need to restart the agent after you update this file.

### Configurable Server Visibility Settings

Setting	Description
basicEnabled	Indicates whether the machine agent should report the basic metrics through the SIM extension. Set this to false to use Sigar to report basic metrics. This setting only affects monitoring on Linux.  Default = true

<p>volumeMonitorConfig: maxNumberVolumes</p>	<p>Do not report more than N volumes, where N = maxNumberVolumes. Default = 5</p> <div style="border: 1px solid red; padding: 5px; margin-top: 10px;"> <p>Changing this setting can affect the resource consumption of your deployment. Before you increase this setting, verify that your application environment and Controller can handle the increased resource requirements.</p> </div>
<p>volumeMonitorConfig: whitelistSelectorRegex</p>	<p>Volumes with names that match this regular expression are always reported, up to the maximum specified by maxNumberVolumes. Default = ""</p>
<p>volumeMonitorConfig: blacklistSelectorRegex</p>	<p>Volumes with names that match this regular expression are excluded. This setting is useful for filtering out irrelevant metrics.</p> <p>If a volume name matches both the blacklist and whitelist regexes, metrics for that network are reported (the whitelist takes priority) up to the maximum specified by maxNumberVolumes.</p> <p>The default ServerMonitoring.yml file does not include a blacklistSelectorRegex field for volumes. If you want to add one, use the same indentation and formatting as the volumeMonitorConfig:whitelistSelectorRegex field.</p>
<p>volumeMonitorConfig: samplingInterval</p>	<p>Specifies a custom sampling interval for collecting volume metrics on Linux. Default = 3000</p>
<p>networkMonitorConfig: maxNumberNetworks</p>	<p>Do not report more than N networks, where N = maxNumberNetworks. Default = 5</p> <div style="border: 1px solid red; padding: 5px; margin-top: 10px;"> <p>Changing this setting can affect the resource consumption of your deployment. Before you increase this setting, verify that your application environment and Controller can handle the increased resource requirements.</p> </div>
<p>networkMonitorConfig: whitelistSelectorRegex</p>	<p>Networks with names that match this regular expression are always reported, up to the maximum specified by maxNumberNetworks.</p> <p>To report metrics for one or more virtual networks, specify a regex that matches the virtual network names.</p> <p>When collecting Process metrics on Solaris, the Standalone Machine Agent observes and captures only the first 80 characters of each process name and argument list. This means that the agent considers only the first 80 characters of each process string when it applies whitelists. Default = ""</p>
<p>networkMonitorConfig: blacklistSelectorRegex</p>	<p>Networks with names that match this regular expression are excluded. This setting is useful for filtering out irrelevant metrics.</p> <p>If a network name matches both the blacklist and whitelist regexes, metrics for that network are reported (the whitelist takes priority) up to the maximum specified by maxNumberNetworks.</p> <p>The default regex excludes virtual networks. To monitor a set of one or more virtual networks, edit whitelistSelectorRegex to include the networks of interest. To monitor all virtual networks, change the blacklistSelectorRegex to an empty string.</p> <p>When collecting Process metrics on Solaris, the Standalone Machine Agent observes and captures only the first 80 characters of each process name and argument list. This means that the agent considers only the first 80 characters of each process string when it applies blacklists. Default = "^veth.* ^vnet.*"</p>

defaultProcessClassSelector	<p>The default "class selector" based on a <i>class_selector_regex</i>. If this regex is defined, and a match is found in the process command line, the class name is the first group occurrence of that regex in the command line.</p> <p>Default = ""</p>
processClassSelectorRegexList	<p>A list of <i>class_name:regex</i> mappings. If the command line for a process matches <i>regex</i>, the metrics for that process are assigned to <i>class_name</i>. This setting is useful when you want to ensure that high-priority processes get reported, even if the number of defined classes is higher than the <i>maxNumberMonitoredClasses</i> setting.</p> <p>Here is an example:</p> <pre>processMonitorConfig:     processClassSelectorRegexList :         machineAgentTasks: '.*java.*machineagent.*'         controllerTasks: '.*java.*controller.*'         nextOne: '.*svchost.*'</pre> <p>The Machine Agent assigns a process to a class as follows:</p> <ol style="list-style-type: none"> <li>1. Assign to the first match in <i>processClassSelectorRegexList</i>.</li> <li>2. If there is no match for the <i>processClassSelectorRegexList</i> (step1), apply the <i>defaultProcessClassSelector</i> regex to the command line.</li> <li>3. If there is no match for the <i>defaultProcessClassSelector</i> regex (step 2), use the process name (truncated if the name exceeds the <i>maxClassIdLength</i>).</li> </ol> <p>The following steps outline the recommended workflow for updating this list:</p> <ol style="list-style-type: none"> <li>1. The default <i>ServerMonitoring.yml</i> file includes an example that is commented out. If you are updating the default list for the first time, it is good practice to             <ol style="list-style-type: none"> <li>a. Create a copy of the example.</li> <li>b. Uncomment the copy (remove the <code>&lt;!--</code> and <code>--&gt;</code> comment tags) and edit it as needed.</li> </ol> </li> <li>2. Choose the process(es) that you want to monitor on the host machine.</li> <li>3. Create a regex to match the process name(s) of interest. It is good practice to test the regex using an online regex validator.</li> <li>4. Add the regex to the list. You should order the mappings by priority, highest to lowest. If a command line matches multiple regexes, the first match is used.</li> <li>5. Save the <i>ServerMonitoringConfiguration.yml</i> file.</li> <li>6. Wait 15 minutes or longer for the updated list to take effect, then verify that the matching process(es) appear in the Controller UI.</li> </ol>
samplingInterval	<p>Indicates how often to gather metric data. Units in milliseconds.</p> <p>Default = 60000 (60 seconds)</p>
maxClassIdLength	<p>Specifies the maximum process class name length. Any process class name that is longer than the specified maximum is truncated. The global maximum of the process class name is 100. If this variable is set to be greater than 100, then the process name is truncated at 100.</p> <p>Default = 50</p>
processSelectorRegex	<p>Contains a regular expression that specifies which processes should be monitored by the machine agent. The regular expression is compared against the full command line that was used to start the process.</p> <p>The default regular expression will filter out any processes where the command line ends with a close bracket (']').</p> <p>For Linux, this means that the process arguments could not be found and usually that is a kernel process.</p> <p>For Windows, no processes should end with a bracket character, so consequently the regex should include all processes on Windows.</p> <p>Default = ""^.[^)]\$""</p>

minLiveTimeMillisBeforeMonitoring	<p>Specifies the minimum amount of time a process must be alive before it is monitored by the machine agent. This is useful for preventing the machine agent from being overloaded by monitoring short-lived processes. Units in milliseconds.</p> <p>Default = 60000 (60 seconds)</p>
maxNumberMonitoredClasses	<p>Specifies the maximum number of process classes that the machine agent monitors.</p> <p>The processes that are reported are the processes using the highest CPU and memory that match the regex specified by processSelectorRegex.</p> <p>Default = 20</p> <div style="border: 1px solid red; padding: 5px; margin-top: 10px;"> <p>Changing this setting can affect the resource consumption of your deployment. Before you change this setting verify your application environment and Controller can handle the increased resource requirements.</p> </div>
defaultDiskSectorSize	<p>Specifies the default sector size (in bytes) for each disk if the machine agent cannot determine the sector size. This value is used to calculate the number of bytes read/written for the disk.</p> <p>Default = 512</p>
memoryMonitorConfig	<p>Specifies a custom sampling interval for collecting memory metrics on Linux:</p> <p>Default = 3000</p>
cpusMonitorConfig	<p>Specifies a custom sampling interval for collecting CPU metrics on Linux:</p> <p>Default = 3000</p>
tag	<p>A list of user-defined tags for the individual server. You can use these tags to query, filter, aggregate, and compare related servers. See <a href="#">Server Tagging</a>.</p> <p>Each tag is specified by a key-value pair. You can define tag names up to 127 unicode characters and tag values up to 255 unicode characters. Define each key on a separate line. All key/value strings should be within single quotes. If a key has multiple values, delineate the list with commas:</p> <pre>&lt;key&gt;: [ &lt;value&gt; ] tags:   'Location': ['NYC', 'Data Center', 'Server Room 7']   'Environment': ['preProduction']</pre>

## Docker Visibility Settings

For information about Docker Visibility settings on the Standalone Machine Agent, see [Configuring Docker Visibility](#).

## Process Limits

Other configurable settings are available in the Controller Admin UI, for example:

- Total number of processes displayed in the UI for a single query is 5000 processes per call.
- Total number of processes tracked per account. The default value is 10000 processes per account.

For more details, see [Controller Settings for Server Visibility](#).