



Monitoring Docker

Restricted Rights Legend

The information contained in this document is confidential and subject to change without notice. No part of this document may be reproduced or disclosed to others without the prior permission of eG Innovations, Inc. eG Innovations, Inc. makes no warranty of any kind with regard to the software and documentation, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose.

Trademarks

Microsoft Windows, Windows 2008, Windows 7/8/10, Windows 2012 and Windows 2016 are either registered trademarks or trademarks of Microsoft Corporation in United States and/or other countries.

The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

Copyright

© 2016 eG Innovations, Inc. All rights reserved.

The copyright in this document belongs to eG Innovations, Inc. Complying with all applicable copyright laws is the responsibility of the user.

Table of Contents

INTRODUCTION	1
1.1 How does eG Enterprise Monitor the Docker?	3
1.2 Pre-requisites for Monitoring Docker in an Agent based approach.....	3
1.3 Pre-requisites for Monitoring Docker in an Agentless approach.....	3
MONITORING THE DOCKER	5
2.1 The Docker Engine Layer	5
2.1.1 Docker Service Status Test.....	5
2.1.2 Docker Storage Test.....	7
2.1.3 Docker Events Test	9
2.1.4 Docker Images Test.....	12
2.1.5 Docker Images - Performance Test	15
2.2 The Docker Containers Layer	18
2.2.1 Docker Containers - Status Test.....	19
2.2.2 Docker Containers - Uptime Test.....	22
2.2.3 Docker Containers - Connectivity Test	24
2.2.4 Docker Containers - Performance Test	26
CONCLUSION	34

Introduction

Docker is an open platform for developing, shipping, and running applications. Docker is designed to deliver your applications faster. With Docker you can separate your applications from your infrastructure and treat your infrastructure like a managed application. Docker helps you ship code faster, test faster, deploy faster, and shorten the cycle between writing code and running code. Docker does this by combining a lightweight container virtualization platform with workflows and tooling that help you manage and deploy your applications.

At its core, Docker provides a way to run almost any application securely isolated in a container. The isolation and security allow you to run many containers simultaneously on your host. The lightweight nature of containers, which run without the extra load of a hypervisor, means you can get more out of your hardware.

Docker has two major components:

- **Docker:** the open source container virtualization platform.
- **Docker Hub:** the Software-as-a-Service platform for sharing and managing Docker containers.

Docker uses a client-server architecture. The Docker *client* talks to the Docker *daemon*, which does the heavy lifting of building, running, and distributing the Docker containers. Both the Docker client and the daemon *can* run on the same system, or you can connect a Docker client to a remote Docker daemon. The Docker client and daemon communicate via sockets or through a RESTful API.




Figure 1: The Docker architecture

INTRODUCTION

The Docker daemon runs on a host machine. The user does not directly interact with the daemon, but instead through the Docker client.

The Docker client, in the form of the docker binary, is the primary user interface to Docker. It accepts commands from the user and communicates back and forth with a Docker daemon.

To understand Docker's internals, you need to know about three components:

- **Docker images:** A Docker image is a read-only template. For example, an image could contain an Ubuntu operating system with Apache and your web application installed. Images are used to create Docker containers. Docker provides a simple way to build new images or update existing images, or you can download Docker images that other people have already created. Docker images are the **build** component of Docker.
- **Docker registries:** Docker registries hold images. These are public or private stores from which you upload or download images. The public Docker registry is provided with the Docker Hub. It serves a huge collection of existing images for your use. These can be images you create yourself or you can use images that others have previously created. Docker registries are the **distribution** component of Docker.
- **Docker containers:** Docker containers are similar to a directory. A Docker container holds everything that is needed for an application to run. Each container is created from a Docker image. Docker containers can be run, started, stopped, moved, and deleted. Each container is an isolated and secure application platform. Docker containers are the **run** component of Docker.

Due to the lightweight architecture of the docker and fast accessibility of the applications, Docker is gaining a rapid foothold among IT giants. As continuous access to the applications is the key in such environments, even the smallest slip in the performance of the Docker would result in huge losses. To ensure the 24x7 availability of the Docker and high performance rate, administrators need to closely monitor the performance and status of the Docker and its associated components, promptly detect abnormalities, and fix them before end-users notice.

eG Enterprise offers a specialized *Docker* monitoring model that monitors each of the key indicators of the performance of Docker - such as the images, containers etc. - and proactively alerts administrators to potential performance bottlenecks, so that administrators can resolve the issues well before end-users complain.




Figure 2: The layer model of Docker server

Each layer of Figure 2 above is mapped to a variety of tests, each of which report a wealth of useful information related to the docker server. Using these metrics, administrators can find quick and accurate answers to the following performance queries:

- Is the docker service installed?
- What is the current status of the docker service?

- How well the data space and metadata space of the docker server is utilized?
- How many events were triggered on the containers of the docker server? Which type of event was triggered the maximum and the minimum? – is it the *create* events or *start* events or *stop* events or *die* events?
- What is the total count of the docker images in the server? How many images are used to create the containers?
- What is the disk space utilization of the images that are mapped to the containers?
- How many containers are created from each image of the docker server? How many containers are actually running and how well the CPU, disk and memory resources are utilized by those containers?
- How many containers are available in the server and among them how many containers are currently running? How many containers are added recently and how many are actually removed?
- What is the uptime of each container?
- Is the container available over the network?
- Are the read and write storage processor (SP) caches used optimally? Which storage processor's cache may require right-sizing, and which cache is it - read or write?
- What is the disk space utilization of each container?
- How well data is transmitted and received from each container?
- What is the rate of errors that are transmitted through each container?
- How well the memory is utilized by each container?
- How well data is read from and written to each container?

1.1 How does eG Enterprise Monitor the Docker?

eG Enterprise employs both *agentbased* and *agentless* approach to monitor the Docker.

1.2 Pre-requisites for Monitoring Docker in an Agent based approach

In an agent based approach, the eG agent communicates with the docker using *sudo* commands. Therefore, in the test configuration page, the **USESUDO** parameter is set to **Yes**, by default.

1.3 Pre-requisites for Monitoring Docker in an Agentless approach

In an agentless approach, the eG agent communicates with the docker through the eG agent install user. Provide the credentials of such a user while adding the Docker component. For an eG agent to collect the required metrics from the Docker, the eG agent install user should be within a Docker group.

Alternately, if the eG agent install user is not within the Docker group or if the administrators in your environment do not wish to add a new user to the docker group, then you can do the following to collect the required metrics:

1. Login to the docker host through **putty.exe**.
2. Upon logging in, navigate to the **/etc/sudoers** file of the docker. The default user privileges to the server will be "ALL".
3. Say for example, if **eguser** is your eG agent install user, then, provide the following privileges to the **eguser**:

INTRODUCTION

```
eguser ALL=NOPASSWD: /usr/bin/docker
```

4. Finally, save the file. Once the necessary user privileges are provided, the target Docker server is ready for monitoring by the eG Enterprise suite.

Once the aforesaid requirements are fulfilled, the eG agent will report a host of useful metrics revealing the status and performance of the Docker and present these performance statistics in the eG monitoring model using the hierarchical layer model representation of Figure 2.

The chapter that follows will discuss each layer of Figure 2 in great detail.

Monitoring the Docker

This chapter deep dives into every layer of the *Docker* monitoring model, the tests mapped to each layer, and the measures every test reports. Since the **Operating System**, **Network** and **TCP** layers of the Docker is similar to that of a *Linux* server and had been dealt in extensively in *Monitoring Unix and Windows servers* document, let us now discuss the layers that are exclusive to the Docker in the forthcoming sections.

2.1 The Docker Engine Layer

The tests mapped to this layer reports the current state of the docker service, the utilization of data and metadata space of the docker, events triggered on the containers of the docker server, image available in the docker, the statistics revealing the number of images mapped to containers and those that are not mapped, individual resource utilization of the containers that are created based on each image etc.



Figure 3: The tests mapped to the Docker Engine layer

2.1.1 Docker Service Status Test

A Docker container created from a Docker image holds everything that is needed for an application to run. In order to create a Docker container, the Docker service should be up and running without a glitch. If the Docker containers cannot be created or if the containers become inaccessible, then the applications running on the containers will be rendered inaccessible to the users causing inconvenience. To avoid such inconvenience to the users accessing the Docker containers, administrators need to constantly monitor the current state of the Docker service. The **Docker Service Status** test exactly helps administrators in this regard!

This test helps administrators to continuously monitor the Docker service and reports the current state of the service. In addition, this test helps administrators to figure out whether the Docker service is installed and loaded.

Purpose	Helps administrators to continuously monitor the Docker service and reports the current state of the service. In addition, this test helps administrators to figure out whether the Docker service is
---------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

	installed and loaded								
Target of the test	A Docker server								
Agent deploying the test	An internal/remote agent								
Configurable parameters for the test	<ol style="list-style-type: none"> 1. TEST PERIOD – How often should the test be executed 2. HOST – Specify the host for which the test is to be configured 3. PORT - The port number at which the specified HOST listens. The default is 2375. 4. USESUDO – By default, this flag is set to Yes indicating either of the following: <ul style="list-style-type: none"> ➤ The eG agent is using the sudo commands to collect the required metrics from the target docker server if monitored in an agent based manner or ➤ The target docker server is monitored in an agentless manner and the eG agent install user is within a docker group. If the target docker server is monitored in an agentless manner and the user privileges are altered in the /etc/sudoers file as mentioned in Section 1.2, then, set this flag to No. 								
Outputs of the test	One set of results for the Docker server to be monitored								
Measurements made by the test	Measurement Is service installed?: Indicates whether/not the Docker service is installed.	Measurement Unit	Interpretation <p>The values reported by this measure and their numeric equivalents are available in the table below:</p> <table border="1"> <thead> <tr> <th>Measure Value</th> <th>Numeric Value</th> </tr> </thead> <tbody> <tr> <td>No</td> <td>0</td> </tr> <tr> <td>Yes</td> <td>1</td> </tr> </tbody> </table> <p>Note:</p> <p>This measure reports the Measure Values listed in the table above while indicating whether the Docker service is installed or not. However, in the graph of this measure, this measure is indicated using only the Numeric Values listed in the above table.</p>	Measure Value	Numeric Value	No	0	Yes	1
Measure Value	Numeric Value								
No	0								
Yes	1								

	<p>Is service loaded?: Indicates whether/not the Docker service is loaded.</p>		<p>The values reported by this measure and their numeric equivalents are available in the table below:</p> <table border="1" data-bbox="1003 306 1444 496"> <thead> <tr> <th>Measure Value</th><th>Numeric Value</th></tr> </thead> <tbody> <tr> <td>No</td><td>0</td></tr> <tr> <td>Yes</td><td>1</td></tr> </tbody> </table> <p>Note: This measure reports the Measure Values listed in the table above while indicating whether the Docker service is loaded or not. However, in the graph of this measure, this measure is indicated using only the Numeric Values listed in the above table.</p>	Measure Value	Numeric Value	No	0	Yes	1
Measure Value	Numeric Value								
No	0								
Yes	1								
	<p>Docker service status: Indicates the current state of the Docker service.</p>		<p>If the Docker service is currently running, then the value of this measure is reported as <i>Active</i> and if the Docker service is not running, then the value of this measure is reported as <i>In active</i>.</p> <p>The values reported by this measure and their numeric equivalents are available in the table below:</p> <table border="1" data-bbox="1003 1087 1444 1277"> <thead> <tr> <th>Measure Value</th><th>Numeric Value</th></tr> </thead> <tbody> <tr> <td>In active</td><td>0</td></tr> <tr> <td>Active</td><td>1</td></tr> </tbody> </table> <p>Note: This measure reports the Measure Values listed in the table above while indicating whether the Docker service is active or not. However, in the graph of this measure, the state is indicated using only the Numeric Values listed in the above table.</p>	Measure Value	Numeric Value	In active	0	Active	1
Measure Value	Numeric Value								
In active	0								
Active	1								

2.1.2 Docker Storage Test

When the Docker service runs for the first time in the Docker server, a defined quantity of data space and metadata space is allocated from the Docker volume, using which data can be stored in the Docker container outside of the boot volume (but within the root file system). By default, the Docker offers thin provisioning facility of volumes which means that you have a big pool of available storage blocks within the volume and can create block devices such as virtual disks, of arbitrary size from that pool. The storage blocks will be marked as used or taken from the pool only when you actually write to it. The storage blocks are mapped to the data space and the metadata space for storage

needs. Here, the Docker images are stored in the data space and the definite data about the docker images are stored in the metadata space. Since the Docker images are the basic building blocks of the Docker, if the data space or the metadata space is not sufficient, then new images and new containers based on the images cannot be created. If sufficient images are not available, then users cannot create new containers and simultaneously applications will also become inaccessible or the users will have only a limited choice of applications. To avoid such eventualities, it is inevitable that the data space and metadata space be monitored at regular intervals. The **Docker Storage** test helps you in this regard!

This test reports the utilization of data and metadata space within a Docker server and proactively alerts administrators to potential space crunch, if any.

Note:

This test will run only if the Docker server is installed with the **Device Mapper**, which is a Linux kernels framework for mapping physical block devices onto higher-level virtual block devices.

Purpose	Reports the utilization of data and metadata space within a Docker server and proactively alerts administrators to potential space crunch, if any		
Target of the test	A Docker server		
An internal/remote agent	An internal/remote agent		
Configurable parameters for the test	<ol style="list-style-type: none"> TEST PERIOD – How often should the test be executed HOST – Specify the host for which the test is to be configured PORT - The port number at which the specified HOST listens. The default is 2375. USESUDO – By default, this flag is set to Yes indicating either of the following: <ul style="list-style-type: none"> ➤ The eG agent is using the sudo commands to collect the required metrics from the target docker server if monitored in an agent based manner or ➤ The target docker server is monitored in an agentless manner and the eG agent install user is within a docker group. If the target docker server is monitored in an agentless manner and the user privileges are altered in the /etc/sudoers file as mentioned in Section 1.2, then, set this flag to No. DOCKER USER AND DOCKER PASSWORD- Specify the credentials of the user who has the right to execute commands 		
Outputs of the test	One set of results for the Docker server being monitored		
Measurements made by the test	Measurement	Measurement Unit	Interpretation
	Data space used: Indicates the amount of data space currently utilized by this server.	GB	A low value is desired for this measure. If the value of this measure is closer to the Data space total measure, then it indicates that the Docker server is currently running short of space. Administrators have to ensure that sufficient data space is allocated to the Docker server.

	Data space total: Indicates the total amount of data space allocated to this server.	GB	
	Data space utilization: Indicates the percentage of data space utilized by this server.	Percent	A high value for this measure indicates that the Docker server is currently running out of space.
	Metadata space used: Indicates the amount of metadata space currently utilized by this server.	GB	<p>Metadata is data that describes other data. Meta is a prefix that in most information technology usages means "an underlying definition or description."</p> <p>Metadata summarizes basic information about data, which can make finding and working with particular instances of data easier. For example, <i>author</i>, <i>date created</i> and <i>date modified</i> and <i>file size</i> are examples of very basic document metadata. Having the ability to filter through that metadata makes it much easier for someone to locate a specific document.</p> <p>In addition to document files, metadata is used for images, videos, spreadsheets and web pages. The use of metadata on web pages can be very important. Metadata for web pages contain descriptions of the page's contents, as well as keywords linked to the content. These are usually expressed in the form of metatags.</p>
	Metadata space total: Indicates the total amount of metadata space allocated to this server.	GB	
	Metadata space utilization: Indicates the percentage of metadata space utilized by this server.	Percent	A low value is desired for this measure.

2.1.3 Docker Events Test

By default, the Docker containers can be created, run, started, stopped, moved, and deleted. These actions are called events in the Docker events API. By constantly monitoring the events triggered by the Docker, administrators can identify how many create events were triggered while creating the containers, how many events were triggered to delete the containers etc. The **Docker Events** test helps administrators to figure out the various events triggered by the Docker server.

This test monitors the Docker server and reports the total number of events triggered on the containers of the Docker server. In addition, this test helps administrators to figure out how many times the start and stop events

were triggered to start and stop the containers. Alongside, administrators can also identify the create events, delete events and die events that were triggered on the server. Using this test, administrators can easily identify how well the containers are managed within the Docker server and proactively be alerted to unpleasant eventualities caused due to the slowdowns in application processing as well as processing bottlenecks.

Purpose	Monitors the Docker server and reports the total number of events triggered on the containers of the Docker server. In addition, this test helps administrators to figure out how many times the start and stop events were triggered to start and stop the containers. Alongside, administrators can also identify the create events, delete events and die events that were triggered on the server.			
Target of the test	A Docker server			
Agent deploying the test	An internal/remote agent			
Configurable parameters for the test	<ol style="list-style-type: none"> 1. TEST PERIOD – How often should the test be executed 2. HOST – Specify the host for which the test is to be configured 3. PORT - The port number at which the specified HOST listens. The default is 2375. 4. USESUDO – By default, this flag is set to Yes indicating either of the following: <ul style="list-style-type: none"> ➤ The eG agent is using the sudo commands to collect the required metrics from the target docker server if monitored in an agent based manner or ➤ The target docker server is monitored in an agentless manner and the eG agent install user is within a docker group. If the target docker server is monitored in an agentless manner and the user privileges are altered in the /etc/sudoers file as mentioned in Section 1.2, then, set this flag to No. 5. DD FREQUENCY – Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is 1:1. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying none against DD FREQUENCY. 6. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option. <p>The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:</p> <ul style="list-style-type: none"> ○ The eG manager license should allow the detailed diagnosis capability ○ Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. 			
Outputs of the test	One set of results for the target Docker server being monitored			
Measurements made by the	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center; width: 33.33%;">Measurement</th> <th style="text-align: center; width: 33.33%;">Measurement Unit</th> <th style="text-align: center; width: 33.33%;">Interpretation</th> </tr> </thead> </table>	Measurement	Measurement Unit	Interpretation
Measurement	Measurement Unit	Interpretation		

test	Total events: Indicates the total number of events that occurred during the last measurement period.	Number	The detailed diagnosis of this measure if enabled, lists the Container ID, Container Name, Image, Event type and Time at which each event occurred.
	Create events: Indicates the number Create events that occurred during the last measurement period.	Number	The detailed diagnosis of this measure if enabled, lists the Container ID, Container Name, Image, Event type and Time at which each event was created.
	Delete events: Indicates the number of Delete events that occurred during the last measurement period.	Number	The detailed diagnosis of this measure if enabled, lists the Container ID, Container Name, Image, Event type and Time at which each event was deleted.
	Start events: Indicates the number of Start events that occurred during the last measurement period.	Number	The detailed diagnosis of this measure if enabled, lists the Container ID, Container Name, Image, Event type and Time at which the events were started.
	Stop events: Indicates the total number of Stop events that occurred during the last measurement period.	Number	The detailed diagnosis of this measure if enabled, lists the Container ID, Container Name, Image, Event type and Time at which the events were stopped.
	Die events: Indicates the number of <i>die events</i> during the last measurement period.	Number	The detailed diagnosis of this measure if enabled, lists the Container ID, Container Name, Image, Event type and Time at which the <i>die events</i> occurred.

The detailed diagnosis of the *Total Events* measure lists the containers on which the events occurred.

Event Time	Container ID	Container Name	Image	Event Type
Oct 20, 2015 18:22:58	48c300292f300e5f24fe401b4139b63375a06cf015fc0dd303f08adfd46a76fe	adoring_mclean	tomcat:8	exec_create:date
2015-10-20T18:19:33.000000000+05:30	48c300292f300e5f24fe401b4139b63375a06cf015fc0dd303f08adfd46a76fe	adoring_mclean	tomcat:8	exec_start:date
2015-10-20T18:19:34.000000000+05:30	e6fab132f60d91c764ae3f976e4de9bcb4a66d65631786fa2373034c94b5a2d1	sick_mayer	ubuntu:14.04	exec_create:date
2015-10-20T18:19:34.000000000+05:30	e6fab132f60d91c764ae3f976e4de9bcb4a66d65631786fa2373034c94b5a2d1	sick_mayer	ubuntu:14.04	exec_start:date

Figure 4: The detailed diagnosis of the Total Events measure

2.1.4 Docker Images Test

Images are the basic building blocks of the Docker, and are organized in a layered manner. The images are utilized as *read-only* templates for building multiple Docker containers using layered Docker filesystems. The containers so created will share common files and enhance disk usage and downloads on the containers. Once the containers are created, multiple applications can be installed on them. The applications can also be updated to a new version by simply building a new layer on the existing images rather than replacing the whole image or entirely rebuilding the containers. The containers can be mapped to a single image or to multiple images. If a single image is alone over-utilized in creating the containers or if any image remains unmapped to the containers, then such unmapped images will remain under-utilized but still will occupy unnecessary disk space. If too many such unmapped images exists, then the creation of new images may be impacted and applications cannot be bundled to the existing images/containers resulting in a decreased reusability of disk space, increased disk usage and reduction in the building speed of the containers. This is why, administrators need to frequently check for the images and the disk space occupied by those images. The **Docker Images** test does this check.

This test reports the total number of images that are available on the Docker host/server. In addition, this test helps administrators to compare the numerical statistics of the images that are mapped to the containers and those that are not mapped to any container. Likewise, administrators can also be able to figure out the disk space utilization of images that are mapped to the containers and the disk space utilization of the images that are not mapped to the containers. Using this test, administrators can figure out the images that are sparsely utilized, the images that are consuming too much of disk space etc and take remedial actions to restrict the disk space utilization of the images.

Purpose	Reports the total number of images that are available on the Docker host/server. In addition, this test helps administrators to compare the numerical statistics of the images that are mapped to the containers and those that are not mapped to any container. Likewise, administrators can also be able to figure out the disk space utilization of images that are mapped to the containers and the disk space utilization of the images that are not mapped to the containers
Target of the test	A Docker server
Agent deploying the test	An internal/remote agent

Configurable parameters for the test	<ol style="list-style-type: none"> 1. TEST PERIOD – How often should the test be executed 2. HOST – Specify the host for which the test is to be configured 3. PORT - The port number at which the specified HOST listens. The default is 2375. 4. USESUDO – By default, this flag is set to Yes indicating either of the following: <ul style="list-style-type: none"> ➤ The eG agent is using the sudo commands to collect the required metrics from the target docker server if monitored in an agent based manner or ➤ The target docker server is monitored in an agentless manner and the eG agent install user is within a docker group. <p>If the target docker server is monitored in an agentless manner and the user privileges are altered in the /etc/sudoers file as mentioned in Section 1.2, then, set this flag to No.</p> 5. DD FREQUENCY – Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is 1:1. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying none against DD FREQUENCY. 6. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option. <p>The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:</p> <ul style="list-style-type: none"> ○ The eG manager license should allow the detailed diagnosis capability ○ Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. 						
Outputs of the test	One set of results for the target Docker server that is being monitored						
Measurements made by the test	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center; width: 33.33%;">Measurement</th> <th style="text-align: center; width: 33.33%;">Measurement Unit</th> <th style="text-align: center; width: 33.33%;">Interpretation</th> </tr> </thead> <tbody> <tr> <td style="padding: 5px;">Total Images: Indicates the total number of images available in the Docker server.</td> <td style="padding: 5px; text-align: center;">Number</td> <td style="padding: 5px;"></td> </tr> </tbody> </table>	Measurement	Measurement Unit	Interpretation	Total Images: Indicates the total number of images available in the Docker server.	Number	
Measurement	Measurement Unit	Interpretation					
Total Images: Indicates the total number of images available in the Docker server.	Number						

	Images mapped to containers: Indicates the number of images that are mapped to the Docker containers.	Number	Ideally, the value of this measure should be high. A low value of this measure indicates that more number of images are left idle and more disk space is occupied by unused images. This may cause potential space crunch in the disk. The detailed diagnosis of this measure if enabled, lists the images and the containers to which the images are mapped.
	Images not mapped to any container: Indicates the number of images that are not mapped to any Docker container.	Number	Ideally, the value of this measure should be low. The detailed diagnosis of this measure if enabled, lists the images that are not mapped to any Docker containers.
	Disk space used by images mapped to containers: Indicates the amount of disk space utilized by the images that are mapped to the Docker containers.	MB	The detailed diagnosis of this measure if enabled, lists the images and the containers.
	Disk space used by images not mapped to any container: Indicates the amount of disk space utilized by the images that were not mapped to the Docker containers.	MB	A high value of this measure indicates space crunch in the disk. The detailed diagnosis of this measure if enabled, lists the images, which were not mapped with containers.

The detailed diagnosis of the Images mapped to containers measure lists the Image, Image ID, the containers mapped to the image, the time at which the container was created and size of the image. Using this measure administrators can figure out the image that is widely used to create a container.

IMAGE NAME	IMAGE ID	CONTAINERS NAME	CREATED	SIZE(MB)
debian:8	4a5e6db8c069	grave_kowalevski	9 weeks ago	125.2
ubuntu:14.04	91e54dfb1179	sick_mayer,boring_poincare,insane_einstein,distracted_ardinghelli	9 weeks ago	188.4
tomcat:8	6fe1972e6b08	adoring_mclean	8 weeks ago	347.7

Figure 5: The detailed diagnosis of the Images mapped to containers measure

2.1.5 Docker Images - Performance Test

A Docker server can have multiple images using which containers are created. In large environments where the requirements of users vary from minute to minute, it is important for an administrator to cater to their specific needs by providing the images of their choice so that the users can create the containers and load applications accordingly. Often administrators may be confused while identifying the needs of the users. An image that is needed today may not be needed tomorrow to create a container. In such cases, the images may be unutilized but still hold on valuable disk space. To optimize the disk space and eliminate unused images, administrators are required to analyze the performance of the images as well as identify the images that are most commonly used by the users. The **Docker Images - Performance** test helps administrators in this regard!

This test monitors the images of the Docker server and reports the number of containers created using each image. By analyzing the number of running containers created from each image and the resource utilization of those containers, administrators can figure out the image that is widely used by the users to create the containers and analyze how well the resources are utilized by those containers.

Purpose	Monitors the images of the Docker server and reports the number of containers created using each image. By analyzing the number of running containers created from each image and the resource utilization of those containers, administrators can figure out the image that is widely used by the users to create the containers and analyze how well the resources are utilized by those containers
Target of the test	A Docker server
Agent deploying the test	An internal/remote agent

Configurable parameters for the test	<ol style="list-style-type: none"> 1. TEST PERIOD – How often should the test be executed 2. HOST – Specify the host for which the test is to be configured 3. PORT - The port number at which the specified HOST listens. The default is 2375. 4. USESUDO – By default, this flag is set to Yes indicating either of the following: <ul style="list-style-type: none"> ➤ The eG agent is using the sudo commands to collect the required metrics from the target docker server if monitored in an agent based manner or ➤ The target docker server is monitored in an agentless manner and the eG agent install user is within a docker group. <p>If the target docker server is monitored in an agentless manner and the user privileges are altered in the /etc/sudoers file as mentioned in Section 1.2, then, set this flag to No.</p> 5. DOCKER USER and DOCKER PASSWORD – Specify the credentials of the user who is authorized to execute the docker commands on the target server. 6. CONFIRM PASSWORD – Confirm the DOCKER PASSWORD by retyping it here. 7. WEBPORT – By default, the Docker server listens on 2375. This implies that while monitoring a Docker server via an SSL-enabled environment, the eG agent, by default, connects to port 2375 of the Docker server to pull out metrics. In some environments however, this default port may not apply. In such a case, against the WEBPORT parameter, you can specify the exact port at which the Docker server in your environment listens so that the eG agent communicates with that port. 8. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option. <p>The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:</p> <ul style="list-style-type: none"> ○ The eG manager license should allow the detailed diagnosis capability ○ Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. 						
Outputs of the test	One set of results for each image available on the target Docker server that is being monitored						
Measurements made by the test	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center; width: 33.33%;">Measurement</th><th style="text-align: center; width: 33.33%;">Measurement Unit</th><th style="text-align: center; width: 33.33%;">Interpretation</th></tr> </thead> <tbody> <tr> <td style="vertical-align: top;"> Total containers: Indicates the total number of containers that were mapped to this image. </td><td style="text-align: center; vertical-align: top;"> Number </td><td style="vertical-align: top;"> The detailed diagnosis of this measure if enabled, lists the name of the containers. Compare the value of this measure across the images to identify the image to which maximum number of containers are mapped. </td></tr> </tbody> </table>	Measurement	Measurement Unit	Interpretation	Total containers: Indicates the total number of containers that were mapped to this image.	Number	The detailed diagnosis of this measure if enabled, lists the name of the containers. Compare the value of this measure across the images to identify the image to which maximum number of containers are mapped.
Measurement	Measurement Unit	Interpretation					
Total containers: Indicates the total number of containers that were mapped to this image.	Number	The detailed diagnosis of this measure if enabled, lists the name of the containers. Compare the value of this measure across the images to identify the image to which maximum number of containers are mapped.					

	Running containers: Indicates the number of containers that are created from this image and are currently in <i>Running</i> state.	Number	Ideally, the value of this measure should be high. A low value for this measure is a cause of concern. The detailed diagnosis of this measure if enabled, lists the containers that are currently in <i>Running</i> state and the resource utilization of the containers.
	CPU utilization of containers: Indicates the percentage of CPU utilized by the containers (created from this image) that are currently in <i>Running</i> state.	Percent	This measure is only applicable for the Docker server 1.5 and above. A sudden increase or decrease in the value of this measure could be a cause of concern.
	Memory used by containers: Indicates the amount of memory utilized by the containers (created from this image) that are currently in <i>Running</i> state.	MB	This measure is only applicable for Docker server 1.5 and above. A low value is desired for this measure. A high value of this measure indicates that the Docker is running out of space.
	Memory utilization of containers: Indicates the percentage of memory utilized by the containers (created from this image) that are currently in <i>Running</i> state.	Percent	This measure is only applicable for the Docker server 1.5 and above. A value close to 100% indicates that the memory is running out of space.
	Disk space used by image: Indicates the amount of disk space utilized by this image.	MB	Ideally, the value of this measure should be low. A significant increase in the value of this measure could be a cause of concern.

The detailed diagnosis of the *Total containers* measure lists the containers that were mapped to an image.

Detailed Diagnosis Measure Graph Summary Graph Trend Graph Fix History Fix Feedback

Component	Test	Measured By	Descriptor	Measurement
Docker_10_169:2375:Docker	Docker Images - Perfo	Docker_10_169	ubuntu:14.04	Total containers

Timeline

Latest Submit

Details of containers mapped to image

CONTAINER NAME

Oct 23, 2015 17:31:21

sick_mayer
boring_poincare
insane_einstein
distracted_ardinghelli

Figure 6: The detailed diagnosis of the Total containers measure

The detailed diagnosis of the *Running containers* lists the containers that are currently running and the split up resource utilization (CPU, memory and memory usage) for each of the listed containers.

Detailed Diagnosis Measure Graph Summary Graph Trend Graph Fix History Fix Feedback

Component	Test	Measured By	Descriptor	Measurement
Docker_10_169:2375:Docker	Docker Images - Perfo	Docker_10_169	ubuntu:14.04	Running containers

Timeline

Latest Submit

Details of running containers

CONTAINER NAME %CPU MEMORY USED %MEM

Oct 23, 2015 17:31:21

sick_mayer	0	0.5	0.0128
------------	---	-----	--------

Figure 7: The detailed diagnosis of the Running containers measure

2.2 The Docker Containers Layer

This layer monitors the docker containers and reports the performance of each container, the uptime of the container, the network connectivity to each container and the resource ad memory utilization of each container etc.

Docker Containers

Docker Containers - Status

Docker Containers - Uptime

cadvisor

dd-agent

Figure 1.1: The tests mapped to the Docker Containers layer

2.2.1 Docker Containers - Status Test

Docker containers can be built within seconds based on a Docker image. In large environments where multiple containers are created simultaneously, administrators may actually want to keep track on the containers that are currently running, the containers that are currently stopped, and the containers that were removed etc so that containers that are sparsely used and containers that crashed can be identified easily. The **Docker Containers - Status** test helps administrators to optimize the containers and enhance the overall performance of the Docker server.

This test monitors the containers available in the Docker server and reports the numerical statistics of the total containers, the containers in *Running* state, the containers that were added, the containers that were removed etc. In addition, this test reports the containers that are not running for a time duration above the specified limit and the containers that are utilizing disk space above the specified limit. By analyzing this, administrators can figure out the overall performance of the Docker server and identify bottlenecks if any.

Purpose	Monitors the containers available in the Docker server and reports the numerical statistics of the total containers, the containers in Running state, the containers that were added, the containers that were removed etc. In addition, this test reports the containers that are not running for a time duration above the specified limit and the containers that are utilizing disk space above the specified limit
Target of the test	A Docker server
Agent deploying the test	An internal/remote agent

Configurable parameters for the test	<ol style="list-style-type: none"> 1. TEST PERIOD – How often should the test be executed 2. HOST – Specify the host for which the test is to be configured 3. PORT - The port number at which the specified HOST listens. The default is 2375. 4. USESUDO – By default, this flag is set to Yes indicating either of the following: <ul style="list-style-type: none"> ➤ The eG agent is using the sudo commands to collect the required metrics from the target docker server if monitored in an agent based manner or ➤ The target docker server is monitored in an agentless manner and the eG agent install user is within a docker group. <p>If the target docker server is monitored in an agentless manner and the user privileges are altered in the /etc/sudoers file as mentioned in Section 1.2, then, set this flag to No.</p> 5. SIZE LIMIT IN MB – By default, the disk space utilized by each container may vary according to the size of the container. Sometimes, the containers may be over-utilizing the disk space which when left unnoticed may hamper the creation of new containers. To figure out such containers, you can specify a disk space limit beyond which the containers can be termed as <i>large sized container</i>. For example, if you wish to monitor the number of containers that are utilizing disk space above 50 MB, then specify 50 against this text box. 6. TIME LIMIT IN WEEKS- For this test to report the numerical statistics of the containers that were not started/running, set a valid value against this parameter. For example, if you wish to report the containers that were not started for more than 3 weeks, then set 3 against this text box. 7. DD FREQUENCY – Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is 1:1. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying none against DD FREQUENCY. 8. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option. <p>The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:</p> <ul style="list-style-type: none"> ○ The eG manager license should allow the detailed diagnosis capability ○ Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. 		
Outputs of the test	One set of results for the Docker server being monitored		
Measurements made by the	Measurement	Measurement Unit	Interpretation

test	Total containers: Indicates the total number of containers in this server.	Number	The detailed diagnosis of this measure if enabled, lists the name of the container, the container ID, the image used to create the container, the container creation time, the status of the container and the size of the container.
	Running containers: Indicates the number of containers that are in <i>Running</i> state.	Number	A high value of this measure indicates that more number of containers are in running state as a result more number of applications can be accessed actively. To know which containers are currently running, use the detailed diagnosis capability of this measure.
	Stopped containers: Indicates the number of containers that are currently not running in this server.	Number	To know more about the containers that are currently not running, use the detailed diagnosis capability of this measure (if enabled).
	Added containers: Indicates the number of containers that were created during the last measurement period.	Number	The detailed diagnosis of this measure, if enabled, lists the names of the containers that were created.
	Removed containers: Indicates the number of containers that were removed from this server during the last measurement period.	Number	The detailed diagnosis of this measure, if enabled, lists the containers that were removed. If too many containers are removed, then the users may not be able to access the hosted applications thus leading to a performance bottleneck of the Docker server. Administrators are therefore required to keep a check on the containers that are removed from the Docker server.
	Large size containers: Indicates the number of containers that were using more disk space than the limit configured against the SIZE LIMIT IN MB parameter.	Number	A low value is desired for this measure. The detailed diagnosis of this measure if enabled, lists the names of the containers that are over-utilizing the disk space.
	Containers not started for long time: Indicates the number of containers that were not running for more than the configured TIME LIMIT IN WEEKS .	Number	Ideally, the value of this measure should be low. The detailed diagnosis of this measure if enabled, lists the containers that are not running for longer time.

The detailed diagnosis of the *Total Containers* measure lists all the containers that were created and the images that were used to create those containers.




Figure 8: The detailed diagnosis of the Total Containers measure

2.2.2 Docker Containers - Uptime Test

In environments where Docker server is used extensively, it is essential to monitor the uptime of critical containers within the Docker server. By tracking the uptime of each of the containers, administrators can determine what percentage of time a container has been up. Comparing this value with service level targets, administrators can determine the most trouble-prone areas of the infrastructure.

In some environments, administrators may schedule periodic reboots of their containers. By knowing that a specific container has been up for an unusually long time, an administrator may come to know that the scheduled reboot task is not working on a container.

The **Docker Containers - Uptime** test included in the eG agent monitors the uptime of critical containers in the target Docker server.

Purpose	monitors the uptime of critical containers in the target Docker server
Target of the test	A Docker server
Agent deploying the test	An internal/remote agent

Configurable parameters for the test	<ol style="list-style-type: none"> 1. TEST PERIOD – How often should the test be executed 2. HOST – Specify the host for which the test is to be configured 3. PORT - The port number at which the specified HOST listens. The default is 2375. 4. USESUDO – By default, this flag is set to Yes indicating either of the following: <ul style="list-style-type: none"> ➤ The eG agent is using the sudo commands to collect the required metrics from the target docker server if monitored in an agent based manner or ➤ The target docker server is monitored in an agentless manner and the eG agent install user is within a docker group. <p>If the target docker server is monitored in an agentless manner and the user privileges are altered in the /etc/sudoers file as mentioned in Section 1.2, then, set this flag to No.</p> 5. DD FREQUENCY – Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is 1:1. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying none against DD FREQUENCY. 6. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option. <p>The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:</p> <ul style="list-style-type: none"> ○ The eG manager license should allow the detailed diagnosis capability ○ Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. 		
Outputs of the test	One set of results for each container available in the Docker server being monitored		
Measurements made by the	Measurement	Measurement Unit	Interpretation

test	<p>Has the container been restarted?: Indicates whether/not this container was rebooted.</p>		<p>The values reported by this measure and their numeric equivalents are available in the table below:</p> <table border="1" data-bbox="1000 302 1449 487"> <thead> <tr> <th>Measure Value</th><th>Numeric Value</th></tr> </thead> <tbody> <tr> <td>No</td><td>0</td></tr> <tr> <td>Yes</td><td>1</td></tr> </tbody> </table> <p>Note: This measure reports the Measure Values listed in the table above while indicating whether /not this container was rebooted. However, in the graph of this measure, the measure is indicated using only the Numeric Values listed in the above table.</p>	Measure Value	Numeric Value	No	0	Yes	1
Measure Value	Numeric Value								
No	0								
Yes	1								
	<p>Uptime of the container during the last measure period: Indicates the time duration for which this container has been up since the last time this test ran</p>	Secs	<p>If this container has not been rebooted during the last measurement period and the agent has been running continuously, this value will be equal to the measurement period. If this container was rebooted during the last measurement period, this value will be less than the measurement period of the test. The accuracy of this metric is dependent on the measurement period - the smaller the measurement period, greater the accuracy.</p>						
	<p>Total uptime of the container: Indicates the total time that this container has been up since its last reboot.</p>	Mins	<p>Administrators may wish to be alerted if a container has been running without a reboot for a very long period. Setting a threshold for this metric allows administrators to determine such conditions.</p>						

2.2.3 Docker Containers - Connectivity Test

When the Docker starts, it creates a virtual interface called “Ethernet bridge” on the Docker server. The Ethernet bridge automatically forwards packets to/from the Docker to the external network interfaces. Every time, the Docker creates the containers, it creates a pair of “peer” interfaces that are like opposite ends of a pipe. The packets sent through one of the peer interfaces will be received by the other peer interface. By binding the peer interfaces to the Ethernet bridge, the Docker creates a virtual subnet shared between the Docker server and every Docker container so as to ensure uninterrupted communication between the Docker server and the Docker containers. For a connection to be uninterrupted, it is essential to keep track of the network delay as well as the packet loss. The **Docker Containers – Connectivity** test helps administrators to keep track of such delays and loss.

This test monitors the network connectivity to the Docker containers from an external location.

Purpose	Monitors the network connectivity to the containers from an external location
Target of the test	A Docker server

Agent deploying the test	An internal/remote agent		
Configurable parameters for the test	<ol style="list-style-type: none"> TEST PERIOD – How often should the test be executed HOST – Specify the host for which the test is to be configured PORT - The port number at which the specified HOST listens. The default is 2375. USESUDO – By default, this flag is set to Yes indicating either of the following: <ul style="list-style-type: none"> ➤ The eG agent is using the sudo commands to collect the required metrics from the target docker server if monitored in an agent based manner or ➤ The target docker server is monitored in an agentless manner and the eG agent install user is within a docker group. If the target docker server is monitored in an agentless manner and the user privileges are altered in the /etc/sudoers file as mentioned in Section 1.2, then, set this flag to No. PACKETSIZE – Here, specify the size of the packets used for the test. PACKETCOUNT – Here, specify the number of packets to be transmitted during the test. TIMEOUT – Specify the maximum time (in seconds) that the test should wait for the response to a transmitted packet. A response received after the TIMEOUT period is ignored by the test. The default timeout period is 100 seconds. PACKETINTERVAL – Represents the interval (in milliseconds) between successive packet transmissions during the execution of the test for a specific target. 		
Outputs of the test	One set of results for each container available in the target Docker server being monitored		
Measurements made by the test	Measurement	Measurement Unit	Interpretation
	Network availability of container: Indicates whether/not the network connection to this container is available.	Percent	A value of 100 indicates that this container is accessible over the network. The value 0 indicates that the container is inaccessible. Typically, the value 100 corresponds to a <i>Packet loss</i> of 0.
	Avg network delay: Indicates the average delay between transmission of packets to this container and receipt of the response to the packet at the source.	Secs	Ideally, the value of this measure should be low.
	Min network delay: The minimum time between transmission of a packet and receipt of the response back.	Secs	A significant increase in the minimum round-trip time is often a sure sign of network congestion.

	Packet loss: Indicates the percentage of packets lost during transmission from source to target and back.	Percent	Packet loss is often caused by network buffer overflows at a network router or by packet corruptions over the network. The detailed diagnosis for this measure provides a listing of routers that are on the path from the external agent to target server, and the delays on each hop. This information can be used to diagnose the hop(s) that could be causing excessive packet loss/delays.
--	---------------------------------------------------------------------------------------------------------------------	---------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

2.2.4 Docker Containers - Performance Test

This test monitors each container available in the Docker and reports the CPU utilization, I/O processing, memory related statistics such as memory utilization, paging in/paging outs, errors that were detected etc. Using this test, administrators can easily figure out processing/memory bottlenecks and rectify the same before the users complain of slow responsiveness of the containers.

Note:

This test will be supported only for the Docker server version 1.5 and above.

Purpose	Monitors each container available in the Docker and reports the CPU utilization, I/O processing, memory related statistics such as memory utilization, paging in/paging outs, errors that were detected etc
Target of the test	A Docker server
Agent deploying the test	An internal/remote agent

Configurable parameters for the test	<ol style="list-style-type: none"> 1. TEST PERIOD – How often should the test be executed 2. HOST – Specify the host for which the test is to be configured 3. PORT - The port number at which the specified HOST listens. The default is 2375. 4. USESUDO – By default, this flag is set to Yes indicating either of the following: <ul style="list-style-type: none"> ➤ The eG agent is using the sudo commands to collect the required metrics from the target docker server if monitored in an agent based manner or ➤ The target docker server is monitored in an agentless manner and the eG agent install user is within a docker group. <p>If the target docker server is monitored in an agentless manner and the user privileges are altered in the /etc/sudoers file as mentioned in Section 1.2, then, set this flag to No.</p> 5. DOCKER USER and DOCKER PASSWORD – Specify the credentials of the user who is authorized to execute the docker commands on the target server 6. CONFIRM PASSWORD – Confirm the DOCKER PASSWORD by retyping it here. 7. WEBPORT – By default, the Docker server listens on 2375. This implies that while monitoring a Docker server via an SSL-enabled environment, the eG agent, by default, connects to port 2375 of the Docker server to pull out metrics. In some environments however, this default port may not apply. In such a case, against the WEBPORT parameter, you can specify the exact port at which the Docker server in your environment listens so that the eG agent communicates with that port. 8. DETAILED DIAGNOSIS - To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the On option. To disable the capability, click on the Off option. <p>The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:</p> <ul style="list-style-type: none"> ○ The eG manager license should allow the detailed diagnosis capability ○ Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0. 		
Outputs of the test	One set of results for each container available in the Docker server being monitored		
Measurements made by the test	Measurement	Measurement Unit	Interpretation
	Disk space usage: Indicates the amount of disk space utilized by this container.	MB	The detailed diagnosis of this measure if enabled, lists the containers in this server.
	Data received rate: Indicates the rate at which the data was received by this container.	Mbps	A sudden increase or decrease in the value of this measure could be a cause of concern.

	Incoming traffic: Indicates the rate at which packets were received by this container.	Pkts/sec	A significant increase or decrease in the value of this measure may alter traffic condition in the Docker server.
	Errors received: Indicates the number of errors occurred while data was received by this container.	Number	Ideally, the value of this measure should be zero.
	Packets dropped during reception: Indicates the number of packets dropped by this container during reception.	Number	Ideally, the value of this measure should be zero.
	Data transmit rate: Indicates the rate at which the data was transmitted by this container.	Mbps	
	Outgoing traffic: Indicates the rate at which packets were transmitted by this container.	Pkts/sec	
	Errors transmitted: Indicates the number of errors occurred while data was transmitted by this container.	Number	Ideally, the value of this measure should be zero.
	Packets dropped during transmission: Indicates the number of packets dropped by this container during transmission.	Number	Ideally, the value of this measure should be zero.
	CPU utilization: Indicates the percentage of CPU that is currently utilized by this container.	Percent	Comparing the value of this measure across the containers will enable you to accurately identify the container on which CPU-intensive applications are executing.

	<p>CPU utilization in kernelmode: Indicates the percentage of CPU utilized by this container in kernel mode.</p>	Percent	<p>A processor in a server has two different modes: user mode and kernel mode. The processor switches between the two modes depending on what type of code is running on the processor. Applications run in user mode, and core operating system components run in kernel mode.</p> <p>In Kernel mode, the executing code has complete and unrestricted access to the underlying hardware. It can execute any CPU instruction and reference any memory address. Kernel mode is generally reserved for the lowest-level, most trusted functions of the operating system. Crashes in kernel mode are catastrophic; hence they will halt the entire container.</p> <p>A high value for this measure indicates that the container is taking too much of CPU resources to execute processes in kernel mode.</p>
	<p>CPU utilization in usermode: Indicates the percentage of CPU utilized by this container in user mode.</p>	Percent	<p>In User mode, the executing code has no ability to directly access hardware or reference memory. Code running in user mode must delegate to system APIs to access hardware or memory. Due to the protection afforded by this sort of isolation, crashes in user mode are always recoverable. Most of the code running on the containers will execute in user mode. User mode processes communicate and use the kernel through the Kernel API and system calls.</p>
	<p>Memory used: Indicates the amount of memory that is currently utilized by this container.</p>	MB	
	<p>Memory limit: Indicates the maximum amount of memory that is allocated to this container.</p>	MB	
	<p>Memory utilization: Indicates the percentage of memory utilized by this container.</p>	Percent	<p>A high value for this measure indicates that the memory resources of the container is depleting alarmingly.</p>
	<p>Memory max usage: Indicates the maximum amount of memory utilized by this container.</p>	MB	

	Active anonymous memory : Indicates the amount of anonymous memory that has been identified as active by the kernel.	MB	Anonymous memory is the large, zero-filled block of memory that is directly mapped by kernel from the anonymous memory region when large memory with ideally multiples of page sizes is required for the containers. The pages of anonymous memory do not link to any file on the disk, and are the part of program's data area or stack.
	Inactive anonymous memory: Indicates the amount of anonymous memory that has been identified as inactive by the kernel.	MB	All anonymous pages are initially active. In that, some pages will be tagged as inactive when the kernel sweeps over the memory at regular intervals. Whenever the inactive pages are accessed, they are immediately retagged as active. The inactive pages will be swapped when the kernel is almost out of memory, and time comes to swap out to disk.
	Cache memory: Indicates the amount of memory used by processes of the control group that can be associated precisely with a storage block on this container.	MB	When you read from and write to files on the disk, the amount of cache memory will increase. Size of the cache memory depends on the number of read operations and write operations performed on this container.
	Active file: Indicates the amount of cache memory that has been identified as active by the kernel.	MB	Pages in the cache memory can be swapped between active and inactive states similar to the anonymous memory but the exact rules used by the kernel to move memory pages between active and inactive sets are different from the rules used for the anonymous memory. The cache memory pages can be immediately retrieved in a cheaper way when the kernel needs to reclaim memory while the anonymous pages and dirty/modified pages have to be written to disk first before retrieving process.
	Inactive file: Indicates the amount of cache memory that has been identified as inactive by the kernel.	MB	Docker on Linux also makes use of another technology called <code>cgroups</code> or control groups. A key to running applications in isolation is to have them only use the resources you want. This ensures containers are good multi-tenant citizens on a host. Control groups allow Docker to share available hardware resources to containers and, if required, set up limits and constraints. For example, limiting the memory available to a specific container.
	Memory mapped: Indicates the amount of memory mapped by the processes in the control group.	MB	

	Page faults: Indicates the number of times a process in this container triggered a page fault.	Number	A page fault occurs when a process accesses a part of its virtual memory space which is nonexistent or protected. The former can happen if the process is buggy and tries to access an invalid address (it will then be sent a SIGSEGV signal, typically killing it with the famous Segmentation fault message). The latter can happen when the process reads from a memory zone which has been swapped out, or which corresponds to a mapped file: in that case, the kernel will load the page from disk, and let the CPU complete the memory access. It can also happen when the process writes to a copy-on-write memory zone: likewise, the kernel will preempt the process, duplicate the memory page, and resume the write operation on the process' own copy of the page.
	Page major faults: Indicates the number of times a process in this container triggered a major page fault.	Number	The major page faults are regular faults occur when the kernel actually has to read the data from disk, duplicate an existing page, or allocate an empty page.
	Page ins: Indicates the number of pages that were added to the control group of this container.	Number	
	Page outs: Indicates the number of pages that were not billed to the control group of this container.	Number	
	Resident set size: Indicates the amount of memory that doesn't correspond to anything on disk such as stacks and heaps, and anonymous memory maps of this container.	MB	
	Huge resident set size: Indicates the amount of maximum memory that doesn't correspond to anything on disk such as stacks and heaps, and anonymous memory maps of this container.	MB	

	Swap memory: Indicates the amount of swap memory that is currently utilized by the processes in the control group of this container.	MB	An unusually high value for the swap memory can indicate a memory bottleneck.
	Memory unevictable: Indicates the amount of memory that cannot be reclaimed by this container.	MB	Generally, the unevictable memory has been locked with mlock, and is often used by crypto frameworks to make sure that secret keys and other sensitive material never gets swapped out to disk.
	Writeback: Indicates the amount of memory that was written back in this container.	MB	A high value is preferred for this measure. Normally, the data are first written into the cache before writing it into the memory or disk that supports caching. During idle machine cycles, the data are written from the cache into the memory or onto the disk at high speed. In this way, the Write back caches improve performance and writing speed than the normal RAM or disk.
	Failures: Indicates the number of memory failures that occurred in this container.	Number	Ideally, the value of this measure should be low.
	Data reads: Indicates the rate at which the data was read from this container.	Mbps	Compare the values of these measures across the containers to identify the slowest container in terms of processing read and write operations (respectively).
	Data writes: Indicates the rate at which the data were written into this container.	Mbps	
	Data sync: Indicates the rate at which the synchronous I/O operations were performed on this container.	Mbps	In synchronous I/O file, a thread starts a I/O operation and immediately enters a <i>wait</i> state until the I/O request has completed such that the I/O operations of this container are performed in one-by-one manner so as to prevent unwanted data traffic.
	Data async: Indicates the rate at which the asynchronous I/O operations were performed on this container.	Mbps	A thread performing asynchronous file I/O sends an I/O request to the kernel by calling an appropriate function. If the request is accepted by the kernel, the calling thread continues processing another job until the kernel signals to the thread that the I/O operation is complete. Therefore, speed of the I/O operations will be increased.

MONITORING THE DOCKER

	Total data: Indicates the rate at which the total I/O operations were performed on this container.	Mbps	
--	--------------------------------------------------------------------------------------------------------------	------	--

Conclusion

This document has clearly explained how eG Enterprise monitors the **Docker** server. We can thus conclude that eG Enterprise, with its ability to provide in-depth insight into the performance of SAN storage infrastructures, is the ideal solution for monitoring such environments. For more information on eG Enterprise, please visit our web site at www.eginnovations.com or write to us at sales@eginnovations.com.