

GANGLIA INSTALLATION GUIDE

-Prashant Bhamidipati
(prashant@hepmail.uta.edu)

Section-1 : Ganglia installation procedure
Section-2 : GMOND & GMETAD Configuration

Section – 1 Ganglia installation procedure

Ganglia is a complete and real-time monitoring and execution system intended to provide system statistics and important performance metrics for cluster computing systems. It is capable of providing excellent support for clusters as small as a few computers to those as large as 512 systems. One noteworthy feature is that it uses widely popular tools such as RRDtools, XRD, XML and PHP enabled web interfaces for its operation.

The ganglia system basically comprises of three sections:

- The GMOND or Ganglia Monitoring Daemon
The GMOND runs on every machine needing to be monitored
- The GMETAD or the Ganglia Meta Daemon
The GMETAD runs on the main machine of the cluster which is the web server / or is the main machine communicating with the web server
- The Ganglia Web Front-end
The ganglia web front-end is the system that runs the web page for the entire monitoring system. It needs the presence of GMETAD for successful running.

This document intends to provide concise installation instructions for ganglia. Detailed instructions may be found at <http://ganglia.sourceforge.net>. Users of this document are encouraged to signup on their mailing list as it gives an exposure to a wide variety of problems and ways to tackle them.

Important points to note:-

1. In a given cluster, only one machine will run the GMETAD daemon. The others will run only the GMOND. This machine may run the GMOND depending on whether it needs to be monitored in the cluster. Generally, it does run the GMOND.
2. The GMETAD is the machine that aggregates data from the various nodes of the cluster and hence will need to be equipped with the RRDtools.
3. There will be one machine in the cluster that will be the web server. It is mandatory that the GMETAD be running on this machine also. This means that RRDtools also has to be installed here.
4. The machine acting as the web-server will need to be PHP enabled also.

It is recommended that installation be done as **root**.

1.1 Installation of the RRDtools

Here, the procedure is explained for installation from RPMs The RRDtools package is required on every machine running the GMETAD since it supports the functions of the GMETAD by providing a way to efficiently store the data collected.

1. Download the developer package "rrdtool-devel-1.0.38-1.i386.rpm" from the downloads section of the RRDtools web page. i.e.: (<http://www.rrdtool.com/download.html>)
2. Install the developer package in the location decided by the system administrator.
`Your_prompt>rpm-Uvh rrdtool-devel-1.0.38-1.i386.rpm`
3. Download the regular release "rrdtool-1.0.38-1.i386.rpm" from the RRDtools web page <http://www.rrdtool.com/download.html>
4. Install the regular package using
`Your_prompt>rpm-Uvh rrdtool-1.0.38-1.i386.rpm`

This completes the RRD tool installation. On completion, please make sure that the following two files are present in appropriate locations

- `librrd.a` in `/usr/lib/librrd.a`
- `rrd.h` in `/usr/include/rrd.h`

These two files are important for data storage and their absence will result in no data being registered in the main machine.

1.2 Installation of the GMOND from source

The installation procedure explained for installation from source. If RPMs are a preferred means of installation, please refer the main web page. The GMOND has to be running on every machine in the cluster that needs to be monitored.

1. Download the latest tarball (`ganglia-monitor-core-2.5.1.tar.gz`) from <http://ganglia.sourceforge.net/downloads.php> into the location where it is to be installed.

2. Unzip the distribution

```
Your_prompt>gunzip ganglia-monitor-core-2.5.1.tar.gz
```

```
Your_prompt>tar -xvf ganglia-monitor-core-2.5.1.tar
```

3. Move to the distribution directory

```
Your_prompt>cd ganglia-monitor-core-2.5.1/
```

4. Run the configuration script

```
Your_prompt>./configure
```

5. Run the make file

```
Your_prompt>make
```

6. Once all is built, install everything by

```
Your_prompt>make install
```

Hereon, the instructions given are Linux specific. They ensure that the daemon is started at boot up.

7. Copy the `gmond.init` file to `/etc/rc.d/init.d/` (Commands assume that the PWD is the distribution directory)

```
Your_prompt>cp gmond/gmond.init /etc/rc.d/init.d/gmond
```

8. Copy the `gmond.conf` file to `/etc/`

```
Your_prompt>cp gmond/gmond.conf /etc/gmond.conf
```

9. Add GMOND to the list of programs at startup and check if it responds well. Listed here is the sequence of the outputs

```
Your_prompt>chkconfig --add gmond
```

```
Your_prompt>chkconfig --list gmond
```

```
GMOND    0:off    1:off    2:on     3:on     4:on     5:on     6:off
```

10. Start GMOND and observe if the following lines are output by Linux

```
Your_prompt> /etc/rc.d/init.d/gmond start
Starting GANGLIA gmond: [ OK ]
```

11. Test the GMOND installation by running

```
Your_prompt> telnet localhost 8649 | grep "hostname"
```

where 'hostname' can be the IP address of any of the hosts in the cluster with which this host may be communicating. The output should show a listing of the GMOND hosts which have the hostname specified in 'hostname' and which are present in the cluster being monitored

The procedure listed above is for the installation of the GMOND on member machines of the cluster.

1.3 Installation of the GMETAD from source

The installation procedure explained for installation from source. If RPMs are a preferred means of installation, please refer the main web page. These steps are pretty similar to the previous installation of the GMOND. However, there are a few changes in a few steps in the beginning that need to be noted.

1. Download the latest tarball (ganglia-monitor-core-2.5.1.tar.gz) from <http://ganglia.sourceforge.net/downloads.php> into the location where it is to be installed.

1. Unzip the distribution in the preferred directory

```
Your_prompt>gunzip ganglia-monitor-core-2.5.1.tar.gz
```

```
Your_prompt>tar -xvf ganglia-monitor-core-2.5.1.tar
```

2. Move to the distribution directory

```
Your_prompt>cd ganglia-monitor-core-2.5.1/
```

3. Run the configuration script

```
Your_prompt>./configure --with-gmetad
```

4. Run the make file

```
Your_prompt>make
```

5. Install GMOND by running with the following command

```
Your_prompt>make install
```

Hereon, the instructions given are Linux specific. These instructions ensure that the daemon is started at boot up.

6. Copy the gmetad.init file from the ganglia distribution directory to /etc/rc.d/init.d/ (Commands assume that the PWD is the distribution directory)

```
Your_prompt>cp gmetad/gmetad.init /etc/rc.d/init.d/gmetad
```

7. Copy the gmetad.conf file to /etc/ from the Ganglia distribution directory

```
Your_prompt>cp GMETAD/gmetad.conf /etc/gmetad.conf
```

8. Make sure of the existence of the 'rrd' folders. IF the RRDs have been installed but the folders haven't. This should not occur normally, but there have been instances of this

happening. As mentioned earlier, this step is extremely important because they assist the GMETAD with data storage.

```
Your_prompt> mkdir -p /var/lib/ganglia/rrds
Your_prompt> mkdir -p /var/lib/ganglia/rrds
Your_prompt> chown -R nobody /var/lib/ganglia/rrds
```

9. Add GMETAD to the list of programs at startup and check if it responds well. Listed here is the sequence of your outputs

```
Your_prompt> chkconfig --add gmetad
Your_prompt> chkconfig --list gmetad
GMETAD    0:off    1:off    2:on    3:on    4:on    5:on    6:off
```

10. Start GMETAD and observe if the following lines are output by Linux

```
Your_prompt> /etc/rc.d/init.d/gmetad start
Starting GANGLIA gmetad:                [ OK ]
```

11. Test the GMETAD installation by running

```
Your_prompt> telnet localhost 8651 | grep "hostname"
```

where 'hostname' can be the IP address of any of the hosts in the cluster with which this host may be communicating. The output should show a listing of the GMETAD hosts which have the hostname specified and which are present in the cluster being monitored

The procedure listed above is for the installation of the GMETAD on member machines of the cluster.

1.4 Installation of the GANGLIA WEB FRONTEND from source

Here are instructions to install the ganglia web front-end on the web server. Once again, make sure the GMETAD is installed on the same system too, irrespective of whether it is running on another machine

1. Ensure that PHP is installed on the web server. Look for the following lines in the httpd.conf file to make sure

```
<IfDefine HAVE_PHP4>
LoadModule php4_module  extramodules/libphp4.so
AddModule mod_php4.c
</IfDefine>

AddType application/x-httpd-php    .php .php4 .php3 .phtml
AddType application/x-httpd-php-source .phps
```

2. The PHP installation may be tested in the following manner:
In the DocumentRoot, create a file of the name php_testing.php. Have only the following line

```
<?phpinfo()?>
```

Close this file. Attempt to view the URL http://yourwebserver/php_testing.php . This should show up a web page similar to the following.



System	Linux hepfn000.uta.edu 2.4.7-10smp #1 SMP Thu Sep 6 17:09:31 EDT 2001 i686
Build Date	Jan 23 2003 22:13:01
Configure Command	'./configure' '--with-apxs2=/usr/local/apache2/bin/apxs'
Server API	Apache 2.0 Filter
Virtual Directory Support	disabled
Configuration File (php.ini) Path	/usr/local/lib/php.ini
PHP API	20020918
PHP Extension	20020429
Zend Extension	20021010
Debug Build	no
Thread Safety	disabled
Registered PHP Streams	php, http, ftp

This program makes use of the Zend Scripting Language Engine:
Zend Engine v1.3.0, Copyright (c) 1998-2002 Zend Technologies



PHP Credits

Configuration

PHP Core

Directive	Local Value	Master Value
-----------	-------------	--------------

3. If PHP is not installed on the web server, please follow instructions for updating the web server whichever it may be with the appropriate PHP module. **This web-interface will not run without PHP.**
4. Download the latest tarball (ganglia-webfrontend-2.5.3.tar.gz) from <http://ganglia.sourceforge.net/downloads.php> into a location by the system administrator for installing the front end.
5. It is recommended to unzip this in the folder that is the DocumentRoot for the web server. Alternately it may be unzipped into any other folder but the configuration has to be done accurately.

Your_prompt>gunzip ganglia-webfrontend-2.5.1.tar.gz

Your_prompt>tar -xvf ganglia-webfrontend-2.5.1.tar

6. Test the web front-end bu going to <http://localhost/ganglia-web-frontend-2.5.1/>

This should complete the installation.

The next section explains the configuration file settings.

Section – 2 GMOND & GMETAD CONFIGURATION

The following explains the steps involved in the changes to the configuration files for both the GMOND and the GMETAD.

2.1 GMOND Configuration

The GMOND default options work for most clusters. However, It is suggested that the following changes me made to the gmond.conf file located at `/etc/gmond.conf`

Please follow the sample gmond file attached at the end of Section 2.1 along with the following explanation.

`# The name of the cluster this node is a part of`

`# default: "unspecified"`

`name "LINUX_CLUSTER"`

"LINUX_CLUSTER" is the name of the cluster that is taken as an example.

`# The multicast channel for gmond to send/receive data on`

`# default: 239.2.11.71`

`# mcast_channel nnn.nnn.nnn.nnn`

This is the default channel that GMOND communicates with its peer. It uses the multicast channel 239.2.11.71. There is no reason to change it unless this channel is already in use by some other application on that network and it is desired to have another channel for the purpose. If it is desired that another channel be used for this purpose, please make sure that the channel you are using is a multicast channel and use that number. So it is suggested that this parameter be kept at the default value unless circumstances dictate otherwise.

`# The multicast port for gmond to send/receive data on`

`# default: 8649`

`# mcast_port nnnn`

This parameter specifies the port on which the GMOND peers communicate. As in the case of the multicast channel, unless already in use this port number may be kept unchanged. However, if another channel has to be used, the appropriate port number should be specified above

`# The multicast interface for gmond to send/receive data on`

`# default: the kernel decides based on routing configuration`

`# mcast_if eth1`

This value will be decided by GMOND on the basis of the information provided by the kernel. However if it so happens that there are multiple Ethernet ports on the system and a specific Ethernet port has to be used to ganglia operations, then the value above has to be set appropriately.

`# The number of nodes in your cluster. This value is used in the`

`# creation of the cluster hash.`

`# default: 1024`

`# num_nodes nn`

This parameter is used to specify the number of nodes in the cluster that this host is a part of. Please make sure that this value is appropriately configured. Although it wont make much of a difference, a larger number than the actual number of nodes will cause some additional space to be reserved by ganglia for data storage.

```
# Run gmond in "debug" mode. Gmond will not background. Debug messages
# are sent to stdout. Value from 0-100. The higher the number the more
# detailed debugging information will be sent.
# default: 0
# debug_level 10
```

THIS SECTION IS THE MOST IMPORTANT SETTING FOR DEBUGGING ERRORS IN THE INSTALLATION. SET THE DEBUG LEVEL TO ANY VALUE ABOVE 0 TO GET GMETAD TO SPIT OUT ERRORS IN THE INSTALLATION. OF COURSE, THIS STEP NEEDS TO BE TAKEN IF THE REGULAR INSTALLATION PROCEDURE FAILS.

All the other values may be kept according to their default values. They need to be changed only if a customized operation is desired. The configuration details mentioned above are a list of changes required for the basic operation of Ganglia.

Please remember that the location of this file is `/etc/gmond.conf`. It was moved to this location during the GMOND installation.

This completes the GMOND configuration. On completion, run the following command to restart the GMOND daemon with the new configuration parameters.

```
Your_prompt>/etc/rc.d/init.d/gmond restart
Starting GANGLIA gmond: [ OK ]
```

This completes the GMOND installation.

Here is a sample GMOND configuration file :

```
# $Id: gmond.conf,v 1.2 2002/09/19 00:37:18 sacerdoti Exp $
# This is the configuration file for the Ganglia Monitor Daemon (gmond)
# Documentation can be found at http://ganglia.sourceforge.net/docs/
#
# To change a value from it's default simply uncomment the line
# and alter the value
#####
#
# The name of the cluster this node is a part of
# default: "unspecified"
# name "My Cluster"
name "LINUX_CLUSTER"
#
# The owner of this cluster. Represents an administrative
# domain. The pair name/owner should be unique for all clusters
# in the world.
# default: "unspecified"
# owner "My Organization"
#
# The latitude and longitude GPS coordinates of this cluster on earth.
# Specified to 1 mile accuracy with two decimal places per axis in Decimal
# DMS format: "N61.18 W130.50".
# default: "unspecified"
# latlong "N32.87 W117.22"
#
# The URL for more information on the Cluster. Intended to give purpose,
# owner, administration, and account details for this cluster.
# default: "unspecified"
```

```

# url "http://www.mycluster.edu/"
#
# The location of this host in the cluster. Given as a 3D coordinate:
# "Rack,Rank,Plane" that corresponds to a Euclidean coordinate "x,y,z".
# default: "unspecified"
# location "0,0,0"
#
# The multicast channel for gmond to send/receive data on
# default: 239.2.11.71
# mcast_channel 239.2.11.71
#
# The multicast port for gmond to send/receive data on
# default: 8649
# mcast_port 8649
#
# The multicast interface for gmond to send/receive data on
# default: the kernel decides based on routing configuration
# mcast_if eth1
#
# The multicast Time-To-Live (TTL) for outgoing messages
# default: 1
# mcast_ttl 1
#
# The number of threads listening to multicast traffic
# default: 2
# mcast_threads 2
#
# Which port should gmond listen for XML requests on
# default: 8649
xml_port nnnn
#
# The number of threads answering XML requests
# default: 2
# xml_threads 2
#
# Hosts ASIDE from "127.0.0.1"/localhost and those multicasting
# on the same multicast channel which you will share your XML
# data with. Multiple hosts are allowed on multiple lines.
# Can be specified with either hostnames or IP addresses.
# default: none
# trusted_hosts 1.1.1.1 1.1.1.2 1.1.1.3 \
# 2.3.2.3 3.4.3.4 5.6.5.6
#
# The number of nodes in your cluster. This value is used in the
# creation of the cluster hash.
# default: 1024
num_nodes 35
#
# The number of custom metrics this gmond will be storing. This
# value is used in the creation of the host custom_metrics hash.
# default: 16
# num_custom_metrics 16

```

```
#
# Run gmond in "mute" mode. Gmond will only listen to the multicast
# channel but will not send any data on the channel.
# default: off
# mute on
#
# Run gmond in "deaf" mode. Gmond will only send data on the multicast
# channel but will not listen/store any data from the channel.
# default: off
# deaf on
#
# Run gmond in "debug" mode. Gmond will not background. Debug messages
# are sent to stdout. Value from 0-100. The higher the number the more
# detailed debugging information will be sent.
# default: 0
# debug_level 10
#
# If you don't want gmond to setuid, set this to "on"
# default: off
# no_setuid on
#
# Which user should gmond run as?
# default: nobody
# setuid nobody
#
# If you do not want this host to appear in the gexec host list, set
# this value to "on"
# default: off
# no_gexec on
#
# If you want any host which connects to the gmond XML to receive
# data, then set this value to "on"
# default: off
# all_trusted on
```

2.2 GMETAD Configuration

The behavior of the GMETAD is completely controlled by the configuration parameters. The configuration and importance of the most important parameters are listed below.

```
# The data_source tag must immediately be followed by a unique
# string which identifies the source then a list of machines
# which service the data source in the format ip:port, or name:port.
# If a # port is not specified then 8649 (the default gmond port) is
# assumed.
# default: There is no default value
# data_source "my box" localhost my.machine.edu:8655 1.2.3.5:8655
# data_source "another source" 1.3.4.7:8655 1.3.4.8
data_source "LINUX_CLUSTER" 127.0.0.1:8649
```

The data source is the most important parameter in the GMOND configuration and it is absolutely imperative that this parameter be set to the name of the cluster being monitored so that the system can monitor the state of that cluster. If two or more farms are being monitored then there will be an additional line in the data_source configuration that will read

```
data_source "SECOND_CLUSTER_NAME" nnn.nnn.nnn.nnn:8649
```

where nnn.nnn.nnn.nnn is the IP address of the host which houses the GMETAD of the second cluster.

```
# List of machines this gmetad will share XML with
# default: There is no default value
# trusted_hosts 127.0.0.1 169.229.50.165
#
```

This is a configuration setting that is useful when there is a second cluster like the one mentioned above. In the second cluster, the GMETAD will be SHARING its information of the main GMETAD. Hence in the second cluster, this parameter will be directed towards the system housing the GMETAD in the main cluster (i.e:IP ADDRESS)

```
# The port gmetad will answer requests for XML
# default: 8651
# xml_port 8651
#
```

GMETAD communicates its XML data on the port 8651 by default. Hence as explained above, it is unnecessary to change this value unless 8651 is already being used by another application. In that case this value will be set to the port allotted for use by the system administrator.

```
# Where gmetad stores its round-robin databases
# default: "/var/lib/ganglia/rrds"
# rrd_rootdir "/some/other/place"
```

It was mentioned in the RRD installation that the RRDtools are installed at `/var/lib/ganglia/rrds/` by default. There was an option to install RRDs at another location. Hence this parameter directs Ganglia to look in a location other than the default to pick up its information.

```
# Setting the debug_level above zero will make gmetad output
# debugging information and stay in the foreground
# default: 0
# debug_level 10
```

THIS SECTION IS THE MOST IMPORTANT SETTING FOR DEBUGGING ERRORS IN THE INSTALLATION. SET THE DEBUG LEVEL TO ANY VALUE ABOVE 0 TO GET GMETAD TO SPIT OUT ERRORS IN THE INSTALLATION. OF COURSE, THIS STEP NEEDS TO BE TAKEN IF THE REGULAR INSTALLATION PROCEDURE FAILS.

Please remember that the location of this file is `/etc/gmetad.conf`. It was moved to this location during the GMETAD installation.

This completes the GMETAD configuration. On completion, run the following command to restart the GMETAD daemon with the new configuration parameters.

```
Your_prompt>/etc/rc.d/init.d/gmetad restart
Starting GANGLIA gmetad: [ OK ]
```

This completes the GMETAD installation.

Here is a sample GMETAD configuration file:

```
# This is an example of a Ganglia Meta Daemon configuration file
# http://ganglia.sourceforge.net/
#
# $Id: gmetad.conf,v 1.3 2002/09/19 18:56:42 sacerdoti Exp $
#
# Setting the debug_level above zero will make gmetad output
# debugging information and stay in the foreground
# default: 0
# debug_level 10
#
# The data_source tag must immediately be followed by a unique
# string which identifies the source then a list of machines
# which service the data source in the format ip:port, or name:port.
# If a # port is not specified then 8649 (the default gmond port) is
# assumed.
# default: There is no default value
# data_source "my box" localhost my.machine.edu:8655 1.2.3.5:8655
# data_source "another source" 1.3.4.7:8655 1.3.4.8
data_source "LINUX_CLUSTER" 127.0.0.1:8649
data_source "SECOND_CLUSTER_NAME" nnn.nnn.nnn.nnn:8649
#
# List of machines this gmetad will share XML with
# default: There is no default value
# trusted_hosts 127.0.0.1 169.229.50.165
#
# If you don't want gmetad to setuid then set this to off
# default: on
# setuid off
#
# User gmetad will setuid to (defaults to "nobody")
# default: "nobody"
# setuid_username "nobody"
#
# The port gmetad will answer requests for XML
```

```

# default: 8651
# xml_port 8651
#
# The number of threads answering XML requests
# default: 2
# server_threads 4
#
# Where gmetad stores its round-robin databases
# default: "/var/lib/ganglia/rrds"
# rrd_rootdir "/some/other/place"

```

If all the GMETAD and the GMOND are turned on, then the display should show you something similar to what is seen at <http://hepfm000.uta.edu/ganglia-webfrontend-2.5.1/>



This screen is proof that Ganglia has been installed. Customizations may be made according to the needs of the system administrator. For all these, it is advised to refer to the main Ganglia web page at <http://ganglia.sourceforge.net>.

Thank You.