

Red Hat Enterprise MRG 1.2

Grid Installation Guide

Installation information for the Grid
component of Red Hat Enterprise MRG



Lana Brindley

Red Hat Enterprise MRG 1.2 Grid Installation Guide

Installation information for the Grid component of Red Hat Enterprise MRG

Edition 3

Author

Lana Brindley

lbrindle@redhat.com

Copyright © 2008 Red Hat, Inc

Copyright © 2008 Red Hat, Inc. This material may only be distributed subject to the terms and conditions set forth in the Open Publication License, V1.0 or later (the latest version of the OPL is presently available at <http://www.opencontent.org/openpub/>).

Red Hat and the Red Hat "Shadow Man" logo are registered trademarks of Red Hat, Inc. in the United States and other countries.

All other trademarks referenced herein are the property of their respective owners.

1801 Varsity Drive
Raleigh, NC 27606-2072 USA
Phone: +1 919 754 3700
Phone: 888 733 4281
Fax: +1 919 754 3701
PO Box 13588 Research Triangle Park, NC 27709 USA

This book will show you how to download and install the MRG Grid component of the Red Hat Enterprise MRG distributed computing platform. For detailed information on the use of MRG Grid, see the MRG Grid User Guide.

Preface	v
1. Document Conventions	v
1.1. Typographic Conventions	vi
1.2. Pull-quote Conventions	vii
1.3. Notes and Warnings	viii
2. We Need Feedback!	viii
1. Overview	1
2. Installing MRG Grid	3
2.1. Installing MRG Grid on Red Hat Enterprise Linux 5	3
2.2. Installing MRG Grid on Red Hat Enterprise Linux 4	3
3. Getting Started with MRG Grid	5
4. MRG Grid Benefits and Features	7
5. More Information	11
A. Revision History	13

Preface

Red Hat Enterprise MRG

This book contains basic overview and installation information for the MRG Grid component of Red Hat Enterprise MRG. Red Hat Enterprise MRG is a high performance distributed computing platform consisting of three components:

1. *Messaging* — Cross platform, high performance, reliable messaging using the Advanced Message Queuing Protocol (AMQP) standard.
2. *Realtime* — Consistent low-latency and predictable response times for applications that require microsecond latency.
3. *Grid* — Distributed High Throughput Computing (HTC) and High Performance Computing (HPC).

All three components of Red Hat Enterprise MRG are designed to be used as part of the platform, but can also be used separately.

MRG Grid

Grid computing allows organizations to fully utilize their computing resources to complete high-performance tasks. By monitoring all resources - rack-mounted clusters and general workstations - for availability, any spare computing power can be redirected towards other, more intensive tasks until it is explicitly required again. This allows a standard networked system to operate in a way that is similar to a supercomputer.

MRG Grid provides High Throughput and High Performance computing and enables enterprises to achieve higher peak computing capacity as well as improved infrastructure utilization by leveraging their existing technology to build high performance grids. MRG Grid provides a job queueing mechanism, scheduling policy, priority scheme, resource monitoring, and resource management. Users submit their jobs to MRG Grid, where they are placed into a queue. MRG Grid then chooses when and where to run the jobs based upon a policy, carefully monitors their progress, and ultimately informs the user upon completion.

MRG Grid is based on the [Condor Project](http://www.cs.wisc.edu/condor/)¹ developed within the [University of Wisconsin-Madison](http://www.wisc.edu/)². Condor also offers a comprehensive library of freely available documentation in its [Manual](http://www.cs.wisc.edu/condor/manual/)³.

1. Document Conventions

This manual uses several conventions to highlight certain words and phrases and draw attention to specific pieces of information.

In PDF and paper editions, this manual uses typefaces drawn from the [Liberation Fonts](https://fedorahosted.org/liberation-fonts/)⁴ set. The Liberation Fonts set is also used in HTML editions if the set is installed on your system. If not, alternative but equivalent typefaces are displayed. Note: Red Hat Enterprise Linux 5 and later includes the Liberation Fonts set by default.

¹ <http://www.cs.wisc.edu/condor/>

² <http://www.wisc.edu/>

³ <http://www.cs.wisc.edu/condor/manual/>

⁴ <https://fedorahosted.org/liberation-fonts/>

1.1. Typographic Conventions

Four typographic conventions are used to call attention to specific words and phrases. These conventions, and the circumstances they apply to, are as follows.

Mono-spaced Bold

Used to highlight system input, including shell commands, file names and paths. Also used to highlight key caps and key-combinations. For example:

To see the contents of the file **my_next_bestselling_novel** in your current working directory, enter the **cat my_next_bestselling_novel** command at the shell prompt and press **Enter** to execute the command.

The above includes a file name, a shell command and a key cap, all presented in Mono-spaced Bold and all distinguishable thanks to context.

Key-combinations can be distinguished from key caps by the hyphen connecting each part of a key-combination. For example:

Press **Enter** to execute the command.

Press **Ctrl-Alt-F1** to switch to the first virtual terminal. Press **Ctrl-Alt-F7** to return to your X-Windows session.

The first sentence highlights the particular key cap to press. The second highlights two sets of three key caps, each set pressed simultaneously.

If source code is discussed, class names, methods, functions, variable names and returned values mentioned within a paragraph will be presented as above, in **Mono-spaced Bold**. For example:

File-related classes include **filesystem** for file systems, **file** for files, and **dir** for directories. Each class has its own associated set of permissions.

Proportional Bold

This denotes words or phrases encountered on a system, including application names; dialogue box text; labelled buttons; check-box and radio button labels; menu titles and sub-menu titles. For example:

Choose **System > Preferences > Mouse** from the main menu bar to launch **Mouse Preferences**. In the **Buttons** tab, click the **Left-handed mouse** check box and click **Close** to switch the primary mouse button from the left to the right (making the mouse suitable for use in the left hand).

To insert a special character into a **gedit** file, choose **Applications > Accessories > Character Map** from the main menu bar. Next, choose **Search > Find...** from the **Character Map** menu bar, type the name of the character in the **Search** field and click **Next**. The character you sought will be highlighted in the **Character Table**. Double-click this highlighted character to place it in the **Text to copy** field and then click the **Copy** button. Now switch back to your document and choose **Edit > Paste** from the **gedit** menu bar.

The above text includes application names; system-wide menu names and items; application-specific menu names; and buttons and text found within a GUI interface, all presented in Proportional Bold and all distinguishable by context.

Note the **>** shorthand used to indicate traversal through a menu and its sub-menus. This is to avoid the difficult-to-follow 'Select **Mouse** from the **Preferences** sub-menu in the **System** menu of the main menu bar' approach.

Mono-spaced Bold Italic or ***Proportional Bold Italic***

Whether Mono-spaced Bold or Proportional Bold, the addition of Italics indicates replaceable or variable text. Italics denotes text you do not input literally or displayed text that changes depending on circumstance. For example:

To connect to a remote machine using ssh, type **ssh *username@domain.name*** at a shell prompt. If the remote machine is **example.com** and your username on that machine is john, type **ssh *john@example.com***.

The **mount -o remount *file-system*** command remounts the named file system. For example, to remount the **/home** file system, the command is **mount -o remount */home***.

To see the version of a currently installed package, use the **rpm -q *package*** command. It will return a result as follows: ***package-version-release***.

Note the words in bold italics above — username, domain.name, file-system, package, version and release. Each word is a placeholder, either for text you enter when issuing a command or for text displayed by the system.

Aside from standard usage for presenting the title of a work, italics denotes the first use of a new and important term. For example:

When the Apache HTTP Server accepts requests, it dispatches child processes or threads to handle them. This group of child processes or threads is known as a *server-pool*. Under Apache HTTP Server 2.0, the responsibility for creating and maintaining these server-pools has been abstracted to a group of modules called *Multi-Processing Modules (MPMs)*. Unlike other modules, only one module from the MPM group can be loaded by the Apache HTTP Server.

1.2. Pull-quote Conventions

Two, commonly multi-line, data types are set off visually from the surrounding text.

Output sent to a terminal is set in Mono-spaced Roman and presented thus:

```
books      Desktop  documentation  drafts  mss    photos  stuff  svn
books_tests Desktop1  downloads      images  notes  scripts svgs
```

Source-code listings are also set in Mono-spaced Roman but are presented and highlighted as follows:

```
package org.jboss.book.jca.ex1;
import javax.naming.InitialContext;
```

```
public class ExClient
{
    public static void main(String args[])
        throws Exception
    {
        InitialContext iniCtx = new InitialContext();
        Object          ref    = iniCtx.lookup("EchoBean");
        EchoHome        home   = (EchoHome) ref;
        Echo             echo   = home.create();

        System.out.println("Created Echo");

        System.out.println("Echo.echo('Hello') = " + echo.echo("Hello"));
    }
}
```

1.3. Notes and Warnings

Finally, we use three visual styles to draw attention to information that might otherwise be overlooked.



Note

A Note is a tip or shortcut or alternative approach to the task at hand. Ignoring a note should have no negative consequences, but you might miss out on a trick that makes your life easier.



Important

Important boxes detail things that are easily missed: configuration changes that only apply to the current session, or services that need restarting before an update will apply. Ignoring Important boxes won't cause data loss but may cause irritation and frustration.



Warning

A Warning should not be ignored. Ignoring warnings will most likely cause data loss.

2. We Need Feedback!

If you find a typographical error in this manual, or if you have thought of a way to make this manual better, we would love to hear from you! Please submit a report in Bugzilla: <http://bugzilla.redhat.com/bugzilla/> against the product **Red Hat Enterprise MRG**.

When submitting a bug report, be sure to mention the manual's identifier: *Grid_Installation_Guide*

If you have a suggestion for improving the documentation, try to be as specific as possible when describing it. If you have found an error, please include the section number and some of the surrounding text so we can find it easily.

Overview

MRG Grid provides High Throughput and High Performance computing and enables enterprises to achieve higher peak computing capacity as well as higher IT utilization by leveraging their existing infrastructure and building high performance grids. High Throughput Computing (HTC) delivers large amounts of computing power over a sustained period of time (months or years), whereas High Performance Computing (HPC) delivers significant computing power over a short period of time (hours or days).

MRG Grid enables enterprises and research organizations to bring the power of distributed computing across their entire infrastructure to tackle the largest computational problems in a highly efficient and effective manner.

Installing MRG Grid

In order to install MRG Grid you will need to have registered your system with [Red Hat Network](#)¹. This table lists the Red Hat Enterprise MRG channels available on Red Hat Network for MRG Grid.

Channel Name	Operating System	Architecture
Red Hat MRG Grid	RHEL-4 AS	32bit, 64bit
Red Hat MRG Grid	RHEL-4 ES	32bit, 64bit
Red Hat MRG Grid	RHEL-5 Server	32bit, 64bit
Red Hat MRG Grid	non-Linux	32bit

Table 2.1. Red Hat Enterprise MRG Channels Available on Red Hat Network



Important

Before you install Red Hat Enterprise MRG check that your hardware and operating system are supported. A complete list is available on the [Red Hat Enterprise MRG Supported Hardware Page](#)².

2.1. Installing MRG Grid on Red Hat Enterprise Linux 5

1. Install the MRG Grid group using the **yum** command.

```
# yum groupinstall "MRG Grid"
```

2. You can check the installation location and that the components have been installed successfully by using the **rpm -q** command with the name of the package you installed. For example:

```
# rpm -q condor
/etc/condor
/etc/condor/condor_config
/usr/bin/condor
...
[output truncated]
```



Note

If you find that yum is not installing all the dependencies you require, make sure that you have registered your system with [Red Hat Network](#)³.

2.2. Installing MRG Grid on Red Hat Enterprise Linux 4

1. Install the MRG Grid components using the **up2date** command.

¹ <https://rhn.redhat.com/help/about.pxt>

```
# up2date condor mrg-grid-docs
```

2. You can check the installation location and that the components have been installed successfully by using the `rpm -q1` command with the name of the package you installed. For example:

```
# rpm -q1 condor
/etc/condor
/etc/condor/condor_config
/usr/bin/condor
...
[output truncated]
```



Note

If you find that **up2date** is not installing all the dependencies you require, make sure that you have registered your system with *Red Hat Network*⁴.

Getting Started with MRG Grid

1. It is advised that you review the local configuration file before starting MRG Grid. By default, this file is located at `~condor/condor_config.local`. However, some installations might put this file in a different location.

Check the exact location of the local configuration file by running the following command from the shell prompt:

```
$ condor_config_val LOCAL_CONFIG_FILE
/var/lib/condor/condor_config.local
```

2. The default configuration sets up a *Personal Condor*. This is a specific configuration suited for individual users who do not have their own pool of machines. To find out more about Personal Condor, see the [Condor Manual](#)¹.
3. To allow other machines to join your pool additional configuration must be done. Specifically, you will need to customize the **HOSTALLOW_WRITE** option. Open the `~condor/condor_config.local` file in your preferred text editor and locate the section titled **Host/IP Access Levels**.
4. The value for this option should be set to allow machines to join your pool and submit jobs. Any machine that you give write access to using the **HOSTALLOW_WRITE** option should also be given read access using the **HOSTALLOW_READ** option.

```
HOSTALLOW_WRITE = *.your.domain.com
```



Warning

The simplest option is to change the **HOSTALLOW_WRITE** option to **HOSTALLOW_WRITE = ***. However, this will allow anyone to submit jobs or add machines to your pool. This is a serious security risk and therefore not recommended.

5. MRG Grid is usually run as a service. To start MRG Grid run the following command:

```
# service condor start
```



Note

See *Further Reading*, for places to turn for help if you have trouble with installing MRG Grid.

For detailed configuration information, see the *MRG Grid User Guide*.

MRG Grid Benefits and Features

Benefits

MRG Grid provides significant benefits and value for enterprises, including:

Power

MRG Grid can process the largest computational workloads, from massively parallel High Performance Computing jobs to long-running High Throughput Computing jobs

Peak Workload Handling

MRG Grid adds on-demand computational power for handling peak loads. This is achieved through capabilities such as cycle-stealing, scheduling, and the use of remote grids.

Flexibility

MRG Grid provides complete flexibility and can run high-burst or lengthy computations in both centralized and distributed grids. Jobs can be run on various platforms including Linux and Windows. Furthermore, MRG Grid can schedule virtualized environments and workloads for the upmost flexibility in utilizing infrastructure.

Powerful Management Tools

Managing MRG Grid is simplified by leveraging the Red Hat Enterprise MRG unified, browser-based management console. The Red Hat Enterprise MRG integrated management tools enable administrators to manage, configure, provision, deploy, and monitor their grid deployments using the same tools they use for MRG Messaging and MRG Realtime.

Features

MRG Grid provides a broad set of features across both High Throughput Computing and High Performance Computing, including:

Virtualization

Allows for submission of a virtual machine (VM) as a user job, supporting migration of the VM

Dedicated and Undedicated Node Management (Cycle-Stealing)

Allows for dedicated resources (clusters) to be augmented with otherwise undedicated (desktops) using flexible policies

Multiple Standards-Based APIs

The web service interface provides job submission and management functionality; the command line interface provides functionality that is highly scriptable, with consistent output

Security

Authentication using multiple mechanisms

Privacy provided by network encryption

Integrity of network traffic

Authorization through flexible configuration policies

Federated Grids/Clusters

A mechanism known as flocking allows independent pools to use each others' resources, controllable by customizable policies

Management Tools

Powerful browser-based management tools for managing daemons and machines, security, compute jobs, scalability settings, priorities, and more. Also provides sophisticated monitoring capabilities.

Workflow Management

The ability to specify job dependencies, via [DAGMan](#)¹, allows for construction and execution of complex workflows

Accounting

User and group resource utilization is tracked and accessible to administrators

ClassAds

A flexible language for policy and meta-data description

Policies

Flexible, customizable policies specified by jobs and resources via ClassAds

High Availability

The Negotiator and Collector, via the High Availability Daemon (HAD), and the Schedd, via Schedd Fail-over, can have their state replicated to allow for graceful fail-over upon service disruption

Compute On-Demand (COD)

The ability for a node or set of nodes to be claimed by a user in such a way that others may use the claimed nodes until the user needs them

Priority Based Scheduling

Priority scheduling is performed at the granularity of a user

Fair-share scheduling can be performed on groups of users

Priority management is controllable by administrators

Account Remapping

Allows for execution across administrative domains

Enhance security by using a restricted pool of users to run jobs on execute machines

Privilege Separation

Only a single, specialized, audited component requires root/administrator permissions on execute nodes

Parallel Universe

Provides an extensible framework for running parallel - including Message Passing Interface (MPI) - jobs

Co-allocation of compute nodes is done automatically

¹ <http://www.cs.wisc.edu/condor/dagman/>

Framework implementation for MPICH1, [MPICH2](#)², and [LAM](#)³ provided

Java Universe

Explicit support of jobs written in Java

Time Scheduling for Job Execution (Cron)

Allows a job or multiple jobs to be started at specific times, with customizable policy for failures such as missed deadlines

Backfill

Allows otherwise unused nodes to run jobs provided by [BOINC](#)⁴

File Staging

Support for automatic file staging in the absence of a shared file system. Job input and online file streaming from submit to execute nodes using Chirp

Condor-C

Allows for jobs in one queue to be moved to another queue

² <http://www.mcs.anl.gov/research/projects/mpich2/>

³ <http://www.lam-mpi.org/>

⁴ <http://boinc.berkeley.edu/>

More Information

Reporting Bugs

Follow these instructions to enter a bug report:

1. You will need a [Bugzilla](#)¹ account. You can create one at [Create Bugzilla Account](#)².
2. Once you have a Bugzilla account, log in and click on [Enter A New Bug Report](#)³.
3. You will need to identify the product (Red Hat Enterprise MRG), the version (1.1), and whether the bug occurs in the software (component=grid) or in the documentation (component=Grid_Installation_Guide).

Further Reading

- Red Hat Enterprise MRG and MRG Grid Product Information
 - <http://www.redhat.com/mrg>
- MRG Grid User Guide and other Red Hat Enterprise MRG manuals
 - http://redhat.com/docs/en-US/Red_Hat_Enterprise_MRG
- Condor Manual
 - <http://www.cs.wisc.edu/condor/manual/>

Appendix A. Revision History

Revision 3.1	Thu Oct 29 2009	Lana Brindley lbrindle@redhat.com
Final version for 1.2 release		
Revision 3.0	Thu Sep 24 2009	Lana Brindley lbrindle@redhat.com
Moved 'Getting Started' info to its own chapter BZ #482959 - local configuration file location		
Revision 2.6	Mon Jan 19 2009	Lana Brindley lbrindle@redhat.com
Added links to product page		
Revision 2.5	Mon Dec 22 2008	Michael Hideo mhideo@redhat.com
Further changes as per BZ #470847		
Revision 2.3	Mon Nov 24 2008	Lana Brindley lbrindle
Minor updates and changes in preparation for delivery to QE		
Revision 2.2	Thu Oct 30 2008	Lana Brindley lbrindle
Minor updates and changes in preparation for technical review		
Revision 2.1	Thu Oct 30 2008	Lana Brindley lbrindle
Ported information from Deployment Guide to Installation Guide		

