

Red Hat Enterprise MRG 1.1

MRG Release Notes

Release Notes for the Red Hat Enterprise MRG 1.1 Release



Lana Brindley

Red Hat Enterprise MRG 1.1 MRG Release Notes

Release Notes for the Red Hat Enterprise MRG 1.1 Release Edition 2

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 with the restrictions noted below (the latest version of the OPL is presently available at <http://www.opencontent.org/openpub/>).

Distribution of substantively modified versions of this document is prohibited without the explicit permission of the copyright holder.

Distribution of the work or derivative of the work in any standard (paper) book form for commercial purposes is prohibited unless prior permission is obtained from the copyright holder.

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 27606-2072USAPhone: +1 919 754 3700

Phone: 888 733 4281

Fax: +1 919 754 3701

PO Box 13588Research Triangle Park, NC 27709USA

These Release Notes contain important information available at the time of release of Red Hat Enterprise MRG 1.1. Known problems, resources, and other issues are discussed here. Read this document before beginning to use the Red Hat Enterprise MRG distributed computing platform.

1. System Requirements	1
1.1. Supported Hardware and Platforms	1
1.2. Installing and Configuring Red Hat Enterprise MRG	2
1.2.1. Installing and Configuring MRG Messaging	2
1.2.2. Installing and Configuring MRG Realtime	3
1.2.3. Installing and Configuring MRG Grid	3
1.2.4. Installing and Configuring the MRG Management Console	4
2. Known Issues	5
2.1. MRG Messaging Known Issues	5
2.2. MRG Realtime Known Issues	14
2.3. MRG Grid Known Issues	18
A. Revision History	25

System Requirements

This section contains information related to installing Red Hat Enterprise MRG, including hardware and platform requirements.

1.1. Supported Hardware and Platforms

Red Hat Enterprise MRG is highly optimized to run on Red Hat Enterprise Linux 5 and later due to its inclusion of MRG Realtime. The MRG Messaging and MRG Grid capabilities can also run on other platforms, but without the full benefits of running on Red Hat Enterprise Linux 5 and later.

	x86	x86_64	RHEL 4.7	RHEL 5.2	Windows	Solaris
MRG Messaging Native Linux Broker	X	X	X	X		
MRG Messaging Client - Java *	X	X	X	X	X	X
MRG Messaging Client - JMS *	X	X	X	X	X	X
MRG Messaging Client - C++	X	X	X	X		
MRG Messaging Client - Python	X	X	X	X		
MRG Messaging Client - .NET <i>preview</i>	X	X			X	
MRG Messaging Client - Ruby <i>preview</i>	X	X	X	X		
MRG Grid Scheduler	X	X	X	X		
MRG Realtime	X	X		X		

Table 1.1. Supported Hardware and Platforms

**Note**

* The Java and JMS MRG Messaging Clients are supported for use with Java 1.5 and Java 6 JVMs. For Sun JVMs, it is recommended to use Java 1.5.15 or later or 1.6.06 or later.

1.2. Installing and Configuring Red Hat Enterprise MRG

In order to download and install Red Hat Enterprise MRG 1.1 on your system, you need to subscribe to the appropriate channels on the Red Hat Network (RHN).

Channel Name	Platform	Architecture
MRG Grid	RHEL-4 AS	32bit, 64bit
MRG Grid	RHEL-5 Server	32bit, 64bit
MRG Grid	RHEL-4 ES	32bit, 64bit
MRG Grid	non-Linux	32bit
MRG Grid Execute Node	RHEL-4 AS	32bit, 64bit
MRG Grid Execute Node	RHEL-5 Server	32bit, 64bit
MRG Grid Execute Node	RHEL-4 ES	32bit, 64bit
MRG Management Console	RHEL-4 AS	32bit, 64bit
MRG Management Console	RHEL-4 ES	32bit, 64bit
MRG Management Console	RHEL-5 Server	32bit, 64bit
MRG Messaging	RHEL-4 AS	32bit, 64bit
MRG Messaging	RHEL-4 ES	32bit, 64bit
MRG Messaging	RHEL-5 Server	32bit, 64bit
MRG Messaging	non-Linux	32bit
MRG Messaging Base	RHEL-4 AS	32bit, 64bit
MRG Messaging Base	RHEL-4 ES	32bit, 64bit
MRG Messaging Base	RHEL-5 Server	32bit, 64bit
MRG Realtime	RHEL-5 Server	32bit, 64bit

Table 1.2. Red Hat Network Channels

1.2.1. Installing and Configuring MRG Messaging

Installation information for MRG Messaging is available in the *MRG Messaging Installation Guide*. For use and configuration details, see the *MRG Messaging User Guide*. For information on developing your own programs for MRG Messaging, start with the *MRG Messaging Tutorial*.

MRG Messaging version 1.1 also includes a preview of the Qman feature. This feature creates a management bridge for communication between one or more remote brokers.

Red Hat Enterprise MRG documentation is available for download at the [Red Hat Enterprise MRG Documentation Website](http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_MRG/)¹.

¹ http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_MRG/

1.2.2. Installing and Configuring MRG Realtime

Installation information for MRG Realtime is available in the *MRG Realtime Installation Guide*. For information on tuning MRG Realtime, see the *MRG Realtime Tuning Guide*.

MRG Realtime is incompatible with the Xen-based virtualization features in Red Hat Enterprise Linux 5 and is not supported for use with any virtualization technology.

The MRG Realtime kernel may be rebased over the lifetime of a Red Hat Enterprise MRG release, however there are no guarantees of a stable kernel Application Binary Interface (kABI) over the life of Red Hat Enterprise MRG.

Because MRG Realtime provides an updated Linux kernel, it is certified for use on a subset of the hardware systems certified for Red Hat Enterprise Linux. MRG Realtime is certified on x86 and x86_64 architectures only. Furthermore, Red Hat works with hardware vendors to certify systems for use with MRG Realtime based on customer demand. At initial release, the list of certified systems for MRG Realtime includes:

HP

- DL320g5
- DL360g5
- DL380g5
- DL585g5

IBM

- LS21 blade
- HS21 blade

For an updated list of certified systems, see the [Red Hat Hardware Catalog](#)².

Red Hat Enterprise MRG documentation is available for download at the [Red Hat Enterprise MRG Documentation Website](#)³.

1.2.3. Installing and Configuring MRG Grid

The 1.1 release is the first general release for MRG Grid. This version contains support for:

- High Availability
- Dynamic Provisioning
- Low latency scheduling

This version of MRG Grid provides support for MRG Grid Amazon EC2 Execute Node. The product that is run within the EC2 instance will determine the available support level.

Installation and configuration information for MRG Grid is available in the *MRG Grid Installation Guide*. For use and configuration details, see the *MRG Grid User Guide*.

² <https://hardware.redhat.com/>

³ http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_MRG/

Red Hat Enterprise MRG documentation is available for download at the [Red Hat Enterprise MRG Documentation Website](http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_MRG/Documentation_Website)⁴.

1.2.4. Installing and Configuring the MRG Management Console

Installation and configuration information for the MRG Management Console is available in the *MRG Management Console Installation Guide*.

Red Hat Enterprise MRG documentation is available for download at the [Red Hat Enterprise MRG Documentation Website](http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_MRG/Documentation_Website)⁵.

⁴ http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_MRG/

⁵ http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_MRG/

Known Issues

This section describes some of the most important known issues in Red Hat Enterprise MRG.

2.1. MRG Messaging Known Issues

The 1.1 release of MRG Messaging contains updated packages that fix several bugs.

This release relates to the following Errata:

- RHEA-2009:0035
- RHEA-2009:0038

Bug Fixes

The following is a list of bug fixes in the 1.1 release of MRG Messaging:

Bugzilla Number	Description
457922	<p>Performance tests occurring over a loopback connection are much faster than when run over an ethernet connection, where limited bandwidth is not the cause. The error given is</p> <pre>eth0: too many iterations (6) in nv_nic_irq.</pre> <p>To workaround this issue, open <code>/etc.modprobe.conf</code> and add the following line:</p> <pre>options forcedeth max_interrupt_work=30</pre>
468713	<p>The <code>client::AckPolicy</code> class has been removed.</p> <p>The following classes have been changed:</p> <ul style="list-style-type: none"> • <code>client::SubscriptionManager</code> • <code>SessionManager::subscribe()</code> • <code>SessionManager::findSubscriptions()</code> <p>The following classes have been added:</p> <ul style="list-style-type: none"> • <code>client::SubscriptionSettings</code> • <code>client::Subscription</code> <p>Classes were changed for the following reasons:</p> <ul style="list-style-type: none"> • The use of <i>bool</i> for accept and acquire modes was confusing

Bugzilla Number	Description
	<ul style="list-style-type: none"> • Creating individual subscriptions is more flexible and less error prone than setting values on the SubscriptionSettings that apply to all subsequent subscriptions. • AckPolicy is only provided for control of accepting messages, not for acquiring them. • AckPolicy was copied to each new subscription, there was no way to access the instance associated with a particular subscription. • The old API had no class to represent subscription. The Subscription class fills that gap.
470067	Applications that use the SIGPIPE on startup were exiting if they attempted to write to a closed socket. This functionality was disabled in the C++ client in order to prevent this.
433797	If the broker is killed while still connected, the perftest client would sporadically perform a core dump. This was caused by the client exiting without joining the thread that was running the dispatch loop. It was fixed by ensuring that all connection dispatch threads are joined during the connection close function.
442591	When brokers are federated, all routes are completely static. This results in messages being routed to destinations even if there are no bound consumers at that destination. A dynamic binding feature was implemented that ensures that messages are dropped if there is no consumer bound to the destination.
446644	The MRG Management Console was producing graphs that plotted lines even when the data rate dropped to zero.
449756	When first connecting to the MRG Management Console with a browser, it was requesting login credentials more than once. This issue has now been fixed and login occurs as expected.
450749	Errors in JMS URLs were not being detected. URLs are now checked, and an error thrown when a malformed URL is encountered.
452123	In a cluster using transient queues and transient messages, the last surviving node of the cluster should automatically convert all its queues to durable mode and flush the current state of the

Bugzilla Number	Description
	queues to a persistence store. This process now occurs as expected, and when a node is added back to the cluster all messages become transient again.
455087	If a broker running on a 64-bit machine with a store module and any existing store directory is stopped using ctrl+c , the broker performs a core dump.
461816	After a session.detach function was called, the broker sent a message.transfer , which can cause problems for the subsequently attached sessions. This issue was resolved.
462121	When an istream.read() or ignore() operation failed and was followed by a call to tellg() on Red Hat Enterprise Linux 4, an error was being produced. This was occurring because tellg() returned the value -1 instead of the expected offset. This behavior was corrected and the issue resolved.
465213	When an exception is encountered in the Java client, it does not always communicate the exact problem that led to the exception. Some instances of this problem have been addressed, and the issue will be resolved fully in a future release.
466955	When consuming messages using the python client using an accept mode of explicit, causes the broker to consume large amounts of processor capacity. This is due to the broker retaining records of each delivery and keeping them until the message has been completed. This was resolved by causing the broker to discard these records unless it is in windowing mode.
468428	If the <i>perftest</i> tool causes an <i>Enqueue capacity threshold exceeded</i> error, the broker was failing with a signal of SIGSEV . This was caused by a callback function that pointed back to an object in the publishers session. It was resolved by ensuring that the callbacks are all unregistered before the target is deleted.
469561	On Red Hat Enterprise Linux 4, the command service qpid start was failing. This was due to sections of the init script being specific to Red Hat Enterprise Linux 5. The init file was fixed to correct this issue.

Bugzilla Number	Description
469822	The RDMA plugin was failing due to incorrect or unavailable configuration, this is expected behavior however the error message was not explicit. These errors are now handled more gracefully.
469853	Exchanges within a cluster that had an additional state declared were not replicating this state to new cluster members. This was corrected by modifying cluster replication to use the same encode and decode mechanism as the persistent store. This resolved the problem and ensured consistency between clustering and persistence.
470167	The Management Agent daemon was not shutting down cleanly. This behavior has now been corrected and the daemon shuts down as expected.
470228	When the broker was restarted with different sizing parameters after having been started with the persistence store loaded, the broker was aborting. This was due to the number of files specified in the second instance being smaller than those in the original count. This error was resolved.
470286	The MRG Management Console was creating graphs that ran over the right hand edge of the available screen area. This was caused by a clock skew which resulted in statistics being sent with a future time. A check was added so that if the receive time of the data is in the future, the current time will be used instead.
471061	In cases where the journal is small and messages are large, the file header write operation may cause the file to be overwritten by the next write cycle before the header has returned.
471095	Deadlocks were being created in federations where one thread would be waiting for a lock on the same link object as another thread. This issue was resolved and deadlocks no longer occur.
471247	When connections were opened concurrently, running the perftest tool caused the broker to crash. This issue occurred because the result of a call to <code>inet_ntoa()</code> was being written to a statically allocated buffer, and simultaneous calls created the crash. This behavior was changed and the broker no longer crashes in this situation.

Bugzilla Number	Description
471319	When the broker was started as a daemon, and then instructed to close, the broker could not find the correct port to enable it to shut down. This confusion was corrected and the broker now behaves as expected.
471584	When the --require-encryption option was used, the broker returns an exception due to it only recognizing brokers that use TCP transport. This was resolved by updating the URL format to include other forms of transport.
471601	When attempting to read libaio returns, if the order is scrambled, the operation sporadically failed with a message saying RHM_IORES_EMPTY when unread records still existed. This is due to a logic error, where the getevents operation occurs but the return value is not checked. This error has not been reproduced on a production system and will be fully addressed in a future release.
471606	Some two phase commit processes were failing intermittently with an illegal read/write error. This was due to a lock that prevented the thread from beginning until after a flush had been completed, and another thread destroyed the journal during this period. A mutex was added to resolve this issue, and the error has not been reproduced.
471612	The journal auto-expand feature is not enabled in version 1.1, however the command line options and internal status flags exist. These have been suppressed until the feature is enabled.
453536	Support was added for a feature that allows the user to set Time To Live (TTL) values for message expiration.
441700	A ring queue feature has been implemented. This is a fixed size queue, where new publications result in the oldest message being dropped.
446938	Support was added for RBAC and ACL, in order to allow multiple message flows to securely use the same broker.
452015	A custom exchange was created that imposes an order on all routed messages, and adds a header to all delivered messages with a queue sequence number.
452017	A Last Value Queue (LVQ) was added. This queue gives each message a given identifier.

Bugzilla Number	Description
	When a new message arrives with the same identifier, the old message is overwritten.
453535	A clustering feature was added that causes a group of brokers to behave like a single broker and provides fault tolerance.
462046	Upon a broker crash, the database will be automatically recovered. Additionally, if the database is from a previous version of DB4, it will be upgraded.
467735	The character set and encoding is now determined automatically by the system.
453530	The XML exchange is now built as a separate module that can be loaded for use with Red Hat Messaging.
460664	The broker now authenticates the user ID of any published messages. An error is thrown and the message will not be delivered if the user ID is different to that of the authenticated session.
461292	A wait() method was added to the public interface of the SubscriptionManager and Dispatcher modules, to allow threads to be correctly shutdown.
446213	Transactions have been optimized so that they no longer cause a write to disk when all the consumed messages are transient.
453527	Infiniband is now built as a separate module that can be loaded for use with Red Hat Messaging.
431874	A new client management agent is now included.
467883	TX transactions within a cluster will now replicate their transactional states.
468859	A feature was added that allows the user to specify the order in which URLs should be tried for failover.
470261	When a message is sent to a cluster, and the message is received but is not accepted, other members of the cluster need to replicate the credit information associated with that transaction.
467882	The cluster protocol was extended to ensure consistent completion between cluster members.
467014	Support was added for multi-byte character sets with UTF-8.
468173	The cluster now supports the Application Interface Specification (AIS) Redundant Ring

Bugzilla Number	Description
470288	The MRG Management Console now displays a shorted version of the command property (showing the full path to the job's executable) in the job pane.
470315	The MRG Management Console was showing that it had at least one Scheduler and one Collector, even when the tabs did not contain any. This was occurring because the queries were assuming that corresponding system data existed. This has now been fixed to handle the case where the system data is missing.
472215	qpidd rmgr::get_events() was producing a JERR__AIO: AIO error. This was caused by a logic error and resolved by ensuring the readsize is floored to the closest sblk boundary
473122	Incorrect records being enqueued or dequeued as part of a recovered transaction were resulting in either a shortfall or excess of records. txtest was modified to use a 36 character string representation of UUIDs for xids.
472813	If the broker was forced into external mode and all messages were stored in the BDB rather than the journal itself, then txtest failed with a broker exception. This was caused by a logic error, and fixed by causing the read pipeline to be able to handle out-of-order RIDs.
473496	A performance regression occurred in broker management. This regression was resolved,
472145	The ruby-qpidd MRG Ruby Client package was included.
471290	A feature was added to monitor CMan quorum status.
466822	When using RDMA, topic exchange can cause deadlocks. This was resolved by rearranging the locking order.
459949	Producer performance dropped as subscribers were increased to more than 2 on a single queue. This issue was resolved.
472242	Cluster race conditions were discovered and rectified.
470287	The MRG Management Console was displaying incorrect behavior when the Grid tab was clicked. This behavior was corrected.
470299	The number of slots reported was multiplied by the number of HA Collectors in a pool. This behavior was rectified.

Bugzilla Number	Description
472456	Spurious errors were being logged by the FailoverListener. These errors are now suppressed.
472526	Non-functional components of qpidd-perftest were removed
472806	Setting a negotiator limit was resulting in an error. This has now been corrected
472983	A feature was added to Mint that purges the statistic update and job records. The time between purges is configurable and defaults to 24 hours
470588	The primary group of the qpidd user is ais when installing cluster module
471467	Loaded modules could not be enabled in the MRG Management Console due to a capitalization error. This was corrected.
471507	Default Scheduler statistics default timing changed from 10 minutes to 60 minutes.
472928	Threads were hanging waiting for an enqueue completion notification that never arrives. This was resolved by removing the check for a stopped condition
474054	The chkconfig metadata was added in to the sesame init scripts
474373	When an agent connected to the QMF Broker crashes the broker can retain partial schema information. The behavior was changed so that the information was completed when the agent reconnects.
474173	A deadlock involving ConnectionImpl::close() and ConnectionImpl::shutdown() was resolved
474371	qpidd+store exits on journal recovery with the message Timeout waiting for AIO in MessageStoreImpl::recoverMessages() . This was corrected by adjusting the sblk offset
474516	QueueOptions.setOrdering(FIFO) never returns if LVQ was not set previously. This has been fixed
474657	Status updates are no longer published for objects with no status
470285	Slots were individually publishing system information. This behavior was changed so that only the broker publishes system information.

Bugzilla Number	Description
470323	The MRG Management Console Slot Utilization graph displays all slots published to all brokers. This behavior was changed to show only those slots associated with the system

Table 2.1. MRG Messaging Bug Fixes

Enhancements

The following is a list of enhancements in the 1.1 release of MRG Grid:

Bugzilla Number	Description
431874	A new client management agent is now included.
441700	A ring queue feature has been implemented. This is a fixed size queue, where new publications result in the oldest message being dropped.
446213	Transactions have been optimized so that they no longer cause a write to disk when all the consumed messages are transient.
446938	Support was added for RBAC and ACL, in order to allow multiple message flows to securely use the same broker.
452015	A custom exchange was created that imposes an order on all routed messages, and adds a header to all delivered messages with a queue sequence number.
452017	A Last Value Queue (LVQ) was added. This queue gives each message a given identifier. When a new message arrives with the same identifier, the old message is overwritten.
453527	Infiniband is now built as a separate module that can be loaded for use with Red Hat Messaging.
453530	The XML exchange is now built as a separate module that can be loaded for use with Red Hat Messaging.
453535	A clustering feature was added that causes a group of brokers to behave like a single broker and provides fault tolerance.
453536	Support was added for a feature that allows the user to set Time To Live (TTL) values for message expiration.
460664	The broker now authenticates the user ID of any published messages. An error is thrown and the message will not be delivered if the user ID is different to that of the authenticated session.

Bugzilla Number	Description
461292	A <code>wait()</code> method was added to the public interface of the SubscriptionManager and Dispatcher modules, to allow threads to be correctly shutdown.
462046	Upon a broker crash, the database will be automatically recovered. Additionally, if the database is from a previous version of DB4, it will be upgraded.
467014	Support was added for multi-byte character sets with UTF-8.
467735	The character set and encoding is now determined automatically by the system.
467882	The cluster protocol was extended to ensure consistent completion between cluster members.
467883	TX transactions within a cluster will now replicate their transactional states.
468173	The cluster now supports the Application Interface Specification (AIS) Redundant Ring
468859	A feature was added that allows the user to specify the order in which URLs should be tried for failover.
470261	When a message is sent to a cluster, and the message is received but is not accepted, other members of the cluster need to replicate the credit information associated with that transaction.
470288	The MRG Management Console now displays a shorted version of the command property (showing the full path to the job's executable) in the job pane.
470315	The MRG Management Console was showing that it had at least one Scheduler and one Collector, even when the tabs did not contain any. This was occurring because the queries were assuming that corresponding system data existed. This has now been fixed to handle the case where the system data is missing.

Table 2.2. MRG Grid Enhancements

2.2. MRG Realtime Known Issues

The 1.1 release of MRG Realtime contains updated kernel packages that fix several security issues and several bugs.

This update has been rated as having important security impact by the Red Hat Security Response Team.

This release relates to the following Errata:

- RHSA-2009:0009
- RHSA-2009:0053
- RHBA-2009:0037

Security Fixes

The following is a list of security fixes in the 1.1 release of MRG Realtime:

CVE Number	Rating	Description
CVE-2008-0598	Important	Tavis Ormandy discovered a deficiency in the Linux kernel 32-bit and 64-bit emulation. This could allow a local, unprivileged user to prepare and run a specially-crafted binary which used this deficiency to leak uninitialized and potentially sensitive data
CVE-2008-3831	Important	Olaf Kirch reported a flaw in the i915 kernel driver that only affects the Intel G33 series and newer. This flaw could, potentially, lead to local privilege escalation
CVE-2008-4554	Important	Miklos Szeredi reported a missing check for files opened with <code>O_APPEND</code> in the <code>sys_splice()</code> . This could allow a local, unprivileged user to bypass the append-only file restrictions
CVE-2008-4576	Important	A deficiency was found in the Linux kernel Stream Control Transmission Protocol (SCTP) implementation. This could lead to a possible denial of service if one end of a SCTP connection did not support the AUTH extension
CVE-2008-4618	Important	Wei Yongjun reported a flaw in the Linux kernel SCTP implementation. In certain code paths, <code>sctp_sf_violation_paramlen()</code> could be called with a wrong parameter data type. This could

CVE Number	Rating	Description
		lead to a possible denial of service
CVE-2008-5029	Important	The __scm_destroy function in the Linux kernel could make indirect recursive calls to itself when the fput function was called to close a socket. This could lead to a possible local denial of service attack
CVE-2008-3528	Low	The ext2 and ext3 filesystem code failed to properly handle corrupted data structures, leading to a possible local denial of service issue when read or write operations were performed

Table 2.3. MRG Realtime Security Fixes

Bug Fixes

The following is a list of bug fixes in the 1.1 release of MRG Realtime:

Bugzilla Number	Description
460102	Several System on Chip (SoC) audio drivers allocated memory in the platform device probe function but did not free this memory in the event of an error. Instead, the memory was freed in the device probe function's error path. This could result in a <i>double free</i> error. With this update, errors cause memory to be freed correctly
462281 & 462282	When a check was made to see if the netlink attribute fitted into available memory, the value returned, <i>remaining</i> , could become negative due to alignment in nla_next() . GCC set <i>remaining</i> to unsigned when testing against the sizeof(*nla) , however. As a consequence, the test would always succeed and the function nla_for_each_attr() could, potentially, access memory outside the received buffer. With this update, sizeof is cast to an integer, ensuring sizeof(*nla) does a signed test and preventing the illegal memory de-reference
465744 & 465745	If a user-space process used a <i>SIGIO</i> notification and did not disable it before closing the file descriptor, a stale pointer was left in the async_queue of the real-time clock. When a different user-space process subsequently used a <i>SIGIO</i> notification, the kernel de-referenced this pointer and crashed. With this updated

Bugzilla Number	Description
	kernel, <i>SIGIO</i> notifications are disabled when the file descriptor is closed, preventing this
458797	The real-time kernel included with Red Hat Enterprise MRG did not randomize <i>exec</i> , <i>heap</i> or <i>libc</i> addresses. This update corrects this error: <i>exec</i> , <i>heap</i> or <i>libc</i> addresses are now randomized
465862	The Infiniband driver unlocked rwlocks in the same order they were received. This creates kernel errors as rwlocks are expected to unlock in reverse order. This has been fixed by allowing rwlocks to release out of order
469186	An extra function call was being made, which was causing latency regression problems. The call was made an inline function which returned the latency to normal levels
455095	The events_trace latency tracer interface requires assembly-linked annotation. It also requires that the edx register is used before it can be stored on the stack. When used on i386 architectures, these can result in the events_trace latency tracer giving bad parameters to syscalls. This was corrected by adding asmLinkage annotation to the sys_call function, and changing the order in which data was passed to the edx register
466153, 466557, 466558, 466563, 466564, 466565, 466568, 466569 & 466571	Network drivers update: e1000e, tg3, bnx2, igb, ixgb, ixgbe, bnx2x, cxgb3 and solarflare
461330, 466554 & 466572	Storage drivers update: qla2xxx and lpfc
466750	select() and poll() have an effective resolution of 1 to 10 milliseconds (1000 - 100Hz). Applications cannot specify more accurate delays. The internals of select() and poll() interfaces have been changed to use hrtimers, which allows applications to operate at the nanosecond level
467739	Added amd64_edac driver
467783	Added infrastructure for SCSI device handlers with scsi_dh .
468205	The OCFS2 filesystem has not been tested for use with MRG Realtime. The feature was disabled until further testing can be undertaken
468231	New Red Hat Enterprise MRG drivers enabled in the MRG Realtime kernel: amd64_edac and bnx2x

Bugzilla Number	Description
468988	Disable KVM in the MRG Realtime kernel configuration
467781	A bug that caused corruption in e1000e chips was found in the dynamic ftrace tracer, which is used in the Red Hat Enterprise MRG kernel. Protection against this bug was added by preventing code from being processed that might be freed later on
466341	The MRG Realtime kernel failed to boot on Intel Canelands processors (16 and 24 cores). This was fixed by making the mem_map allocations continuous.
460217	A spurious 40ms latency was found in messaging workloads. This was corrected by making the tcp_delack_min able to be tuned using /proc .

Table 2.4. MRG Realtime Bug Fixes

2.3. MRG Grid Known Issues

The 1.1 release of MRG Grid contains updated packages that fix several known bugs.

This release relates to the following Errata:

- RHBA-2009:0037
- RHBA-2009:0036

Bug Fixes

The following is a list of bug fixes in the 1.1 release of MRG Grid:

Bugzilla Number	Description
471958	Jobs submitted to the Xen VM universe were not starting. This was caused by the <code>SendAliveToParent()</code> function failing. This was resolved by disabling CHILDALIVE functions for the gahp servers.
457496	When running condor_version on a platform other than Red Hat Enterprise Linux 5, the command returned an error stating that the platform was UNKNOWN . The configure.ac file was modified and the condor_version command now recognizes Red Hat Enterprise Linux 4, Fedora and CentOS platforms.
463042	When jobs are submitted, the submitter information is given in the form owner@uid_domain . Jobs were recording only the owner information, which caused problems with matching jobs to owners. This was changed

Bugzilla Number	Description
	so that jobs record both the owner and the user ID domain information, which allowed jobs and submitters to be matched correctly.
467128	In order to implement the MRG/EC2 Enhanced feature, the job router hooks feature and the parse-Classad-offset fix were included.
467353	The Amazon gahp code in gsoap_commands.C was attempting to fill a buffer using only a single read() call and when it failed, a critical error occurred. The functionality was changed to use a full_read() instead, which performs multiple reads to fill a buffer if necessary, and thereby eliminating the error.
469388	When a node was reconfigured to run with different plugins and then restarted, the master node would fail to recognise the new plugins and subsequently perform a core dump when it was shut down. This was circumvented by causing the problematic Qpid Management Function (QMF) structures to not attempt to clean up. The core issue will be addressed in a future release.
470320	Although Red Hat Enterprise Linux 4 machines were loading plugins correctly, they were not publishing their existence in the pool. This was caused by a call to <code>failover.reset</code> never returning because it existed in a different thread. It was not fixed, but a workaround was implemented until a fix is found. The workaround involved defining FieldTable::clear in a .cpp file, in addition to making qmf-gen generate separate blocks.
470326	A debug command was causing options to be printed as they were parsed, resulting in unexpected output upon command execution. The command was removed, and output was tested to be as expected.
470862	The condor_schedd was logging information about failure to change the ownership of non-existent job spool directories when the logging level was set to <i>D_ALWAYS</i> . In large pools, this created significant latencies as useless entries were logged and caused the file to rotate too often. This behaviour was changed to only occur at the <i>D_FULLDEBUG</i> debug level.
470870	A bug in the Qpid Management Framework (QMF) was causing the condor_schedd daemon to run out of file descriptors after a certain amount of jobs were submitted, and an exception

Bugzilla Number	Description
	would be thrown. The QMF bug was resolved, fixing this issue.
471270	When submitting jobs, an error occurred that references the University of Wisconsin's Condor Manual. The error message was reworded, and the reference removed.
471332	When submitting jobs larger than 4GB in the VM universe, an error occurred that stated that the file could not be modified. This was because a function call was suffering a size overflow. The overflow limit was changed to rectify this problem.
471604	The \$(LIB) macro in the /etc/condor/condor_config file always points to /usr/lib . When installed on 64-bit architectures, files are stored in /usr/lib64 . The macro was updated so that it points to the correct location when used on 64-bit machines.
471620	A dependency was added for the mkisofs package.
468221	The condor keyboard daemon detects events occurring in X. This information is used to vacate jobs when a user returns. This feature was not working under Linux. A new X_CONSOLE_DISPLAY parameter was created to fix this problem and contains the name of the display that the keyboard daemon should monitor.
458889	After the job hooks have been run, the starter was entering a mode with incorrect privileges and was unable to access files created by a job. This was fixed by causing the hooks to run with the correct privileges.
472084	Jobs submitted to the VM universe were creating files that Xen could need due to SELinux permissions. This was fixed by creating a policy that defines a type and context for Condor.
471902	When the condor_master was shut down, the PID file was not being cleaned up. This was fixed and the PID file is now cleaned up as expected.
471886	Condor was not on by default. The init script was edited to resolve this.
471888	The low-latency feature was not on by default. The init script was edited to have carod enabled to resolve this.

Bugzilla Number	Description
471889	The MRG/EC2 Enhanced feature was not on by default. The init script was edited to have the caroniad daemon enabled to resolve this.
472129	A series of unexplained errors were occurring in the VMGaHPLog file, followed by a shutdown. Additional debugging information was added, which resolved the problem.
471957	The caro daemon needed to have the universe specified, even in situations where the universe had already been set. This was fixed by causing the daemon to check for the vanilla universe every time. If no job universe is specified, it will set it to vanilla before submitting the job.
472419	Errors were occurring when Condor was installed on a Red Hat Enterprise Linux 5 system with SELinux disabled. This was fixed by adding a check and only setting up the SELinux context if SELinux is enabled.
471628	The vmgahp startup was failing on systems where mkisofs was not installed. This is correct behaviour, but the error was not logged. The output of /usr/sbin/condor_vm_xen.sh check was included in the log to resolve this.
472421	Errors were occurring when Condor was installed on a Red Hat Enterprise Linux 4 system with SELinux disabled. This was fixed by causing the condor_startd daemon to run by default in an unconfined context on Red Hat Enterprise Linux 4.
471901	The condor_master daemon should wait for all daemon it is managing to exit before it shuts down. However, the Condor init script uses killproc to shutdown the condor_master which does not allow it to exit gracefully. The init script was changed so that it uses SIGQUIT instead.
472837	Jobs were not being scheduled correctly with disk constraints were specified. This was fixed by changing the way disk constraints were calculated.
454456	If the carod daemon was stopped it could not be immediately restarted. It was discovered that a socket was not being closed correctly and this issue was resolved.
474071	Hook scripts were converting regular expression lines to lowercase, which was causing problems with the values of those regular

Bugzilla Number	Description
	expressions being lowercased also. The logic was restructured to avoid this problem.
474725	The security session cache was not handling privilege-switching accurately when read access required authentication. This was fixed by causing an authenticated security session to be created only when the read level access requires authentication.
474749	The mailer was causing errors when finalizing MRG/EC2 Enhanced jobs. This was caused by a misconfiguration that was subsequently fixed.
474931	The header of the global condor_config file was changed to strongly discourage users from editing it. It now instructs users to edit the local_config_file instead.
475235	When two schedulers are configured for high availability, and a failover occurs, duplicate jobs are created. This issue was resolved by adding a QMF_STOREFILE option to the QMF plugin. All high availability schedulers should configure SCHEDD.QMG_STOREFILE to be in a shared location they can all access, in order to resolve this issue.
474862	File permission errors were experienced on Red Hat Enterprise Linux 4 systems. These errors have been fixed.

Table 2.5. MRG Grid Bug Fixes

Enhancements

The following is a list of enhancements in the 1.1 release of MRG Grid:

Bugzilla Number	Description
429787	The management console is now able to view statistics about components within a pool.
429788	The management console can now act on components of a pool, such as starting components and changing policies.
447069	The condor_gridmanager was changed to enable resource limits to be applied differently depending on the grid type. This was achieved through the new configuration parameters GRIDMANAGER_MAX_SUBMITTED_JOBS_PER_RESOURCE_<gridtype> and GRIDMANAGER_MAX_PENDING_SUBMITS_PER_RESOURCE_<gridtype>
453131	The MRG/EC2 Enhanced feature was added to allow for the transparent running of vanilla jobs

Bugzilla Number	Description
	using the Enhanced Cloud Computing (EC2) feature. It builds on the MRG/EC2 Basic feature already in Condor.
454430	Support was added for scheduling very small jobs using the AMQP and MRG Messaging functionality.
454450	The low-latency feature was packaged into an RPM that depends on condor.
465167	A function was added that allows remote configuration and disabling of features on machines in the pool, using puppet.
468196	By default, condor has the kbdd feature disabled on Linux. This was changed so that the kbdd package is built and included when run on Linux platforms.
462471	A backup of the job queue log is maintained so that jobs are not lost in a failure.

Table 2.6. MRG Grid Enhancements

Other changes

The following is a list of other changes in the 1.1 release of MRG Grid:

Bugzilla Number	Description
470138	Glide-in is not supported in version 1.1.
470385	Deployment tools are not supported in version 1.1.
470394	The standard universe is not supported in version 1.1.
470399	The condor_configure and condor_install commands have been removed, as other existing tools perform the same functions.
470400	The condor_master_off and condor_schedd_reconfig commands have been deprecated from version 1.1.
470401	The condor_static package is not supported in version 1.1.
470406	Globus and the components that rely on it has been removed from version 1.1.
470937	Previous versions of Condor introduced dependancies on qpdc and qmf in order to provide the QMF plugins. The condor-qmf-plugins package dependancy was removed from version 1.1.

Bugzilla Number	Description
471747	As the standard universe is not supported in this version, the default universe has been changed from standard to the vanilla universe.
474690	All bugs in the RHBA-2008:0990 errata were built for Red Hat Enterprise Linux 4

Table 2.7. MRG Grid other changes

Appendix A. Revision History

Revision 1.0	Tue Feb 3 2009	Lana Brindley lbrindle@redhat.com
Removed BZ#483238		
Revision 0.14	Mon Feb 2 2009	Lana Brindley lbrindle@redhat.com
BZ#483238		
Revision 0.13	Thu Jan 22 2009	Lana Brindley lbrindle@redhat.com
BZ#481054		
Revision 0.12	Mon Jan 19 2009	Lana Brindley lbrindle@redhat.com
Updated errata numbers		
Revision 0.11	Mon Jan 19 2009	Lana Brindley lbrindle@redhat.com
Added links to product page		
Revision 0.10	Fri Jan 16, 2009	Lana Brindley lbrindle@redhat.com
BZ #479466 - reworded		
Revision 0.9	Mon Jan 12, 2009	Lana Brindley lbrindle@redhat.com
BZ #479466		
Revision 0.8	Fri Jan 9, 2009	Lana Brindley lbrindle@redhat.com
RHEA-2008:0994 Updated System Req's		
Revision 0.7	Wed Jan 7, 2009	Lana Brindley lbrindle@redhat.com
BZ#479015		
Revision 0.6	Fri Dec 12, 2008	Lana Brindley lbrindle@redhat.com
New bugs in RHBA-2008:0990 (BZ#474725, BZ#474749, BZ#474931, BZ#475235) RHBA-2008:1033 RHEA-2008:0994		
Revision 0.5	Tue Dec 9, 2008	Lana Brindley lbrindle@redhat.com
Updates from RHBA-2008:0990		
Revision 0.4	Mon Dec 8, 2008	Lana Brindley lbrindle@redhat.com

Appendix A. Revision History

BZ #457922

BZ #468713

BZ #470067

Revision 0.3 Fri Dec 5, 2008

Updated System Req's

Lana Brindley lbrindle@redhat.com

Revision 0.2 Wed Dec 3, 2008

Added information from RHBA-2008:0990

Lana Brindley lbrindle@redhat.com

Revision 0.1 Tue Dec 2, 2008

Initial document creation

Lana Brindley lbrindle@redhat.com