

RED HAT® ENTERPRISE LINUX®

DATASHEET

AVAILABILITY ADD-ONS

Resilient Storage

RED HAT ENTERPRISE LINUX RESILIENT STORAGE ADD-ON

Gain improved resiliency to ensure that data is available when you need it



The Red Hat Enterprise Linux Resilient Storage Add-On offers improved resiliency to ensure that data is available to meet the needs of the business by enabling a shared storage or clustered file system to access the same storage device over a network.

By providing consistent storage across a cluster of servers, Red Hat's Resilient Storage Add-On creates a pool of data that is not only available to each server in the group, but also protected if any one server fails. This Add-On works in parallel with the Red Hat Enterprise Linux High Availability Add-On.

KEY FEATURES

Red Hat's Resilient Storage Add-On includes: a clustered file system Global File System 2(GFS2), a clustered Logical Volume Manager (CLVM), clustered mirroring capabilities(cmirror) and clustering extensions for Samba (CTDB).

GFS2 Shared Storage File System

GFS2 is a clustered file system which ensures that data is shared among GFS2 nodes with a single, consistent, and coherent view of the file system name space. GFS2 is generally used for targeting high availability clusters of 2-16 nodes.

Ensure Data and Cache Consistency

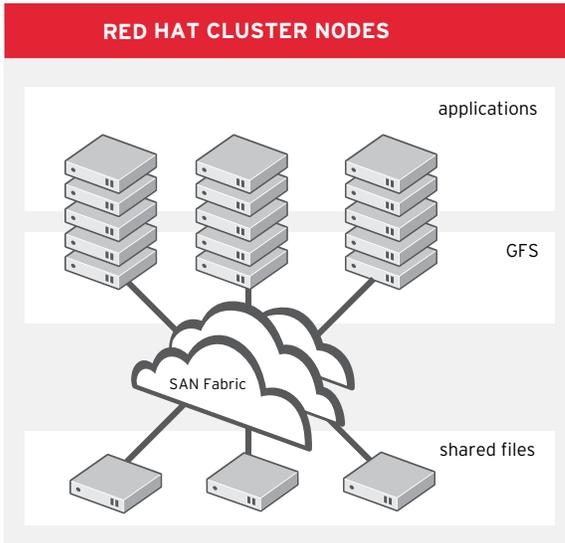
When using the GFS2 file system, a single version of all files in a cluster is visible to all nodes within that cluster. Each server in a cluster has direct access to a shared block device over a local storage area network (SAN) of up to 100 terabytes. Data and cache consistency is ensured using a cluster-wide locking mechanism called distributed lock manager (DLM) to arbitrate access to the storage. Each member of the cluster thus has direct access to the same storage device, and all cluster nodes access the same set of files.

READ ABOUT OTHER
ADD-ON OPTIONS
AVAILABLE

www.redhat.com/rhel/add-ons/

In order to prevent members of the cluster from updating the same storage block at the same time – causing data corruption – GFS2 uses a cluster-wide locking mechanism to arbitrate access to the storage as a concurrency-control mechanism. For example, before creating a new file or writing to a file that is opened on multiple servers, the file system component on the server must obtain the correct lock. (See Figure 1.)

Figure 1



GFS2's sharing style file system ensures data consistency and allows for the most complete coverage of your data.

GFS2 performs well with workloads where each node writes almost exclusively to a particular set of files that are not shared with other nodes, or where a set of files is to be shared in an almost exclusively read-only approach across a set of nodes. This results in a minimum of cross-node cache invalidation; thus maximizing performance.

GFS2 file systems provide broader and more mainstream support in that it:

- Supports SELinux extended attributes
- Supports nanosecond time-stamp
- Introduces Dynamic journal addition
- Is part of the upstream kernel (integrated into 2.6.19)

While a GFS2 file system can be implemented either in a standalone system or as part of a cluster configuration, Red Hat Enterprise Linux 6 does not support the use of GFS2 as a single-node file system. Red Hat does support a number of other high-performance single-node file systems to reduce the overhead as compared to a cluster file system.

Cluster Logical Volume Manager (CLVM)

The Clustered Logical Volume Manager (CLVM) is a set of clustering extensions to Linux Volume Manager. These extensions allow a cluster of computers to manage shared storage (for example, on a SAN) using LVM. While CLVM is part of the Resilient Storage Add-On, you also need the High Availability Add-On to leverage the functionality of CLVM.

Consider using CLVM based on your system requirements and the following guidelines:

- If only one node of your system requires access to the storage you are configuring as logical volumes, then you can use LVM without the CLVM extensions and the logical volumes created with that node are all local to the node.
- If you are using a clustered system for failover where only a single node that accesses the storage is active at any one time, you should use High Availability Logical Volume Management agents (HA-LVM).
- If more than one node of your cluster will require access to your storage which is then shared among the active nodes, then you must use CLVM. CLVM allows a user to configure logical volumes on shared storage by locking access to physical storage while a logical volume is being configured, and uses clustered locking services to manage the shared storage.

Clustered SAMBA

Clustered SAMBA (CTDB) is a cluster implementation of the TDB database used by Samba. To use CTDB, a clustered file system must be available and shared on all nodes in the cluster. GFS2 provides that functionality. CTDB provides clustered features on top of GFS2. As of the Red Hat Enterprise Linux 6.2 release, CTDB also runs a cluster stack in parallel to the one provided by Red Hat Enterprise Linux clustering. CTDB manages node membership, recovery/failover, IP relocation and Samba services.

Classic sharing methods are required to share Samba's Trivial Database (TDB) – a database API that allows multiple simultaneous writers and uses locking internally to keep writers from trampling on each other.

Cluster-Wide Mirroring (Cmirror)

Cmirror enables creation of a mirrored logical volume in a cluster. In order to create a mirrored LVM volume in a cluster the cluster and cluster mirror infrastructure must be running, the cluster must be quorate, and the locking type in the lvm.conf file must be set correctly to enable cluster locking.

FEATURE SUMMARIES

GFS2	GFS2 is a clustered file system in which data is shared among GFS2 nodes with a single, consistent, and coherent view of the file system name space. Processes on different nodes may work with GFS2 files in the same way that processes on one node can share files in a local file system.
CLVM	The CLVM provides a cluster-wide version of LVM2. CLVM provides the same capabilities as LVM2 on a single node, but makes the volumes available to all nodes in a Red Hat cluster. The logical volumes created with CLVM are available to all nodes in a cluster in a single, consistent, view.
CTDB	CTDB is a specialized distributed database used to clusterize Samba on top of a distributed file system like GFS2. By using CTDB Samba is able to properly share files from multiple nodes at the same time (active-active) maintaining full SMB semantics without degrading performances.
Cmirror	Cmirror provides the capability to create mirrored clustered logical volumes. It is an extension of logical volume mirroring as applicable in a stand-alone environment to a clustered environment.

COMPATIBILITY OF SERVERS AND RED HAT ENTERPRISE LINUX VERSIONS

The Load Balancer Add-On is fully compatible with other Red Hat Enterprise Linux Add-Ons, including the High Availability Add-On (for load balancing applications across redundant servers).

SERVER AND VERSION COMPATIBILITY

Red Hat Enterprise Linux Version	Variants	Releases
Red Hat Enterprise Linux 5	Red Hat Enterprise Linux Server Red Hat Enterprise Linux AP	Red Hat Enterprise Linux 5.5 and later
Red Hat Enterprise Linux 6	Red Hat Enterprise Linux Serve Red Hat Enterprise Linux for SAP Business Apps	Red Hat Enterprise Linux 6.0 and later

In conjunction with Red Hat Enterprise Linux, Red Hat offers a portfolio of Add-Ons to extend the features of your Red Hat Enterprise Linux subscription. Add-Ons allow you to tailor your application environment to suit your particular computing requirements. With increased flexibility and choice, you can select the availability, scalability, and management features required by your organization when and where they are needed.

DATASHEET**EXECUTIVE SUMMARY**

Red Hat's Resilient Storage Add-On enables a shared storage or clustered file system to access the same storage device over a network. By providing consistent storage across a cluster of servers, Red Hat's Resilient Storage Add-On creates a pool of data that is available to each server in the group, but that is also protected in case any one server fails. You can read more about our portfolio of Add-On offerings by visiting <http://www.redhat.com/rhel/add-ons/>.

HOW TO ORDER THE RESILIENT STORAGE ADD-ON

The Red Hat Enterprise Linux Resilient Storage Add-On is delivered through Red Hat Network, similar to the way regular Red Hat Enterprise Linux content is provided. Customers using the RHN Satellite can make use of Satellite to manage Load-Balancer systems.

To order Red Hat Enterprise Linux Resilient Storage Add-On, please contact your local Red Hat Account Representative.

ABOUT RED HAT

Red Hat, the world's leading provider of open source solutions and an S&P 500 company, is headquartered in Raleigh, NC with more than 70 offices spanning the globe. Red Hat provides high-quality, affordable technology with its operating system platform, Red Hat Enterprise Linux, together with cloud, virtualization, management, storage and service-oriented architecture (SOA) solutions, including Red Hat Enterprise Virtualization and JBoss Enterprise Middleware. Red Hat also offers support, training and consulting services to its customers worldwide.

SALES AND INQUIRIES

NORTH AMERICA 1-888-REDHAT1 www.redhat.com	EUROPE, MIDDLE EAST AND AFRICA 00800 7334 2835 www.europe.redhat.com europe@redhat.com	ASIA PACIFIC +65 6490 4200 www.apac.redhat.com apac@redhat.com	LATIN AMERICA +54 11 4329 7300 www.latam.redhat.com info-latam@redhat.com
---	--	--	--