

Release Notes for SUSE Linux Enterprise High Availability Extension 11 SP3




Release Notes for SUSE Linux Enterprise High Availability Extension 11 SP3

Version 11.3.13 (2013-05-27)

Abstract

These release notes apply to all SUSE Linux Enterprise High Availability Extension 11 SP3 based products (e.g. for x86, x86_64, Itanium, Power and System z). Some sections may not apply to a particular architecture or product. Where this is not obvious, the respective architectures are listed explicitly in these notes. Instructions for installing SUSE Linux Enterprise High Availability Extension can be found in the README file on the CD.

Manuals can be found in the docu directory of the installation media. Any documentation (if installed) can be found in the /usr/share/doc/ directory of the installed system.

This SUSE product includes materials licensed to SUSE under the GNU General Public License (GPL). The GPL requires SUSE to provide the source code that corresponds to the GPL-licensed material. The source code is available for download at <http://www.suse.com/download-linux/source-code.html> . Also, for up to three years after distribution of the SUSE product, upon request, Novell will mail a copy of the source code. Requests should be sent by e-mail to mailto:sle_source_request@novell.com  or as otherwise instructed at <http://www.suse.com/download-linux/source-code.html> . Novell may charge a reasonable fee to recover distribution costs.

Contents

1 Purpose	5
2 Features and Versions	6
2.1 hawk's Cluster Simulator Now Allows Exploration of Configuration Changes	6
2.2 Pacemaker Initial Monitor (probe) Timeout Now Uses Different Default	6
2.3 Pacemaker Utilization Is Now Also Supported for Groups	6
2.4 SBD takes quorum and pacemaker state into account during self-fencing decision	7
2.5 Cluster File System—Oracle Cluster File System 2 (OCFS2)	7
2.6 Clustered Logical Volume Manager2—cLVM2	8
2.7 High Availability Cluster Resource Manager (Pacemaker)	8
2.8 Cluster and Distributed Systems Infrastructure (corosync and OpenAIS)	9
2.9 Server	9
2.9.1 RAID10 support in cmirrord	9
2.10 Network	10
2.10.1 Linux Virtual Server Load Balancer (ipvs) Extends Support for IPv6	10
2.11 Resource Management	10
2.11.1 Data Replication—Distributed Remote Block Device (DRBD)	10
2.11.2 IP Load Balancing—Linux Virtual Server (LVS)	11
2.11.3 Relax and Recover (ReaR)	11
2.11.4 Distributed Lock Manager (DLM)	11
2.12 Systems Management	11
2.12.1 hawk status display now allows for direct transfer to the history explorer	11
2.12.2 Improve usability and accessibility of the cluster's web interface hawk	12
2.12.3 YaST2 configuration module provided for Relax and Recover (ReaR) bare metal recovery tool	12
3 Changed Functionality in SUSE Linux Enterprise High Availability Extension 11 SP3	13
3.1 pacemaker-pygui Rename and Split	13
3.2 New Packages from SP2 not Installed Automatically on Update	13
3.3 Non-production Fencing/STONITH Agents Moved	13
4 Deprecated Functionality in SUSE Linux Enterprise High Availability Extension 11 SP3	14
4.1 OCFS2's O2CB Stack.	14
4.2 Samba Clustered Trivial Database (CTDB)	14
4.3 DRBD Resource Agent	14
4.4 Heartbeat	15
4.5 EVMS2 Replaced with LVM2	15
5 Infrastructure, Package and Architecture Specific Information	16
5.1 Architecture Independent Information	16
5.1.1 Changes in Packaging and Delivery	16
5.2 Systems Management	16
5.2.1 Dashboard for multiple clusters	16
6 Other Updates	17
6.1 crm Shell:	17
7 Update-Related Notes	18
7.1 Update from SUSE Linux Enterprise High Availability Extension 11 SP 2	18
7.1.1 Pacemaker Initial Monitor (probe) Timeout Now Uses Different Default	18
8 Supported Deployment Scenarios SUSE Linux Enterprise High Availability Extension 11 SP3.....	19
8.1 Local Cluster	19

8.2 Metro Area Cluster	19
8.3 Geographical Clustering	19
9 Known Issues in SUSE Linux Enterprise High Availability Extension 11 SP3	21
9.1 Linux Virtual Server Tunnelling Support	21
9.2 Samba CTDB Cluster Rolling Update Support	21
10 Further Notes on Functionality	22
10.1 Cluster-concurrent RAID1 Resynchronization	22
10.2 Quotas on OCFS2 Filesystem	22
11 Support Statement for SUSE Linux Enterprise High Availability Extension 11 SP3	24
11.1 Long-term Strategy for SLE HA's Graphical User Interface	24
12 Miscellaneous	26
13 Technical Information	27
13.1 OCF Resource Agent for oracle no longer creates audit logs	27
13.2 IPv6 Implementation and Compliance	27
13.2.1 Linux Virtual Server Load Balancer (ipvs) Extends Support for IPv6	27
14 More Information and Feedback	28

1 Purpose

SUSE Linux Enterprise High Availability Extension is an affordable, integrated suite of robust open source clustering technologies that enable enterprises to implement highly available Linux clusters and eliminate single points of failure.

Used with SUSE Linux Enterprise Server, it helps firms maintain business continuity, protect data integrity, and reduce unplanned downtime for their mission-critical Linux workloads.

SUSE Linux Enterprise High Availability Extension provides all of the essential monitoring, messaging, and cluster resource management functionality of proprietary third-party solutions, but at a more affordable price, making it accessible to a wider range of enterprises.

It is optimized to work with SUSE Linux Enterprise Server, and its tight integration ensures customers have the most robust, secure, and up to date high availability solution. Based on an innovative, highly flexible policy engine, it supports a wide range of clustering scenarios.

With static or stateless content, the High Availability cluster can be used without a cluster file system. This includes web-services with static content as well as printing systems or communication systems like proxies that do not need to recover data.

Finally, its open source license minimizes the risk of vendor lock-in, and its adherence to open standards encourages interoperability with industry standard tools and technologies.

In Service Pack 3, a large number of improvements have been added, some of which are called out explicitly here. For the full list of changes and bugfixes, refer to the change logs of the RPM packages. Note that these changes are in addition to those that have already been added with Service Pack 1 and 2.

2 Features and Versions

This section includes an overview of some of the major features and new functionality provided by SUSE Linux Enterprise High Availability Extension 11 SP3.

2.1 hawk's Cluster Simulator Now Allows Exploration of Configuration Changes

hawk's cluster simulator allowed the exploration of the cluster's reaction to status events such as node down or resource monitoring failures only. Examining the predicted result for changes to the cluster configuration was only possible via the crm shell.

hawk now supports the creation of a shadow configuration, and can be asked to explore the effects of applying this new configuration to the cluster.

2.2 Pacemaker Initial Monitor (probe) Timeout Now Uses Different Default

Pacemaker executes an initial monitor for each service on every node ("probe"). The timeout for this monitor operation could be specified by configuring a monitor operation with interval set to 0, or would otherwise be inherited from the cluster-wide default operation timeout.

To ease the need to configure a dedicated monitor timeout for the initial probe, pacemaker will now use the timeout from the most frequent configured monitor automatically. If no monitor operation is configured at all, the cluster-wide default continues to apply.

Users which relied on the probe inheriting the timeout from the cluster-wide timeout in the presence of other monitor operations are advised to configure an explicit monitor operation instead.

2.3 Pacemaker Utilization Is Now Also Supported for Groups

Pacemaker supports utilization-based resource placement in the cluster. Previously, utilization from individual resources was not summarized to the whole group. This could lead to scenarios where only a few resources in the group could be placed, but the rest would remain stopped.

The utilization configured is now summed up for the whole group so that the whole group will be placed correctly.

2.4 SBD takes quorum and pacemaker state into account during self-fencing decision

Some scenarios experience temporary outages of the SBD partitions. If the majority of fencing devices were lost, this would cause a node fence, even if the cluster was otherwise healthy.

SBD now takes into account if the cluster partition remains quorate and whether the local node is healthy from pacemaker's state. This functionality must be explicitly enabled by adding "-P" to /etc/sysconfig/sbd. Please review the updated sbd manpage for a full discussion of this functionality.

2.5 Cluster File System—Oracle Cluster File System 2 (OCFS2)

Cluster file systems are used to provide scalable, high performance, and highly available file access across multiple instances of SUSE Linux Enterprise High Availability Extension servers. Oracle Cluster File System 2 (OCFS2) is a POSIX-compliant shared-disk cluster file system for Linux. OCFS2 is developed under a GPL open source license.

New features included in OCFS2 with this product release are:

- Support for repquota has been added.

Beyond this, OCFS2 continues to deliver the functionality provided in the previous releases:

- Indexed directories, delivering high performance regardless of number of files per directory
- Meta data checksumming, detects all on-disk corruption and capable of correcting some errors transparently
- Improved performance for deletion
- Improved allocation algorithms reduce fragmentation for large files
- Access Control Lists (ACL)
- Quota support
- POSIX conforming file locking
- expand file system during operation

With these features OCFS2 can be used as generic file system for common use without previous limitations to specific workloads. Workloads for OCFS2 included in this product are, but are not limited to:

- Central storage area for virtual machine images
- Central storage area for file servers
- Shared file system for High Availability
- Oracle Database
- All applications using a cluster file system (e.g. Tibco)

The full functionality of OCFS2 is only available in combination with the OpenAIS, corosync, and Pacemaker-based cluster stack.

2.6 Clustered Logical Volume Manager2—cLVM2

The Clustered Volume Manager allows multiple nodes to read and write volumes on a single storage device at the block layer level. It features creation and reallocation of volumes on a shared storage infrastructure like SAN or iSCSI, and allows moving volumes to a different storage device during operation. It can be used for volume snapshots for later recovery if needed.

cLVM2 has been updated to the latest upstream version.

2.7 High Availability Cluster Resource Manager (Pacemaker)

Pacemaker orchestrates the cluster's response to change events such as node failures, resource monitoring failures, permanent or transient administrative changes, and ensures that service availability is recovered.

New features introduced by Pacemaker and included with this product release are:

- Dynamic utilization-based resource placement
- A powerful web console (hawk) for management and monitoring

hawk has been enhanced to support full cluster administration, including access control lists, the cluster test drive, and a new graphical history explorer.

hawk now supports an extensible UI wizard with a set of templates already included that can be used to guide users through standard set up tasks.

- Support for resource templates to simplify and reduce CIB complexity
- Log file query tools available from the commandline shell. The CIB now also records the source and user of a given change to ease history analysis.
- Passwords can now be saved external to the CIB for secure storage.
- The number of concurrent live migrations can now be limited to avoid overloading node capacity.
- The CIB now supports cluster-wide, in addition to node, attributes.

With unified command line support system setup, managing and integration is made easier. To extend High Availability to all types of applications, resource agent templates and templates for configuration examples are provided for customization.

2.8 Cluster and Distributed Systems Infrastructure (corosync and OpenAIS)

The Corosync Cluster Engine is an OSI certified implementation of a complete cluster engine. This component provides membership, ordered messaging with virtual synchrony guarantees, closed process communication groups, and an extensible framework.

- Support unicast communication in addition to multicast and broadcast, with full YaST support.

2.9 Server

2.9.1 RAID10 support in cmirrord

Previous versions only supported one physical volume per mirror side and a total of two mirrors. If more PVs were added to grow the RAID, the logical volumes were possibly allocated on the same side of the mirror.

In addition to growing the single physical volume per mirror side on the SAN level, which remains the recommended solution for growing LVM2 mirrors, it is now possible to add multiple physical volumes per mirror leg.

This can be configured by setting `mirror_legs_require_separate_pvs` to 1 in `/etc/lvm/lvm.conf`, and add tags to the PVs using `"pvchange --addtag"`; PVs with the same tags will be assumed to be on the same mirror leg, and consequently LVs be allocated correctly.

2.10 Network

2.10.1 Linux Virtual Server Load Balancer (ipvs) Extends Support for IPv6

The LVS/ipvs load balancing code did not fully support RFC2460 and fragmented IPv6 packets which could lead to lost packets and interrupted connections when IPv6 traffic was fragmented.

The load balancer has been enhanced to fully support IPv6 fragmented extension headers and is now RFC2460 compliant.

2.11 Resource Management

2.11.1 Data Replication—Distributed Remote Block Device (DRBD)

Data replication is part of a disaster prevention strategy in most large enterprises. Using network connections data is replicated between different nodes to ensure consistent data storages in case of a site failure.

Data replication is provided in SUSE Linux Enterprise High Availability Extension 11 SP3 with DRBD. This software based data replication allows customers to use different types of storage systems and communication layers without vendor lock-in. At the same time, data replication is deeply integrated into the operating system and thus provide ease-of-use. Features related to data replication and included with this product release are:

- YaST setup tools to assist initial setup
- Fully synchronous, memory synchronous or asynchronous modes of operation
- Differential storage resynchronization after failure
- Bandwidth of background resynchronization tunable
- Shared secret to authenticate the peer upon connect
- Configurable handler scripts for various DRBD events
- Online data verification

With these features data replication can be easier configured and used. And with improved storage resynchronization recovery times will be decreased significantly.

The distributed replicated block device (DRBD) version included supports active/active mirroring, enabling the use of services such as cLVM2 or OCFS2 on top.

2.11.2 IP Load Balancing—Linux Virtual Server (LVS)

Linux Virtual Server (LVS) is an advanced IP load balancing solution for Linux. IP load balancing provides a high-performance, scalable network infrastructure. Such infrastructure is typically used by enterprise customers for web servers or other network related service workloads.

With LVS network requests can be spread over multiple nodes to scale the available resources and balance the resulting workload. By monitoring the compute nodes, LVS can handle node failures and redirect requests to other nodes maintaining the availability of the service.

2.11.3 Relax and Recover (ReaR)

New in SUSE Linux Enterprise High Availability Extension 11 SP3

On the x86 and x86-64 architectures, a disaster recovery framework is included. ReaR allows the administrator to take a full snapshot of the system and restore this snapshot after a disaster on recovery hardware.

2.11.4 Distributed Lock Manager (DLM)

The DLM in SUSE Linux Enterprise High Availability Extension 11 SP3 supports both TCP and SCTP for network communications, allowing for improved cluster redundancy in scenarios where network interface bonding is not feasible.

2.12 Systems Management

2.12.1 hawk status display now allows for direct transfer to the history explorer

Failure messages in hawk's status display now directly link to the history explorer to ease issue analysis.

2.12.2 Improve usability and accessibility of the cluster's web interface hawk

A major focus area of SLE HA 11 SP3 has been to improve the accessibility of the web interface and to also make the interface reflect the cluster status better by adding icons to all relevant status messages.

2.12.3 YaST2 configuration module provided for Relax and Recover (ReaR) bare metal recovery tool

The product now includes a YaST2 module for configuring the most common scenarios of ReaR.

Please install the yast2-rear package if you wish to use this component.

3 Changed Functionality in SUSE Linux Enterprise High Availability Extension 11 SP3

3.1 pacemaker-pygui Rename and Split

In response to customer demand, the Python-based GUI component (formerly packaged as `pacemaker-pygui`) has been split into a server and client package, allowing server installs without client software.

The new packages are called `pacemaker-mgmt` for the server, and `pacemaker-mgmt-client` for the client.

After an update from GA, the client package may not automatically be installed, depending on installer settings, and thus the `hb_gui` and `crm_gui` commands unavailable.

The resolution is to install the package manually.

3.2 New Packages from SP2 not Installed Automatically on Update

New functionality provided via new packages is not automatically installed by an update, which strives to preserve the existing functionality. It is recommended to install the HA pattern manually to take advantage of all new functionality in SUSE Linux Enterprise High Availability Extension 11 SP3.

3.3 Non-production Fencing/STONITH Agents Moved

The `ssh`, `external/ssh`, and `null` STONITH agents have been moved to the `libglue-devel` package, and are no longer installed by default.

These fencing agents are not suitable for production environments and should only be used for limited functionality demo setups. This move clarifies their intended use case.

4 Deprecated Functionality in SUSE Linux Enterprise High Availability Extension 11 SP3

4.1 OCFS2's O2CB Stack.

The legacy O2CB in-kernel stack of OCFS2 is only supported in combination with Oracle RAC. Oracle RAC, due to its technical limitations, cannot be combined with the pacemaker-based cluster stack.

4.2 Samba Clustered Trivial Database (CTDB)

SUSE Linux Enterprise High Availability Extension 11 SP3 includes the Samba CTDB extension, including an OCF-compliant resource agent to orchestrate fail-over. This is fully supported, together with exporting Samba CTDB from OCFS2.

Due to technical limitations, this also includes the CTDB internal fail-over functionality for IP address take-over. Please note that this part is not supported by Novell. Only Pacemaker clusters are fully supported.

The `smb_private_dir` parameter for the CTDB resource agent is now deprecated and has been made optional. Existing installations using CTDB should remove this parameter from their configuration at their next convenience.

Several new parameters have been added to the CTDB resource agent in this release - run `"crm ra info CTDB"` for details. Two of these parameters, `ctdb_manages_samba` and `ctdb_manages_winbind`, default to "yes" for compatibility with the previous releases. Existing installations should update their configuration to explicitly set these parameters to "yes", as the defaults will be changed to "no" in a future release.

4.3 DRBD Resource Agent

The new version of DRBD included in SUSE Linux Enterprise High Availability Extension 11 SP3 also supplies a new, updated Open Clustering Framework resource agent from the provider `linbit`.

It is recommended that setups are converted from `ocf:heartbeat:drbd` to use the new `ocf:linbit:drbd` agent. Some new features, such as dual-primary support for master resources, is only available in the new version.

4.4 Heartbeat

Whereas SUSE Linux Enterprise Server 10 clusters utilized heartbeat as the cluster infrastructure layer, providing messaging and membership services, SUSE Linux Enterprise 11 High-Availability Extension uses corosync and openais. heartbeat is no longer included with the product.

Please use the `hb2openais.sh` tool for migrating your SUSE Linux Enterprise Server 10 environment to SUSE Linux Enterprise High Availability Extension 11 SP3.

4.5 EVMS2 Replaced with LVM2

Since EVMS2 has been depreciated in SUSE Linux Enterprise Server 11, the clustered extensions are also no longer available in SUSE Linux Enterprise High Availability Extension 11 SP3. A conversion tool is supplied as part of the `lvm2-clvm` package. After the conversion, the former C-EVMS2 segments can be used as regular, full-featured LVM2 logical volumes.

For more details, refer to </usr/share/doc/packages/lvm2-clvm/README.csm-converter>.

5 Infrastructure, Package and Architecture Specific Information

5.1 Architecture Independent Information

5.1.1 Changes in Packaging and Delivery

5.1.1.1 New or Removed Packages

New Packages (compared with version 11 SP2 GA):

- `ibqb-devel`
- `libqb0`
- `rubygem-bundler`
- `rubygems`
- `sbd`
- `yast2-rear`

Removed Packages (compared with version 11 SP2 GA):

- None

5.2 Systems Management

5.2.1 Dashboard for multiple clusters

To ease the management of multiple clusters, hawk can now display a dashboard showing a summary of several clusters in parallel. This feature can be accessed using the "Dashboard" link in the upper right corner on hawk's login screen.

6 Other Updates

6.1 crm Shell:

The group command within the configuration submenue has been extended. The group command is now able to extend existing groups by adding new resources. Resources can also be deleted.

The group command within the configuration submenue has been extended. The group command is now able to extend existing groups by adding new resources. Resources can also be deleted.

7 Update-Related Notes

This section includes update-related information for this release.

7.1 Update from SUSE Linux Enterprise High Availability Extension 11 SP 2

7.1.1 Pacemaker Initial Monitor (probe) Timeout Now Uses Different Default

Pacemaker executes an initial monitor for each service on every node ("probe"). The timeout for this monitor operation could be specified by configuring a monitor operation with interval set to 0, or would otherwise be inherited from the cluster-wide default operation timeout.

To ease the need to configure a dedicated monitor timeout for the initial probe, pacemaker will now use the timeout from the most frequent configured monitor automatically. If no monitor operation is configured at all, the cluster-wide default continues to apply.

Users which relied on the probe inheriting the timeout from the cluster-wide timeout in the presence of other monitor operations are advised to configure an explicit monitor operation instead.

8 Supported Deployment Scenarios SUSE Linux Enterprise High Availability Extension 11 SP3

The SUSE Linux Enterprise High Availability Extension stack supports a wide range of different cluster topologies.

Local and Metro Area (stretched) clusters are supported as part of a SUSE Linux Enterprise High Availability Extension subscription. Geographical clustering requires an additional Geo Clustering for SUSE Linux Enterprise High Availability Extension subscription.

8.1 Local Cluster

In a local cluster environment, all nodes are connected to the same storage network and on the same network segment; redundant network interconnects are provided. Latency is below 1 millisecond, and network bandwidth is at least 1 Gigabit/s.

Cluster storage is fully symmetric on all nodes, either provided via the storage layer itself, mirrored via MD Raid1, cLVM2, or replicated via DRBD.

In a local cluster all nodes run in a single corosync domain, forming a single cluster.

8.2 Metro Area Cluster

In a Metro Area cluster, the network segment can be stretched to a maximum latency of 15 milliseconds between any two nodes (approximately 20 miles or 30 kilometers in physical distance), but fully symmetric and meshed network inter-connectivity is required.

Cluster storage is assumed to be fully symmetric as in local deployments.

As a stretched version of the local cluster, all nodes in a Metro Area cluster run in a single corosync domain, forming a single cluster.

8.3 Geographical Clustering

A Geo scenario is primarily defined by the network topology; network latency higher than 15 milliseconds, reduced network bandwidth, and not fully interconnected subnets. In these scenarios, each site by itself

must satisfy the requirements of and be configured as a local or metropolitan cluster as defined above. A maximum of three sites are then connected via Geo Clustering for SUSE Linux Enterprise High Availability Extension; for this, direct TCP connections between the sites must be possible, and typical latency should not exceed 1 second.

Storage is typically asymmetrically replicated by the storage layer, such as DRBD, MD Raid1, or vendor-specific solutions.

DLM, OCFS2, and cLVM2 are not available across site boundaries.

9 Known Issues in SUSE Linux Enterprise High Availability Extension 11 SP3

9.1 Linux Virtual Server Tunnelling Support

The LVS TCP/UDP load balancer currently only works with Direct Routing and NAT setups. IP-over-IP tunnelling forwarding to the real servers does not currently work.

9.2 Samba CTDB Cluster Rolling Update Support

The CTDB resource should be stopped on all nodes prior to update. Rolling CTDB updates are not supported for this release, due to the risk of corruption on nodes running previous CTDB versions.

10 Further Notes on Functionality

10.1 Cluster-concurrent RAID1 Resynchronization

To ensure data integrity, a full RAID1 resync is triggered when a device is re-added to the mirror group. This can impact performance, and it is thus advised to use multipath IO to reduce exposure to mirror loss.

Due to the need of the cluster to keep both mirrors uptodate and consistent on all nodes, a mirror failure on one node is treated as if the failure had been observed cluster-wide, evicting the mirror on all nodes. Again, multipath IO is recommended to reduce this risk.

In situations where the primary focus is on redundancy and not on scale-out, building a storage target node (using `md raid1` in a fail-over configuration or using `drbd`) and reexporting via iSCSI, NFS, or CIFS could be a viable option.

10.2 Quotas on OCFS2 Filesystem

To use quotas on ocfs2 filesystem, the filesystem has to be created with appropriate quota features: 'usrquota' filesystem feature is needed for accounting quotas for individual users, 'grpquota' filesystem feature is needed for accounting of quotas for groups. These features can be also enabled later on an unmounted filesystem using `tunefs.ocfs2`.

For quota-tools to operate on the filesystem, you have to mount the filesystem with 'usrquota' (and/or 'grpquota') mount option.

When a filesystem has appropriate quota feature enabled, it maintains in its metadata how much space and files each user (group) uses. Since ocfs2 treats quota information as a filesystem internal metadata, there is no need to ever run `quotacheck(8)` program. Instead, all the needed functionality is built into `fsck.ocfs2` and the filesystem driver itself.

To enable enforcement of limits imposed on each user / group, run `quotaon(8)` program similarly as for any other filesystem.

Commands `quota(1)`, `setquota(8)`, `edquota(8)` work as usual with ocfs2 filesystem. Commands `repquota(8)` and `warnquota(8)` do not work with ocfs2 because of a limitation in the current kernel interface.

For performance reasons each cluster node performs quota accounting locally and synchronizes this information with a common central storage once per 10 seconds (this interval is tunable by `tunefs.ocfs2` using options 'usrquota-sync-interval' and 'grpquota-sync-interval'). Thus quota information need not be

exact at all times and as a consequence user / group can slightly exceed their quota limit when operating on several cluster nodes in parallel.

11 Support Statement for SUSE Linux Enterprise High Availability Extension 11 SP3

Support requires an appropriate subscription from Novell; for more information, please see: http://www.novell.com/products/server/services_support.html.

A Geo Clustering for SUSE Linux Enterprise High Availability Extension subscription is needed to receive support and maintenance to run geographical clustering scenarios, including manual and automated setups.

Support for the DRBD storage replication is independent of the cluster scenario and included as part of the SUSE Linux Enterprise High Availability Extension product and does not require the addition of a Geo Clustering for SUSE Linux Enterprise High Availability Extension subscription.

General Support Statement

The following definitions apply:

- L1: Installation and problem determination - technical support designed to provide compatibility information, installation and configuration assistance, usage support, on-going maintenance and basic troubleshooting. Level 1 Support is not intended to correct product defect errors.
- L2: Reproduction of problem isolation - technical support designed to duplicate customer problems, isolate problem areas and potential issues, and provide resolution for problems not resolved by Level 1 Support.
- L3: Code Debugging and problem resolution - technical support designed to resolve complex problems by engaging engineering in patch provision, resolution of product defects which have been identified by Level 2 Support.

Novell will only support the usage of original (unchanged or not recompiled) packages.

11.1 Long-term Strategy for SLE HA's Graphical User Interface

SLE HA customers desire a long-term statement about support and active development for the cluster stack's graphical user interface. The supported GUI should be flexible, usable without client-side installation of additional packages, firewall-friendly and also be usable from non-Linux desktops as well as mobile devices.

SUSE is committed to providing an easy to use yet powerful GUI for SLE HA. Because it best meets customer requirements, effort is being focused on the web-based frontend (hawk); new features and

functionality will primarily be developed and improved in this user interface. hawk already provides the best GUI experience on SLE HA today.

The X11-based "hb_gui", which required a Linux client or an X11 server, is now in maintenance mode and is not scheduled to receive new functionality during the SLE HA 11 lifecycle, and will no longer be part of SLE HA 12.

12 Miscellaneous

13 Technical Information

13.1 OCF Resource Agent for oracle no longer creates audit logs

The "ocf:heartbeat:oracle" resource agent has been improved to no longer trigger audit logs to be written, reducing the logging overhead of the cluster.

13.2 IPv6 Implementation and Compliance

13.2.1 Linux Virtual Server Load Balancer (ipvs) Extends Support for IPv6

The LVS/ipvs load balancing code did not fully support RFC2460 and fragmented IPv6 packets which could lead to lost packets and interrupted connections when IPv6 traffic was fragmented.

The load balancer has been enhanced to fully support IPv6 fragmented extension headers and is now RFC2460 compliant.

14 More Information and Feedback

- Read the READMEs on the CDs.
- Get detailed changelog information about a particular package from the RPM:

```
rpm ---changelog --qp <FILENAME>.rpm
```

<FILENAME>. is the name of the RPM.

- Check the ChangeLog file in the top level of CD1 for a chronological log of all changes made to the updated packages.
- Find more information in the docu directory of CD1 of the SUSE Linux Enterprise High Availability Extension CDs. This directory includes PDF versions of the SUSE Linux Enterprise High Availability Extension startup and preparation guides.
- <http://www.novell.com/documentation/sles11/> contains additional or updated documentation for SUSE Linux Enterprise High Availability Extension.
- Visit <http://www.novell.com/linux/> for the latest Linux product news from SUSE and Novell, and <http://www.novell.com/linux/source/> for additional information on the source code of SUSE Linux Enterprise products.

Thanks for using SUSE Linux Enterprise High Availability Extension in your business.

The SUSE Linux Enterprise High Availability Extension Team.