Release Notes

SUSE Linux Enterprise High Availability Extension 12

SUSE Linux Enterprise High Availability Extension is a suite of clustering technologies that enable enterprises to implement highly available Linux clusters and eliminate single points of failure. This document gives an overview of features of SUSE Linux Enterprise High Availability Extension and their limitations. Some sections do not apply to a particular architecture or product, this is explicitly marked.

Manuals can be found in the `docu` directory of the installation media, or in the directory `/usr/share/doc/` on the installed system (if installed).

Publication date: 2014-10-08 , Version: 12.0.20140929

Contents

1 SUSE Linux Enterprise High Availability Extension

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 it's adherence to open standards encourages interoperability with industry standard tools and technologies.

2 What's New?

Cluster File System

GFS2 cluster file system with read/write support, to complement the SUSE recommended OCFS2 cluster file system.

Load balancer

HAProxy as layer 4 load balancer added, to complement the linux virtual server load balancer.

History Explorer

HAWK history explorer now includes off-line analysis capabilities.

Resource Agents

Resource agents got multiple updates, including a resource agent to handle SCSI reservations

See https://www.suse.com/releasenotes/x86_64/SLE-HA/12-GEO/ for what's new in the Geo clustering option for the SUSE Linux Enterprise High Availability Extension 12.

3 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 12.

4 Infrastructure, Package and Architecture Specific Information

4.1 Systems Management

4.1.1 Improved Help in crm Shell

Some commands were undocumented or did not cover all use cases.

The interactive help has been overhauled, and the online help for the crm shell has been updated and improved.

4.2 Performance Related Information

4.2.1 Faster Unlink Performance

An unlink would generate unnecessary journalling requests. This would hinder any write operations that would follow.

The unlink path has been changed to avoid the unnecessary journalling.

5 Packages and Functionality Changes

5.1 New Packages

5.1.1 Include fence agent using SCSI reservations

Support fencing via SCSI reservations.

The fence-agents package is now included which includes an agent capable of fencing via SCSI reservations.

5.1.2 Provide haproxy as an alternative to LVS

Haproxy is useful in some scenarios where LVS cannot be used.

Haproxy 1.5 is now provided.

5.2 Updated Packages

5.2.1 Relax-and-Recover: Upgrade to Latest Stable Upstream Version

The current version of Relax-and-Recover (abbreviated as 'rear', formerly 'ReaR') comes with many new features and fundamental changes. In particular the file system creation is totally different (via the so called "layout workflow" instead of the old "dr workflow" that has been removed in version 1.14).

Warning: Users who already use it and want to upgrade (in particular when upgrading from SUSE Linux Enterprise 11), must carefully re-validate whether their particular disaster recovery procedure still works.

A new and different RPM package name is used that contains the version (e.g. "rear116"). It conflicts with all other rear packages to avoid that an installed version could get accidentally replaced with another version. This means "just a normal package installation" that would do a version change of an already installed Relax-and Recover software is intentionally prohibited.

Furthermore this version upgrade supersedes the formerly provided package rear-SUSE.

5.2.2 Support MariaDB as well as MySQL

MariaDB is an alternative to MySQL, and the resource agent for HA should support both.

The resource agent now supports MariaDB.

5.2.3 Hawk: Limit access to wizards based on ACL permissions

A user with limited ACL access may not have the correct permissions to apply the configuration created by a wizard. Hawk should detect this and warn before allowing the wizard to be applied.

Hawk warns if the user lacks necessary write permissions to apply a configuration wizard.

5.2.4 crmsh: Configure all nodes in the cluster simultaneously using cluster scripts

Ensuring that all nodes in the cluster have the correct packages and file system configuration can be difficult. Sometimes, a configuration wizard could be improved if it could make system-level changes or inspect the state of all cluster nodes before applying a new CIB configuration.

The crm shell now includes cluster scripts, which are able to validate and modify the configuration of cluster nodes.

5.2.5 Allow the configuration of corosync via crmsh

The crm shell is an excellent tool for configuring pacemaker, but the cluster stack consists of additional software. It should be possible to leverage the abilities of crmsh for configuring these as well.

crmsh can now configure the cluster infrastructure directly, via the new cluster and corosync subcommands.

5.2.6 Enable the use of reports generated by crm_report in crmsh and the history explorer

A cluster report generated by crm_report contains much of the information needed to debug issues with a cluster. It should be possible to open these reports in the history explorer and crmsh.

Both hawk and crmsh can now open reports generated by crm_report. However, reports generated by hb_report contain additional information and are preferred when possible.

5.2.7 Improved Help in crm Shell

Some commands were undocumented or did not cover all use cases.

The interactive help has been overhauled, and the online help for the crm shell has been updated and improved.

5.2.8 Simplified OCFS2 configuration with wizard in Hawk

The configuration of OCFS2 has been simplified, and the Hawk configuration wizard as well as the configuration template in crm shell should reflect this.

The Hawk wizard for cLVM2/OCFS2 has been updated.

5.2.9 Hawk wizard for configuring NFS

To make the configuration of an NFS server easier, Hawk should include a configuration wizard for this purpose.

Hawk includes a configuration wizard that configures an NFS server.

5.2.10 Allow resource sets and templates for location constraints

Enabling the use of resource sets and templates in location constraints makes configuration easier and can simplify the CIB.

Both crmsh and hawk now supports creating location constraints that reference resource sets. The syntax for crmsh is described in the Admin Guide.

5.2.11 crmsh: Manage multiple resources as one using resource tags

It may be desired to start and stop multiple resources all at once, without having explicit dependencies between those resources.

This feature adds resource tags to Pacemaker. Tags are collections of resources that do not imply any colocation or ordering constraints, but can be referenced in constraints or when starting and stopping resources.

5.2.12 crmsh: Enable anonymous shadow CIBs

When scripting cluster configuration changes, the scripts are more robust if changes are applied with a single commit. Creating a shadow CIB to collect the changes makes this easier, but has previously required naming the shadow CIB.

crmsh now allows the creation of shadow CIBs without explicitly specifying a name. The name will be determined automatically and will not clash with any other previously created shadow CIBs.

5.2.13 Allow use of IP addresses without reliance on DNS in csync2 configuration

Allowing the use of plain IP addresses can make it easier to configure csync2.

csync2 allows the use of IP addresses for connectivity via a special syntax. This is documented in Procedure 3.8: Configuring Csync2 with YaST in the Admin Guide.

5.2.14 Enable colocating resources without further dependency between them

Sometimes, it is desired that two resources should run on the same node. However, there should be no further dependency implied between the two resources, so that if one fails, the other one can keep running.

This can be accomplished using a third, dummy resource which both resources depend on in turn. To make it easier to create a configuration like this, the command "assist weak-bond" has been added to crmsh.

5.2.15 hawk: Wizards Can Use Multiple Instances of a Template Within a Workflow

Simplify wizards with multiple identical config instances.

In a wizard workflow, multiple steps to reference the same wizard template with different sets of parameters are supported.

This is done by adding a type attribute to the template instance in the workflow, which can be separate from the name identifying the step in the workflow. Each step in the workflow can also override any attribute of the wizard template, including the description.

This change is backwards compatible.

5.2.16 Hawk off-line history explorer

Sometimes, a system adminstrator may want to use the Hawk history explorer to examine a cluster report on a machine that is not part of the cluster.

This feature allows hawk to be installed and run without requiring pacemaker to be running, and the history explorer can function as expected.

5.2.17 Cluster health evaluation tool built in to crmsh

Configuring a cluster node correctly is complicated, and problems with hardware, network connectivity and system software can cause issues when configuring cluster resources.

The crm shell now includes a cluster health script, which can run both on an already-configured cluster node as well as a potential cluster node, to perform various system-level health checks. It then presents a list of potential issues, such as not having NTP configured correctly or warning about low disk space.

6 Update-Related Notes

This section includes update-related information for this release.

7 Support Statement for SUSE Linux Enterprise High Availability Extension 12

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

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.

8 How to Obtain Source Code

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@suse.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.

9 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 a PDF version of the High Availability Guide.
- http://www.suse.com/documentation/sle_ha/ contains additional or updated documentation for SUSE Linux Enterprise High Availability Extension 12.
- Visit http://www.suse.com/products/ for the latest product news from SUSE and http://www.suse.com/download-linux/source-code.html 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.