

Release Notes for release 1.0.11 of Linuxha.net

The following notes describe the changes since 1.0.10.

Fix: Await DRBD Initialization

When forming the cluster without DRBD modules loaded there was the possibility of the module not registering the new devices with udev prior to an auto-started application attempting to make use of them. This potential has now been removed - if Linuxha.net loads the DRBD kernel module it will await the presence of /dev entries prior to allowing an application to attempt to make use of such devices.

Improvement: Allow "--forcesync" in "clbuildapp" at all times

In previous Linuxha.net versions the use of --forcesync when a normal --sync was all that was necessary caused a DRBD failure to be shown. As of 1.0.11 this has been fixed,

Improvement: Application Start/Stop entries containing multiple commands

Previous versions of Linuxha.net allowed an application to include parameters to the start/stop commands. As of 1.0.11 it is now possible to specify complete shell statements. For example the following are examples of commands which would fail under previous versions but should now work;

```
/path/script1; /path/anotherscript
```

or even;

```
/path/script1 && /path/script2
```

Improvement: Earlier application abort if node unsuitable

If a node was not deemed suitable to host an application in previous versions of Linuxha.net it would still start the DRBD devices before aborting - leaving them active and the application down! For 1.0.11 the application will query the cluster daemon on the node status earlier in the start-up avoiding this potential problem.

Fix: Abort start/stop script on timeout

If a start/stop script failed caused a timeout event and failed to heed the SIGTERM sent to the process it could eventually fall into an endless loop - since although the "clstartapp" and "clhaltapp" scripts to eventually call SIGKILL they never abort after sending it. Now fixed.

Fix: Correctly handle LEAST_APP_LOAD & LEAST_CPU_LOAD

When an application set a "preferred_node" attribute to either "LEAST_APP_LOAD" or "LEAST_CPU_LOAD" it would previously always start the application on the local node. This has been fixed to check the applications currently running on a node or the 15 minute load average respectively.

Release Notes for release 1.0.10 of Linuxha.net

The following notes describe the changes since 1.0.9.

Improvement: DRBD 0.7.18

This releases of Linuxha.net simply updates DRBD to the latest version which should ensure the software will install/build on even the very newest distributions.

Release Notes for release 1.0.9 of Linuxha.net

The following notes describe the changes since 1.0.8.

New Feature: IP Range Support

The biggest change is the support of IP address ranges when specifying the IP addresses a particular application should present on a particular network. Consider the example of attempting to allocate 4 IP addresses leading to a element in the “appconf.xml” similar to:

```
<network net="main"
  ip="172.16.177.200,172.16.177.201,172.16.177.202,172.16.177.203"
  netmask="255.255.255.0"/>
```

It has always been possible to use the following “short-hand” method of writing the IP addresses:

```
<network net="main"
  ip="172.16.177.200,201,202,203"
  netmask="255.255.255.0"/>
```

However even this can become cumbersome, so version 1.0.9 allows ranges to be defined, for example:

```
<network net="main"
  ip="172.16.177.200-203,210,50-60"
  netmask="255.255.255.0"/>
```

The above would define 16 IP addresses for the application.

Fix: Support Dynamic additional/removal of Volume Groups

Although it was theoretically possible to add and remove complete volume groups from running applications this level of dynamic change had never actually been tried. In testing this facility a small bug with 1.0.8 and previous versions was found. This has now been corrected.

This is a powerful feature and users are encouraged to seek further help from the mailing lists before using such a feature prior to this being fully documented as part of the 1.2.0 user guide.

Improvement: Added the “useful” directory

An additional directory has been added - called “useful”. This directory will contain various items that are of use to 3rd party software. At present the directory simply contains a logrotate configuration file suitable for managing the log files generated by Linuxha.net.

Release Notes for release 1.0.8 of Linuxha.net

The following notes describe the changes since 1.0.7.

New Feature: “Arp Flux” Handling

The Cluster Network Daemon - “cldnetd” now makes use of the “arp_filter” flag in “/proc” to ensure that when multiple network cards are used for a network, they should work when spread across multiple switches. This allows the cluster architect greater flexibility for the network design and allows redundancy to be added by using multiple switches.

No additional changes are required to make use of this feature - it is automatically turned on when 1.0.8 is installed. No existing network topology will “break” by the introduction of this feature.

Improvement: Timestamp Addition

When an application generates log information during the start-up or shutdown of the application these entries are now prefixed by a line with the following format:

```
TIMESTAMP: YYYY.MM.DD HH:MM:SS
```

This allows the administrator to correct associated log entries with particular times - which helps when performing problem resolution.

Fix: Improved Route Definitions following IP Fail-over

When an IP fail-over occurs the gateway portion was not being correctly configured when adding back the routes on the new interface. This often led to routing problems associated with the default route. This has been corrected.

Improvement: Improved Lems Resilience to Module problems

Previously if a particular module that a Lems daemon called caused an issue it could result in the Lems daemon for that particular application dying. The code calling these modules has been improved to catch and log such errors - making it now impossible for a badly written module to stop all other monitoring for the application.

Release Notes for release 1.0.6/1.0.7 of Linuxha.net

The following notes describe the changes since 1.0.5 - essentially just bug fixes.

Fix: Improved Application Fail-over Reliability

It was found that the handling of child processes during an application fail-over was influenced on occasion by the machine load, sometimes causing applications to fail to start correctly. The underlying child signal handling has been improved to ensure that this no longer occurs.

Fix: Improved Cluster Abort

When the cluster is shutdown via `clhalt(1M)` it now does so cleanly. It also ensures that the other cluster daemons stop at the same time as well. Shutting down the cluster when no applications are running typically takes less than a second now.

Fix: Cluster Daemons Signal Handling

It was found that under certain conditions a broken socket connection (or a time-out) might cause a cluster daemon to receive a "SIGPIPE" signal - which was not handled. Now configured to expect such signals and handle appropriately.

Fix: "drbd_tool" Support for "disconnect"

Although not used by the automated tools, support for the "disconnect" functionality for the underlying DRBD devices is occasionally useful for rolling upgrades. This functionality was supported on the command line but the logic for performing the necessary actions was missing. Now fully working.

Fix: Lock Daemon Error

A certain logic path (not often encountered) in the lock daemon had a fatal error which caused the daemon to abort. Now fixed.

Fix: Manual Page Inaccuracies

The `clstartapp(1M)` manual page has been modified to ensure the correct information for the "--reallyforce" option is included. Also included for the first time are details for the "--fsonly" and "--ignoreapp" flags.

Fix: LVM2 Volume Group Activation

The LVM sub-system will now perform a `vgscan(1)` prior to returning LVM details. This fixes some errors under certain application start-up or rebuild - though this was only seen on Slackware but might have affected other distributions.

Fix: RPM Package Upgrades work Reliably

Newer versions use a different algorithm to determine whether or not symbolic links need to be

removed. This fixes the problem of “missing links” following an RPM package upgrade.

Release Notes for release 1.0.5 of Linuxha.net

The following notes describe the changes since 1.0.4 and are mainly concerned with fixing some recently found bugs, though some additional features are also present.

New Feature: Network Attribute Support

It has been found that if a network consists of a single card, and that card is simply a cross-over cable, loss of the other node will be seen as a double-failure; a failure of the network and then remote host. This double-failure scenario is not currently dealt with - and hence monitoring of such networks for physical link failure is not advisable.

Hence 1.0.5 supports the ability to indicate network attributes - and currently the only attribute supported is cross-over. Consider the following example networks section from a cluster configuration file:

```
<node>
    <name>servera</name>
    <network name="a" cards="eth0,eth1"/>
    <network name="b" cards="eth2" attributes="crossover"/>
</node>
<node>
    <name>serverb</name>
    <network name="a" cards="eth0,eth1"/>
    <network name="b" cards="eth2" attributes="crossover"/>
</node>
```

The above configuration indicates that network “b” consists of a cross-over cable and so no monitoring of this network occurs.

New Feature: Application startup “--ignoreapp” Flag

The “clstartapp” command now supports an “--ignoreapp” flag. When this flag is specified the return code for the application start-up simply generates a warning if non-zero - the application will then continue to complete initialization into the cluster.

This flag is typically used during debugging of application configuration issues and complements the “--fsonly” flag (which mounts all file systems but does not register the application as actually started in the cluster).

New Feature: Cluster Status Shown when cluster Down

Previous versions of Linuxha.net simply indicated the cluster was not running when “clstat” was run and no cluster daemons were running. This limitation has been resolved and the cluster and all applications will be shown even if not running, for example:

```
Cluster: slack10c1 - DOWN
```

Node	Status
slack10s1	DOWN
slack10s2	DOWN

Application	Node	State	Started	Monitor	Stale	Fail-over?
test	N/A	DOWN	N/A	N/A	N/A	N/A

New Feature: Cluster Status Shows “Unregistered” Applications

Previously if an application configuration file was changed and the cluster started without the application being rebuilt the “clstat” command would not show it. Such applications are now shown as follows.

Application	Node	State	Started	Monitor	Stale	Fail-over?
test	N/A	DOWN	N/A	N/A	N/A	N/A [Rebuild]

The presence of the “[Rebuild]” is used to indicate that the application is not currently registered and will require a rebuild (via “clbuildapp”) before it can be used.

New Feature: Validation Online Growth of Ext3 File Systems

The support for the “ext2online” tool for online growth of Ext3 file systems has now been validated. It should be noted that this facility tool has not yet reached version 1.0 and thus it is recommended that use of this feature is considered “emergency only” at present.

New Feature: Auto Activation of Volume Groups

Previously it was expected that all volume groups would be activated when the node was booted. If this was not the case there was the possibility that the “volume group build” phase of defining an application would fail. The volume group management layer now automatically scans all volumes when necessary to ensure this possibility can not now occur.

Fix: Online Growth of Reiserfs File systems

The command generated to grow the file system was incorrect - this should now work with-out problems.

Fix: Correctly Handle unresolved Hostname

During the cluster build phase if a hostname could not be resolved to an IP address the error generated contained an error itself - the result of which was that the hostname that could not be resolved to an IP address was not actually shown.

Fix: Cluster Build abort on Fatal Errors

Some conditions regarding the “global” definitions for a cluster reported an error, but allowed the

build to continue. The result was an incorrectly built cluster configuration that might cause problems later. All such errors are now trapped and the cluster build aborts in the expected manner.

Fix: Cluster formation of Secondary Node Only

It was found that if the cluster was formed with only the secondary node the no applications would be registered! This now works as expected.

Release Notes for release 1.0.4 of Linuxha.net

The following notes describe the changes since 1.0.3 and are mainly concerned with fixing some recently found bugs.

Installation Process Changes

If the user space DRBD tools have been installed the post-installation summary will include lines similar to the following:

```
Installation Summary
DRBD Kernel module      : OK [User space tools only]
```

If the MII fallback binary fails then the new message generated now will be:

```
MII fallback binary     : FAILED [Warning only]
```

This is classed as a warning since it is unlikely the tool will be needed anyway with recent hardware. This tool is likely to fail to compile on all modern Linux 2.6 variants and has been superseded with the widely available “ethtool”.

Stricter Network Checks during Cluster Build

Consider the following network configuration for the two nodes:

```
<node>
    <name>suse10s1</name>
    <network name="prod" cards="eth1,eth2"/>
    <network name="drbd" cards="eth0"/>
</node>
<node>
    <name>suse10s2</name>
    <network name="prod" cards="eth2,eth3"/>
    <network name="drbd" cards="eth2"/>
</node>
```

This is an incorrect configuration since “suse10s2” defines “eth2” as part of two different networks - however the cluster build fails to spot that the “suse10s2” server is using different interfaces and errors are passed over during the build stage, (shown below in bold).

```
INFO 02/12/2005 23:49:34 Checking ssh cabability between Primary IP addresses (node
names).
INFO 02/12/2005 23:49:35 Networks defined for suse10s1: drbd,prod
INFO 02/12/2005 23:49:35 Networks defined for suse10s1: drbd,prod
INFO 02/12/2005 23:49:35 Validated network IP for drbd on suse10s1: eth0 -
192.100.100.1
INFO 02/12/2005 23:49:35 Validated network IP for prod on suse10s1: eth1 - 190.10.1.1
INFO 02/12/2005 23:49:35 Validated network IP for drbd on suse10s2: eth2 - 190.10.1.2
INFO 02/12/2005 23:49:35 Validated network IP for prod on suse10s2: eth2 - 190.10.1.2
INFO 02/12/2005 23:49:35 Validated unique network/interfaces for suse10s1:
INFO 02/12/2005 23:49:35 Interface eth1 is on network 190.10.1.0
INFO 02/12/2005 23:49:35 Interface eth0 is on network 192.100.100.0
INFO 02/12/2005 23:49:35 Interface wlan0 is on network 192.168.1.0
INFO 02/12/2005 23:49:36 Validated unique network/interfaces for suse10s2:
INFO 02/12/2005 23:49:36 Interface eth2 is on network 190.10.1.0
INFO 02/12/2005 23:49:36 Interface eth0 is on network 192.100.100.0
```

```
INFO 02/12/2005 23:49:36 Successfully copied network configuration to suse10s2.
INFO 02/12/2005 23:49:36 Validated data replication network is ok (drbd).
INFO 02/12/2005 23:49:36 Able to send ping to DRBD Ip address 192.100.100.1 for
suse10s1
INFO 02/12/2005 23:49:36 Able to send ping to DRBD Ip address 190.10.1.2 for suse10s2
INFO 02/12/2005 23:49:36 Node suse10s1 is not running a cldaemon process (good)
INFO 02/12/2005 23:49:36 node suse10s2 is not running a cldaemon process (good)
```

As of version 1.0.4 such a configuration now produces an error:

```
INFO 03/12/2005 00:45:17 Networks defined for suse10s1: drbd,prod
INFO 03/12/2005 00:45:17 Networks defined for suse10s1: drbd,prod
ERROR 03/12/2005 00:45:17 Duplicated network cards across networks for node suse10s2.
ERROR 03/12/2005 00:45:17 Card eth2 in prod already defined for network drbd.
```

It is expected that this lack of checking in earlier versions might account for some unexpected failures particularly during fail-over conditions.

XML Dump Utility

The “dump.pl” routine which has not been previously distributed is now present as “/sbin/cluster/tools/dump.pl”. This is not normally used but might be requested to be run against a configuration file for debugging purposes if a problem arises.

Cluster Daemon Improvements

Several changes have been made to the “cldaemon” process - the main cluster daemon.

- The “reload_modules.pl” have a typo which resulted in it failing to load correctly when a cluster configuration daemon started.
- The cluster daemon now should generate files of the following names, irrespective of the name of any log file generated:

```
/var/log/cluster/cldaemon-CLUSTER.stderr
/var/log/cluster/cldaemon-CLUSTER.stderr
```

If an error causes the cluster daemon to abort these files may contain information the developers can use to resolve the problem in a more timely manner.

Lems Daemon Improvement

The lems daemon will also generate a log file which contains the any output from standard error. By default this file will be called:

```
/var/log/cluster/lems/application.stderr
```

The standard log file which captures standard output and general messages has also been renamed to:

```
/var/log/cluster/lems/application.log
```

The “--file” and new “--errfile” can be used to specially indicate the full paths of the files to create to override any default.