

# Release Notes for 1.2.9 of Linuxha.net

## ***Improvement: Net Known Connections Validation***

When the cluster configuration is built warnings are generated if any entries in the “net\_known\_connections” settings can not be pinged. Also every time a cluster is formed information on the number of valid hosts that can be contacted is shown.

## ***Improvement: “clbuild” error Output Improved***

Previously there were certain conditions that could result in a failure but no evidence for the cause of the failure were presented to this screen. Additional error messages have been added to reduce the likelihood of this occurring.

## ***Fix: Node removal from Cluster***

The use of the “clhalt –node name” now works exactly as it should do.

## ***Improvement: Online File system resize***

When attempting to resize an “ext3” based file system, Linuxha.net will now attempt to use “resize2fs” or the newer “ext3online” if found.

## ***Improvement: Log file Management***

The software also includes a sample of a series of rules suitable for Skulker 2 to manage the log files generated by the software.

## ***Improvement: Cluster Status Output***

If the Lems daemon for a particular application fails it will now show as “Down” in the “clstat” output. Previously this was not often the case.

## ***Fix: Application build safety Improvements***

Previously it was possible to use the “clbuildapp –sync” option whilst the application was running which might potentially result in loss of data if run on the incorrect node, or force application unavailability since it would attempt to un-mount the application file systems! This is now detected and if the application is running, the “--sync” option will abort with a suitable error.

### ***Fix: Application Rebuild Detection***

Previously if the cluster was up and running and an application needed rebuilding the “clstat” output might not show it if the application was valid at the time of the last cluster formation. This has been fixed – as soon as an application requires a rebuild it is shown in the “clstat” command output.

### ***Fix: Automatic DRBD Upgrade Feature***

If later versions of DRBD 0.7.x are added to the correct directory then cluster formation will correctly find it, compile it and install it on the next cluster formation. Previously the pattern matching used meant this might appear to randomly work.

### ***Improvement: More robust cluster formation for single node cluster***

If just a single node is available occasionally cluster formation would fail in the past with time-out errors – this condition has been removed.

### ***Improvement: Application IP Probing***

If an application defines multiple IP addresses the probing for addresses use to fail; it now correctly works with patterns allowing detection of such applications when the software is restarted with applications currently running.

## Release Notes for 1.2.8 of Linuxha.net

### ***Improvement: Migration to DRBD 0.7.25***

The default version of DRBD distributed with Linuxha.net has been increased to the latest available for DRBD 0.7 series. This introduces some bug fixes and latest distribution support.

### ***Improvement: Email Alert Support***

The cluster configuration can optionally include an “alert\_to” attribute. In this case certain conditions will result in an email being sent to the specified address to ensure the administrator is aware of any conditions that might be impacting the cluster.

## Release Notes for 1.2.7 of Linuxha.net

## Release Notes for release 1.2.6 of Linuxha.net

The following notes describe changes compared to release 1.2.5.

### ***Improvement: net\_settle Support***

It has been found that some environments require a little time for the physical link state checks on network cards to start functioning correctly. Whether this is down to operating system, NIC firmware or particular switches is not known. However to combat the problem a net attribute in the globals section is now available, “net\_settle”. It is optional and defaults to 20 seconds. For most configurations setting it to 1 should be fine however.

The example belows the attribute in the “globals” section of the cluster configuration.

```
<global>
  <name>cluster1</name>
  <version>0.1</version>
  <data>replicated</data>
  <datadetail>drbd</datadetail>
  <dataprotocol>C</dataprotocol>
  <logdir>/var/log/cluster</logdir>
  <key>mykey</key>
  <portpool>9901,9999</portpool>
  <maxblockdevs>50</maxblockdevs>
  <drbd_network>a</drbd_network>
  <net_reuse>2.5</net_reuse>
  <net_settle>20</net_settle>
  <parallel>0</parallel>
</global>
```

### ***Improvement: Link Failure Detection Confirmation***

Before the network daemon assumes that a network link has failed it will attempt both supported link detection methods rather than just relying on the one previously used.

### ***Improvement: NIC Status Temporarily Unavailable***

In some situations the tests to validate network link status can result in Linux returning 'resource temporarily unavailable' errors. This is now treated as a warning and will not cause a IP fail-over to occur.

## ***Improvement: Automated Network Recovery***

Previously when all network connections for a given network were down on a node then it was the responsibility of the administrator to inform the cluster daemons to re-attempt use of that network. Now the Linuxha.net network daemon automatically re-probes previously failed networks every two minutes and enables them if link connectivity can be established.

## ***Improvement: Flexible Cluster Partitioning Handling***

In previous versions Linuxha.net would identify a node as being partitioned from any networking [when possible] and then sit ideally awaiting the partitioning to resolve. In practice it was found that a more sensible method to handle this scenario was to simply reset the server in question. This is now optionally available and will likely become the default in later versions. To enable this functionality a new optional attribute in the globals section can be added:

```
<global>
  <name>cluster1</name>
  <version>0.1</version>
  <data>replicated</data>
  <datadetail>drbd</datadetail>
  <dataprotocol>C</dataprotocol>
  <logdir>/var/log/cluster</logdir>
  <key>mykey</key>
  <portpool>9901,9999</portpool>
  <maxblockdevs>50</maxblockdevs>
  <drbd_network>a</drbd_network>
  <net_reuse>2.5</net_reuse>
  <on_partition>die</no_partition>
  <parallel>0</parallel>
</global>
```

## ***Improvement: More Descriptive Error Text***

Several conditions that result in messages being logged in various files have been improved so they are more readable and are more descriptive.

## ***Fix: DRBD Kernel Module Rebuild***

When Linuxha.net automatically rebuilds the kernel module [either due to a later drbd tar file found, or a kernel upgrade having been performed], it now cleans the source build environment to ensure the build/install is correctly performed.

### ***Fix: Cluster form "--nodrbdcheck" Argument***

This was not correctly implemented in previous versions - it made no difference to the cluster formation actions when specified. Starting with 1.2.6 it now correctly does not attempt to look for updated versions of DRBD. On machines where the DRBD module can take several minutes to load this results in much speedier cluster formations.

## Release Notes for release 1.2.5 of Linuxha.net

The following notes describe changes compared to release 1.2.4.

### ***Improvement: Dynamic Netmask Updates***

The last remaining IP configuration that was not truly dynamic has been fixed: it is now possible to change the netmasks of application IP addresses whilst the application is running.

### ***Improvement: DRBD v8 Preparation***

This version of Linuxha.net **does not** support DRBD and relies instead on the versions in the v0.7.x series. However some work in preparation to including support on the newer version 8 have been included in this release.

### ***Improvement: DRBD 0.7.23 bundled***

The included version of DRBD has been bumped up to 0.7.23 compared to 0.7.21 which was found in version 1.2.4 of Linuxha.net.

### ***Improvement: Improved Application Build Script Data Synchronization***

The support for data synchronization in the “clbuildapp” script has been improved. It now shows the estimated number of minutes the synchronization attempt will complete in. Further it is now possible to specify a “--rate” option which overrides the built in synchronization rate for this application - something that can be very useful during the initial data synchronization.

### ***Fix: Log Rotate Script***

This auxiliary script supplied works in a more standard manner to fit in with other Log Rotate scripts. The permissions on it have also been corrected to fall in line with other configuration files.

### ***Improvement: Over-mount Handling***

The “clhaltapp” to stop an application now successfully detects and handles the condition where one or more file systems have been mounted on sub-directories on file systems for this application. Prior to 1.2.5 this conditions would result in the “clhaltapp” routine failing to stop the application.



## Release Notes for release 1.2.4 of Linuxha.net

The following notes describe changes in since 1.2.2 [covers 1.2.3 and 1.2.4].

### ***Improvement: Multiple Address Family Support***

Although support for many IP addresses for a single application was available, only IP addresses on the same network address were officially supported. For example, although it might be possible to configure the following IP addresses, it was never tested:

```
192.168.1.22/24
192.168.100.22/24
```

Now that functionality has been tested and works with the improvements to the netmask specification support [see next improvement]. Indeed the flexibility of the configuration of IP addresses has been changed enough to provide an additional tool to allow administrator's to check support for various IP configurations. This tool is “show\_ip\_map” and is given an IP entry and a default netmask. An example use might be:

```
# /sbin/cluster/tools/show_ip_map 192.168.1.22-24,192.167.2.21/16,24/16 \
255.255.255.0
```

The output is the IP addresses and netmasks that would be configured if the above was used for an application. In this case:

```
192.168.1.22      255.255.255.0
192.168.1.23      255.255.255.0
192.168.1.24      255.255.255.0
192.167.2.21      255.255.0.0
192.167.2.24      255.255.0.0
```

### ***Improvement: Netmask flexibility***

In previous releases it was only possible to specify a single netmask for all IP addresses associated with a particular application, for example:

```
<networks>
    <network net="prod"
```

```
        ip="192.168.1.221,192.168.1.222"  
        netmask="255.255.0.0"/>  
</networks>
```

For 1.2.4 the above will still work, but each individual IP address can override the default, for example the below entry would configure 192.168.1.221 with a netmask of 255.255.0.0, whilst 192.168.1.222 would have a netmask of 255.255.255.0.

```
<networks>  
    <network net="prod"  
        ip="192.168.1.221,192.168.1.222/255.255.255.0"  
        netmask="255.255.0.0"/>  
</networks>
```

Note: Changing the netmask of an existing IP address associated with the running application will not take affect until the application is restarted. Adding new addresses or removing old address with the default or specified netmask will be added dynamically.

Netmasks that override the default can also be specified in the standard shorthand form. For example the following is equivalent to the above:

```
<networks>  
    <network net="prod"  
        ip="192.168.1.221,192.168.1.222/24"  
        netmask="255.255.0.0"/>  
</networks>
```

### ***Improvement: Network attribute “nossh”***

By default Linuxha.net will attempt to use SSH across different networks to communicate between the hosts. If certain networks do not support SSH between the two machines then this can result in slower than normal operations whilst timeouts on that network occur.

It is now possible to indicate that use of SSH should be avoided on a particular network by using a new “nossh” attribute on that network. For example consider the following network section:

```
<node>  
    <name>centos43s1</name>  
    <network name="prod" cards="eth2,eth1"/>  
    <network name="drbd" cards="eth0" attributes="nossh"/>  
</node>  
<node>  
    <name>centos43s2</name>
```

```
<network name="prod" cards="eth0,eth2"/>
<network name="drbd" cards="eth1" attributes="nossh"/>
</node>
```

This means that only the “prod” network will be used for SSH communication between the two nodes. It should be noted that SSH communication is an key component of Linuxha.net and disabling it across all networks will result in an un-usable cluster! Only use this attribute if strictly necessary!

## ***Improvement: Pseudo Network Resource “MONITOR”***

When an application makes use of IP addresses on a network, loss of that network will result in the application failing over to the other node. However if the application does not present an IP address then the application will continue running. This situation can be avoided by using a pseudo IP on the network in question, for example:

```
<networks>
  <network net="prod" ip="192.168.1.221" netmask="255.255.0.0"/>
  <network net="internal" ip="MONITOR"/>
</networks>
```

In the above configuration loss of the “internal” network on the currently running node for the application in question will cause it to fail-over, even though it does not require any IP addressing on that network.

## ***Improvement: Running “clstat” against unvalidated Cluster***

If the cluster configuration has been changed “clstat” will now issue a warning that the results might not reflect actual conditions, rather than just showing an error and requesting the use of the “--nochecksums” argument. For example:

```
ERROR 28/11/2006 14:37:13 Unvalidated Cluster configuration - clstat might
ERROR 28/11/2006 14:37:13 return inconsistent results until 'clbuild' is run.
Cluster: cl01 - DOWN
```

Node	Status
centos43s1	DOWN
centos43s2	DOWN

  

Application	Node	State	Started	Monitor	Stale	Fail-over?
nfs	N/A	DOWN	N/A	N/A	N/A	N/A
nfs2	N/A	DOWN	N/A	N/A	N/A	N/A
apache	N/A	DOWN	N/A	N/A	N/A	N/A

## ***Improvement: DRBD 0.7.21 Bundled***

The default installation offered by Linuxha.net is now DRBD 0.7.21. If administrators have a later tarball this will automatically be compiled and used - see the Web page resources for details.

## ***Improvement: Syntax Handling for Process Monitors***

When checking the command specified to restart a failed process in the process monitor, the monitor will treat white space and a semi-colon as the command name delimiters, thus allowing

arguments to be handled correctly, or even a sequence of commands to be executed.

### ***Fix: Additional clbuildapp Error Handling***

A condition where “clbuildapp” was run against a cluster which did not have a cluster configuration file (“/etc/cluster/clconf.xml”) is now handled in a more sensible manner with a suitably worded error message produced.

### ***Fix: IP Address Removal Failure Checked***

In previous versions when an IP address was removed there could be occasions when the return code from the “ifconfig” command indicated an error, but the IP address associated was actually removed. This condition is now fixed - thus a potential problem of IP addresses for a package not clearing down when that package is stopped has now been removed.

### ***Fix: Data Synchronization Display***

The data synchronization display incorrectly displayed multi-gigabyte synchronization sizes as mega-bytes. This has now been resolved.

### ***Update: Manual Page updates for Configuration Files***

The manual pages describing the cluster and application configuration file formats have been updated to reflect the latest supports elements and attributes of these XML files.

## Release Notes for release 1.2.2 of Linuxha.net

The following notes describe the changes since 1.2.1.

Note: Due to the major improvement implemented it is necessary to perform a “migration” to 1.2.2. This is an on line process that typically takes 5-10 minutes only. The steps needed are shown following descriptions of the Improvements/Fixes found in this release.

### ***Improvement: Automated Port Allocation***

One of the biggest causes of configured issues has now been removed - the administrator is no longer responsible for choosing port numbers for the various daemons! All ports are allocated from the “port pool” as necessary. As stated the “Migration process” documented below should be followed when 1.2.2 is installed.

### ***Fix: Lems Daemon Process Monitor pattern matching***

Previously when handling more complex pattern matching the results might not be as expected since leading and trailing spaces are not removed and may be treated as part of the pattern. This has now been resolved.

### ***Fix: “clrunapp” uses Application configured Start Time***

Prior to 1.2.2 the “clrunapp” command would default to a 30 second application start-up timeout if no parameter was provided on the command line. As of 1.2.2 it will use the value configured as part of the application settings, and only default to 30 if the application does not define such a setting.

### ***Migration process for 1.2.1 to 1.2.2***

Although only a minor increment version change the main reason for the update is greatly improved handling of network ports used by the application. In previous versions ports for all daemons were explicitly defined by the administrator whenever the cluster was built or new applications added.

*With version 1.2.2 this is no longer the case.*

From this version onwards the administrator is still responsible for defining a “port pool” when building the cluster - but nothing else. All port allocations, whether for data replication or cluster daemons, will be allocated automatically when application or cluster requires them. This removes a common source of faults - especially when cluster contain multiple applications.

As always it is recommended that this document is used as a baseline only; you should document the upgrade steps necessary for each cluster and also ideally test it [even in software - such as on VMWare(R)] before attempting the steps on any production environments.

It is recommended that all the steps documented here are performed *in the same session* since for much of the change the cluster software will not actually be running!

The steps here can be performed if the cluster and applications are running or not.

## Sample Cluster Configuration

The sample cluster is running version 1.2.1 with a single application:

```
[root@centos41s2 ~]# clstat
Cluster: cluster1 - UP

      Node      Status
centos41s1      UP
centos41s2      UP

Application      Node      State  Started  Monitor  Stale  Fail-over?
      apache centos41s2  STARTED  0:00:01  Running    0      Yes
```

The first step of course is to simply install 1.2.2 on both nodes in the usual manner, for example if the RPM is to be installed use:

```
# rpm --force -Uvh linuxha12-1.2.2-1.noarch.rpm
```

Once installed kill off all cluster daemons on both nodes:

```
[root@centos41s1 tmp]# ps -ef | grep cluster1
root      5822      1  0 06:00 ?        00:00:00 cldaemon-cluster1
root      5825      1  0 06:00 ?        00:00:00 cllockd-cluster1
root      5829      1  0 06:00 ?        00:00:00 clnetd-cluster1
root      5832      1  0 06:00 ?        00:00:00 clhbd-cluster1
root      5833     5832  0 06:00 ?        00:00:00 clhbd2-cluster1
```

Note: Alter the above command to grep for the name of cluster in question [“cluster1” in the example above].

Now the processes should be removed:

```
[root@centos41s1 tmp]# kill 5822 5825 5829 5832 5833
```

Repeat this on the other node in the cluster.

Now remove the bold lines in the “/etc/cluster/clconf.xml” file on one of the nodes:

```
<global>
  <name>cluster1</name>
  <version>0.1</version>
  <data>replicated</data>
  <datadetail>drbd</datadetail>
  <dataprotocol>C</dataprotocol>
  <logdir>/var/log/cluster</logdir>
  <key>hello_fred</key>
  <port>9900</port>
  <clport>9849</clport>
  <clnetdport>9850</clnetdport>
  <clhbdport>9851</clhbdport>
  <portpool>9901,9999</portpool>
  <maxblockdevs>50</maxblockdevs>
  <drbd_network>a</drbd_network>
  <echotype>ICMP</echotype>
</global>
```

Then on that node run:

```
[root@centos41s2 tmp]# clbuild
```

For each application configured in the cluster simply edit the lems.local.xml file, for example “/etc/cluster/apache/lems.local.xml” - and remove the highlighted line:

```
<globals modules="/sbin/cluster/lems/modules"
  programs="/sbin/cluster/lems/programs"
  logs="/var/log/cluster/lems"
  port="8800"
/>
```

For each application in the cluster run the clbuildapp against them, for example:

```
[root@fc5s1 oracle-xe]# clbuildapp -A apache
```

Now form the cluster again:

```
[root@centos41s2 tmp]# clform --force
```

Check if any “lems” daemons are running for any applications and kill them off.

```
[root@centos41s1 tmp]# ps -ef | grep lem
root      6452      1  0 06:02 ?          00:00:09 lems-apache
```

On each node where an application is running manually re-start the Lems daemon for each application, for example:

```
[root@centos41s1 tmp]# lems.pl --application apache \  
--config /etc/cluster/apache/lems.local.xml --detach --verbose
```

## 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 3<sup>rd</sup> 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 `c1halt (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.