

Introduction

This document describes an approach that will allow users of the 1.0.x series of releases for Linuxha.net to migrate to the new 1.2.x series. Given that the software will obviously be used on platforms that must be available 24x7 the upgrade process describes is slightly longer than necessary since it performs the upgrade online.

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!

Sample Cluster Configuration

The sample cluster is running version 1.0.11 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
```

This is running DRBD 0.7.18:

```
[root@centos41s2 ~]# head -2 /proc/drbd
version: 0.7.18 (api:78/proto:74)
SVN Revision: 2176 build by root@centos41s2, 2006-06-01 21:26:59
```

Recording Initial Configuration

On each host create a file called “/tmp/copy_drbd_bins” with the following contents:

```
#!/usr/bin/perl

sub get_installed_drbdsetup_version {
    if(! -x "/sbin/drbdsetup") {
        return undef;
    }
    my $r=`/sbin/drbdsetup --version 2>&1`;
    my ($x) = ($r =~ /^Version:\s+(\d+\.\d+\.\d+)\s+/m);
    return $x;
}

if(! -d "/sbin/cluster/drbd-bin") {
    mkdir "/sbin/cluster/drbd-bin";
}

$v=get_installed_drbdsetup_version();
exit(0) if !defined($v);
exit(0) if -x "/sbin/cluster/drbd-bin/drbdsetup-$v";
system("cp /sbin/drbdsetup /sbin/cluster/drbd-bin/drbdsetup-$v");
chmod 0555,"/sbin/cluster/drbd-bin/drbdsetup-$v";
exit(0);
```

Once the file has been created it should be run on both machines:

```
[root@centos41s1 ~]# /tmp/copy_drbd_bins  
[root@centos41s2 ~]# /tmp/copy_drbd_bins
```

It should return immediately [having created a copy of the current version of the “drbdsetup” installed].

Software Installation

It is assumed that a version of “linuxha12” [version 1.1.05 - though ideally 1.2.0 or higher] is available in “/tmp” on both nodes. The first step, on node 1 [centos41s1] is to remove the “linuxha” package and install the “linuxha12” package:

```
[root@centos41s1 tmp]# rpm -e linuxha  
[root@centos41s1 tmp]# rpm -Uvh linuxha12-1.2.0-1.noarch.rpm  
Preparing... ##### [100%]  
 1:linuxha12 ##### [100%]
```

Following the installation almost all commands will fail, for example:

```
[root@centos41s1 tmp]# clstat  
ERROR 02/06/2006 01:34:38 Missing following attributes from globals section:  
ERROR 02/06/2006 01:34:38 clhbdport
```

On “centos42s1” edit the /etc/cluster/clconf.xml and add a “clhbdport” entry into the “globals” section, for example:

```
<global>  
  <name>cluster1</name>  
  <version>0.1</version>  
  <data>replicated</data>  
  <datadetail>drbd</datadetail>  
  <dataprotocol>C</dataprotocol>  
  <logdir>/var/log/cluster</logdir>  
  <key>jhwjehrkjwhrkrhkjw</key>  
  <port>9900</port>  
  <clport>9849</clport>  
  <clhbdport>9848</clhbdport>  
  <clnetdport>9850</clnetdport>  
  <portpool>9901,9999</portpool>  
  <maxblockdevs>50</maxblockdevs>  
  <drbd_network>a</drbd_network>  
  <echotype>ICMP</echotype>  
</global>
```

Also it is recommended that a “net_known_connections” section is added to the end of the file. For example:

```
<known_net_connections>  
  www.bbc.co.uk  
  www.demon.net  
</known_net_connections>
```

This should provide IP address or qualified host names of other machines that will be used to indicate whether a network problem is occurring, or whether the remote node is actually down.

At this point “clstat” should work again, though it will require the “--nochecksums” option:

```
[root@centos41s1 tmp]# clstat --nochecksums  
Cluster: cluster1 - UP
```

Node	Status
centos41s1	UP
centos41s2	UP

Application	Node	State	Started	Monitor	Stale	Fail-over?
apache		STARTING	0:00:28	Down	0	Yes

Note that the application definition is shown incorrectly - that is expected until the upgrade process completes!

Now perform the software installation on that node as well:

```
[root@centos41s2 tmp]# rpm -e linuxha
[root@centos41s2 tmp]# rpm -Uvh linuxha12-1.2.0-1.noarch.rpm
warning: linuxha12-1.2.0-1.noarch.rpm: V3 DSA signature: NOKEY, key ID 98b47b3f
Preparing... ##### [100%]
 1:linuxha12 ##### [100%]
```

Software reconfiguration

On “centos41s1” run the following:

```
[root@centos41s1 tmp]# clbuild -F
```

After a short while it should say:

```
INFO 02/06/2006 21:57:49
INFO 02/06/2006 21:57:49 Clbuild has completed without errors or warnings
INFO 02/06/2006 21:57:49
```

At this point check the running daemons on this node:

```
[root@centos41s1 tmp]# ps -ef | grep cluster1
root      8098      1  0  01:08 ?        00:00:02 cldaemon-cluster1
root      8100      1  0  01:08 ?        00:00:00 cllockd-cluster1
root      8102      1  0  01:08 ?        00:00:02 clnetd-cluster1
```

At this point the cluster daemons on “centos41s1” should be killed off manually ;

```
[root@centos41s1 tmp]# kill 8098 8100 8102
```

Check that they have died:

```
[root@centos41s1 tmp]# ps -ef | grep cluster1|grep -v grep
```

Also perform the same steps on the other node:

```
[root@centos41s2 tmp]# ps -ef | grep cluster1
root      3783      1  0  Jun01 ?        00:00:37 cldaemon-cluster1
root      3786      1  0  Jun01 ?        00:00:00 cllockd-cluster1
root      3790      1  0  Jun01 ?        00:00:33 clnetd-cluster1
root      5418    4401  0  18:40 pts/0    00:00:00 grep cluster1
[root@centos41s2 tmp]# kill 3783 3786 3790
[root@centos41s2 tmp]# ps -ef | grep cluster1|grep -v grep
```

At this point all cluster daemons will have been killed, but any applications should remain running without problems.

Reform the Cluster

At this point reform the cluster - ensuring the “--force” and “--noapps” options are given:

```
[root@centos41s1 tmp]# clform --force --noapps
INFO 02/06/2006 22:05:17 Validated checksum for cluster configuration
INFO 02/06/2006 22:05:17 SSH communication to centos41s2 will be:
INFO 02/06/2006 22:05:17 192.168.1.61 ("a" network)
INFO 02/06/2006 22:05:17 Checking that the cluster is not already running...
INFO 02/06/2006 22:05:17 *** ATTEMPTING TO FORM CLUSTER cluster1 ***
INFO 02/06/2006 22:05:18 Forced Start of cldaemon on centos41s1...
INFO 02/06/2006 22:05:19 Forced Start of cldaemon on centos41s2...
INFO 02/06/2006 22:05:19 Waiting for cluster to form...
INFO 02/06/2006 22:05:25 Cluster cluster1 started successfully.
```

Now check the cluster state:

```
[root@centos41s1 tmp]# clstat
Cluster: cluster1 - UP

      Node      Status
centos41s1      UP
centos41s2      UP

Application      Node      State  Started  Monitor  Stale  Fail-over?
-----
apache centos41s2  STARTED  0:03:20  Running    0      Yes
```

Also check the running cluster daemons on each node:

```
[root@centos41s1 tmp]# ps -ef | grep cluster1
root      9968      1  0  22:05 ?          00:00:00 cldaemon-cluster1
root      9971      1  0  22:05 ?          00:00:00 cllockd-cluster1
root      9975      1  0  22:05 ?          00:00:00 clnetd-cluster1
root      9978      1  0  22:05 ?          00:00:00 clhbd-cluster1
root      9979     9978  0  22:05 ?          00:00:00 clhbd2-cluster1
```

and:

```
[root@centos41s2 tmp]# ps -ef | grep cluster1
root      5532      1  0  18:45 ?          00:00:00 cldaemon-cluster1
root      5534      1  0  18:45 ?          00:00:00 cllockd-cluster1
root      5536      1  0  18:45 ?          00:00:00 clnetd-cluster1
root      5538      1  0  18:45 ?          00:00:00 clhbd-cluster1
root      5539     5538  0  18:45 ?          00:00:00 clhbd2-cluster1
```

Application Changes

The other change to be aware of is to remove the following attributes from any application “network” settings;

- checklist
- checkpercent
- pingtype
- pingtimeout

For example if a network section for the sample “apache” application contained:

```
<networks>
  <network net="a" ip="192.168.1.202,203"
    checklist="192.168.1.1" checkpercent="100"
    pingtype="icmp" pingtimeout="3"/>
</networks>
```

Then it should instead be altered to:

```
<networks>
  <network net="a" ip="192.168.1.202,203"/>
</networks>
```

If any application is altered, run the “clbuildapp” tool against it:

For example:

```
[root@centos41s1 tmp]# clbuildapp -A apache
```

The new style of output will be shown:

```
INFO 03/06/2006 13:14:07 Backups directory defaulted to /clbackup
INFO 03/06/2006 13:14:07
INFO 03/06/2006 13:14:07 Validation of Application 'apache' started.
INFO 03/06/2006 13:14:07 ['/var/log/cluster/build/apache-check-300606031314.log']
INFO 03/06/2006 13:14:09 Initial Validation of Application successful.
INFO 03/06/2006 13:14:09
INFO 03/06/2006 13:14:09 NOTE: Application previously defined - being reconfigured.
INFO 03/06/2006 13:14:09
INFO 03/06/2006 13:14:09 Host Environment Validation started.
INFO 03/06/2006 13:14:09 ['/var/log/cluster/build/apache-envcheck-300606031314.log']
INFO 03/06/2006 13:14:11 Host Environment Validation successful.
INFO 03/06/2006 13:14:11
INFO 03/06/2006 13:14:11 Cluster state      : RUNNING
INFO 03/06/2006 13:14:11 Application state: RUNNING
INFO 03/06/2006 13:14:11
INFO 03/06/2006 13:14:11 Volume Group Configuration started.
INFO 03/06/2006 13:14:11 ['/var/log/cluster/build/apache-lvm-300606031314.log']
INFO 03/06/2006 13:14:13 Volume Group Configuration successful.
INFO 03/06/2006 13:14:13
INFO 03/06/2006 13:14:13 Application Resource Allocation started.
INFO 03/06/2006 13:14:13 ['/var/log/cluster/build/apache-build-300606031314.log']
INFO 03/06/2006 13:14:18 Application Resource Allocation successful.
INFO 03/06/2006 13:14:18
```

And that should be it!