
Subject: DRBD+Heartbeat clustering question!
Posted by [zenny](#) on Sun, 04 May 2008 08:45:04 GMT
[View Forum Message](#) <> [Reply to Message](#)

I have gone through http://wiki.openvz.org/HA_cluster_with_DRBD_and_Heartbeat and other relevant links there thoroughly.

1) The setup mentioned there seems to trigger to serve the secondary HN node in case either primary HN Node or any containers inside primary stops working, if I understand right.

But in OVZ case, what I have seen is mostly the containers stops working due to resource problem. In that case, the explained setup in the wiki may not help.

Let us assume that we have a primary HN 1 and a secondary HN 2 as pictured as in attachment (ovz-clustering.txt) with equal resources (CPU power and memory). Let us say one of the containers in primary HN 1 (say Container 1db) failed, then the primary HN 1 stops working and the secondary starts. But since the secondary also has the same resources, it will also stop as soon as it will be failed over from primary to secondary. Then the entire failover clustering setup collapses.

What I want to achieve is, in case Container 1a fails then Container 2a alone will be activated in secondary server whereas other containers in HN 1 and physical HN 1 remains running in the primary server. Thus, the resource allocation will not be consumed at once so that both of the server will not collapse at the same time.

Similarly, any instance of Container 1db failure will be load balanced by Container 2db without affecting any other containers in the primary HN 1 as well as HN1 itself.

Once the physical HN1 collapses then the entire secondary HN 2 will start serving.

Is this possible from the setup mentioned in the wiki or there are specific tweaks that is needed to accomplish this setup?

2) And what is the use of two ethernet cards in the setup? Is that merely for crossover cables to transfer data realtime between the primary and secondary hardware nodes? Can't it be achieve with a single card?

Thanking you in advance for your attention. Kindly make advise if I misunderstood the entire DRBD-Heartbeat concept. I wholeheartedly welcome any comments.

File Attachments

1) [ovz-clustering.txt](#), downloaded 600 times

Subject: Re: DRBD+Heartbeat clustering question!
Posted by [dim](#) on Sun, 04 May 2008 11:44:34 GMT
[View Forum Message](#) <> [Reply to Message](#)

The problem is that drbd is about remote replication and not about clustering. So, the options are:
1) to have per-Container drbd device.
2) to setup cluster filesystem (GFS, for example) over drbd device.

Subject: Re: DRBD+Heartbeat clustering question!
Posted by [zenny](#) on Sun, 04 May 2008 13:14:20 GMT
[View Forum Message](#) <> [Reply to Message](#)

Do I have to deploy both options or either one?

Subject: Re: DRBD+Heartbeat clustering question!
Posted by [dim](#) on Sun, 04 May 2008 15:04:20 GMT
[View Forum Message](#) <> [Reply to Message](#)

either one.

in the first case you can remount this device on dst node without affecting other Containers, while in the second one you can keep Containers mounted on both nodes and no remount is required.

Subject: Re: DRBD+Heartbeat clustering question!
Posted by [zenny](#) on Sun, 04 May 2008 19:47:49 GMT
[View Forum Message](#) <> [Reply to Message](#)

Dim: Thanks. But does the default GFS package that comes with CentOS5.1 conflicts with OpenVZ? Or there are special packages compatible with OpenVZ?

GFS is new to me. Do I need to install it only on hardware node or in the containers, too?

And what about the GFS licensing? Can GFS be used to production servers like CentOS? Any idea about free alternatives to GFS?

Subject: Re: DRBD+Heartbeat clustering question!
Posted by [dim](#) on Sun, 04 May 2008 20:22:57 GMT
[View Forum Message](#) <> [Reply to Message](#)

zenny wrote on Sun, 04 May 2008 15:47Dim: Thanks. But does the default GFS package that comes with CentOS5.1 conflicts with OpenVZ? Or there are special packages compatible with OpenVZ?

You may use default one.

Quote:

GFS is new to me. Do I need to install it only on hardware node or in the containers, too?

You need to install Containers' private over GFS partition.

Subject: Re: DRBD+Heartbeat clustering question!
Posted by [zenny](#) on Wed, 07 May 2008 10:04:32 GMT
[View Forum Message](#) <> [Reply to Message](#)

Dim:

Thanks for your reply and time. It was indeed enlightening.

After going through some posts in this forum, GFS does not seem to be an ideal solution in terms of speed.

1) However, as you advised to have per-Container drdb device instead of GFS, do I need to follow the instructions without formatting the /vz volume and then make it a DRBD + ext3 volume later as explained in http://wiki.openvz.org/HA_cluster_with_DRBD_and_Heartbeat ? I meant to install cluster manager in HN node?

2) And then again do I have to install the cluster manager in the containers also to achieve what I thought in the first post? Please elaborate.

3) I am yet to figure out two ethernet cards in each node in the above wikipage setup? Is that for the crossover cables to backup data or is it as a failover ethernet card? Just simply wondering!

Thanks!

Subject: Re: DRBD+Heartbeat clustering question!
Posted by [dim](#) on Wed, 07 May 2008 10:32:09 GMT
[View Forum Message](#) <> [Reply to Message](#)

zenny wrote on Wed, 07 May 2008 06:04Dim:

1) However, as you advised to have per-Container drdb device instead of GFS, do I need to follow the instructions without formatting the /vz volume and then make it a DRBD + ext3 volume later as explained in http://wiki.openvz.org/HA_cluster_with_DRBD_and_Heartbeat ? I meant to install cluster manager in HN node?

I didn't advise, I just presented available choices. DRBD is not ideal solution as well, you know. Yes, you need to follow instructions and install cluster manager in HN. But its configuration should

differ - each Container should be a resource, not vz service itself.

Quote:

2) And then again do I have to install the cluster manager in the containers also to achieve what I thought in the first post? Please elaborate.

No, this doesn't help. The problem that in case of failure you need to mount Container's private on failover node in order to start Container, while cluster manager inside it can't do this.

Quote:

3) I am yet to figure out two ethernet cards in each node in the above wikipage setup? Is that for the crossover cables to backup data or is it as a failover ethernet card? Just simply wondering!

I think it is for performance mostly. DRBD is network replication and its health depends on transfer rate greatly.

Subject: Re: DRBD+Heartbeat clustering question!
Posted by [zenny](#) on Wed, 07 May 2008 13:35:25 GMT
[View Forum Message](#) <> [Reply to Message](#)

DRBD is not ideal solution as well, you know.

It gave me more confusion now. What do you think ideal choice, then? If you were me, what would you go for in such a situation? Just want to know expert advise!

PS: Indeed, I am learning a whole lot from you with each of your reply, though they are short but worths millions, so is that of Kir

Subject: Re: DRBD+Heartbeat clustering question!
Posted by [dim](#) on Thu, 08 May 2008 06:53:05 GMT
[View Forum Message](#) <> [Reply to Message](#)

zenny wrote on Wed, 07 May 2008 09:35DRBD is not ideal solution as well, you know.

It gave me more confusion now. What do you think ideal choice, then?
SAN storage, replication by hardware. But it is very expensive