
Subject: vzmigrate and shared filesystems
Posted by [Zirafarafa](#) on Fri, 25 Jul 2008 13:44:08 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi all

I have a set of openvz servers which can access a shared GFS filesystem.

Some of my vz guests are hosted on the GFS system, and they are running fine.

However, if I want to migrate the vz online to a different host, how would I accomplish this, as all hosts can see the gfs share - no rsync is needed.

Any pointers for me?

Subject: Re: vzmigrate and shared filesystems
Posted by [kenjy](#) on Sat, 26 Jul 2008 16:29:11 GMT
[View Forum Message](#) <> [Reply to Message](#)

I don't understand the question, if you have a GFS who shares with a lot of hosts then why you just simple point your VE in the other HW node to the correct path? really if you wanna move a VE the way is too easy, you just need to copy/move it to a new place and rsync its a good tool

Subject: Re: vzmigrate and shared filesystems
Posted by [Zirafarafa](#) on Sun, 27 Jul 2008 06:29:15 GMT
[View Forum Message](#) <> [Reply to Message](#)

I need to be able to online migrate the running VE to a different HW node to be able to optimise memory/processing load across my HW nodes.

I know I can start up the VE on a different node, but I want to be able to migrate with no downtime.

Subject: Re: vzmigrate and shared filesystems
Posted by [kenjy](#) on Mon, 28 Jul 2008 05:15:06 GMT
[View Forum Message](#) <> [Reply to Message](#)

Just rsync from the first HN to the second one and do vzctl stop VE in the first HN and then vzctl start VE in the second one.

Subject: Re: vzmigrate and shared filesystems
Posted by [Zirafarafa](#) on Mon, 28 Jul 2008 08:04:17 GMT
[View Forum Message](#) <> [Reply to Message](#)

kenjy wrote on Mon, 28 July 2008 07:15 Just rsync from the first HN to the second one and do vzctl stop VE in the first HN and then vzctl start VE in the second one.

Hi

There are two problems with that suggestion.

1. I dont need to rsync, as both nodes can see the same filesystem over GFS
2. Stopping and starting the VE defeats the original intent of an online migration.

My problem is that vzmigrate will try to rsync the data, which is unnecesary, and potentially damaging to the data.

Subject: Re: vzmigrate and shared filesystems
Posted by [dowdle](#) on Tue, 29 Jul 2008 01:17:55 GMT
[View Forum Message](#) <> [Reply to Message](#)

I do not have any OpenVZ setups with shared storage so I'm not speaking from experience here... but... are you aware that vzmigrate is a shell script? Examine it, copy it... and modify it as needed to do whatever is necessary to do the operation.

I would imagine you could just checkpoint the system on the source node, transfer the checkpoint to the destination node, restore from checkpoint on the destination machine. I'm not sure what you'd need to do with the checkpointed system on the source machine... but a close look at how vzmigrate handles the checkpointing should expose at least one method.

My goal here isn't to give you the perfect answer but to get you on a path to discovery of your own solution. After you put enough effort into it and get it working as desired, you are strongly encouraged to create a wiki page on the topic.

I have run into at least one person who has a custom built OpenVZ kernel patched against a non-stable branch... and they were able to get NFS based guests and live migration working. I'd love to see a page on that too... but I'd really rather stick with the stable OpenVZ kernel branches.

An alternative that would require a total reworking of your setup, so it isn't recommended, is to use DRBD and heartbeat to create a high availability solution. There is a wiki page on that setup but again, it isn't something I've tried myself.

Subject: Re: vzmigrate and shared filesystems
Posted by [Zirafarafa](#) on Tue, 29 Jul 2008 06:51:49 GMT
[View Forum Message](#) <> [Reply to Message](#)

dowdle wrote on Tue, 29 July 2008 03:17An alternative that would require a total reworking of your setup, so it isn't recommended, is to use DRBD and heartbeat to create a high availability solution. There is a wiki page on that setup but again, it isn't something I've tried myself.

Currently my shared iscsi from which the GFS is mounted is from a pair of DRBD servers, shared to 6 openvz servers.

I will check the vzmigrate script. Probably all I will need to do is to remove the rsync commands, but leave most of the rest the same.