Subject: OpenVZ cluster
Posted by devonblzx on Tue, 17 Jul 2007 19:06:27 GMT
View Forum Message <> Reply to Message

Has anyone had any success with clustering OpenVZ nodes? Anyone have any ideas on how to approach this other than using vzmigrate?

Subject: Re: OpenVZ cluster
Posted by pringleso on Wed, 01 Aug 2007 15:37:27 GMT
View Forum Message <> Reply to Message

i havent had any luck with it either, i just use my own firewall.

Subject: Re: OpenVZ cluster
Posted by devonblzx on Wed, 01 Aug 2007 15:39:38 GMT
View Forum Message <> Reply to Message

Did you mean to reply to my other thread? (iptables firewall)

Subject: Re: OpenVZ cluster
Posted by pringleso on Wed, 01 Aug 2007 15:51:26 GMT
View Forum Message <> Reply to Message

woops, yeap!

but regurading this thread, i have a feeling that VEs can't be part of a "true" cluster since clustering is kernel based, of course it's possible with a lot of hacking. However, I'm sure that cluster apps such as the folding project or seti would work on your ve.

Subject: Re: OpenVZ cluster
Posted by rickb on Wed, 01 Aug 2007 21:23:21 GMT
View Forum Message <> Reply to Message

Hi, cluster is a broad term and can mean many things from marketing to technical definitions. tell us what you are wanting to accomplish.

I have had good success using openVZ 2.6.9 + GFS + fiber channel for shared storage, using this, a VE can be stopped on one node and fired up another instantly without file transfers between nodes (they both share the same fs). If your application can handle this, you could run the same VE on each server and balance the traffic between them.- Add an LVS router in front of the two HNs. Its not plug and play but if you are looking to accomplish a specific task it is the best solution I have come across for HA and load distribution.

Rick Blundell

Subject: Re: OpenVZ cluster
Posted by devonblzx on Wed, 01 Aug 2007 21:29:15 GMT
View Forum Message <> Reply to Message

Yes, I was a little vague. I was just trying to hear some stories of what people have done.

I was thinking more of a way to combine all the nodes you have to have them to balance the load and then have some sort of shared storage, sort of like your idea. I supposed we could have some sort of redundant NFS solution and have all the nodes sharing the storage incase of one going down or being overloaded we could always start it up on another one and stop it on the original but would a gigabit NIC be fast enough to give me the disk speed I require?

Subject: Re: OpenVZ cluster
Posted by rickb on Wed, 01 Aug 2007 21:38:39 GMT
View Forum Message <> Reply to Message

i would avoid nfs like the plague. when you read about what nfs can do, and use it in its most basic form, it does work great. its quick and cheap to deploy. when you start doing tons of concurrent and random io operations to your nfs box, from my experience, even with all of the tweaks on google and oreilley books, you are going to get dog slow performance, I am talking about xx or xxx KB/sec.

If you want good performance for remote storage there are some new options available, iscsi and ata over ethernet. to use these effectively you need to be using a switch supporting jumbo frames and configure your networ interface appropriately. using standard mtu will seriously degrade performance due to 1500byte fragmentation. from my experience, 2Gbit fiber channel blows the doors off everythign else in terms of performance, but it may be pricey depending on what business application you are using it for. I am using apple xserve xraid units, which are 2X 6x @raid10 sata. I regularly see throughput rates of XX MB/sec writes and 5000-10000 trans/sec writes. if you can manage this over nfs, I would be glad to purchase your consulting services . check out apple's fiber channel units or if you find another solution, let me know as I am interested in this area of computing.

Rick