system toolbox - Podrška #21285

glusterfs, cloud hosting

07.10.2010 16:39 - Ernad Husremović

Status:	Zatvoreno	**Početak:**	07.10.2010
Prioritet:	Normalan	**Završetak:**	
Odgovorna osoba:	Ernad Husremović	**% završeno:**	0%
Kategorija:		**Procjena vremena:**	0.00 sat
Ciljna verzija:			

Opis

Historija

#1 - 07.10.2010 16:39 - Ernad Husremović

http://rackerhacker.com/ - ekstra blog

#2 - 07.10.2010 16:40 - Ernad Husremović

- Naslov promijenjeno iz glusterfs u glusterfs, cloud hosting

http://rackerhacker.com/redundant-cloud-hosting-configuration-guide/

#3 - 07.10.2010 16:53 - Ernad Husremović

http://www.mail-archive.com/gluster-users@gluster.org/msg02016.html

```
Hello,
DRBD is faster because, among other things, it replicate byte blocks on disks,
whatever they are. A corrupted master will be simply replicated...
In order to try to improve glusterfs (metadata file level) performance one could
try NUFA first at client side and then server side replication.
NUFA is _very_ good to accelerate local writes and reads at distributed
configuration.
Maybe, such concepts could be developed to replication in future.
Regards.
Andre Felipe Machado
http://www.techforce.com.br
```

#4 - 07.10.2010 17:01 - Ernad Husremović

http://serverfault.com/questions/146677

've finally managed to get this solved using GlusterFS in both boxes. Some things learned in the process:

- First I've tried a generic RAID 1 setup. The main problem with this is that the client always use **tcp to contact both servers**, even when one of them is in the same machine. So I've to change client configurations to replace the tpc 'local' volume with a direct access (storage/posix) volume
- To avoid stressing the network link, **every client read** use the local storage with directive option read-subvolume. Off course to keep the RAID1 integrity GlusterFS always check other volumes as well, but the actual file is retrieved directly from disk
- Performance is good, but client process seems like memory hug. I think is related to quickread volume, I need to investigate further

Modified client configuration:

```
# Server1 configuration (RAID 1)
volume server2-tcp
    type protocol/client
    option transport-type tcp
    option remote-host server2
    option transport.socket.nodelay on
    option transport.remote-port 6996
    option remote-subvolume brick1
end-volume

volume posix-local
    type storage/posix
    option directory /shared
```

```
end-volume

volume locks-local
    type features/posix-locks
    subvolumes posix-local
end-volume

volume brick-local
    type performance/io-threads
    option thread-count 8
    subvolumes locks-local
end-volume

volume mirror-0
    type cluster/replicate
    option read-subvolume brick-local
    subvolumes brick-local server2-tcp
end-volume

.....
```

#5 - 07.10.2010 17:10 - Ernad Husremović

http://groups.drupal.org/node/61908

http://github.com/evolvingweb/cookbooks/tree/master/glusterfs/

#6 - 07.10.2010 17:13 - Ernad Husremović

Peek is Cooking with Chef

#7 - 26.11.2010 20:10 - Ernad Husremović

- Status promijenjeno iz Dodijeljeno u Zatvoreno