


Gluster performance testing

Notes from  **NDS-860** - Gluster performance problems CLOSED :

I ran the following command to reproduce the disk workload created by launching the postgresql container:

```
for j in `seq 0 4`; do time for i in `seq 0 1024`; do dd if=/dev/zero of=/path/to/storage/temp/test$i.tmp bs=1 count=1024 status=none && sync ; done ; done
```

(print the time it takes to write 1024 1k file x5)

The baseline time for writing to the non-gluster filesystem was 24 seconds.

I verified the CoreOS instances had the proper mount options for the gluster filesystem.

The baseline time for writing to the gluster file system before any modifications was 93 seconds.

I followed tuning recommendations from:

<http://lists.gluster.org/pipermail/gluster-users/2016-January/024865.html>
<http://blog.gluster.org/2016/10/gluster-tiering-and-small-file-performance/>

After following the above turning recommendations the improvement was only 2 seconds.

The gluster profiler did not provide any useful information, WRITE and FSYNC are the source of the latency issues.

There is no significant CPU, memory, disk io or network usage during the tests.