

D4Science Infrastructure - Task #1221

Assess the difference in the performance of Thredds with and without nginx

Oct 22, 2015 07:29 PM - Gianpaolo Coro

Status:	Closed	Start date:	Oct 22, 2015
Priority:	Normal	Due date:	
Assignee:	Gianpaolo Coro	% Done:	100%
Category:	Application	Estimated time:	0.00 hour
Target version:	D4Science Infrastructure Upgrade to gCube 3.9.0		
Infrastructure:	Development, Pre-Production		
Description			
Thredds streams a lot of information to visualisation clients. I'm going to evaluate if nginx is a bottleneck for our use case.			

History

#2 - Oct 26, 2015 11:25 AM - Gianpaolo Coro

- Assignee changed from Gianpaolo Coro to Tommaso Piccioli
- % Done changed from 0 to 50

I detected around 60% of performance decrease in visualising maps using thredds with nginx. These results could also depend on the different performance of the portals I used to make the experiments, but the difference is very large and I suppose nginx plays a role in this delay. I used a very high resolution GIS layers to do the experiments.

Furthermore, here is a test you can run to get a geoTiff from this service:

http://thredds-d-d4s.d4science.org/thredds/wcs/public/netcdf/WOA2005TemperatureAnnual_CLIMATOLOGY_METEOROLOGY_ATMOSPHERE__nc?request=GetCoverage&version=1.0.0&service=WCS&format=GeoTIFF&coverage=t00an1&bbox=0,-80,360,80&time=0001-07-01T12:00:00Z&vertical=0

At the present status, if not better explored, I would avoid using nginx in front of Thredds.

#3 - Oct 26, 2015 03:49 PM - Tommaso Piccioli

- Assignee changed from Tommaso Piccioli to Gianpaolo Coro
- % Done changed from 50 to 0

I really doubt that the performance differences are due to the proxy nginx.

However, to conduct a proper test we can choose one of these ways, or both:

- Temporarily remove the proxy nginx server development
- Put both virtual machines on the same server xen with the same storage server

At the moment tomcat port 8180 on the dev server is temporarily open, you can test thredds on dev without going through the nginx proxy

#4 - Oct 26, 2015 05:48 PM - Gianpaolo Coro

I have done more precise tests visualising two large layers on port 80 and 8180 alternatively:

http://thredds-d-d4s.d4science.org:8180/thredds/wms/public/netcdf/global-reanalysis-phys-001-004-b-ref-fr-mjm95-gridt__ENVIRONMENT_OCEANS_ELEVATION_1366210702774.nc?service=wms&version=1.3.0&request=GetMap&layers=votemper&styles=&srs=EPSG:4326&CRS=EPSG:4326

http://thredds-d-d4s.d4science.org/thredds/wms/public/netcdf/global-reanalysis-phys-001-004-b-ref-fr-mjm95-gridt__ENVIRONMENT_OCEANS_ELEVATION_1366210702774.nc?service=wms&version=1.3.0&request=GetMap&layers=votemper&styles=&srs=EPSG:4326&CRS=EPSG:4326

The average reponse time is variable because it depends on statistical fluctuations in network performance, thus I made several measurements. The 8180 link results to be 1% faster on average.

If this threshold is neglectable with respect to the requirement of having a proxy in front of the service then we can go forward.

#5 - Nov 27, 2015 04:15 PM - Gianpaolo Coro

- *Status changed from New to Closed*

- *% Done changed from 0 to 100*