

D4Science Infrastructure - Task #536

Thredds: tomcat runs as root

Sep 04, 2015 12:33 PM - Roberto Cirillo

Status:	Closed	Start date:	Sep 23, 2015
Priority:	High	Due date:	
Assignee:	Tommaso Piccioli	% Done:	100%
Category:	System Application	Estimated time:	0.00 hour
Target version:	improve tomcat security		
Infrastructure:	Production		
Description The tomcat on "thredds.d4science.org" has been installed under the user tomcat but it runs as root. Now it is configured to run under the port 80. I guess, it runs as root because thredds needs to run under the port 80. We should try to install a nginx as frontend (as done for geoservers) and change the tomcat port to 8080. Do you think is possible to do it in the next release?			
Subtasks: Support # 740: Create a Thredds instance for the development environment Closed			
Related issues: Related to D4Science Infrastructure - Task #183: tomcat running as root user ... Closed May 30, 2015 Related to D4Science Infrastructure - Task #1283: Thredds layers colors not v... Closed Nov 23, 2015 Related to D4Science Infrastructure - Task #2946: Install ganglia and the loc... Closed Mar 23, 2016			

History

#1 - Sep 04, 2015 12:42 PM - Andrea Dell'Amico

- Related to Task #183: tomcat running as root user under port 80 added

#2 - Sep 04, 2015 12:44 PM - Andrea Dell'Amico

As usual, I'd prefer to automate all the VM configuration. Is it possible? Is there a devel instance where we can test the new configuration?

#3 - Sep 04, 2015 12:51 PM - Roberto Cirillo

Andrea Dell'Amico wrote:

As usual, I'd prefer to automate all the VM configuration. Is it possible? Is there a devel instance where we can test the new configuration?

Unfortunately there isn't a dev instance for thredds

#4 - Sep 04, 2015 12:52 PM - Andrea Dell'Amico

Is it possibile to create one and perform some basic tests on it?

#5 - Sep 04, 2015 12:56 PM - Roberto Cirillo

Andrea Dell'Amico wrote:

Is it possibile to create one and perform some basic tests on it?

I guess it is possible. [Gianpaolo Coro](#) could you help us and tell us what is the minimum set of resources for a minimal thredds installation in dev, please?

#6 - Sep 04, 2015 12:57 PM - Gianpaolo Coro

Thredds is an old service installation that needs larger file system capacity with respect to other services. It needs to store geospatial files locally to the FS and publishes them as maps.

Currently, there is only one service instance running, because at "that time" we had not enough space to create two machines. Thredds hosts especially environmental maps (and thus files) which need to be both on dev and prod environments.

The best could be to replicate the Thredds machine, creating a dev machine hosting the same files as the current machine.
Please, note that one geospatial file (NetCDF, GeoTiff, ASC) could be as large as 10 GBs.

#7 - Sep 04, 2015 01:01 PM - Andrea Dell'Amico

Gianpaolo Coro wrote:

Thredds is an old service installation that needs larger file system capacity with respect to other services. It needs to store geospatial files locally to the FS and publishes them as maps.

Currently, there is only one service instance running, because at "that time" we had not enough space to create two machines. Thredds hosts especially environmental maps (and thus files) which need to be both on dev and prod environments.

The best could be to replicate the Thredds machine, creating a dev machine hosting the same files as the current machine. Please, note that one geospatial file (NetCDF, GeoTiff, ASC) could be as large as 10 GBs.

Well, it could be temporary. So we could rsync (or get from a backup) the data on a storage disk and mount it on a new machine, that would be deployed in a way similar to the geoservers: recent distribution, newer tomcat, firewall rules, etc.

#8 - Sep 04, 2015 01:08 PM - Gianpaolo Coro

It's OK to me.

#9 - Oct 20, 2015 11:11 AM - Andrea Dell'Amico

- Assignee changed from Andrea Dell'Amico to *_InfraScience Systems Engineer*

The configuration with tomcat running as an unprivileged user has been successfully tested in dev, see [#740](#)

#10 - Oct 20, 2015 01:14 PM - Tommaso Piccioli

- Status changed from New to In Progress

I'm going to install a new production host configured as the dev one, then we will switch to it the production environment.

#11 - Nov 23, 2015 01:28 PM - Tommaso Piccioli

- Assignee changed from *_InfraScience Systems Engineer* to Gianpaolo Coro

thredds-p-d4s.d4science.org up and running, tomcat instance on port 8180 running with gcube user and nginx proxy on port 80

thredds must be configured

#12 - Dec 01, 2015 03:28 PM - Gianpaolo Coro

- Related to Task #1283: Thredds layers colors not visualised on GisViewer added

#13 - Dec 01, 2015 03:31 PM - Gianpaolo Coro

First, I would like to finish the configuration in the development environment for the related ticket [#1283](#). I will then report it to the new service before deploying it in prod.

#14 - Dec 18, 2015 01:39 PM - Gianpaolo Coro

- Assignee changed from Gianpaolo Coro to Roberto Cirillo

The Thredds instance thredds-p-d4s.d4science.org has been configured and the catalog can be accessed here

<http://thredds-p-d4s.d4science.org/thredds/catalog/public/netcdf/catalog.html>

I configured the web application and cleaned-up the content folder under /data/content/

Next actions should be:

@roberto.cirillo@isti.cnr.it to update and check the GHN running on that machine endowed with DataTransfer
@tommaso.piccioli@isti.cnr.it to (i) align the "/data/content/thredds/public/netcdf/" folder on thredds-p-d4s to the corresponding folder on thredds.d4science, (ii) assign the name thredds.d4science as alias to thredds-p-d4s
@gianpaolo.coro@isti.cnr.it double-check the consistency of layers display and use of data by computations

#15 - Mar 03, 2016 04:31 PM - Roberto Cirillo

- Status changed from In Progress to Closed

#16 - Mar 23, 2016 02:29 PM - Andrea Dell'Amico

- Related to Task #2946: Install ganglia and the local nagios checks on thredds-p-d4s.d4science.org added

#17 - Mar 23, 2016 02:53 PM - Andrea Dell'Amico

- Status changed from Closed to In Progress

It seems that the switch did not happen, the old machine is still active.

#18 - Mar 23, 2016 04:06 PM - Roberto Cirillo

Could anyone add another user called gcube1 on thredds-p-d4s?

The gcube user is used by tomcat. I think it's better to have a separated user dedicated to the ghn

#19 - Mar 23, 2016 04:53 PM - Tommaso Piccioli

The ghn will be installed with the user gcube, as it is on the dev machine thredds-d-d4s.d4science.org

#20 - Mar 23, 2016 05:25 PM - Roberto Cirillo

- Status changed from In Progress to Feedback

- Assignee changed from Roberto Cirillo to Gianpaolo Coro

The data-transfer service has been installed (under gcube user) and it is running on devsec scope for testing purpose. Now the tomcat is running on port 8180, the ghn is running on port 8080

#21 - Apr 01, 2016 04:06 PM - Andrea Dell'Amico

Any other progress?

#22 - Apr 01, 2016 04:34 PM - Gianpaolo Coro

Tests for this issue are quite long and require attention. I'm currently managing more urgent issues, whose tickets have been indicated as having very high priority (updating Readme files in 120 components of mine). I will go back to this issue asap.

#23 - Apr 05, 2016 02:50 PM - Gianpaolo Coro

- Assignee changed from Gianpaolo Coro to Andrea Dell'Amico

The thredds-p-d4s.d4science.org installation works very well. It can be substituted to thredds.d4science.org in the prod. environment, after assigning this alias to that machine.

Please, be sure that the RAM enhancements just requested for the other thredds machines have been applied also to thredds-p-d4s.d4science.org. The previous thredds.d4science.org machine can be dismissed after aligning the /data/content/thredds/public/netcdf/ folders with the new machine's.

#24 - May 05, 2016 01:02 PM - Andrea Dell'Amico

- Assignee changed from Andrea Dell'Amico to Tommaso Piccioli

#25 - May 18, 2016 10:41 AM - Roberto Cirillo

Any update on this?

The new vm should be in production in the next week otherwise we have to postpone again this activity

#26 - May 23, 2016 05:30 PM - Roberto Cirillo

- Status changed from Feedback to In Progress

The new VM has been deployed in production

#27 - May 23, 2016 05:31 PM - Roberto Cirillo

- Status changed from In Progress to Closed