

D4Science Infrastructure - Task #8875

Enable VPN for external gCube Liferay Portal instance (EDISON)

Jun 08, 2017 02:30 PM - Massimiliano Assante

Status:	Closed	Start date:	Jun 08, 2017
Priority:	Normal	Due date:	Jun 16, 2017
Assignee:	_InfraScience Systems Engineer	% Done:	100%
Category:	Other	Estimated time:	0.00 hour
Target version:	EDISON Migration outside CNR		
Infrastructure:	Development		
Description <p>The current VM edison.d4science.org (dev3.d4science.org) hosting one instance of gCube Liferay Portal for the EDISON Project has been successfully migrated to an external site part of the EGI Infrastructure.</p> <p>I'm not sure what is the best solution here, whether a VPN or allowing specific exceptions in our firewall for their static public IP</p> <p>They now need to be able to contact some database within CNR Premises.</p> <ul style="list-style-type: none">• jdbc.url: jdbc:postgresql://postgresql-srv-dev.d4science.org:5432/survey-db <p>And successively</p> <ul style="list-style-type: none">• ldap.d4science.org (for writing here they will use our script to export the users). <p>Also, the node must be authorised to run in the Development Infrastructure, @lucio.lელი@isti.cnr.it can you take care of this as soon as they provide us with the IP.</p>			
Related issues:			
Copied to D4Science Infrastructure - Task #9203: Enable access for production...			Closed Aug 21, 2017 Nov 03, 2017

History

#1 - Jun 08, 2017 02:38 PM - Massimiliano Assante

- Description updated

#2 - Jun 09, 2017 02:57 PM - Andrea Dell'Amico

A VPN is overkill for this kind of activity. As soon as we have the IP addresses that need access, I'll grant access to them.

#3 - Jun 09, 2017 03:04 PM - Massimiliano Assante

Andrea, they are going to need to contact also D4Science clusters such as Cassandra and MongoDB and jackrabbit repository and ElasticSearch

#4 - Jun 09, 2017 03:14 PM - Andrea Dell'Amico

Massimiliano Assante wrote:

Andrea, they are going to need to contact also D4Science clusters such as Cassandra and MongoDB and jackrabbit repository and ElasticSearch

If we have a list of resources I still prefer to explicitly open the single services. A VPN would give open access to all our datacenter network.

#6 - Jun 19, 2017 05:57 PM - Andrea Dell'Amico

- Status changed from New to In Progress

- % Done changed from 0 to 100

All the required services are now accessible by dev.datasciencepro.eu

#7 - Jun 19, 2017 05:58 PM - Andrea Dell'Amico

- Status changed from In Progress to Feedback

#8 - Jun 20, 2017 11:36 AM - Boro Jakimovski

Do we need to setup the same for the staging.datasciencepro.eu (at least the db access)

and

www.datasciencepro.eu (194.149.136.227) for production site. Does it means that we will have to connect to another DB instance, and do we need to setup ldap sync with another ldap instance or the same is enough?

Regards,
Boro.

#9 - Jun 23, 2017 03:03 PM - Massimiliano Assante

@andrea.dellamico@isti.cnr.it it seems that the elastic search node (elastic0-d-d4s.d4science.org) is not contactable from 194.149.136.225 (dev.datasciencepro.eu)

```
PORTAL] 1049431 [elasticsearch[Molly Hayes][transport_client_boss][T#1]{New I/O boss #17}} WARN org.elasticsearch.transport.netty - [Molly Hayes] exception caught on transport layer [[id: 0x99a98814]], closing connection
java.net.NoRouteToHostException: No route to host
    at sun.nio.ch.SocketChannelImpl.checkConnect(Native Method)
    at sun.nio.ch.SocketChannelImpl.finishConnect(SocketChannelImpl.java:717)
    at org.jboss.netty.channel.socket.nio.NioClientBoss.connect(NioClientBoss.java:150)
    at org.jboss.netty.channel.socket.nio.NioClientBoss.processSelectedKeys(NioClientBoss.java:105)
    at org.jboss.netty.channel.socket.nio.NioClientBoss.process(NioClientBoss.java:79)
    at org.jboss.netty.channel.socket.nio.AbstractNioSelector.run(AbstractNioSelector.java:312)
    at org.jboss.netty.channel.socket.nio.NioClientBoss.run(NioClientBoss.java:42)
    at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1142)
    at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:617)
    at java.lang.Thread.run(Thread.java:748)
```

#10 - Jun 23, 2017 03:24 PM - Andrea Dell'Amico

The problem is not on our side:

```
# iptables -nvL | grep 194.149.136.225
0      0 ACCEPT    tcp  --  *      *      194.149.136.225      0.0.0.0/0      state NEW tcp dpt:9200
```

#11 - Jun 23, 2017 03:29 PM - Massimiliano Assante

Boro Jakimovski wrote:

Do we need to setup the same for the staging.datasciencepro.eu (at least the db access)

I understand staging.datasciencepro.eu will be your preproduction instance. I think this should still point to the devEDISON VRE on D4Science PreProduction Infrastructure so the same nodes we authorised for dev.datasciencepro.eu

and

www.datasciencepro.eu (194.149.136.227) for production site.

In this case, this portal should point to the EDISON Virtual Organisation on D4Science Production Infrastructure

Does it means that we will have to connect to another DB instance, and do we need to setup ldap sync with another ldap instance or the same is enough?

Yes in the latter case you have to connect to totally different nodes, we will take care of opening the ports on our firewall for this node.

About LDAP

The script that we provide for exporting users on D4Science automatically retrieves the LDAP instance based on the infra. I can enable it on dev.datasciencepro.eu (and it will use our LDAP dev) but you **should first perform a clean** of the users in dev.datasciencepro.eu Liferay DB according to the rules we agreed in our last telCo.

Hope this helps,
M.

#12 - Jun 23, 2017 07:02 PM - Andrea Dell'Amico

I just gave staging.datasciencepro.eu the same accesses that dev.datasciencepro.eu has.

#13 - Jul 05, 2017 08:13 PM - Andrea Dell'Amico

- *Status changed from Feedback to Closed*

I'm closing it. Open another one for the production configuration when you'll be ready.

#14 - Jul 06, 2017 12:42 PM - Boro Jakimovski

- *Copied to Task #9203: Enable access for production EDISON Portal instance added*