

Wikiprint Book

Title: ESGF Local Node Deployment Tutorial

Subject: TracMeteo - ESGFNodeTutorial

Version: 47

Date: 07/07/2022 04:02:43 AM

Table of Contents

ESGF Local Node Deployment Tutorial	3
Index	3
0. Prerequisites	3
1. Previous installation clean up	3
2. Installation from scratch	4
Configuration for publishing	7
Publish the test dataset	9
Publish CORDEX datasets	9
Known issues during installation	9
#error "Pycopg requires PostgreSQL client library (libpq) >= 9.1	9
Failed building wheel for Pillow	10
Installing a custom certificate in the ESGF Node	10
References	10


```

Sending stop command to Solr running on port 8984 ... waiting 5 seconds to allow Jetty process 16554 to stop gracefully.
Stopping Globus Services for Data-Node... (GridFTP) stop_globus_services for datanode
globus-gridftp-server: unrecognized service
Stopping Globus Services for Index-Node... (MyProxy server) stop_globus_services for gateway
Stopping myproxy-server:                                [ OK ]
No MyProxy Process Currently Running...
Tomcat (jsvc) process is running...

stop tomcat: /usr/local/tomcat/bin/jsvc -pidfile /var/run/tomcat-jsvc.pid -stop org.apache.catalina.startup.Bootstrap
(please wait)
postmaster (pid 16024) is running...
Stopping postgresql service:                            [ OK ]
Stopping httpd:                                         [ OK ]
Running shutdown hooks...

-----
Running Node Services...
node type: [ data index idp compute ] (60)
-----
-----

```

Execute source /usr/local/bin/esg-purge.sh && esg-purge all

2. Installation from scratch

Change directory to /usr/local/bin/

```
[root@spock ~]# cd /usr/local/bin/
```

```

[root@spock bin]# wget -O esg-bootstrap http://distrib-coffee.ipsl.jussieu.fr/pub/esgf/dist/devel/esgf-installer/2.5/esg-b
[root@spock bin]# chmod 555 ./esg-bootstrap
[root@spock bin]# ./esg-bootstrap

```

Your directory should look like this:

```

[root@spock bin]# ls
esg-bootstrap  esg-functions  esg-init  esg-node  esg-purge.sh  jar_security_scan  setup-autoinstall

```

Check your node's version:

```

[root@spock bin]# ./esg-node --version
[VERIFIED]
Version: v2.5.9-master-release
Release: Midgard
Earth Systems Grid Federation (http://esgf.llnl.gov)
ESGF Node Installation Script

```

Set node's type:

```

[root@localhost bin]# ./esg-node --set-type data index
node type set to: [ data index ] (12)

```

Install the node:

```

[root@localhost bin]# ./esg-node --install
Please select the ESGF distribution mirror for this installation (fastest to slowest):
-----
[1] http://dist.ceda.ac.uk/esgf

```

```
[2] http://distrib-coffee.ipsl.jussieu.fr/pub/esgf
[3] http://esg-dn2.nsc.liu.se/esgf
[4] http://aims1.llnl.gov/esgf
```

```
-----
select [1] >
```

```
What is the fully qualified domain name of this node? [localhost.localdomain]: spock.meteo.unican.es
What is the admin password to use for this installation? (alpha-numeric only) []:
Please re-enter password:
What is the name of your organization? [localhost]: Unican
Please give this node a "short" name: []: Unican
Please give this node a more descriptive "long" name []: Unican
What is the namespace to use for this node? (set to your reverse fqdn - Ex: "gov.llnl") []: es.unican.meteo.spock
What peer group(s) will this node participate in? (esgf-test|esgf-prod) [esgf-test]:
What is the default peer to this node? [spock.meteo.unican.es]: VESGINT-IDX.IPSL.UPMC.FR
What is the hostname of the node do you plan to publish to? [VESGINT-IDX.IPSL.UPMC.FR]:
What email address should notifications be sent as? []: YOUR-EMAIL
Is the database external to this node? [y/N]:
What is the database connection string? [postgresql://dbsuper@localhost:5432/esgctet]: postgresql://
What is the (low priv) db account for publisher? [esgctet]:
What is the db password for publisher user (esgctet)? []:
```

```
Starting Postgress...
```

```
Starting postgresql service:
```

```
[ OK ]
```

```
0 S postgres 5614 1 6 80 0 - 53982 poll_s 17:52 ? 00:00:00 /usr/bin/postmaster -p 5432 -D /var/lib/pgsql/d
1 S postgres 5631 5614 0 80 0 - 44735 poll_s 17:52 ? 00:00:00 postgres: logger process
1 S postgres 5633 5614 0 80 0 - 53982 poll_s 17:52 ? 00:00:00 postgres: writer process
1 S postgres 5634 5614 0 80 0 - 53982 poll_s 17:52 ? 00:00:00 postgres: wal writer process
1 S postgres 5635 5614 0 80 0 - 54015 poll_s 17:52 ? 00:00:00 postgres: autovacuum launcher process
1 S postgres 5636 5614 0 80 0 - 44734 poll_s 17:52 ? 00:00:00 postgres: stats collector process
```

```
Enter password for postgres user dbsuper:
```

```
Re-enter password for postgres user dbsuper:
```

```
Please Enter PostgreSQL port number [5432]:>
```

```
Would you like a "system" or "user" publisher configuration:
```

```
-----
*[1] : System
```

```
[2] : User
```

```
-----
[C] : (Custom)
-----
```

```
select [1] >
```

```
You have selected: 1
```

```
Publisher configuration file -> [/esg/config/esgctet/esg.ini]
```

```
Is this correct? [Y/n]
```

```
Your publisher configuration file will be: /esg/config/esgctet/esg.ini
```

```
What is your organization's id? [Unican]:
```

```
Would you like to configure this node for CMIP6 publishing (additional project dependencies will be installed)? [y/N]
```

```
[VERIFIED]
```

```
Looking for keystore [/esg/config/tomcat/keystore-tomcat]... (don't see one)...
```

```
Keystore setup:
```

```
Launching Java's keytool:
```

```
store_password = *****
```

Would you like to use the DN: (OU=ESGF.ORG, O=ESGF) ? [Y/n]:

Please enter the password for this keystore :

Please re-enter the password for this keystore:

Enter a single ip address which would be cleared to access admin restricted pages.

You will be prompted if you want to enter more ip-addresses

Do you wish to allow further ips? y/n

n

[VERIFIED]

Create user credentials

Please enter username for tomcat [dnode_user]:

Please enter password for user, "dnode_user" [*****]: 653e78d101f9105fd65755249edf849411e70657814acf06667d910709f3ea

Would you like to add another user? [y/N]:

Please Enter the public (i.e. routable) IP address of this host [10.0.2.15]:> YOUR-PUBLIC-ROUTABLE-IP

Using IP: 10.0.2.15

Do you wish to use an external IDP peer?(N/y):y

Please specify your IDP peer node's FQDN:VESGINT-IDX.IPSL.UPMC.FR\

Server sent 2 certificate(s):

```
1 Subject CN=vesgint-idx.ipsl.upmc.fr, OU=ESGF.ORG, O=ESGF
  Issuer CN=IPSL Simple CA, OU=simpleca.ipsl.upmc.fr, OU=ESGF.ORG, O=ESGF
  sha1   8f 1c 62 12 3a 3f 88 be 12 26 c8 f8 f9 3b da 73 a7 f6 f2 04
  md5    1f fb bd cd 25 cd f5 9d 39 42 d4 c3 ef 2d 98 20

2 Subject CN=IPSL Simple CA, OU=simpleca.ipsl.upmc.fr, OU=ESGF.ORG, O=ESGF
  Issuer CN=IPSL Simple CA, OU=simpleca.ipsl.upmc.fr, OU=ESGF.ORG, O=ESGF
  sha1   15 1d a1 c3 b9 0d 9a 62 3f 99 24 9e 0d 53 6a 23 3b cd c2 19
  md5    cc 08 18 d6 c6 31 1b 91 f7 51 78 04 a5 18 14 50
```

Enter certificate to add to trusted keystore or 'q' to quit: [1] >

Automatic peer with super-node (if you administer one, ensure that it is running) [Y/n]

Do you want to continue with the Globus installation and setup? [Y/n]

Do you want to register the GridFTP server with Globus? [Y/n]:

Please provide a Globus username []: YOUR-GLOBUS-USER

Globus password []: YOUR-GLOBUS-PASSWORD

Server sent 2 certificate(s):

```
1 Subject CN=vesgint-idx.ipsl.upmc.fr, OU=ESGF.ORG, O=ESGF
  Issuer CN=IPSL Simple CA, OU=simpleca.ipsl.upmc.fr, OU=ESGF.ORG, O=ESGF
  sha1   8f 1c 62 12 3a 3f 88 be 12 26 c8 f8 f9 3b da 73 a7 f6 f2 04
  md5    1f fb bd cd 25 cd f5 9d 39 42 d4 c3 ef 2d 98 20

2 Subject CN=IPSL Simple CA, OU=simpleca.ipsl.upmc.fr, OU=ESGF.ORG, O=ESGF
  Issuer CN=IPSL Simple CA, OU=simpleca.ipsl.upmc.fr, OU=ESGF.ORG, O=ESGF
  sha1   15 1d a1 c3 b9 0d 9a 62 3f 99 24 9e 0d 53 6a 23 3b cd c2 19
  md5    cc 08 18 d6 c6 31 1b 91 f7 51 78 04 a5 18 14 50
```

Enter certificate to add to trusted keystore or 'q' to quit: [1] >

```

Finished!...
In order to see if this node has been installed properly you may direct your browser to:
http://spock.meteo.unican.es/thredds
http://spock.meteo.unican.es/esg-orp
http://spock.meteo.unican.es/

Your peer group membership -- : [esgf-test]
Your specified "default" peer : [VESGINT-IDX.IPSL.UPMC.FR]
Your specified "index" peer - : [VESGINT-IDX.IPSL.UPMC.FR] (url = http://VESGINT-IDX.IPSL.UPMC.FR/)
Your specified "idp" peer --- : [VESGINT-IDX.IPSL.UPMC.FR] (name = VESGINT-IDX.IPSL.UPMC.FR)
Your temporary certificates have been placed in /etc/tempcerts
You can install them by executing this : esg-node --install-keypair /etc/tempcerts/hostcert.pem /etc/tempcerts/hostkey.pem
When prompted for the chainfile, specify: /etc/tempcerts/cacert.pem

[Note: Use UNIX group permissions on /esg/content/thredds/esgset to enable users to be able to publish thredds catalogs fr
%> chgrp -R <appropriate unix group for publishing users> /esg/content/thredds

-----
Administrators of this node should subscribe to the
esgf-node-admins@lists.llnl.gov by sending email to: majordomo@lists.llnl.gov
with the body: subscribe esgf-node-admins
-----

v2.5.9-master-release

```

Execute the following:

```

[root@spock bin]# ./esg-node --install-keypair /etc/tempcerts/hostcert.pem /etc/tempcerts/hostkey.pem
...
Please set the password for this keystore :
Please re-enter the password for this keystore:
...
certfile> /etc/tempcerts/cacert.pem
certfile>
...
Is the above information correct? [Y/n]
Is the above information correct? [Y/n]

```

Restart the node:

```

[root@spock bin]# ./esg-node restart

```

Check that everything works ([?https://github.com/ESGF/esgf-installer/wiki/ESGF-Post-Installation-Tests](https://github.com/ESGF/esgf-installer/wiki/ESGF-Post-Installation-Tests)).

If the CoG portal does not work follow the instructions on [?https://www.earthsystemco2.org/projects/cog/install_or_upgrade](https://www.earthsystemco2.org/projects/cog/install_or_upgrade).

Now you should be able to log in the CoG portal using the openid "[?https://spock.meteo.unican.es/esgf-idp/openid/rootAdmin](https://spock.meteo.unican.es/esgf-idp/openid/rootAdmin)" and the password that you chose in the installation process. If you have not installed the idp profile, you can administer your CoG portal from "[?https://YOUR-DOMAIN/login2](https://YOUR-DOMAIN/login2)".

Configuration for publishing

The installation process should have created a user in the postgres database, named rootAdmin. You can check it by running `psql -U dbsuper -d esgset` (to access the postgres cli) and visualizing the table `esgf_security.user`.

```

esgset=# select * from esgf_security.user;
id | firstname | middlename | lastname | email | username | password | dn |
ion_token | notification_code
-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | Admin | | User | emailOfTheAdmin | rootAdmin | hashOfThePassword | | https://domain/esgf-idp

```



```

File "/usr/local/uvcdat/2.2.0/lib/python2.7/site-packages/setuptools-1.4-py2.7.egg/setuptools/dist.py", line 263, in fetch_metadata
File "/usr/local/uvcdat/2.2.0/lib/python2.7/site-packages/setuptools-1.4-py2.7.egg/pkg_resources.py", line 568, in resolve_dependencies
File "/usr/local/uvcdat/2.2.0/lib/python2.7/site-packages/setuptools-1.4-py2.7.egg/pkg_resources.py", line 806, in best_match
File "/usr/local/uvcdat/2.2.0/lib/python2.7/site-packages/setuptools-1.4-py2.7.egg/pkg_resources.py", line 818, in obtain_metadata
File "/usr/local/uvcdat/2.2.0/lib/python2.7/site-packages/setuptools-1.4-py2.7.egg/setuptools/dist.py", line 313, in fetch_metadata
File "/usr/local/uvcdat/2.2.0/lib/python2.7/site-packages/setuptools-1.4-py2.7.egg/setuptools/command/easy_install.py", line 154, in run
File "/usr/local/uvcdat/2.2.0/lib/python2.7/site-packages/setuptools-1.4-py2.7.egg/setuptools/command/easy_install.py", line 154, in run
File "/usr/local/uvcdat/2.2.0/lib/python2.7/site-packages/setuptools-1.4-py2.7.egg/setuptools/command/easy_install.py", line 154, in run
File "/usr/local/uvcdat/2.2.0/lib/python2.7/site-packages/setuptools-1.4-py2.7.egg/setuptools/command/easy_install.py", line 154, in run
File "/usr/local/uvcdat/2.2.0/lib/python2.7/site-packages/setuptools-1.4-py2.7.egg/setuptools/command/easy_install.py", line 154, in run
distutils.errors.DistutilsError: Setup script exited with error: command 'gcc' failed with exit status 1

```

Sorry...

This action did not complete successfully

Please re-run this task until successful before continuing further

Also please review the installation FAQ it may assist you

<https://github.com/ESGF/esgf.github.io/wiki/ESGFNode%7CFAQ>

Failed building wheel for Pillow

This error seems unavoidable but it also seems that it doesn't affect the esgf functionality.

Installing a custom certificate in the ESGF Node

You should own your certificate file (hostcert.crt) and your private key (hostkey.key). Your /etc/httpd/conf/esgf-httpd.conf must reference your certificate and key:

```

228     SSLVerifyClient optional
229     SSLVerifyDepth 10
230     SSLCertificateFile /etc/certs/hostcert.crt
231     #SSLCACertificateFile /etc/certs/esgf-ca-bundle.crt
232     SSLCertificateKeyFile /etc/certs/hostkey.key
233     #SSLCertificateChainFile /etc/certs/cachain.pem
234     SSLOptions +StdEnvVars +ExportCertData

```

Then you have to import your certificate and your key into your tomcat keystore (located in /esg/config/tomcat/ and named esg-truststore.ts and keystore-tomcat). They are configurated in /usr/local/tomcat/conf/server.xml.

1. If the self-signed certificate is installed in keystore-tomcat, remove it with `keytool -delete -alias ALIAS -keystore keystore-tomcat`, where alias can be obtained with `keytool -v -list -keystore keystore-tomcat`.
1. Execute `# openssl pkcs12 -export -in /etc/certs/hostcert.crt -inkey /etc/certs/hostkey.key -out server.p12 -name my-esgf-node -CAfile /etc/certs/hostcert.crt -caname root and keytool -importkeystore -deststorepass PASSWORD -destkeypass PASSWORD -destkeystore keystore-tomcat -srckeystore server.p12 -srcstoretype PKCS12 -srcstorepass PASSWORD -alias my-esgf-node`
1. Ensure it has been correctly installed with `keytool -v -list -keystore keystore-tomcat`.
1. Restart the node: `esg-node restart`
1. More info in [?Stackoverflow](#)

References

- [?ESGF Installation From Scratch](#)