

Wikiprint Book

Title: ESGF Local Node Deployment Tutorial

Subject: TracMeteo - ESGFNodeFullAutoInstall

Version: 8

Date: 08/11/2022 05:16:48 AM

Table of Contents

| | |
|---|----------|
| ESGF Local Node Deployment Tutorial | 3 |
| Index | 3 |
| Prerequisites | 3 |
| 1. Previous installation clean up | 3 |
| Installation from scratch using autoinstaller | 4 |

ESGF Local Node Deployment Tutorial

This page shows how to deploy an ESGF Node that provides data, index, compute and idp services and belongs to the esgf-test federation. The purpose of this node is to test publication and other processes before applying them into production.

UNICAN belongs to the Tier 2 group of nodes so the production node MUST be data only.

This page assumes that command are executed by the root user (or sudo -s).

Index

1. Prerequisites
2. Previous installation clean up
3. Installation from scratch using autoinstaller
4. Configuration for publication
5. Publish the test dataset
6. Publish CORDEX datasets
7. Known issues
8. References

Prerequisites

1. You have to create a globus account - [?https://www.globusid.org/create](https://www.globusid.org/create)
2. Open ports: [Required open ports](#)

```
[root@spock ~]# iptables -L
Chain INPUT (policy DROP)
target     prot opt source                destination
ACCEPT     all  -- anywhere              anywhere
ACCEPT     all  -- anywhere              anywhere           state RELATED,ESTABLISHED
ACCEPT     tcp  -- anywhere              anywhere           tcp dpt:http state NEW
ACCEPT     tcp  -- anywhere              anywhere           tcp dpt:https state NEW
ACCEPT     tcp  -- 192.168.202.0/24      anywhere           tcp dpt:ssh
ACCEPT     tcp  -- 193.144.202.0/24      anywhere           tcp dpt:ssh
ACCEPT     tcp  -- anywhere              anywhere           state NEW tcp dpt:gsiftp
ACCEPT     tcp  -- anywhere              anywhere           state NEW tcp dpt:7512
ACCEPT     tcp  -- anywhere              anywhere           state NEW tcp dpts:50000:51000
ACCEPT     icmp -- 192.168.202.0/24      anywhere           icmp echo-request state NEW,RELATED,ESTABLISHED
```

```
Chain FORWARD (policy DROP)
target     prot opt source                destination
```

```
Chain OUTPUT (policy ACCEPT)
target     prot opt source                destination
```

```
Chain SSHSCAN (0 references)
target     prot opt source                destination
```

1. Previous installation clean up

Execute `esg-node stop` in order to stop the current ESGF services (in case they are running).

```
[root@spock ~]# esg-node stop

EEEEEEEEEEEEEEEEEEEEEEEEEE SSSSSSSSSSSSSSSSS           GGGGGGGGGGGGGGFFFFFFFFFFFFFFFFFFFFFFFFFFFF
E:::::::::::::::::::::E SS::::::::::::::::::::S         GGG:::::::::::::GF:::::::::::::::::::::F
```

```

E::::::::::::::::::ES::::SSSSS::::S  GG::::::::::::GF::::::::::::F
EE::::::::EEEEEEEE::::ES::::S  SSSSSS  G::::GGGGGGG::::GFF:::::FFFFFFF::::F
 E::::E      EEEEEES::::S      G::::G      GGGGGG  F::::F      FFFFFF
 E::::E      S::::S      G::::G      F::::F
 E::::EEEEEEEE  S::::SSSS  G::::G      F:::::FFFFFFF
 E:::::E      SS:::::SSSS  G::::G      GGGGGGGGGG  F:::::FFFFFF
 E:::::E      SSS:::::SS  G::::G      G:::::G  F:::::FFFFFF
 E:::::EEEEEEEE  SSSSS::::S  G::::G      GGGGG::::G  F:::::FFFFFFF
 E::::E      S::::SG:::::G      G::::G  F::::F
 E::::E      EEEEE  S::::S  G::::G      G::::G  F::::F
EE::::EEEEEEEE::::ESSSSSS  S::::S  G::::GGGGGGG::::GFF:::::FF
E:::::ES:::::SSSSS::::S  GG::::::::::::GF:::::::::FF
E:::::ES:::::SS  GGG::::GGG::GF:::::::::FF
EEEEEEEEEEEEEEEEEEEE SSSSSSSSSSSSS  GGGGGG  GGGGFFFFFFFFFFFF.llnl.gov

```

Checking that you have root privs on spock.meteo.unican.es... [OK]

Checking requisites...

Using IP: 193.144.184.40

Stopping Globus Services for Data-Node... (GridFTP) stop_globus_services for datanode

Stopping globus-gridftp-server: OK

Stopping dashboard provider...

Shutting down ESGF Information Provider... [OK]

Shutting down : esgfnode Manager process not present. Attempting to wipe PIDFILE and LOCKFILE.

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)

```

```

5 Z tomcat      7111      1 0 80  0 -      0 do_exe  2017 ?          00:24:39 [jsvc] <defunct>

```

postmaster (pid 30159) is running...

Stopping postgresql service: [OK]

Stopping httpd: [OK]

Running shutdown hooks...

```

-----
Running Node Services...

```

```

node type: [ data ] (4)
-----

```

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

Installation from scratch using autoinstaller

The autoinstaller allows the installation of data only and data+idp+compute+index nodes. Any other combinations should be installed manually. See <https://github.com/ESGF/esgf-installer/wiki/ESGF-Installation-Using-Autoinstaller>.

As root (not sudo!):

```

cd /usr/local/bin
wget -O esg-bootstrap http://distrib-coffee.ipsl.jussieu.fr/pub/esgf/dist/esgf-installer/2.5/esg-bootstrap --no-check-cert
chmod 555 esg-bootstrap
./esg-bootstrap

```

Edit /usr/local/etc/esg-autoinstall.conf. It should look like:

```

#!/usr/bin/expect -f
# -*- mode:shell-script -*-
#
# Configuration file for the esg-autoinstall Expect script.

```

```

#
# If you are opening this file as 'esg-autoinstall.template', either
# from /usr/local/etc or from a Git repository, DO NOT MODIFY IT
# DIRECTLY. It must be copied to /usr/local/etc/esg-autoinstall.conf
# before it will be read.
#
# THIS FILE CONTAINS PASSWORDS -- ensure that when you copy it you set
# the permissions to 600 and the ownership to root, e.g.
#   chown root /usr/local/etc/esg-autoinstall.conf
#   chmod 600 /usr/local/etc/esg-autoinstall.conf
#
### Local Node Settings ###

# Set this to "data" or "data compute" for a data-only node
# Set this to "index" for an index-only node
set NODETYPE "index idp data compute"

# Optional install flags
# Possibilities include "--devel" or "--debug"
set INSTALLFLAGS ""

# Fully qualified domain name of the node
set FQDN "spock.meteo.unican.es"

# Shortname of the node
set SHORTNAME "unican test"

# Long descriptive name of the node
set LONGNAME "unican test"

# Admin password (alphanumeric-only)
set ADMINPASS "PASSWORD"

# IP address cleared for admin-restricted pages
set ADMINIP "0.0.0.0"

# Password for low-privilege publisher database account (alphanumeric-only)
set DBLOWPASS "PASSWORD"

# Globus Online username and password
# (you must have set up this account in advance)
# If your password contains URL special characters, installation may fail.
set GLOBUSUSER "YOUR-GLOBUS-USER"
set GLOBUSPASS "YOUR-GLOBUS-PASSWORD"

# Organization name
set ORGNAME "unican"

# Namespace (reverse FQDN, i.e. "gov.llnl")
set NAMESPACE "es.unican.meteo.spock"

# Peer Group (valid values are usually "esgf-test" and "esgf-prod")
set PEERGROUP "esgf-test"

# Default Peer (set this to empty for the default peergroup member, or to your own FQDN)
set DEFAULTPEER "spock.meteo.unican.es"

# Hostname of the node to publish to
set PUBLISHNODE "spock.meteo.unican.es"

```

```

# Use an external IDP Peer?
# (Unless you know you need to run your own, put y)
set EXTERNALIDP "n"

# FQDN of the external IDP Peer
set IDPPEER "spock.meteo.unican.es"

# E-mail address notifications will be sent as
set ADMINEMAIL "ezequiel.cimadevilla@unican.es"

# Default answer for installing Globus
# Should be be N for upgrade and Y for fresh install
set GLOBUS "y"

# Default answer for automatic peering
# Should be be N for upgrade and Y for fresh install
set AUTOMATICPEER "y"

set THREDDS "N"

### Special-use variables (you probably don't want to change these) ###

# Which install script to run
set ESGNODESCRIPT /usr/local/bin/esg-node

# Disable timeouts (the installer can take a long time)
set timeout -1

# Database connection string
# (form: [username]@[host]:[port]/[database])
# Note: all elements are mandatory, and the "postgresql://" header is
# automatically prepended by the esg-node script!
# Currently, the only allowed database name is "esgcet"
## DO NOT CURRENTLY SET THIS TO ANYTHING BUT ""
#set DBSTRING "esg@localhost:5432/esgcet"
set DBSTRING ""

# Is this an externally managed database?
# (if DBSTRING isn't pointed at localhost, the answer is yes)
# DO NOT CURRENTLY SET THIS TO YES
set DBEXTERNAL no

# Low-privilege database account (this may need to be "esgcet")
set DBLOWUSER "esgcet"

# PostgreSQL port number
# DO NOT CURRENTLY CHANGE THIS
set PGPORT 5432

# Force recreation of SOLR replica indexes
set SOLRREINDEX ""

```