Living Document: Version 0.1, Last Update: 2016/05/23

Earth System Grid Federation Policies & Guidelines

ESGF Executive Committee

Dean N. Williams (Chair), Michael Lautenschlager (Co-Chair), Luca Cinquini, Sebastien Denvil, Martin Juckes, Robert Ferraro, Daniel Duffy, Cecilia DeLuca, V. Balaji, Ben Evans, and Claire Trenham

ESGF Working Team Leads

Sasha Ames, Rachana Ananthakrishnan, Katharina Berger, Nicolas Carenton, Luca Cinquini, Eli Dart, Cecelia DeLuca, Charles Doutriaux, Daniel Duffy, Prashanth Dwarakanath, Sandro Fiore, Sam Fries, Stephan Kindermann, Matthew Harris, Mary Hester, Lukasz Lacinski, Philip Kershaw, Paola Nassisi, Bibi Raju, Torsten Rathmann, George Rumney, Martina Stockhause, and Tobias Weigel

THE EARTH SYSTEM GRID FEDERATION REPRESENTS A MULTINATIONAL EFFORT TO SECURELY ACCESS, MONITOR, CATALOG, TRANSPORT, AND DISTRIBUTE PETABYTES OF DATA FOR CLIMATE CHANGE RESEARCH EXPERIMENTS AND OBSERVATIONS.

ABSTRACT

The Earth System Grid Federation (ESGF) is composed of groups and institutions that have elected to work together to operate a global infrastructure in support of climate science research. Although anybody is always welcome to download, install and run the ESGF software stack as a standalone node, joining the global federation requires understanding and abiding by the established ESGF policies, including:

- Overall governance (http://esgf.llnl.gov/governance.html),
- Security (http://esgf.llnl.gov/media/pdf/ESGF-Software-Security-Plan-V1.0.pdf), and
- Operations (http://esgf.llnl.gov/media/pdf/ESGF-Implementation-Plan-V1.0.pdf).

Ultimately, these policies are in place to provide the best possible experience to the community (in terms of reliability, performance, scalability, efficiency), ensure security and stability, and facilitate the job of the staff administering the ESGF nodes. The ESGF-XC and the ESGF Working Group leads will work with all Nodes in the federation to guarantee that the policies listed in this document are properly understood, monitored, and enforced. If a Node fails to comply, after reasonable communication and consensus efforts have been established, it will be excluded from the federation by all other Nodes.

This document lists the current operational policies and will be published at http://esgf.llnl.gov/media/pdf/ESGF-Polices-Guidelines-V1.0.pdf. The document will be updated as new policies are adopted. For any questions pertaining to this document or the aforementioned policy documents, please email the ESGF Executive Committee (ESGF-XC) @ esgf-xc@llnl.gov. Stability

TABLE OF CONTENTS

1. Setting up an Index Node

1.1 Motivation

An ESGF Index node runs the collection of services and applications necessary to publish and search metadata. At this time, the Index node software stack includes a set of Solr search engines running within Jetty containers, the underlying metadata indexes, the ESGF publishing and search web services running within Tomcat and front-ended by an Apache *httpd* server, and the CoG user interface (see https://www.earthsystemcog.org/ for a description of CoG).

When a client application executes a search request against an ESGF Index node, the client expects to obtain results that cover the whole federation, not just that specific Index node. This is accomplished in one of two possible ways, depending on how the Index node is configured:

1. *Remote shards search:* The targeted index node distributes the search request to all other remote Index nodes in the federation, and then assembles the results before returning them to the client. In this scenario, the total request time is largely dominated by the time it takes the query the slowest remote Index.
2. *Local replica shards search:* The targeted Index node distributes the search request to *local* replicas (aka "copies") of the remote Index nodes, and then returns the results to the client. This configuration is typically much more efficient than the previous one, but it requires that the local Index node runs an additional Solr instance for each remote Index node in the federation.

Clearly, in order to maintain acceptable search performance within the federation, and/or to control the amount of computing resources needed to replicate other Indexes, the number of Index nodes within ESGF *cannot be allowed to grow without control*. Therefore, groups and institutions setting up a new ESGF node and wishing to join the federation are asked to avoid running their own Index node, and to leverage instead the services provided by one (or more) of the existing nodes: publishing their data to that node, and setting up a project on that CoG site to search their data. It is also possible to setup a local CoG instance without the rest of the Index node infrastructure, and configure it to search another Index node (see https://www.earthsystemcog.org/projects/cog/installation_install_upgrade for instructions to install CoG standalone).

In general, setting up a new Index node might make sense if at least one of the following criteria is met:

1. The institution is going to publish a significant amount of data as part of a large international project (such as the Sixth Phased of the Coupled Model Inter-comparison Project [CMIP6])
2. The institution is going to act as a major hub for offering ESGF services to a large geographic region (for example, a continent)
3. The institution is going to serve the needs of a major sponsoring agency (e.g. DOE, NASA, NOAA, NERC, etc.)

Setting up a separate Index node is *not* recommended for nodes that intend to publish only small amounts of data, or data collections that are meant to serve only a restricted community.

In order to provide acceptable performance to search clients, ESGF is also strongly recommending that each site operating an Index node be able to allocate enough computing resources to replicate locally all other Index nodes within the federation (i.e. to execute "local replica searches" as opposed to "remote shard searches"). At this time, this involves running approximately a dozen Solr/Jetty instances, each with a RAM of at least 1 GB, and managing an index size of approximately 10 GB. Preferably, the Index node services should be installed on one or more servers that are kept separate from the rest of the ESGF infrastructure (Data, Compute and IdP nodes). On the Index server, ports 80 and 443 must be opened to the world for connection by search clients; all the Jetty/Solr ports (8983, 8984, etc.) should instead either be restricted to *localhost* only, or at most to specific hosts within the institution firewall. The Solr and Django admin interfaces should also be restricted to specific hosts within the Apache httpd configuration file.

A list of current ESGF Index nodes is maintained at: https://github.com/ESGF/esgf.github.io/wiki/ESGF-Index-and-IdP-nodes.

1.2 Policy

The ESGF-XC shall review all requests that a new Index node be included in federation, and either allow it, or require that the new node be associated with one of the existing Index nodes. For the latter and more common case, new comers to ESGF will be appropriately placed with the right existing Index node for efficiency and connectivity within the federation. Institutions that do run an Index node must commit to guaranteeing an uptime of at least 95%, and to keep the ESGF software stack and underlying operating system up to date with the latest releases and security patches.

2. Setting up an Identity Provider

2.1 Motivation

An ESGF Identity Provider (IdP) operates web services for user registration, authentication, and management of access control attributes. At this time, it includes a back-end Postgres database, the ESGF IdP web application (running within Tomcat and front-ended by an Apache *httpd* server), and a MyProxy server for issuing short-lived X.509 certificates.

The IdP is a linchpin in the operations of the ESGF federation for its users. It's the crux of the security trust model and any downtime for an IdP means that its users cannot authenticate within the federation and therefore cannot access secured resources. Installing and maintaining an Identity Provider involves considerable staff time, especially because the node administrator must constantly be aware of the latest known security risks and promptly install any ESGF patch release and appropriate operating system patches as soon as it is available. Additionally, because ESGF enables Single Sign-On (SSO) across all nodes in the federation, whenever a new ESGF IdP comes online, all other ESGF node administrators must perform configuration steps to include the new IdP in the list of trusted providers. Operators need to commit to maintaining the maximum availability and uptime for this service.

For all the above reasons, ESGF is committed to keep the number of Identity Providers to a small set, which includes a few selected sites for each agency (DOE, NASA, NOAA, NERC, etc.) and continent (America, Asia, Australia, Europe, ...). All other sites are asked to leverage the identity services provided by one of the existing IdPs, by configuring their CoG front-end to redirect user registration and authentication to that site.

A list of current ESGF Identity Provider nodes is maintained at:
https://github.com/ESGF/esgf.github.io/wiki/ESGF-Index-and-IdP-nodes.

Those institutions that do operate an Identity Provider should strive to deploy it on a dedicated server, to provide as much security isolation for the node. On this server, ports 80, 443 and 7512 must be opened outside of the institutional firewall for HTTP and MyProxy clients. A dedicated server also enables fast patching and shorter downtimes whenever OS patches are released to the public.

2.2 Policy

New ESGF node installations shall by default use an existing Identity Provider, and avoid installing a new federated IdP. In those rare cases when a new IdP is really needed, ESGF node administrators must first obtain approval from the ESGF-XC, and then coordinate their installation with the ESGF IdEA (Identity, Entitlement and Access Management) working team @ **esgf-idea@llnl.gov**.

3. Publishing New Data Collections

3.1 Motivation

The ESGF search services are backed up by a distributed federated metadata space. That is, metadata published at one ESGF node are searched and/or replicated by all other nodes, and affect the global view returned to clients in two ways:

1. *Global size:* The performance of search requests is clearly affected by the global size of the metadata indexes that need to be searched.
2. *Search facets:* Even more significantly, the number and values of search facets that are exposed to the clients cover all data collections within the federation.

In particular, when human users use an ESGF CoG portal to browse and search for data of interest, it is very important that each search facet contain a reasonable amount of scientifically correct values that the user can select. For example, the "project" facet should not expose "test" projects, or projects with obscure names that only serve a very restricted community; or the "model" facet should not present values that are not really climate models, but rather represent some other characteristic of a data collection. Another unfortunate example would consist in an institution publishing data that is erroneously flagged as CMIP6 or obs4MIPs, when the data was produced outside of the CMIP6 activities, or does not follow the Obs4MIPs conventions. Finally, it has happened that different projects have used the same named facet (for example: "model") with different meaning, therefore generating confusion for users wishing to use that facet to conduct meaningful searches.

Therefore, ESGF must exercise some form of control over which group or institutions publish into the global metadata space, and collections. It is important to note that the ESFG infrastructure always allows publishing metadata into a *local shard*, which is a metadata index that is *not* replicated across the federation. These metadata are still searchable (and the corresponding data downloadable) by any user connecting to the CoG on that specific node—they are simply not shared across nodes.

In general, publishing data collections into the global federation space makes sense only if one of the two following criteria is met:

1. *Global federation across nodes:* The data are stored across several ESGF nodes, which maintain their separate Indexes. In this case, a client searching for the data will want to obtain complete results, which span the whole data collection
2. *Global federation across projects:* The data from that project need to be searched and retrieved together with data from other projects, maintained at other nodes. For example, a typical use case would be users searching for the same physical fields across CMIP6, Obs4MIPs, and Ana4MIPs.

Vice versa, following are examples of data collections that should *not* be published into the replicated Index nodes:

1. Data collections that are only interesting to a very specific community or project, especially if they are very large collections with hundreds of thousands of records. Instead, that collection should be published to the local shard, and a special project setup on the local CoG to search those metadata.
2. A small data collection used to test publication, or specific search features or functionality

3.2 Policy

ESGF node administrators wishing to publish a new data collection into the common metadata space must first obtain approval from the ESGF-XC, then coordinate their activities with the ESGF publishing team @ **esgf-pwt@llnl.gov**. As part of this process, search facet names, types and values will be reviewed for consistency with facets from existing ESGF projects. Moreover, if the publishing node in question is a data-only node, then the administrator performing the publication will coordinate activity specifically with the administrator of a remote index node, who will have to enable the appropriate access controls for the new publication to occur.