

Operational Practice Statement - SAML profile

This is a living document, minor changes can be expected at any time.

- Introduction
- eduGAIN services
 - Core Services
 - Supplementary services
- Operational Team tasks (SAML Profile)
 - Management of core eduGAIN services
 - Management or supervision of supplementary eduGAIN services
- eduGAIN operational model and availability of services
- Operational Team procedures
 - Registration and modification of SAML profile related federation information
 - Security levels
 - Introduction of new requirements for federation metadata feeds
 - Introduction of new best current practices for federation metadata feeds
 - Metadata aggregation related procedures
 - Aggregation, signing and publishing
 - Handling of aggregation alerts
 - System maintenance
 - System updates
 - Aggregator software updates
 - Backups
 - Disaster recovery
 - Technical details
 - eduGAIN database description
 - eduGAIN Metadata Distribution Service (MDS)
 - Organisation and management of services
 - Security considerations
- References

Introduction

This document describes operational procedures implemented to support eduGAIN SAML services. The procedures regarding eduGAIN membership are described in [eduGAIN Operational Practice Statement](#). The Operational Practice Statement is required by the eduGAIN SAML Profile document [eduGAIN-Profile] and in addition to the [Metadata Aggregation Practice Statement](#) must be seen as complementary to eduGAIN SAML Profile.

eduGAIN services

Under the term services listed are utilities as perceived by external users. The internal organisation of services, flow of information and dependencies are not relevant in this view, but are described in sections further down.

Core Services

Name	Access location	Description	Managed by
MDS	http://mds.edugain.org/edugain-v1.xml	eduGAIN Metadata Distribution Service (MDS) is the central component of the eduGAIN service as a whole. For the detailed description and procedures used in the eduGAIN metadata aggregate distributed by MDS see [eduGAIN-meta]. The eduGAIN metadata aggregate is produced on a separate, secured host (mds-feed) and it is copied to the distribution hosts and served from there by the http server. The file is updated hourly.	OT
The technical site	https://technical.edugain.org	The technical site is directed primarily at the federation level technical personnel. It provides information about eduGAIN members, details about their participation. The technical site is also the distribution point of documentation and the home for several core and supplementary services.	OT
Validator	https://technical.edugain.org/validator	The eduGAIN validator is a service designed for validating metadata adherence to standards and eduGAIN requirements. The software has been created primarily as a component for the eduGAIN metadata aggregation and the details of validation rules are given in [eduGAIN-meta]. The same software enriched by a GUI is used as a tool for manual validation of metadata and serves as a support tool for federation operators.	OT
eduGAIN status information	https://technical.edugain.org/status	This status page provides a view of the eduGAIN database in the part relevant to membership information and the current status of metadata aggregation. The page also displays short summary information about numbers of entities in eduGAIN. The interface provides links to scans of the eduGAIN declaration documents signed by federations, direct links to metadata validation, links to contacts, metadata sources etc.	OT
Entities database GUI	http://technical.edugain.org/entities	This service is an interface to the part of the eduGAIN database which stores information about entities themselves. The interface has many filtering mechanisms and also allows for CSV download for further processing in a spreadsheet.	OT

eduGAIN database API	https://technical.edugain.org/api	The API provides access to most of information stored in the database. In particular, the API may be used by the federations to monitor the eduGAIN aggregation process. Other uses are statistics of various sorts or even download of membership maps.	OT
----------------------	---	--	----

Supplementary services

Name	Access location	Description	Managed by
ECCS	https://technical.edugain.org/eccs/	eduGAIN Connectivity Check Service is a monitoring service for IdPs listed in eduGAIN, testing if they are actually ready for eduGAIN, i.e. if they consume eduGAIN metadata	OT
isFederated Check	https://technical.edugain.org/isFederatedCheck/	This tool searches all known academic identity federations for matching organisations and then displays the results.	OT
CoCo monitor	http://monitor.edugain.org/coco/	Monitoring service testing for REFEDS Code of Conduct compliance	SRCE
Technical testing platform	http://technical-test.edugain.org	This host serves as a playground for software development done by the operational team. All extensions are applied, tested and presented at this platform and then transferred to production using the git mechanism	OT
WIKI		The WIKI is maintained as a part of the GEANT WIKI space. The content is provided by many members of the community. WIKI serves as technical documentation, formal documentation (meeting minutes, documentation of operational procedures) and various guides on joining and making most of eduGAIN	GEANT core
Support	support@edugain.org	eduGAIN support mail contact	

Operational Team tasks (SAML Profile)

Management of core eduGAIN services

- management of virtual machines (access management, system maintenance - installation and updates, global backups, status monitoring)
- management of eduGAIN core services (maintenance of any software tools required by the services, monitoring of services, specialised backups)
- supervision of the aggregation function - reacting to aggregation errors, supporting federations in location of problems
- technical documentation - maintenance of user documentation of eduGAIN services
- user support - done in cooperation with the eduGAIN support team
- management of the development platform (based on the GEANT git)
- service development - configuration changes and extensions of existing services, in particular any development work within the eduGAIN MDS, validator, database

Management or supervision of supplementary eduGAIN services

- eduGAIN OT directly manages:
 - ECCS
 - isFederated check
- eduGAIN OT supervises
 - CoCo monitor
 - WIKI

eduGAIN operational model and availability of services

eduGAIN core function is the metadata exchange point. Federations supply their own metadata and download aggregated metadata to supplement their own and redistribute them within their federation members. Federations are strongly discouraged from pointing any of their members directly to the eduGAIN MDS. The operational baseline of MDS availability is set at 99% for any given month. The unavailability details are provided at [eduGAIN Services Status](#) and system changes are listed at https://technical.edugain.org/system_updates.

While every care is taken that all eduGAIN services function reliably, the selected operational model allows that services updates and modifications can be done at a short-term notice allowing for a small risk of a downtime required to restore the system snapshot.

Operational Team procedures

Registration and modification of SAML profile related federation information

information type	security level
federation SAML policy URL	1

registration practice statement URL	1
federation SAML metadata aggregate access URL	3
federation metadata signing key	4
registrationAuthority attribute value	3

Security levels

security level	description
S	special - delegating representatives requires contact with the federation management
1	informational, not requiring special vetting
2	important contact information (while not currently used it may be introduced in the future)
3	information of eduGAIN operational relevance, requires special care
4	crucial for eduGAIN trust, requires utmost care

Introduction of new requirements for federation metadata feeds

Metadata requirements have direct operational impact. As described in the [eduGAIN-meta] document, a violation of the requirements results in rejection of a federation feed. The requirements as understood in this documents are the rejection rules implemented in the eduGAIN metadata validator and automatically applied in the aggregation process.

As a principle requirements for federation feeds must be based on either general standards to which eduGAIN SAML profile adheres or on the eduGAIN SAML profile. In the case of standards, the experience shows that certain violations are only discovered when reported by participating federations - not all such violations are reported by standard schema validation tools, and in fact not all are just schema errors. Whenever a new problem is reported, the OT makes an assessment whether it in fact violates a required standard and if so then:

- the OT implements a new validator rule initially as a warning;
- the OT informs the SG about adding a new validation rule together with an assessment of which federations may be affected by it and suggests a grace period, after which the new rule will start generating an aggregation error;
- the SG members will be given the opportunity to request a longer time-frame, and eduGAIN Support will work with any participants that are currently breaching this requirement to fix the issues before the grace period ends.

Every rule is documented in the [eduGAIN-meta] .

When raising an error, the validator points to the specific rule in [eduGAIN-meta].

Arising problems which cause actual interoperability issues need to be handled immediately, as described in the metadata aggregation related procedures section below.

Introduction of new best current practices for federation metadata feeds

Additions to metadata best current practices need to be decided by the eduGAIN SG. Each such good practice needs to be implemented as an eduGAIN validator warning by the eduGAIN OT. Each good practice rule needs to be added to [eduGAIN-BCP].

Metadata aggregation related procedures

The technical details of the aggregation process are described in [eduGAIN-meta]. Here we only present the operational implementation of this process.

Aggregation, signing and publishing

The aggregation, signing and publishing of the eduGAIN metadata aggregate is done on an hourly basis.

All information about the system status, federation metadata channel information, federation public keys etc. is kept in the eduGAIN database and taken from there as required within the aggregation process.

- Half past every hour metadata acquisition is started on mds-feed and is performed in the following steps:
 - mds-feed downloads federation metadata feeds using conditional GET
 - if the conditional GET resulted in a download of a new metadata file, such file is passed through the local validator instance, if validation succeeds the downloaded file is used as an input for aggregator, if it fails, the previous correct feed copy is used instead
 - the newest available validated copy of the federation metadata feed is kept for future use
 - the validated metadata files are passed to a pyFF flow, see also [eduGAIN-meta] [Metadata combination and collision handling](#)
 - pyFF aggregates and then signs the resulting feed; currently the signing is done with key files stored at the mds-feed host
 - the resulting file is analysed, split into entities and used to update the edugain-db
 - the final output is uploaded with sftp to the technical host using a dedicated user account on the the technical host
- At 45 minutes past every hour the new copy of eduGAIN metadata aggregate is copied to the final destination directory and when the copy is completed the mv action is performed in order to substitute the production file in an atomic mode
- Finally the new eduGAIN metadata aggregate file is copied to the history repository and compressed

- At midnight (CET) hourly copies of metadata are deleted from the repository, leaving only a single daily file. These daily files can then be used as a source of various data analysis.

Handling of aggregation alerts

As described in [eduGAIN-meta], under certain conditions aggregation alerts are raised. The current practice is that these alerts are sent as e-mails to the eduGAIN OT and to the contact address of the Federation that is affected. There is one exception to that rule - metadata unavailability alerts are sent to the OT only to avoid possible false positives caused by temporary network problems. The OT then makes its own decision when to contact the Federation.

It must be realised that the case of all entities supplied by a large federation being deleted from eduGAIN has heavy consequences - other participating federations will naturally have to drop these entities. When the federation metadata feed becomes available again, other federations may be forced into running emergency regeneration of their metadata, service providers may observe limited breaks in their service. Therefore the eduGAIN OT is making all possible effort to avoid such situations. If the eduGAIN OT realises a very special situation it is allowed to temporarily stop aggregation in order to avoid the deletion of the federation but it MUST immediately notify the eduGAIN SG that such measures have been taken.

One example of such a special situation may be a real case of an introduction of a valid but a very short-lived metadata file followed by a metadata error causing a aggregation reject. That situation, leaving no time for normal procedures to take place, was caused by a configuration error on the federation side and was rectified in a short time while the eduGAIN aggregation was suspended.

System maintenance

System updates

- All virtual machines running eduGAIN services are regularly updated.
- Before an update is planned, the local personnel at PSNC are notified in the case of an update failure and immediate restore. An update forward notice is sent to the eduGAIN SG.
- In the case of large configuration changes, like moving services to new hosts, applying large infrastructure changes etc., a notice at least 7 days in advance is sent to the eduGAIN SG.
- All changes are documented in the log available for inspection at: https://technical.edugain.org/system_updates.
- The unavailability details are provided at [eduGAIN Services Status](#).

Aggregator software updates

Updates to crucial aggregator elements, in particular pyFF, may result in a changed format of resulting metadata aggregate. Any such change will be announced to the eduGAIN SG mailing list. If the OT observes that the update indeed introduces changes to metadata, a beta feed will be created and announced to the SG and a change on the production will be delayed by a two-week testing period. A reminder will be issued a week before the actual change of the production feed.

Backups

- system backups are performed daily as a part of the standard PSNC backup routine
- virtual machine snapshots are performed prior to system updates
- four times a year a full virtual machine dump is performed

Disaster recovery

In the case of an unexpected problem resulting from metadata aggregation (which may result from an unusual error on a federation side or some software bug in one of the aggregation process steps) the eduGAIN OT has access to hourly copies reaching 24 hours back and to several years of daily copies .

Restoration of snapshots or full virtual machines is possible (and has been performed several times not as disaster recovery but in order to get access to some old files for statistics reasons).

Technical details

eduGAIN database description

The eduGAIN database is central to all eduGAIN core services. The database stores:

- general and contact information about participating and candidate federations,
- operational information about participating and candidate federations like metadata URLs, signing keys, registrationAuthority values,
- operational information about the metadata aggregation process including details about metadata acquiring from participating federations, results of metadata validation, cache timers for individual participant federations,
- operational information about entities published through eduGAIN derived from the metadata,
- statistics derived from metadata aggregation, like numbers of entities published by individual federations and much more,
- information collected from supporting monitoring services like ECCS, CoCo.

The database is placed on a host separated from the external network, accessible only through a limited number of secure hosts. Database access is realised via dedicated user accounts with access rights crafted to minimize the possibility of unauthorized changes.

The database is managed mostly via a web interface secured with access passwords. Modification of data on security levels S, 1, 2 can be done without any additional protection. Management of data with security level 3 is protected with on-time passwords mailed to an external mail account of the managing administrator. Management of data with security level 4 requires direct access to the database host.

For security reasons signing keys can be present only for federations which have been approved to be a member of the eduGAIN SAML Profile.

eduGAIN Metadata Distribution Service (MDS)

eduGAIN Metadata Distribution Service (MDS) is the central component of the eduGAIN service as a whole. For the detailed description and procedures used in the eduGAIN metadata aggregate distributed by MDS see [eduGAIN-meta] and the **Aggregation, signing and publishing** subsection within this document. The eduGAIN metadata aggregate is produced on a separate, secured host (mds-feed). Metadata signing is also performed on mds-feed currently with a key file located on the host itself.

In order to minimise risks of exposing a high permissions account on the mds host the resulting aggregate file is transferred from mds-feed to the mds host using a dedicated low permissions account. The aggregate is then moved to the final place on the mds host in a process initiated within the mds host.

Organisation and management of services

Main access host - technical, validator, mds	
DNS names	www.edugain.org , technical.edugain.org ; validator.edugain.org ; mds.edugain.org All these are CNAMEs for massonia.man.poznan.pl
Function	<ul style="list-style-type: none">• serves the eduGAIN aggregate file (updated hourly)• serves the information pages at https://technical.edugain.org which includes the status pages, the eduGAIN database WEB GUI and WEB API interfaces, formal documentation• provides the validator service at https://validator.edugain.org• provides the ECCS WEB interface at https://technical.edugain.org/eccs/• provided the isFederatech Check interface at https://technical.edugain.org/isFederatedCheck/
eduGAIN database - edugain-db	
Function	store all data for services directly managed by the eduGAIN OT
The aggregation host - mds-feed	
Function	acquire and validate federation metadata feeds, create, sign and publish the eduGAIN metadata aggregate.

Security considerations

The security of the eduGAIN SAML services is essentially the security of the eduGAIN aggregate. This in turn depends on:

1. validation of federation metadata input data - their originality and integrity - this depends on the safety of federation certificates (stored in the database) and the safety of the signature verification process itself
2. aggregation process - it is crucial that the resulting aggregate contains exactly the data provided by participating federations (after modifications described in the [eduGAIN-meta])
3. aggregation signature - the eduGAIN signing key and the signing process are the key factors here.

In order to maintain the high security standards the following operational procedures are in place,

1. The edugain-db and mds-feed hosts are located in a secured private network and can be accessed only from a single host in the PSNC network.
2. The access to this single host is only available over SSH and only from a limited list of IP addresses.
3. The federation signing keys can only be stored in the edugain-db with a process that needs to be run directly on the db host. This requires that the OT copy the key to the database host and run the process manually. The key is added to the database only when the decision to actually admit the federation metadata to eduGAIN has been taken. This is an additional security procedure guarding against a mistake in assigning the level of participation for a federation.
4. The eduGAIN signing key is stored on the mds-feed host, where the whole process of aggregation and signing is run hourly. This host cannot be reached from the external network. The resulting signed aggregate is then moved to the distribution host as described in the Metadata aggregation related procedures section.

References

[eduGAIN-CONST] <https://technical.edugain.org/doc/eduGAIN-Constitution-v3ter-web.pdf>

[eduGAIN-Profile] <https://technical.edugain.org/doc/eduGAIN-saml-profile.pdf>

[eduGAIN-OPS] [eduGAIN Operational Practice Statement](#)

[eduGAIN-BCP] [Best Current Practice](#)

[eduGAIN-meta] [Metadata Aggregation Practice Statement](#)

