

# Collabora and NextCloud SAML federation pilot


## Introduction

Documents are at the heart of most projects: people create and share text, spreadsheets, presentations and images. Federated access management has been widely used to control access to published documents, but its use in collaborative creation has usually been limited to wikis and content management systems. Office documents are usually created offline. This pilot explored the use of two new applications that can be used together to provide federated access to both file management and office document creation and collaborative editing.

## Components

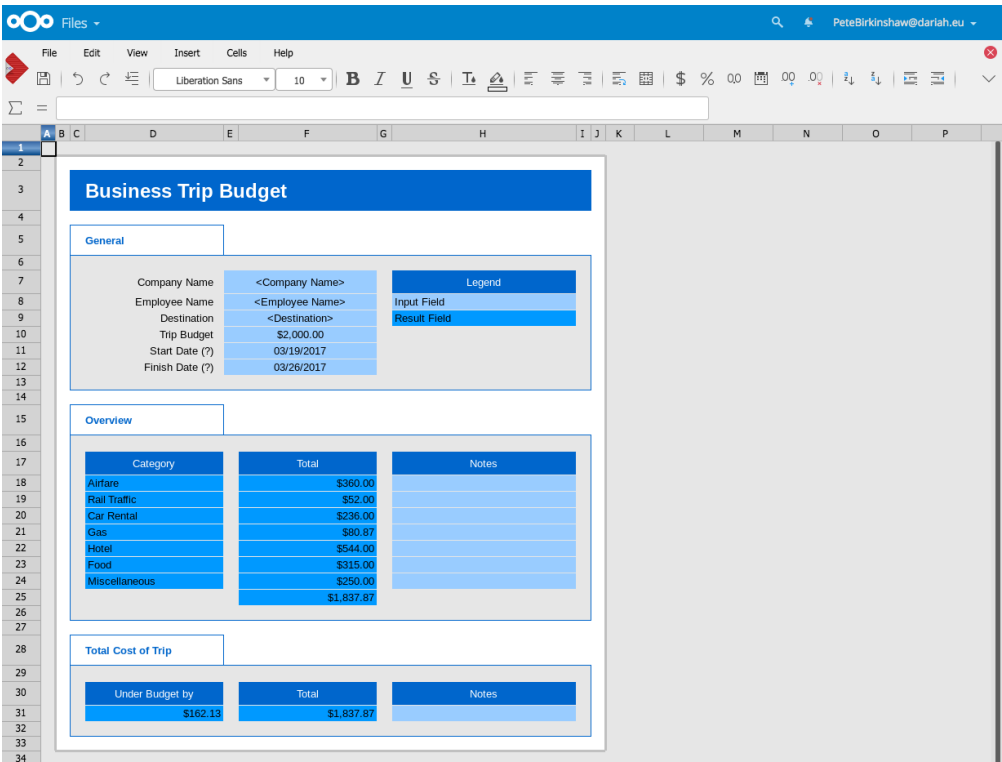
### NextCloud

NextCloud is a web-based document management service. Documents can be uploaded and downloaded via a web interface, or synchronised with local files. Files can be shared with other users. NextCloud was recently forked from OwnCloud. They remain very similar, but NextCloud offers new built-in federated access features using SAML in its free version. We used NextCloud's free edition in this pilot, but the commercially supported edition of OwnCloud may be used in a similar way.



### Collabora Online - LibreOffice for the web

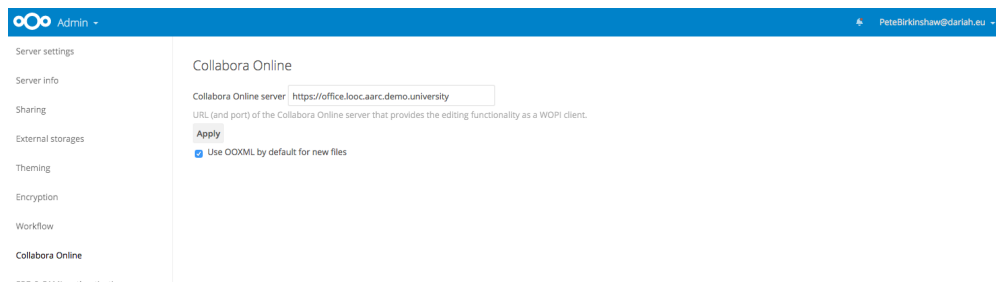
Collabora Online is new software that allows the LibreOffice/OpenOffice office suite to run as a shared web application. Most of the desktop LibreOffice's functionality is available to users in a web page, with the added feature of simultaneous collaborative editing - users can work on the same document at the same time.



## Integrating NextCloud and Collabora Online

Collabora Online provides a basic administration user interface, minimal configuration files, and a generic WOPI API (as used by Microsoft Office). It runs as a web application behind an Apache or NGINX reverse-proxy with SSL. Authentication and authorisation are handled via WOPI and resemble Oauth. As such it requires a second application to act as the primary user interface and authentication provider.

NextCloud is a PHP application and performs the role of , user interface, document store and WOPI authentication server. Collabora Online compatibility is now available via a bundled plugin and the only configuration needed is the URL for the Collabora service. When users open a file in NextCloud control is passed to Collabora Online, and the office application interface appears.



The screenshot shows the NextCloud Admin interface. The top navigation bar is blue with the NextCloud logo and 'Admin' text. On the right, it shows the user 'PeteBirkshaw@dariah.eu'. The left sidebar contains a list of settings categories: Server settings, Server info, Sharing, External storages, Theming, Encryption, Workflow, and Collabora Online. The 'Collabora Online' category is selected, and the main content area displays the configuration for Collabora Online. It includes a text input for 'Collabora Online server' with the value 'https://office.looc.aarc.demo.university', a description of the URL, an 'Apply' button, and a checkbox labeled 'Use OOXML by default for new files' which is checked.

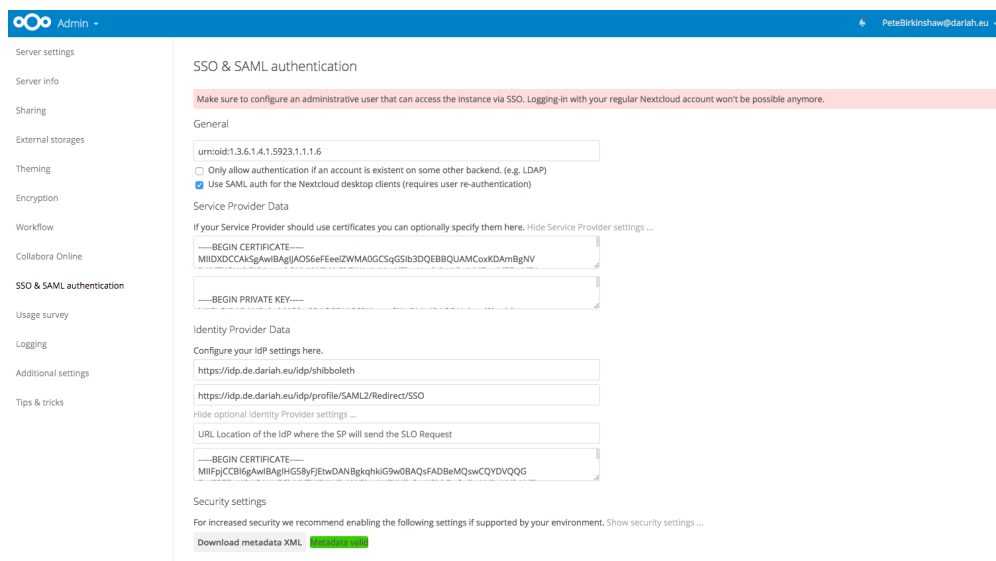
A single Collabora Online service can be used by more than one NextCloud front-end.

## AAI Integration Options

NextCloud needs a means of handling its own authentication. Two options are currently available for federated authentication: a built-in SAML SP, and a external SSO option that relies on the web server handling authentication.

### Built-In SAML

NextCloud's SAML implementation is currently rather limited, and only supports one IdP. This means that a proxy IdP would be needed to provide fully federated access to users.



The screenshot shows the NextCloud Admin interface for 'SSO & SAML authentication'. The top navigation bar is blue with the NextCloud logo and 'Admin' text. On the right, it shows the user 'PeteBirkshaw@dariah.eu'. The left sidebar contains a list of settings categories: Server settings, Server info, Sharing, External storages, Theming, Encryption, Workflow, Collabora Online, SSO & SAML authentication, Usage survey, Logging, Additional settings, and Tips & tricks. The 'SSO & SAML authentication' category is selected, and the main content area displays the configuration for SSO & SAML authentication. It includes a warning message: 'Make sure to configure an administrative user that can access the instance via SSO. Logging-in with your regular Nextcloud account won't be possible anymore.' Below this, there are sections for 'General', 'Service Provider Data', 'Identity Provider Data', and 'Security settings'. The 'General' section has a text input for 'urnoid:1.3.6.1.4.1.5923.1.1.6' and two checkboxes: 'Only allow authentication if an account is existent on some other backend. (e.g. LDAP)' (unchecked) and 'Use SAML auth for the Nextcloud desktop clients (requires user re-authentication)' (checked). The 'Service Provider Data' section has a text input for 'If your Service Provider should use certificates you can optionally specify them here. Hide Service Provider settings ...' and two text areas for 'BEGIN CERTIFICATE' and 'BEGIN PRIVATE KEY'. The 'Identity Provider Data' section has a text input for 'Configure your IdP settings here.' and two text inputs for 'https://idp.de.dariah.eu/idp/shibboleth' and 'https://idp.de.dariah.eu/idp/profile/SAML2/Redirect/SSO'. The 'Security settings' section has a text input for 'URL Location of the IdP where the SP will send the SLO Request' and two text areas for 'BEGIN CERTIFICATE' and 'BEGIN PRIVATE KEY'. There is also a 'Download metadata XML' button.

More advanced SAML features are planned and development seems to be active.

### External authentication

The second SSO option is to rely on the web server (such as Apache or NGINX) and use basic user information passed as environment variables such as REMOTE\_USER.

Admin

PeteBirkshaw@dariah.eu

Server settings

Server info

Sharing

LDAP / AD integration

Theming

Encryption

Workflow

SSO & SAML authentication

SSO & SAML authentication

Make sure to configure an administrative user that can access the instance via SSO. Logging in with your regular Nextcloud account won't be possible anymore.

General

REMOTE\_USER

Only allow authentication if an account is existent on some other backend. (e.g. LDAP)

This means that the Shibboleth SP software, OpenID Connect or even Kerberos can be used, but integration is weak - only the username is available.

## Aggregating attributes from LDAP

It is possible to combine federated authentication with LDAP for additional attributes, and to require presence in the LDAP directory for authentication to succeed.

This would hinder usage across a federation (since users would not exist in the LDAP directory) but may help research organisations, as group membership and access control can be handled by a community LDAP server.

Users

PeteBirkshaw@dariah.eu

+ Add group

Everyone

Admins

Chill\_admins

Chill\_contributors

DARIAH-DE-Cloudshare-Contr...

DARIAH-DE-Sonderheft-Biblioth...

DARIAH-Preservation-users

ERROR

LiKuRe219-II

ReDesign-TextGrid

TextGridDokumentation-adm...

TextGridDokumentation-cont...

TextGridDokumentation-users

admins













no admin role available

Username

Password

Groups

Create

Username	Full name	Password	Email	Groups	Group admin for
 admin	admin	*****		admin	no group
 @dariah.eu		*****	rmstadt.de	jira-developers, dariah-de-co...	no group
 @dariah.eu		*****	ids-mannheim.de	textgrid-contributors, tf-dari...	no group
 @dariah.eu		*****	uni-goettingen.de	jira-developers, dariah-de-co...	no group
 @indid.net		*****		admin	no group
 @dariah.eu		*****	ngit.tu-darmstadt.de	textgrid-contributors, textgr...	no group
 @dariah.eu		*****	uni-wuerzburg.de	jira-developers, dariah-de-co...	no group
 @dariah.eu		*****	aw.de	collection-registry-admins, c...	no group
 @dariah.eu		*****	dirom.de	textgrid-users, li_DARIAH-U...	no group
 @dariah.eu		*****	hab.de	dariah-de-contributors, dhfv...	no group
 @dariah.eu		*****	mpwg-berlin.mpg.de	dariah-de-contributors, dari...	no group
 @dariah.eu		*****	de	dariah-de-contributors, daria...	no group

## Pilot Implementations

Three different demonstrations were set up, so that different features and integration combinations could be explored.

### Demonstration 1: Integrated SAML with one IdP

Built-in SAML

<https://cloud.looc.aarc.demo.university/nextcloud/index.php/>

Admin

PeteBirkshaw@darjah.eu

Server settings

Server info

Sharing

External storages

Theming

Encryption

Workflow

Collabora Online

SSO & SAML authentication

Usage survey

Logging

Additional settings

Tips & tricks

SSO & SAML authentication

Make sure to configure an administrative user that can access the instance via SSO. Logging-in with your regular Nextcloud account won't be possible anymore.

General

urn:oid:1.3.6.1.4.1.5923.1.1.1.6

☐ Only allow authentication if an account is existent on some other backend. (e.g. LDAP)
 ☒ Use SAML auth for the Nextcloud desktop clients (requires user re-authentication)

Service Provider Data

If your Service Provider should use certificates you can optionally specify them here. Hide Service Provider settings ...

-----BEGIN CERTIFICATE-----

MIDXDCCN5gAwIBAgIJAOS6eFEEelZWMAOGCSgG5b3DQEBBQUANCOvKDAmBgNV

-----BEGIN PRIVATE KEY-----

Identity Provider Data

Configure your IdP settings here.

https://idp.de.darjah.eu/idp/shibboleth

https://idp.de.darjah.eu/idp/profile/SAML2/Redirect/SSO

Hide optional Identity Provider settings ...

URL Location of the IdP where the SP will send the SLO Request

-----BEGIN CERTIFICATE-----

MlIfpjCCBtGgAwIBAgIHG5B8yFJEtwDANBgkqhkiG9w0BAQsFAADBeMQswCQYDVQ

Security settings

For increased security we recommend enabling the following settings if supported by your environment. Show security settings ...

Download metadata XML

Configured using NextCloud's built-in SAML to connect to one IdP as if used by a single university or research community using a proxy IdP.

Keys were created using the Shibboleth SP keygen tool and pasted into the configuration form in NextCloud. EPPN was used for NextCloud usernames (identified by URN)

NextCloud generated its own metadata, but the expiry date was only for a few days and so was removed before sharing. Various combinations of encryption and signing can be set.

#### Security settings

For increased security we recommend enabling the following settings if supported by your environment. Hide security settings ...

##### Signatures and encryption offered

- ☐ Indicates that the nameID of the <samlp:logoutRequest> sent by this SP will be encrypted.
- ☒ Indicates whether the <samlp:AuthnRequest> messages sent by this SP will be signed. [Metadata of the SP will offer this info]
- ☐ Indicates whether the <samlp:logoutRequest> messages sent by this SP will be signed.
- ☐ Indicates whether the <samlp:logoutResponse> messages sent by this SP will be signed.
- ☐ Whether the metadata should be signed.

##### Signatures and encryption required

- ☐ Indicates a requirement for the <samlp:Response>, <samlp:LogoutRequest> and <samlp:LogoutResponse> elements received by this SP to be signed.
- ☒ Indicates a requirement for the <saml:Assertion> elements received by this SP to be signed. [Metadata of the SP will offer this info]
- ☒ Indicates a requirement for the <saml:Assertion> elements received by this SP to be encrypted.
- ☐ Indicates a requirement for the NameID element on the SAMLResponse received by this SP to be present.
- ☐ Indicates a requirement for the NameID received by this SP to be encrypted.
- ☐ Indicates if the SP will validate all received XMLs.

##### General

- ☐ ADFS URL-Encodes SAML data as lowercase, and the toolkit by default uses uppercase. Enable for ADFS compatibility on signature verification.

Download metadata XML

## Demonstration 2: External authentication (SAML) plus LDAP

### External SSO

<https://cloud.aarc.federated-example.website/nextcloud/index.php/apps/files/>

## LDAP

Server

Users

Login Attributes

Groups

Advanced

Expert

When logging in, Nextcloud will find the user based on the following attributes:

LDAP / AD Username:

☐

LDAP / AD Email Address:

☐

Other Attributes:

eduPersonPrincipalName

⌵ Edit LDAP Query

LDAP Filter:

(&([objectclass=dariahPerson])([eduPersonPrincipalName=%uid]))

PeteBirkinshaw@dariah.eu

Verify settings

Configuration OK

Back

Continue

Help

Configured to use a single IDP, using the external SSO plugin and a conventional Shibboleth SP. NextCloud was configured to search an LDAP directory for records matching the SAML-authenticated user's `EduPersonPrincipalName`. LDAP was also used to discover which groups a user was a member of. These groups can be used for access control.

Session lifespans for the external authentication service (Shibboleth SP) and Nextcloud's own sessions can become out-of-sync, and require some adjustments to work together consistently.

Internal Username

By default the internal username will be created from the UUID attribute. It makes sure that the username is unique and characters do not need to be converted. The internal username has the restriction that only these characters are allowed: [ a-zA-Z0-9\_@- ]. Other characters are replaced with their ASCII correspondence or simply omitted. On collisions a number will be added/increased. The internal username is used to identify a user internally. It is also the default name for the user home folder. It is also a part of remote URLs, for instance for all \*DAV services. With this setting, the default behavior can be overridden. Leave it empty for default behavior. Changes will have effect only on newly mapped (added) LDAP users.

Internal Username Attribute:

eduPersonPrincipalName

Override UUID detection

## Demonstration 3: Integrated SAML with a federated IdP Proxy

Built-in SAML

<https://files.looc.aarc.federated-example.website>

(Work-in-progress) A similar pilot to the [Demonstration 1](#) but configured to use the [Terena Proxy](#) so that it's easier for a wider range of people to log in.

## An Aside: Federated data storage

NextCloud supports "federated sharing", which permits users to share files between different NextCloud services, and browse user directories on other services. Users are given a globally unique scoped identifier that resembles `EduPersonPrincipalName`. If EPPN is used as the NextCloud username then a user's identifier is scoped twice.

### Federated Cloud

Your Federated Cloud ID: `PeteBirkinshaw@dariah.eu@cloud.aarc.federated-example.website/nextcloud`

The External Storage plugin allows remote data storage to be used, including other NextCloud or OwnCloud services, Windows shares and NFS.

## Caveats

### Speed

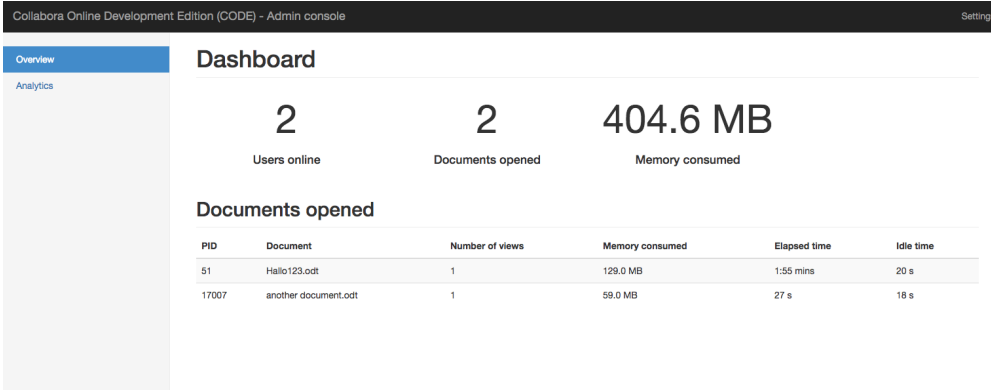
The display of Collabora is generated by sending many tile-like images over the web as individual files, and is rather slow. Over a normal broadband internet connection the display is not quite fast enough to keep up with typing.

Over a much faster connection such as a LAN the speed is greatly improved.

## The Open Source Collabora Online package has restrictions

The CODE (Collabora Online Developer Edition) Docker container used in these pilots is limited to 10 concurrent users. Collabora offer a commercial edition with no limits.

However, there is an unofficial project to help with installing an alternative open source version of the Collabora Online software without these limitations, and without Docker.



## Admin accounts must be created before switching to SSO

Admin users (who are able to configure the service) must be created in Nextcloud, using EPPNs, in advance, before SSO is enabled.

## Cannot easily change SSO methods

NextCloud's SSO plugin offers a choice between the built-in SAML and using external authentication, and it does not seem to be possible to easily switch from one to the other.

## Application passwords are still required

Users will need to create their own passwords in NextCloud to use for syncing files and other non-web access.

## Further Information

- [NextCloud admin manual](#)
- [Nextcloud SAML documentation](#)
- [CODE edition of Collabora](#)
- [LibreOffice Online installer project \(a community alternative to CODE\)](#)
- [NextCloud app passwords, for non-web devices](#)