

KNIME Business Hub Admin Guide

KNIME AG, Zurich, Switzerland
Version 1.9 (last updated on 2024-04-11)



Table of Contents

Introduction	1
Create and manage teams	2
Create a team.	2
Allocate resources to a team	3
Manage team members	5
Delete a team.	5
Execution resources.	7
Create a shared execution context	8
Manage shared execution contexts	9
Advanced configuration of execution contexts.	10
Users management	11
Delete a user	11
Make a user Hub admin	12
Expose external groups inside KNIME Business Hub	16
External OIDC provider	16
LDAP federation	21
Enable external groups.	25
Docker executor images	26
Add extensions to an existing Docker image.	27
Python and Conda in Docker images	28
Customization profiles	31
Create a customization profile	31
Upload a customization profile	35
Apply a customization profile	40
Update a customization profile	46
Detach a customization profile	50
Delete a customization profile.	56
Advanced configuration	60
Configure networking	60
Configure Browser Security	64
Node affinity	66
Create a collection	68
Administrator workflows	71
Workflows overview	71

Requirements and prerequisites	71
Discard Failed Jobs.	72
List All Jobs	75
Delete Old Versions.	78
Scheduled Workflows Kick-Off Times	81
Count Workflows Running Per Day	84
Workflows' Run Time	86
Monitor Users' Usage	88
Monitoring K-AI in KNIME Business Hub.	94
Data access	94
Data structure	94
Deploying the data application	95
Data application overview	96
KNIME Business Hub API documentation	101
Support Bundles and Troubleshooting	102
Generating a support bundle (GUI)	102
Generating a support bundle (CLI)	103
Configuring redaction in support bundles	103
Inspecting support bundles	103
Backup and restore with Velero Snapshots and Kotsadm	107
Creating snapshot backups	107
Backup Troubleshooting	109
Restoring a snapshot	110
Migration of versioning and compatibility	112
General compatibility between KNIME Hub, KNIME Analytics Platform client and executor	112
Migration from space level versioning to item level versioning	112
Changelog (KNIME Business Hub 1.9)	114
KNIME Business Hub 1.9.0	114
Changelog (KNIME Business Hub 1.8)	114
KNIME Business Hub 1.8.3	114
KNIME Business Hub 1.8.2	115
KNIME Business Hub 1.8.1	116
KNIME Business Hub 1.8.0	117
Changelog (KNIME Business Hub 1.7)	119
KNIME Business Hub 1.7.0	119

Changelog (KNIME Business Hub 1.6)	120
KNIME Business Hub 1.6.0	120
Changelog (KNIME Business Hub 1.5)	122
KNIME Business Hub 1.5.2	122
KNIME Business Hub 1.5.1	123
KNIME Business Hub 1.5.0	123
Changelog (KNIME Business Hub 1.4)	125
KNIME Business Hub 1.4.2	125
KNIME Business Hub 1.4.1	126
KNIME Business Hub 1.4.0	126
Changelog (KNIME Business Hub 1.3.0).....	128
KNIME Business Hub 1.3.0	128
Changelog (KNIME Business Hub 1.2.0).....	129
KNIME Business Hub 1.2.0	129

Introduction

KNIME Business Hub is a customer-managed Hub instance. Once you have a license for it and proceed with installation you will have access to Hub resources and will be able to customize specific features, as well as give access to these resources to your employees, organize them into Teams and give them the ability to manage specific resources.

This guide provides information on how to administrate a KNIME Business Hub instance.

To install a KNIME Business Hub please refer to the [KNIME Business Hub Installation Guide](#).

A user guide is also available [here](#), which contains instructions on how to perform team administration tasks. Team admins are designated by the global Hub admin, and have control over their team's allocated resources, can add users to their team, create execution contexts and have an overview of the team's deployments. In this way the Hub administration tasks are distributed and reallocated to those users that have a better overview of their own team necessities.

Create and manage teams

A team is a group of users on KNIME Hub that work together on shared projects. Specific Hub resources can be owned by a team (e.g. spaces and the contained workflows, files, or components) so that the team members will have access to these resources.

Sign in to the KNIME Business Hub instance with the admin user name by visiting the KNIME Business Hub URL.

Then click your profile picture on the right upper corner of the page and select *Administration* to go to the KNIME Business Hub Administration page. Click *Teams* in the menu on the left. Here you will be able to see an overview of the existing teams and you will be able to manage them.

Create a team

To create a new team click the yellow plus button on the right.




Figure 1. Create a new team in the KNIME Business Hub Administration page

After you create a new team you will be redirected to the new team's page. Here you can change the name of the team. To do so click the name of the team under the team logo on the left side of the page. The name of the team can also be changed at any point in time by

the team administrator.

From the team's page you can:

- Add members to the team
- Change their role to, for example, promote a user to team administrator role

Here you might for example want to assign the team to a team administrator. To do so click *Manage team* and enter the user name of the user you want to assign as a team administrator for the current team. Then click on the role and select *Member* and *Admin*. At the same time you might want to delete the global admin user name from the team members list. To do so click the bin icon corresponding to that user. Click *Save changes* to finalize the setting.

Allocate resources to a team

To allocate resources to a team navigate to the KNIME Business Hub Administrator page and select *Teams* from the menu on the left.

Here you can see an overview of the teams available, their allocated resourced, and of their current usage. Click the three dots on the right upper corner of the card corresponding to the team you want to allocate resources to.

KNIME Business Hub Administration

Standard Edition

License

Users

Teams

Teams

Team

CT CDDS Team

Users: 1 of 1

Core Tokens: 8 of unlimited

Team

CD Continuous Deployment Setup

Users: 2 of 3

Core Tokens: 0 of 3

Team

FT Finance Team

Users: 2 of 10

Core Tokens: 0 of unlimited

Team

KA knime_admins team

Users: 1 of 1

Core Tokens: 0 of unlimited

Manage Resources

Manage Members

Figure 2. Manage resources of a team

Select *Manage resources* from the menu. A panel on the right will open where you can select the resources you want to allocate.




Figure 3. Allocate resources to a team

Here you can change:

- The maximum number of members allowed in that team
- The maximum number of execution vCore tokens allowed for that team

Click *Save changes* when you have set up the resources for the current team.

Manage team members

From the KNIME Business Hub Administration page you can also manage the team members.

Click the three dots on the right upper corner of the card corresponding to the team. From the menu that opens select *Manage members*. In the side panel that opens you can add members to a team, or change the team members role.

Delete a team

From the KNIME Business Hub Administration page you can also delete a team.

Click the three dots on the right upper corner of the card corresponding to the team. From the menu that opens select *Delete*. Be aware that this operation will delete also all the team resources, data and deployments.

Execution resources

As mentioned in the previous section you as an Hub admin can assign execution resources to each team.

Team admins will then be able to build execution contexts according to the execution resources that you assigned to their team. These execution contexts will then be dedicated specifically to that team.

As an Hub admin you can also create a shared execution context. Once you create one you can share it with multiple teams.

For an overview of all the available execution contexts click your profile icon on the top right corner of the KNIME Hub and select *Administration* from the drop-down.

You will be then redirected to the KNIME Business Hub administration page.

Here, select *Execution resources* from the menu on the left.

In this page you can see an overview of *All* the execution contexts available on the Hub.

From the toggle at the top you can filter to see only a specific type of execution contexts available in the Hub instance.

Select:

- *Shared*: Shared execution contexts are created by the Hub admin. They can be made available to multiple teams.
- *Dedicated*: Dedicated execution contexts are created by the team admins for their team. Dedicated execution contexts are exclusively used by a single team.




Each team can by default have a maximum of 10 execution contexts and 10000 jobs. As a KNIME Hub admin you can change these limits via a REST API call like the following:

```
``PUT https://api.<base-url>/execution/limits/{scopeId}/{limitKey}``
```

where `{scopeId}` is the team account ID and the `{limitKey}` is account-execution-contexts or account-jobs respectively.

Create a shared execution context

As an Hub admin you can create a shared execution context and make it available to multiple teams. To do so click the  button. A side panel opens where you can set up the new shared execution context.




Figure 4. Create a shared execution context

Here you can give the execution context a name, set up which **Docker image** to use as the executor, give a list of blacklisted nodes, and assign the resources that the execution context will be able to use.




Find more information about how to set up the execution context in the **KNIME Business Hub User Guide**.

Finally, you can configure whether you want the execution context to automatically start and stop. To do so click *Set* under *Configure start and stop behavior* and select *On* (the setting is *Off* by default) from the toggle on top. Then you can indicate the desired inactivity time (in minutes) for the execution context to stop.

The execution context will start automatically when a queued workflow needs to run and stop automatically when there are no more active or queued workflows.

Click *Submit* to create the execution context. A notification will appear where you can click *Manage access* to share the execution context with the teams.

At any time you can also manage the access to a shared execution context by clicking the  icon in the corresponding tile and selecting *Manage access* from the menu.





Figure 5. Manage access for a shared execution context

Manage shared execution contexts

Also from the *Execution resources* page you can have an overview about the current status of an execution context, which teams have access to it, how many jobs are running and also manage the execution context performing the following operations:

- *Start* and *Stop* an execution context by clicking the Start/Stop button in the tiles
- Click the  icon in the tile and from the menu that opens you can:
 - *Edit*: You can change the parameters and configurations in the right side panel that opens.
 - *Manage access*: Manage the access of teams to a shared execution context.
 - *Enable/Disable*: You will need first to delete the jobs associated to the execution context then proceed with disabling it.
 - *Delete*: As a Hub administrator you can delete a shared execution context. You will need to first, delete the jobs associated to the execution context then proceed with disabling it. Finally, you can delete the shared execution context.
 - *Show details*: Selecting this option will open a new page with a list of all the jobs that are running on that execution context, the usage of the execution context

(e.g. how many vCores are in use) and other information. You can also switch to the specific *Executor* tab to see more details about the executor.

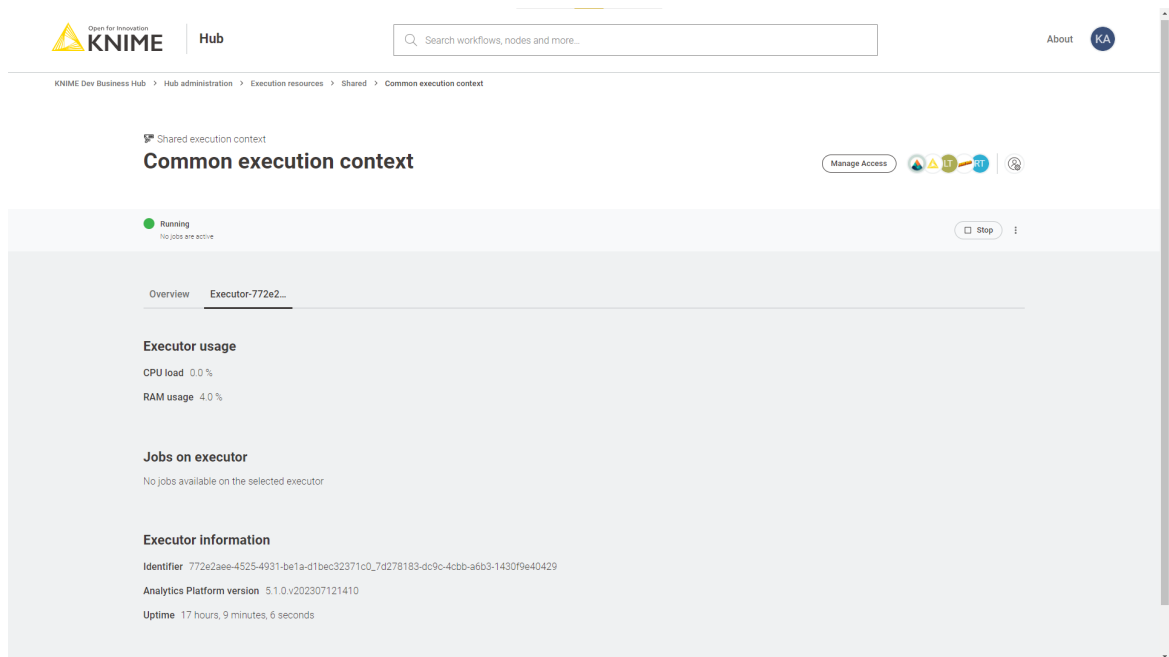


Figure 6. Additional executor information page

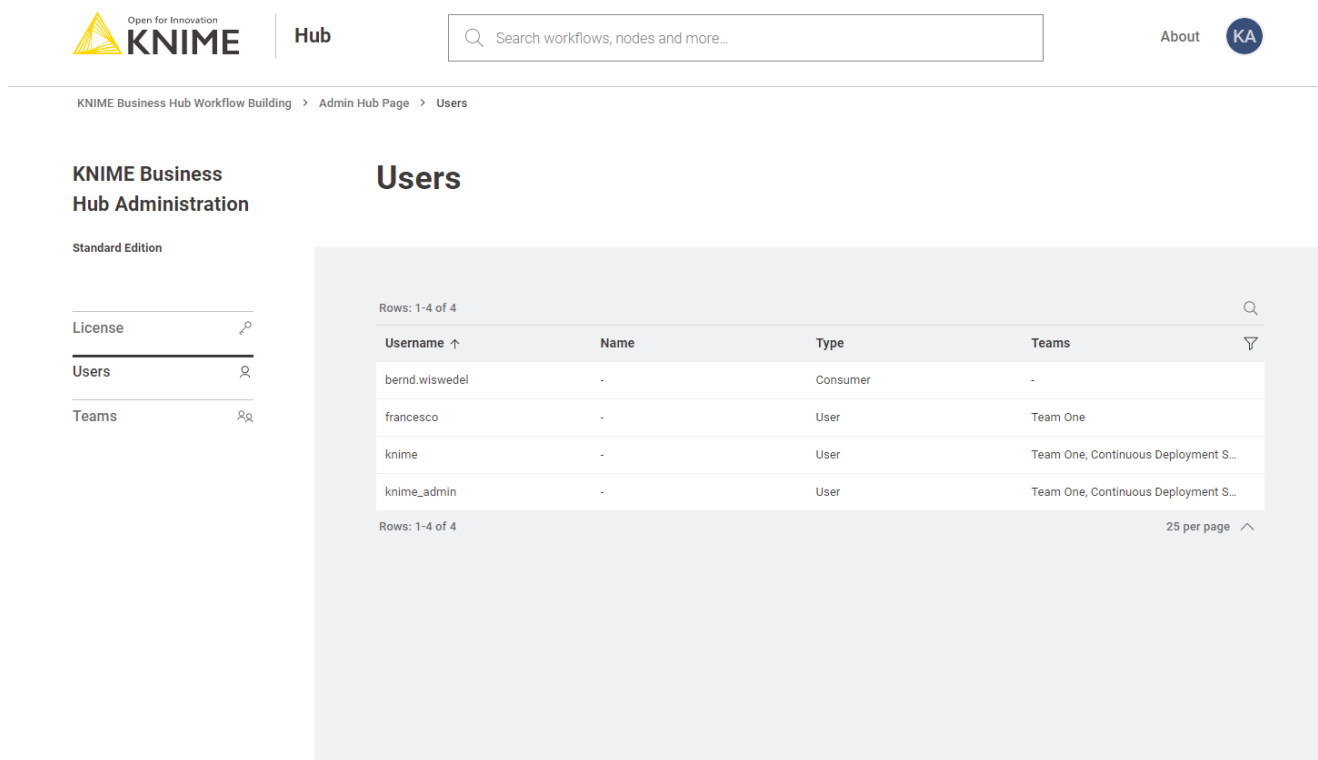
Advanced configuration of execution contexts

Execution contexts can be created and edited also via the Business Hub API.

Find more information on the available configurations in the [Advanced configuration of execution contexts](#) section in KNIME Business Hub User Guide.

Users management

To see a list of all the users that have access to your KNIME Business Hub instance you can go to the KNIME Business Hub Administration page and select *Users* from the menu on the left.



The screenshot shows the KNIME Business Hub Administration interface. On the left, there is a sidebar with the KNIME logo and navigation links: License, Users, and Teams. The main content area is titled 'Users' and displays a table of users. The table has columns for Username, Name, Type, and Teams. There are 4 rows of data. Above the table, there is a search bar and a filter icon. Below the table, there is a pagination control showing '25 per page'.

Username ↑	Name	Type	Teams
bernd.wiswedel	-	Consumer	-
francesco	-	User	Team One
knime	-	User	Team One, Continuous Deployment S...
knime_admin	-	User	Team One, Continuous Deployment S...

Figure 7. Manage users on KNIME Business Hub Administration page

Here you can filter the users based on their team, the type of users and their username and name. To do so click the funnel icon in the users list. You can also search the users by using the magnifier icon and typing the keyword in the field that appears.

Delete a user

You can delete a user from the KNIME Business Hub Administration page. Click the three dots and select *Delete*. You will need to confirm the user deletion in the window that opens by clicking *Delete user*. Be aware that this action will also delete all data from the deleted user and it will not be possible to restore the user.

Note that this user will continue to exist in Keycloak itself and you may want to delete it from there as well.

Make a user Hub admin

Users are managed in the backend via the Keycloak instance embedded in KNIME Business Hub. Therefore, the operation of promoting a registered user to the role of Hub admin is done in Keycloak.

To do so follow these steps:

1. First you will need to access the keycloak admin console. To do so you will need the credentials that are stored in a kubernetes secret called `credential-knime-keycloak` in the `knime` namespace. To get the required credentials, you need to access the instance the Business Hub is running on and run the following command:

```
kubectl -n knime get secret credential-knime-keycloak -o yaml
```

This will return a file that contains the `ADMIN_PASSWORD` and the `ADMIN_USERNAME`. Please notice that they are both base64 encrypted. In order to get the decrypted username and password, you can run the following commands:

```
echo <ADMIN_PASSWORD> | base64 -d  
echo <ADMIN_USERNAME> | base64 -d
```

2. Then go to <http://auth.<base-url>/auth/> and login.
3. In the top left corner click the dropdown and select the "Knime" realm, if you are not there already.

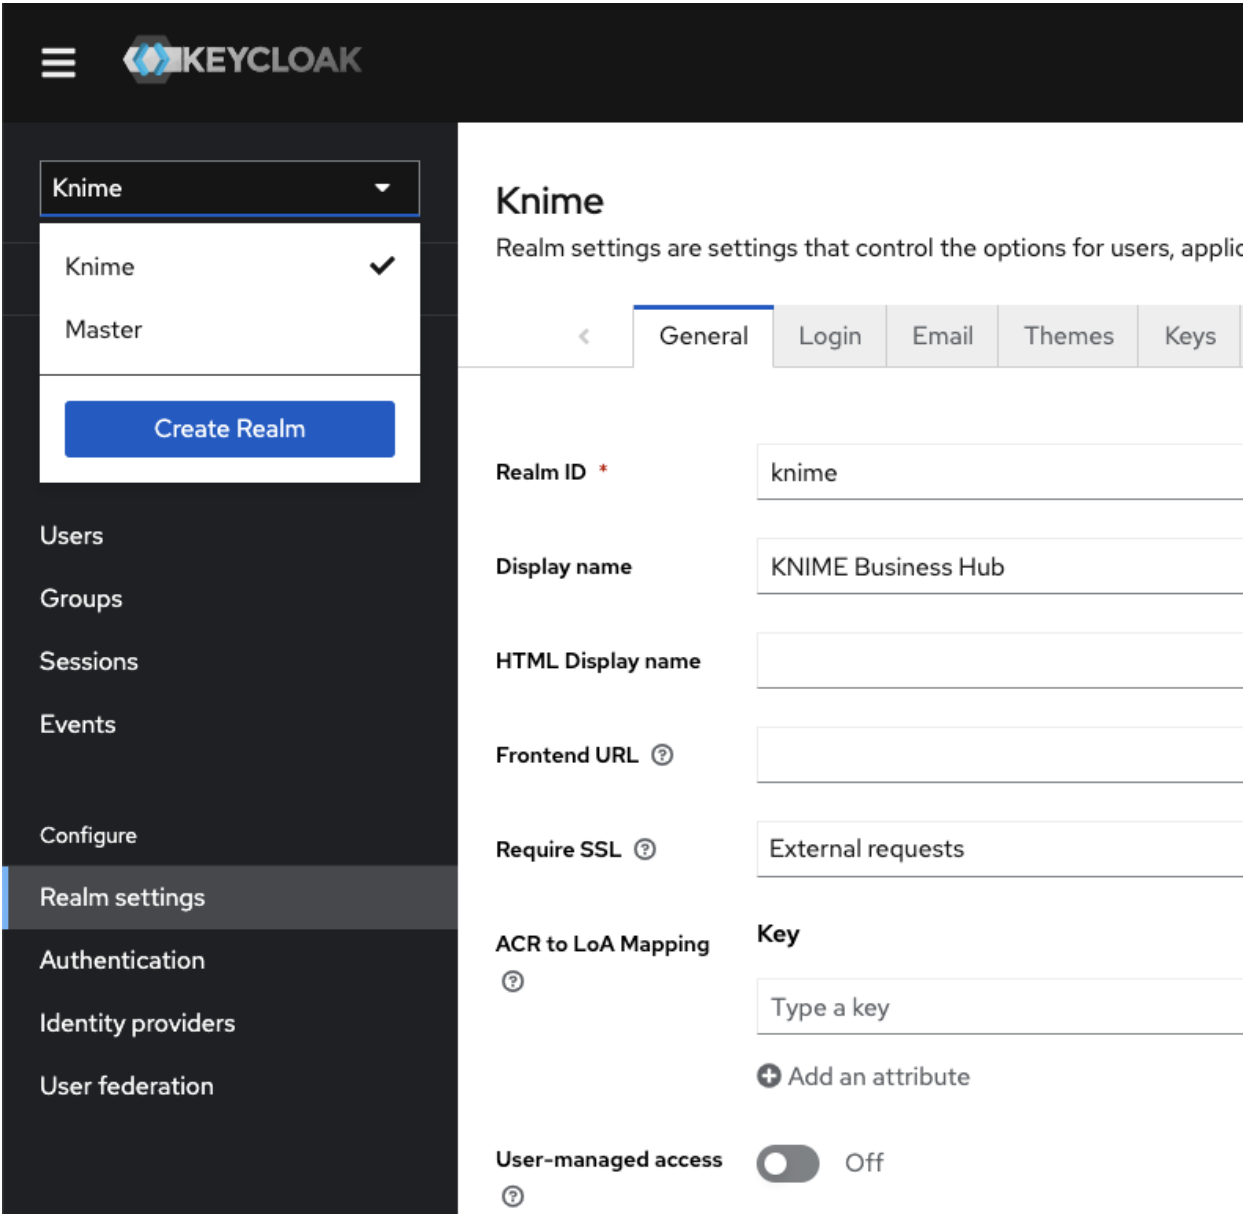


Figure 8. Select the "Knime" realm

- 4. Navigate to the *Users* menu and search for the user by name or email:

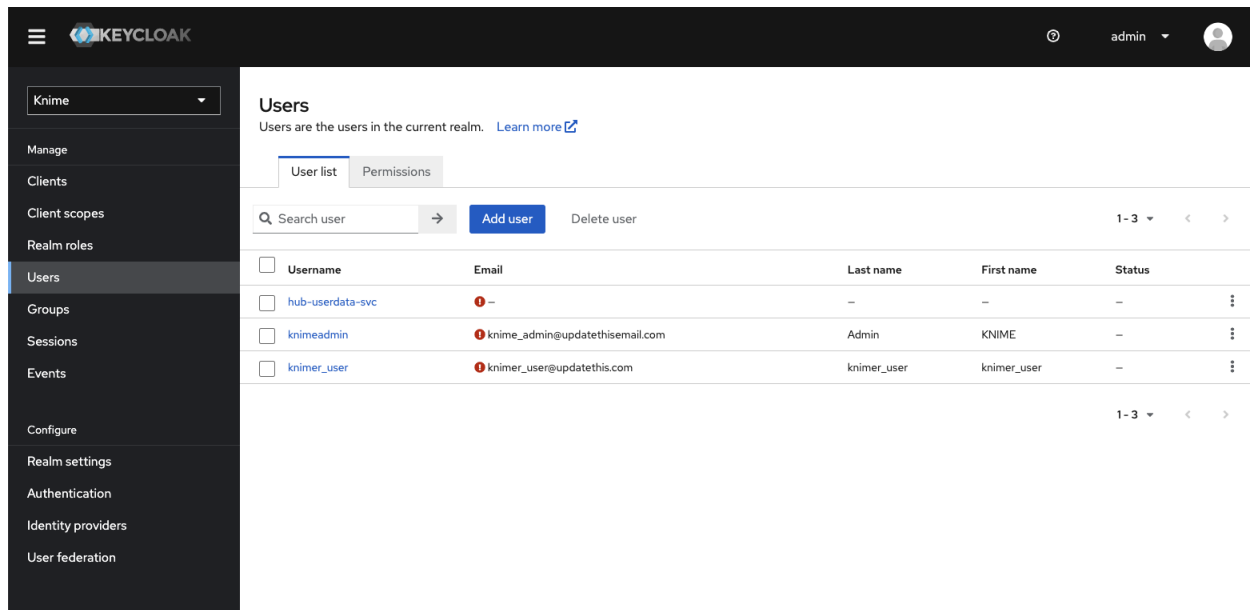


Figure 9. The Keycloak users menu



In order for a user to appear in this list, it is necessary that they have logged into your KNIME Business Hub installation at least once.

- Click the user and go to the *Groups* tab. Click *Join Group* and either expand the *hub* group by clicking it, or search for "admin". Select the admin group and click *Join*:

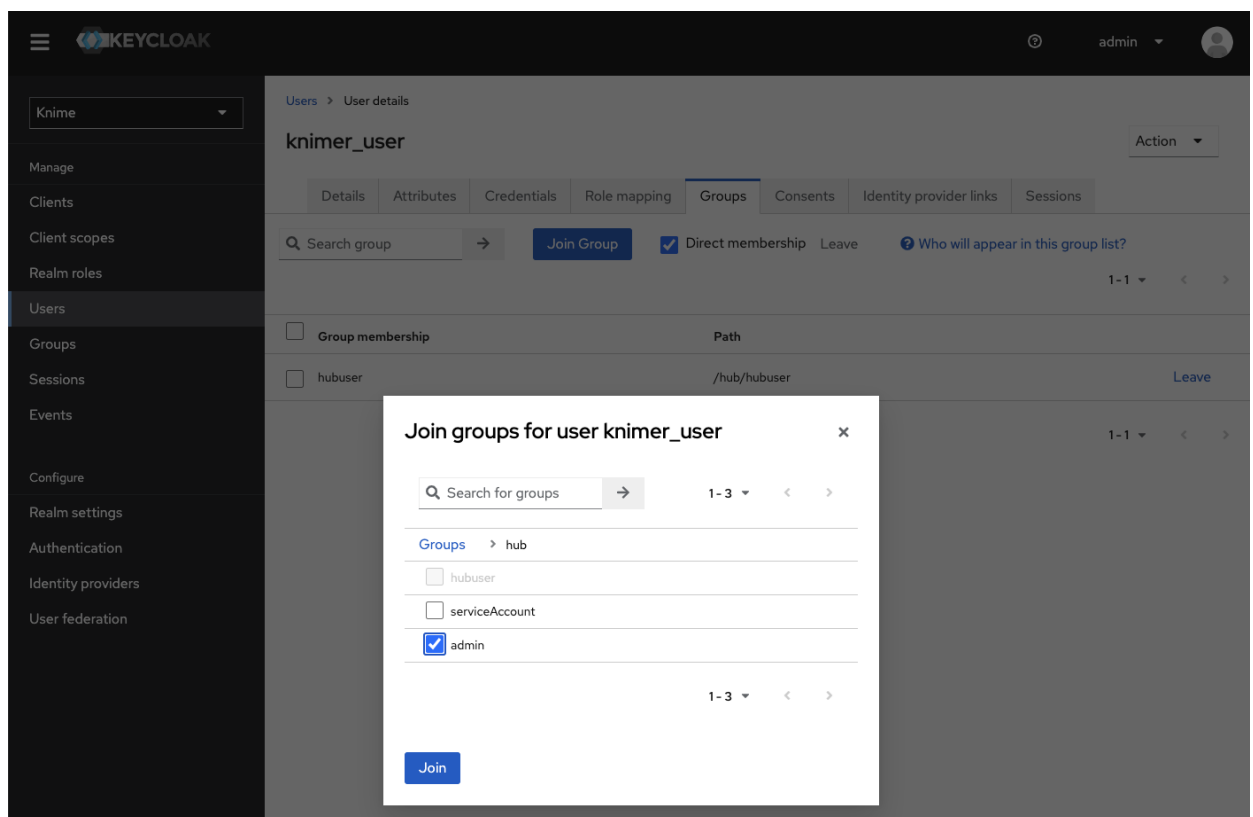


Figure 10. Making a user a Hub admin in Keycloak. If you are searching for the group then the group might show up under its full path "/hub/admin"

6. Done. The user now has the role of Hub admin, and can access the admin pages from within the Hub application to e.g., create teams or delete users.



Please notice that right now there are some operations that can be performed only by the global admin user that was created contextually to the KNIME Business Hub instance installation. Therefore, it is recommended to not delete the corresponding global admin user even when more users have been promoted to global admin users.

Expose external groups inside KNIME Business Hub

As a Global KNIME Hub administrator you can configure groups that are provided via an external identity provider to be exposed inside the KNIME Business Hub instance.

Two possible sources for your external groups are:

1. Groups are provided within the access token of your OIDC provider.
2. Groups are imported from LDAP by federating the login.


External OIDC provider

Assume you have an identity provider that provides groups through a **groups** claim in the access token.

```
{
  "...",
  "groups": [
    "finance",
    "marketing",
    "data"
  ]
}
```



First you need to configure Keycloak in such a way that it can map these groups to a **user attribute**. The second step is to add a mapper that maps these user attributes into the Keycloak's tokens.

Your third-party identity provider should have been set up already. Keycloak has to be configured as follows:




First step is to add an Attribute Importer Mapper.

1. In Keycloak select realm *Knime* in the top left dropdown menu
2. On the left tab select *Identity Providers*
3. Select your third-party provider
4. Switch to the tab *Mappers* and click on *Add mapper*
5. Provide a name for the mapper and set the *Sync mode override* to *Force* to ensure that the user's group memberships are updated upon every login
6. Set *Mapper type* to *Attribute importer*
7. Enter the *Claim* that contains the external groups in the original token (in our example *groups*)
8. In the *User Attribute Name* field enter *external-groups*
9. Click on *Save*



Now, every user in Keycloak who logged in after the mapper has been added will have an *external-groups* attribute associated like in the following picture:



Now, the external groups are known to Keycloak. To expose them inside KNIME Business Hub they need to be mapped into the access tokens issued by Keycloak. For this a second mapper needs to be added, that maps the user attribute *external-groups* to a claim in the user's access token.

To do this you need to add a client scope, which includes a mapper for the user attribute.

1. On the left tab select *Client scopes*
2. Select *groups*
3. Switch to the tab *Mappers*
4. Click on *Add mapper* > *By configuration* and select *User Attribute* from the list
5. Provide a name, e.g. *external-groups-attribute-mapper*
6. Set both fields *User Attribute* and *Token Claim Name* to *external-groups*
7. Ensure that *Add to ID token*, *Add to access token*, *Add to userinfo*, and *Multivalued* are turned on and that *Aggregate attribute values* is turned off
8. Click on *Save*

KEYCLOAK

admin

KNIME

Manage

Clients

Client scopes

Realm roles

Users

Groups

Sessions

Events

Configure

Realm settings

Authentication

Identity providers

User federation

Client scopes

Client scopes are a common set of protocol mappers and roles that are shared between multiple clients. [Learn more](#)

▼ Name

Search for client scope

→

Create client scope

Change type to

1-10

<

>

<input type="checkbox"/>	Name	Assigned type	Protocol	Display order	Description	
<input type="checkbox"/>	acr	Default	OpenID Connect	–	OpenID Connect scope for add acr (authentication context class reference) to the token	⋮
<input type="checkbox"/>	address	Optional	OpenID Connect	–	OpenID Connect built-in scope: address	⋮
<input type="checkbox"/>	email	Default	OpenID Connect	–	OpenID Connect built-in scope: email	⋮
<input type="checkbox"/>	external-groups-scope	Default	OpenID Connect	–	scope for external groups	⋮
<input type="checkbox"/>	groups	None	OpenID Connect	–	–	⋮
<input type="checkbox"/>	microprofile-jwt	Optional	OpenID Connect	–	Microprofile - JWT built-in scope	⋮
<input type="checkbox"/>	offline_access	Optional	OpenID Connect	–	OpenID Connect built-in scope: offline_access	⋮
<input type="checkbox"/>	phone	Optional	OpenID Connect	–	OpenID Connect built-in scope: phone	⋮
<input type="checkbox"/>	profile	Default	OpenID Connect	–	OpenID Connect built-in scope: profile	⋮
<input type="checkbox"/>	role_list	Default	SAML	–	SAML role list	⋮

1-10<>

KEYCLOAK

admin

KNIME

Manage

Clients

Client scopes

Realm roles

Users

Groups

Sessions

Events

Configure

Realm settings

Authentication

Identity providers

User federation

Client scopes > Client scope details

Groups [openid-connect](#)

Action


SettingsMappersScope

Search for mapper→Add mapper

1-1<>

Name	Category	Type	Priority	
groups	Token mapper	User Realm Role	40	⋮

1-1<>



With both mappers in place, the external groups are part of the access tokens issued by Keycloak. By this, the external groups are exposed inside KNIME Business Hub. In order to enable external groups to be used for permissions and access management they need to be configured separately through the admin REST API as described in [Enable external groups](#).

LDAP federation

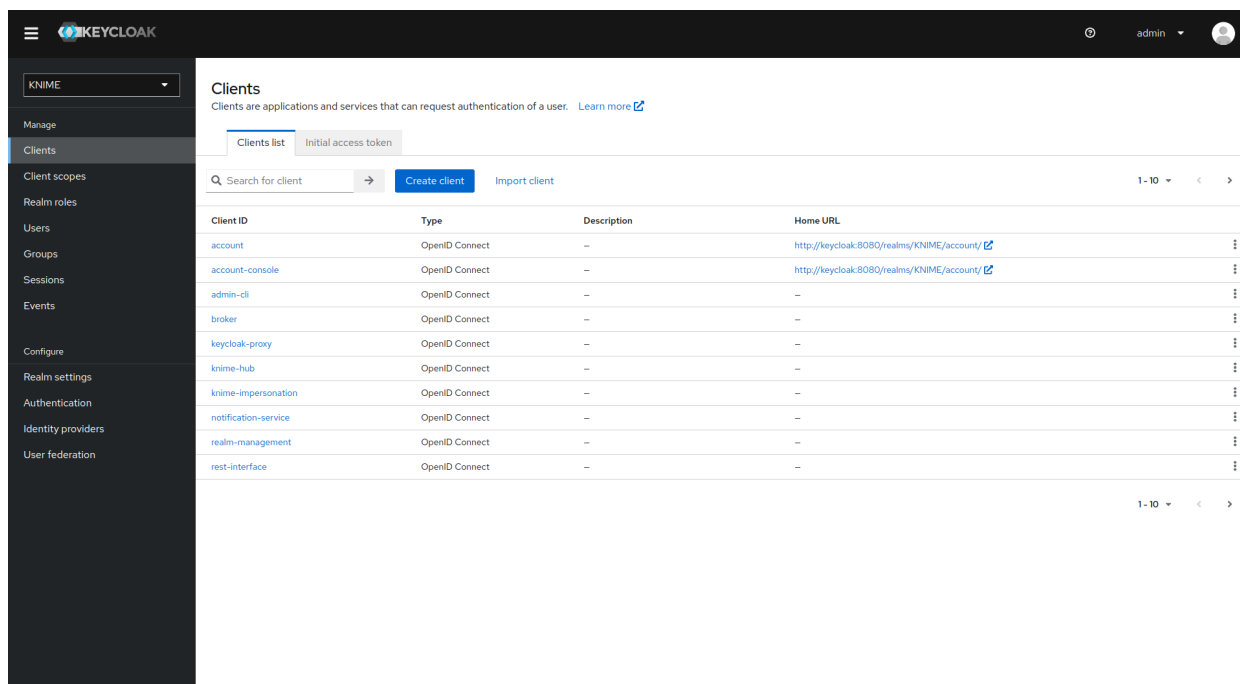
If you have user federation configured for an LDAP instance that also supplies external group names you need to configure mappers that map these groups into the access tokens used inside the Hub instance.

To ensure that groups from Keycloak groups and groups from LDAP are not mixed we recommend to treat external groups as realm roles.

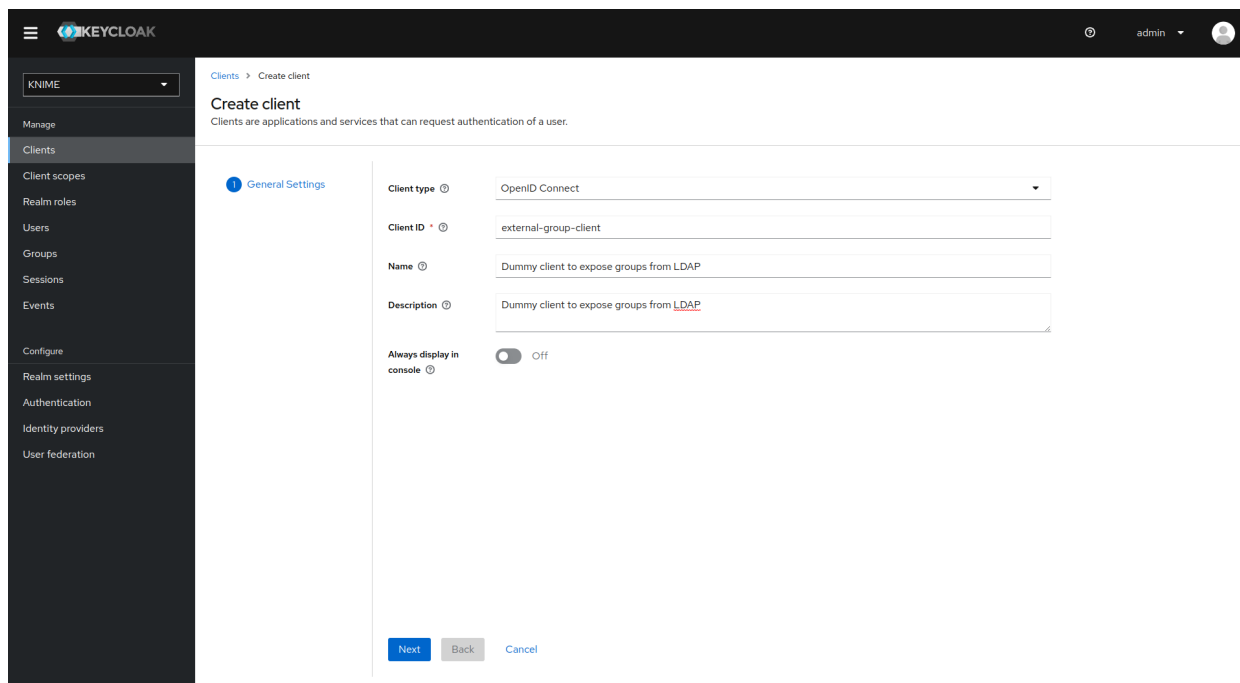
In order to do this we recommend to first create a dummy client for which roles can be created based on the LDAP groups. This will guarantee that any changes will be compatible with future changes to the KNIME Hub client in Keycloak.

To create a new client follow these steps:


1. In Keycloak select realm *Knime* in the top left dropdown menu
2. On the left tab select *Clients* and click *Create client*



3. Set *Client type* to *OpenID Connect*
4. Enter a *Client ID* (in our example *external-group-client*), and a useful *Name* and *Description*
5. Click on *Next*



6. De-select all checkboxes of *Authentication flow* in the *Capability config* section, since this client will not require any capabilities
7. Enable *Client authentication*
8. Click on *Save*



Now that the dummy client is set up, you can proceed to create a mapper that maps the user groups from LDAP to roles inside the dummy client:

1. On the left tab select *User federation* and click on your LDAP configuration
2. Switch to the tab *Mappers*
3. Click on *Add mapper*
4. Provide a name, e.g. *ldap-group-to-dummy-client-role-mapper*
5. Set *Mapper type* to *role-ldap-mapper*
6. Setup the mapper according to your LDAP
7. Disable *User Realm Roles Mapping*
8. Set *Client ID* to the previously created dummy client (in our example *external-group_client*)
9. Click on *Save*

The screenshot shows the Keycloak Admin Console interface. On the left, a sidebar contains a navigation menu with categories like 'Manage' and 'Configure'. The 'User federation' option is highlighted. The main content area is titled 'Create new mapper' and displays the following configuration fields:

- Name: ldap-group-to-dummy-client-role-mapper
- Mapper type: role-ldap-mapper
- LDAP Roles DN: OU=groups, DC=knime, DC=com
- Role Name LDAP Attribute: CN
- Role Object Classes: groupOfNames
- Membership LDAP Attribute: member
- Membership Attribute Type: DN
- Membership User LDAP Attribute: CN
- LDAP Filter: (empty)
- Mode: READ_ONLY
- User Roles Retrieve Strategy: LOAD_ROLES_BY_MEMBER_ATTRIBUTE
- Member-Of LDAP Attribute: memberOf
- Use Realm Roles Mapping: Off
- Client ID: external-group-client

At the bottom of the form are 'Save' and 'Cancel' buttons.

Now if a user logs in with the LDAP credentials the user's groups will be mapped to ad-hoc created client roles inside the 'external-group-client'.

Next, you need to create a mapper that maps a user's realm roles from the dummy realm to the access tokens:

1. On the left tab select *Client scopes*
2. Select *groups*
3. Switch to the tab *Mappers*
4. Click on *Add mapper* > *By configuration* and select *User Client Role* from the list
5. Provide a name, e.g. *external-ldap-client-role-mapper*
6. Set *Client ID* to the previously created dummy client (in our example *external-group_client*)
7. Set *Token Claim Name* to *external-groups*
8. Set *Claim JSON Type* to *String*
9. Ensure that *Add to ID token*, *Add to access token*, *Add to userinfo*, and *Multivalued* are turned on

10. Click on Save

The screenshot shows the Keycloak Admin Console interface. On the left is a dark sidebar with navigation links: Manage, Clients, Client scopes (selected), Realm roles, Users, Groups, Sessions, Events, Configure, Realm settings, Authentication, Identity providers, and User federation. The main content area has a breadcrumb trail: Client scopes > Client scope details > Mapper details. The title is 'Add mapper' with a subtitle: 'If you want more fine-grain control, you can create protocol mapper on this client'. The 'Mapper type' is 'User Client Role'. The 'Name' field contains 'external-ldap-client-role-mapper'. The 'Client ID' dropdown is set to 'external-group-client'. The 'Client Role prefix' field is empty. The 'Multivalued' toggle is turned 'On'. The 'Token Claim Name' field contains 'external_groups'. The 'Claim JSON Type' dropdown is set to 'String'. There are three checkboxes: 'Add to ID token' (On), 'Add to access token' (On), and 'Add to userinfo' (On). At the bottom are 'Save' and 'Cancel' buttons.

Enable external groups

Once you have configured the external groups in Keycloak you need to create the groups that you want to be available inside KNIME Business Hub.

To do so you have to make a PUT request to the corresponding endpoint:

```
PUT https://api.<base-url>/accounts/hub:global/groups/<external-group-name>
```

where <external-group-name> is the name of the group and it must match the group name in the external identity provider.

You can use the following JSON type body in order to set a display name for the group to be shown on KNIME Business Hub.

```
{
  "displayName": "My Group Display Name"
}
```

Docker executor images

In order to create execution contexts for their teams, team admins will need to reference the Docker image of the KNIME Executor that they want to use.

Public Docker executor images are made available by KNIME which correspond to the full builds of KNIME Executor versions 4.7.0 and higher.

The currently available executor images have the following docker image name:

- `registry.hub.knime.com/knime/knime-full:r-4.7.4-179`
- `registry.hub.knime.com/knime/knime-full:r-4.7.5-199`
- `registry.hub.knime.com/knime/knime-full:r-4.7.6-209`
- `registry.hub.knime.com/knime/knime-full:r-4.7.7-221`
- `registry.hub.knime.com/knime/knime-full:r-4.7.8-231`
- `registry.hub.knime.com/knime/knime-full:r-5.1.0-251`
- `registry.hub.knime.com/knime/knime-full:r-5.1.1-379`
- `registry.hub.knime.com/knime/knime-full:r-5.1.2-433`
- `registry.hub.knime.com/knime/knime-full:r-5.1.3-594`
- `registry.hub.knime.com/knime/knime-full:r-5.2.0-271`
- `registry.hub.knime.com/knime/knime-full:r-5.2.1-369`
- `registry.hub.knime.com/knime/knime-full:r-5.2.2-445`
- `registry.hub.knime.com/knime/knime-full:r-5.2.3-477`



Please note that with the release of 5.2.2 KNIME executors will have HTML sanitization of old JavaScript View nodes and Widget nodes turned on by default. This should ensure that no malicious HTML can be output. For more information check the [KNIME Analytics Platform 5.2.2 changelog](#).

However you might want to add specific extensions to the KNIME Executor image that is made available to team admins to create execution contexts.

The following section explains how to do so.

Add extensions to an existing Docker image

In order to install additional extensions and features to the KNIME Executor image, you will need to first create a `Dockerfile`. The file is named `Dockerfile` with no file extension. You can use the example `Dockerfile` below which demonstrates how to extend the base image with a custom set of update sites and features.

i

If you need to install Docker please make sure not to install it on the same virtual machine (VM) where the KNIME Business Hub instance is installed, as it might interfere with `containerd`, which is the container runtime used by Kubernetes.

```
# Define the base image
FROM registry.hub.knime.com/knime/knime-full:r-4.7.4-179

# Define the list of update sites and features
# Optional, the default is the KNIME Analytics Platform update site (first entry in the
# list below)
ENV KNIME_UPDATE_SITES=https://update.knime.com/analytics-
platform/4.7,https://update.knime.com/community-contributions/trusted/4.7
# Install a feature from the Community Trusted update site
ENV KNIME_FEATURES="org.knime.features.geospatial.feature.group"

# Execute extension installation script
RUN ./install-extensions.sh
```

The `KNIME_UPDATE_SITES` environment variable determines the update sites that will be used for installing KNIME Features. It accepts a comma-delimited list of URLs. The `KNIME_FEATURES` environment variable determines the extensions which will be installed in the KNIME Executor. It accepts a comma-delimited list of feature group identifiers. A corresponding update site must be defined in the `KNIME_UPDATE_SITES` list for feature groups to be successfully installed. You can get the necessary identifiers by looking at *Help → About KNIME → Installation Details → Installed Software* in a KNIME instance that has the desired features installed. Take the identifiers from the "Id" column and make sure you do not omit the `.feature.group` at the end (see also screenshot on the next page). The base image contains a shell script `install-extensions.sh` which lets you easily install additional extensions in another `Dockerfile`.

Once the `Dockerfile` has been customized appropriately, you can build a Docker image from it by using the following command:

```
# Replace <image_name> and <tag_name> with actual values
docker build -t <image_name>:<tag_name> .
```

Python and Conda in Docker images

When you create an Execution Context on KNIME Business Hub based on a **full build** you will have KNIME Python bundled available. If you need additional libraries you would need to make them available on the Hub instance.

You can do this in two ways:

1. Use the Conda Environment Propagation node.
2. Customize the Executor image used.

To get started with Conda environment propagation, check out [KNIME Python Integration Guide](#). However, any libraries installed using Conda environment propagation will be removed when the executor restarts and installed again next time, so libraries that are used often should be installed as part of the executor Docker image.

In order to do so you need to:

1. Install Python in the executor Docker image
2. Declare to the execution context the path to the Python installation folder so that the executor can execute Python nodes

Install Python in the executor Docker image

The first step is the installation of Python and an environment manager for instance miniconda on a Docker image.

To do so, first you need the Docker Project to hold a miniconda installer near the Dockerfile, for example:

```
python-image/  
  container/  
    |-Miniconda3-py310_23.3.1-0-Linux-x86_64.sh  
  dockerfiles/  
    |-Dockerfile
```

You will also need to provide a `.yaml` file that will contain all the modules, packages and Python version that you need to install in order to execute the Python scripting nodes.

The `.yaml` file could look like the following:

```

name: py3_knime # Name of the created environment
channels: # Repositories to search for packages
  - defaults
  - anaconda
  - conda-forge
dependencies: # List of packages that should be installed
# - <package>=<version> # This is an example of package entry structure
  - python=3.6 # Python
  - scipy=1.1 # Notebook support
  - numpy=1.16.1 # N-dimensional arrays
  - matplotlib=3.0 # Plotting
  - pyarrow=0.11 # Arrow serialization
  - pandas=0.23 # Table data structures

```

Then you need to pull any **available executor Docker image**, install miniconda in batch mode on the image and define the environment variable for conda, as in the following example. Also you will need to create your environments, that you specified in the `.yaml` files.

```

# getting recent knime-full image as a basis
FROM registry.hub.knime.com/knime/knime-full:r-4.7.3-160

# getting Miniconda and install in batch mode
COPY container/Miniconda3-py310_23.3.1-0-Linux-x86_64.sh /home/knime/miniconda-latest.sh
RUN bash miniconda-latest.sh -b

# adding path to Miniconda bin folder to system PATH variable
ENV PATH="/home/knime/miniconda3/bin:$PATH"

# copy default conda environments into container
COPY --chown=knime envs/ ./temp_envs
RUN conda env create -f ./temp_envs/py3_knime.yaml && \
    rm -rf ./temp_envs

```

When installing conda and creating the environments you will obtain the following paths that will need to be added in the `.epf` file of the customization profile during the **set up of the execution context**.

For example based on the above Dockerfile:

```

<path to conda installation dir>=/home/knime/miniconda3/
<path to default conda environment dir>=<path to conda installation dir>/envs/<name of the env>

```

Now you can build the new Docker image, for example:

```
docker build . -f /dockerfiles/Dockerfile -t knime-full:4.7.3-with-python
```

Finally retag the image to make it useable for your embedded registry on Business Hub:

```
docker tag knime-full:4.7.3-with-python registry.<hub-url>/knime-full:4.7.3-with-python
```

Once you have created the Docker image with Python installed [create an execution context](#) that uses the newly created Docker image.

Set up the execution context

Now you need to set up and customize the execution context.

In order to declare to the execution context the path to the Python installation you will need to build a dedicated [customization profile](#) and apply it to the execution context.

1. Build the .epf file by following the steps in [KNIME Python Integration Guide](#) and exporting the .epf file. To export the .epf file from KNIME Analytics Platform go to *File > Export Preferences...*
2. Open the file and use only the parts related to Python/conda.

The .epf file could look like the following:

```
/instance/org.knime.conda/condaDirectoryPath=<path to conda installation dir>  
/instance/org.knime.python3.scripting.nodes/pythonEnvironmentType=conda  
/instance/org.knime.python3.scripting.nodes/python2CondaEnvironmentDirectoryPath=<path  
to default conda environment dir>  
/instance/org.knime.python3.scripting.nodes/python3CondaEnvironmentDirectoryPath=<path  
to default conda environment dir>
```



Find more details on how to set-up the .epf file in the [Executor configuration](#) section of the KNIME Python Integration Guide.

Now follow these steps to customize the execution context:

1. Build the .zip file containing the [customization profile](#) using the .epf file you just created.
2. Upload the [customization profile](#) .zip file to KNIME Business Hub.
3. Apply the [customization profile](#) to the execution context.

Customization profiles

Customization profiles are used to deliver KNIME Analytics Platform configurations from KNIME Hub to KNIME Analytics Platform clients and KNIME Hub executors.

This allows to define centrally managed:

- Update sites
- Preference profiles (such as Database drivers, Python/R settings)

A profile consists of a set of files that can:

- Be applied to the client during startup once the **KNIME Analytics Platform client is configured**. The files are copied into the user's workspace.
- Be applied to the **KNIME Hub executors** of the execution contexts.

Customization profiles can be:

- Global customization profiles that need to be uploaded and managed by the Global Admin. These can be applied across teams via shared execution context or to specific teams.
- Team's customization profiles, which are scoped to the team, can be uploaded either by a Global Admin or a team admin.

Every user of the KNIME Business Hub instance can then download and use customization profiles.



Currently, customization profiles can be managed on KNIME Hub via REST or using the dedicated Customization Profile data application available [here](#).

Create a customization profile

A customization profile is a file with the `.epf` extension, that contains the needed profile information. So the first step is to create such a file.

1. Create an `.epf` file that contains the needed profile information:
 - a. Build the file from a local KNIME Analytics Platform installation. Set up the needed configuration in a local KNIME Analytics Platform installation, then go to *File → Export Preferences*. Open the created `.epf` file, and look for the entries related to your needed settings. Remove all other settings (as some contain e.g. local paths on your machine, which will inevitably cause issues when applying to

another installation).

- b. Place the `.epf` file in a folder, together with any additional files that need to be distributed along the profile (e.g. database drivers)
- c. Create a `.zip` from that folder



When creating a zip file on macOS using the built in functionality, two files are automatically added that cause the next steps (i.e. applying the profile in Analytics Platform) to fail. There is a way to prevent creation of these files if creating the `.zip` via command line, see [here](#). If in doubt, use a Windows or Linux machine to create the `.zip` file.



The customization profiles on the KNIME Hub instance are going to be accessible without user authentication therefore they shouldn't contain any confidential data such as passwords.

For further details and an example on how to distribute JDBC driver files go to the Hub-managed customization profiles section of the [KNIME Database Extension Guide](#).

Variable replacement

It is possible to use variables inside the preferences files (only those files ending in `.epf`) which are replaced on the client right before they are applied. This makes the Hub-managed customizations even more powerful. These variables have the following format:

`${prefix:variable-name}`. The following prefixes are available:

- **env**: the variable is replaced with the value of an environment value. For example, `${env:TEMP}` will be replaced with `/tmp` under most Linux systems.
- **sysprop**: the variable is replaced with a Java system property.
For example, `${sysprop:user.name}` will be replaced with the current user's name. For a list of standard Java system properties see [the JavaDoc](#). Additional system properties can be defined via `-vmargs` in the `knime.ini`.
- **profile**: the variable will be replaced with a property of the profile in which the current preference file is contained in. Currently `location` and `name` are supported as variable names. For example, `${profile:location}` will be replaced by the file system location of the profile on the client. This can be used to reference other files that are part of the profile, such as database drivers:
`org.knime.workbench.core/database_drivers=${profile:location}/db-driver.jar`

In case you want to have a literal in a preference value that looks like a variable, you have to

use two dollar signs to prevent replacement. For example `$$${env:HOME}` will be replaced with the plain text `${env:HOME}`. If you want to have two dollars in plain text, you have to write three dollars (`$$$${env:HOME}`) in the preference file.



Once you use variables in your preference files they are not standard Eclipse preference files anymore and cannot be imported as they are.

Client customization

Besides the preferences that are exportable by KNIME Analytics Platform there are additional settings that can be added to the preference files to customize clients:

```
/instance/org.knime.workbench.explorer.view/defaultMountpoint/defaultMountpoints
=<mount id1>,<mount id2>,...
```

A comma separated list of default mountpoints that should be loaded, e.g. LOCAL, EXAMPLES, My-KNIME-Hub. Changes to this list only affects new workspaces, i.e. workspaces which already contain default mountpoints will still contain them even though they haven't been defined here. If this option is absent and `defaultMountpoint/enforceExclusion` isn't set to true then all default mountpoints will be added. The current default mountpoints are LOCAL, EXAMPLES, and My-KNIME-Hub.

```
/instance/org.knime.workbench.explorer.view/defaultMountpoint/enforceExclusion=<
true|false>
```

If set to true then all default mountpoint not defined by `/instance/org.knime.workbench.explorer.view/defaultMountpoint/defaultMountpoints` will be removed on start up. Please note that if you want to use this option, the default mountpoints you want to include should **only** be listed in `/instance/org.knime.workbench.explorer.view/defaultMountpoint/defaultMountpoints`, and **not** in their full definition like when exporting the preferences.epf file from a KNIME Analytics Platform.

`/instance/com.knime.customizations/helpContact.buttonText=<label, e.g. Contact Support>`

If set together with `/instance/com.knime.customizations/helpContact.address` a button with the provided label will occur under Help in KNIME Analytics Platform. Clicking on the button will, depending on the `helpContact.address`, either open the default mail client or the default browser with the provided address.

`/instance/com.knime.customizations/helpContact.address=<uri, e.g. mailto:support@company or https://company/support>`

Sets the address of the support contact.

This option only takes effect in combination with

`/instance/com.knime.customizations/helpContact.buttonText`.

`/instance/com.knime.customizations/documentation.buttonText=<label, e.g. Documentation>`

Sets the label of the documentation button that can be found under Help in KNIME Analytics Platform. Clicking on the button will open the default browser and navigate to the documentation. If set to - the button will be hidden.

`/instance/com.knime.customizations/documentation.address=<uri, e.g. https://company/documentation or file:///sharedSpace/documentation>`

Sets the address of the documentation.

By default the documentation address points to the KNIME documentation.

`/instance/com.knime.customizations/windowTitle.appendix=<appendix, e.g. sponsored by company>`

Adds the appendix to the window title of KNIME Analytics Platform.

-

`/instance/com.knime.customizations/updateSite.uris=<uri>,<uri>,...`

Adds the provided addresses to the update sites.

-

```
/instance/com.knime.customizations/updateSite.names=<name>,<name>,...
```

The names that are shown under Available Software Sites for the provided update sites. Note that the number of names must match the number of provided URIs.

```
/instance/com.knime.customizations/updateSite.default.disable=<true|false>
```

Disables the default update sites added by KNIME after a fresh installation or update. If a user enables these update sites again they will remain enabled.

```
/instance/com.knime.customizations/updateSite.default.forceDisable=<true|false>
```

Disables the default update sites added by KNIME after a fresh installation or update. If a user enables these update sites again they will be disabled with the restart of their client.

Upload a customization profile

After creating the .zip file, you need to upload it to KNIME Business Hub. You can do it [via REST](#) or [via data application](#).

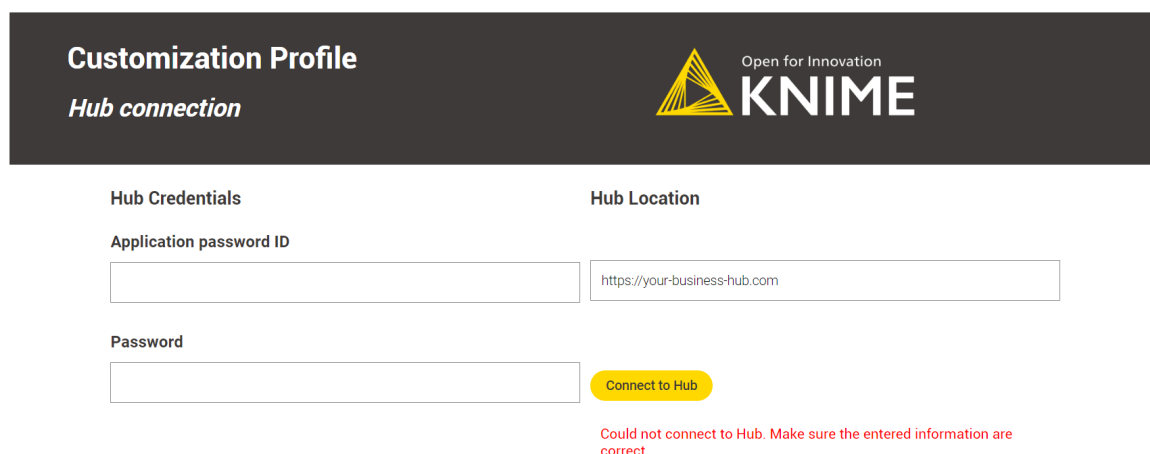
Upload via Data Application

1. If you have not done so already, on KNIME Business Hub, create an [application password](#) for the user uploading the profile.
 - a. For global customization profiles, use a global admin user.
 - b. For team's customization profiles, the user can be either a global admin user or a team admin user.
2. Download the workflow:



Click [here](#) and download the workflow to your KNIME Analytics Platform.

3. Upload the downloaded workflow to your KNIME Business Hub instance via [KNIME Analytics Platform](#).
4. Deploy the workflow as a data application. You can learn more about how to do it [here](#).
5. The first step to begin using the data application is establishing a **connection** with the Hub instance. You must enter your Hub URL and application password in the data app to connect to your KNIME Business Hub instance.



Customization Profile
Hub connection

Hub Credentials

Application password ID

Password

Hub Location

https://your-business-hub.com

Connect to Hub

Could not connect to Hub. Make sure the entered information are correct

Figure 11. The Hub connection step.

6. Select *Upload a New Customization Profile* in the *Action Selection* menu and click *Next*. You can upload multiple customization profiles to your KNIME Business Hub instance.



Customization Profile
Action selection

Please choose the action you wish to perform.

Upload a new customization profile

Figure 12. The upload option is selected in the action selection step.

7. In the next step, select the customization profile **file** previously created, give it a name, and define its scope.
 - a. **Global Admin:** You can choose the customization profile scope if logged in to the Data App as a Global Admin.
 - i. *Global scope:* After being uploaded to KNIME Business Hub, the global customization profile can be applied to a shared or a team execution context.
 - ii. *Team scope:* When a customization profile with team scope is uploaded, it can be applied to the team's executors. It is possible to upload the same customization profile to multiple teams. The Global Admin can upload the customization profile to any team.

Figure 13. Global Admin executing the data app

- b. **Team Admin:** If you're logged in as a team admin, you can only upload the customization profile as a team-scoped type. You can upload the same customization profile for multiple teams. The availability of teams is determined by the permissions of the user who is logged in.
8. In the Upload Results step, a table displays the results of the customization profile upload operation, along with a legend.

If everything goes well, the green light means your customization profile has been successfully uploaded to the KNIME Business Hub. You can now apply it to a **KNIME Analytics Platform installation** or to a **KNIME Business Hub executor**.

9. **Error Handling:** The data application is designed to handle two types of errors:
 - a. Missing customization profile file.

Customization profile details

An error occurred: No .zip file was uploaded. Click "Next" to return to the action selection menu.

Figure 14. This error occurs when no customization profile file is provided.

- b. No team selection (for team-scoped customization profiles).

If you encounter an error message, click *Next* and select the *Turn back to the*

choose *action menu* option. Repeat the operation, fixing the error.



Figure 15. Turn back to the action menu and repeat the operation.

10. Click *Next* to finish the operation. You can return to the *Action Selection* menu to perform additional actions or close the data application directly.

Upload via REST request

1. If you have not done so already, on KNIME Business Hub, create an **application password** for the user uploading the profile.
 - a. For global customization profiles, use a global admin user.
 - b. For team's customization profiles, the user can be either a global admin user or a team admin user.
2. Send a POST request to `https://api.<base-url>/execution/customization-profiles` with the following set up:
 - a. Authorization: select Basic Auth, using username and password from the created application password
 - b. Body: select form-data as request body type and add the following entries:
 - i. Add a new key, set content to "File". Name of the key needs to be "content". The value of the key is the .zip file containing the profile, and the content type is `application/zip`.
 - ii. Add a new key, set content to "Text". Name of the key needs to be "metadata".

A. The value of the key is:

For global customization profiles enter:

```
{
  "name": "<profile_name>",
  "scope": "hub:global"
}
```

For team's customization profiles you first need to obtain the `<team_ID>` of the team you want to upload the customization profile to. To do so you can use the following GET request:

```
GET api.<base-url>/accounts/identity
```

You will find the `<team_ID>` in the body response under `teams`:

```
...
  "teams": [
    {
      "id": "account:team:<team_ID>",
      "name": "<team_name>",
      ...
    }
  ]
...
```

Then the value of the key is:

```
{
  "name": "<profile_name>",
  "scope": "account:team:<team_ID>"
}
```

- B. Set the content type for the second key to `application/json`
 - c. When using Postman there is no need to manually adjust the headers
- 3. Send the request




Figure 16. Set up Postman for a REST call to upload a customization profile to a KNIME Hub instance

- If successful, you will receive a 201 status response. Make note of the created `<profile_ID>`, as this will be used to refer to the profile when requesting it.



You can refer to the API documentation at the following URL for more details about the different calls you can make to the Customization Profiles endpoint.

`http://api.<base-url>/api-doc/execution/#!/Customization%20Profiles`

Apply a customization profile

Apply a customization profile to KNIME Hub executor

Apply via Data App

To apply a customization profile to all executors running in a KNIME Business Hub **execution context**, we can use the data application. However, the type of user logged in affects the operations that can be performed. Refer to the table below for a comprehensive overview.

Table 1. Apply customization profile: user rights

User type	Eligible customization profile	Customization profile type	Eligible Execution Context
Global Admin	All uploaded within the KNIME Business Hub instance	Global	Shared and any team specific execution context
		Team-scoped	Any team execution contexts
Team admin	Only those uploaded in teams in which the user has team admin rights	Team-scoped only	Only the team execution contexts where the customization profile was uploaded. Shared execution contexts are not eligible

To provide a better understanding of the table, here are some examples to demonstrate its functionality:

- A Global Admin can choose one team-scoped customization profile and apply it to any team-scoped execution context within the KNIME Business Hub instance. For instance, the Global Admin can apply *Team A's customization profile* to the execution context of *Team B*, i.e., execution context B.1.
- A Global Admin can also select a **Global** customization profile and apply it to any shared and team-scoped execution context within the KNIME Business Hub instance.
- A team admin can **only** choose team-scoped customization profiles uploaded within the teams with admin rights. For example, they can only apply a customization profile uploaded in *Team A* within the Team A execution contexts.

1. Learn how to download the data app from Community Hub, upload and deploy it in KNIME Business Hub, and authenticate with your application password in the [Upload a customization profile](#) section.
2. Select *Apply an existing Customization Profile* in the *Action Selection* menu to apply a Customization Profile.
3. Choose the desired customization profile from the table and click *Next*. In the table, if

the *Team name* column has a value of *Any*, the customization profile has a global scope and can be applied to any team.

Customization Profile
Action selection

Please choose the action you wish to perform.

Apply an existing customization profile

Select one of the customization profiles that already exists
The customization profile can be applied to a KNIME Business Hub Executor or a KNIME Analytics Platform client.

Show 5 entries Clear Sorting

	Custom profile name	Custom profile type	File uploaded	Created at	Updated at	Team name
<input type="radio"/>	Python profile	Team	windows.zip	2023-11-08	2023-11-08	Dev Team
<input type="radio"/>	Python profile	Team	windows.zip	2023-11-08	2023-11-08	Data Science Team
<input checked="" type="radio"/>	DB configuration	Global	default.zip	2023-11-08	2023-11-08	Any

Previous 1 Next

Filter teams

Excludes

Includes

- Dev Team
- Data Science Team
- Any

Customization profile scope

☒ Global ☐ Team

Figure 17. A Global Admin could select a global or team-scoped customization profile.



It is only possible to apply one customization profile at a time. To apply multiple customization profiles, repeat the operation.


- When you reach the *Apply customization profile* step, choose the option *Apply one customization profile to KNIME Hub Executor*. You'll see a table that shows the available execution contexts based on your user type and the previously selected customization profile scope.



Applying a new customization profile with the same name as an existing one to an execution context causes an internal error. Check the "Existing customization profiles name" column to avoid this.

Select one or more execution contexts from the table where you want to apply the customization profile and click *Next*.

Customization Profile
Apply customization profile



What do you want to do next?

☐ Apply one customization profile to KNIME Analytics Platform client
☒ Apply one customization profile to KNIME Hub executor
☐ None of the above. I am fine for now.

Select to which execution contexts apply the customization profile.

Since the customization profile is team-based, it can only be applied to team-scoped execution contexts.

<input type="checkbox"/>	Team name	Execution context name	Execution context scope	Disabled	Existing customization profiles name
<input type="checkbox"/>	Dev Team	4.7.x stable	Team	true	Missing
<input type="checkbox"/>	Dev Team	Moritz local executor	Team	true	Missing
<input type="checkbox"/>	Dev Team	QA-5.1.1	Team	false	Missing
<input type="checkbox"/>	Dev Team	official-nightly	Team	false	Missing
<input type="checkbox"/>	Dev Team	test EC for docs	Team	false	Missing

Warning: When a new customization profile is added to an execution context with existing ones, it is attached to the execution context along with the existing ones.

Filter by teams

Dev Team

☒ Include disabled execution contexts


Previous 1 Next

Figure 18. The team admin can apply a customization profile uploaded to the Dev Team to the team-scoped execution contexts.

- In the *Application Results* step, a table displays the results of the customization profile application operation, along with a legend.



If everything goes well, the green light indicates that your customization profile has been successfully applied to KNIME Business Hub executors for the selected execution contexts.

Customization Profile
Application results




Customization profile application results

The customization profile "Python profile" has been applied with Team scope to the following execution contexts:

	Team name	Execution context name	Execution context scope	Results
	Dev Team	QA-5.1.1	Team	The customization profile has been applied to the execution context
	Dev Team	official-nightly	Team	The customization profile has been applied to the execution context

Result legend

 The customization profile has been applied to the execution context


 It has not been possible to apply the customization profile to the selected execution contexts. See the results column for further details

Figure 19. The team admin has applied the Python profile to two different execution contexts for the same team.

- Error Handling:** The data application handles two types of errors for the application branch.:
 - No customization profiles were uploaded to the KNIME Business Hub.

- b. Attempting to apply a customization profile to an execution context with a customization profile with the same name.

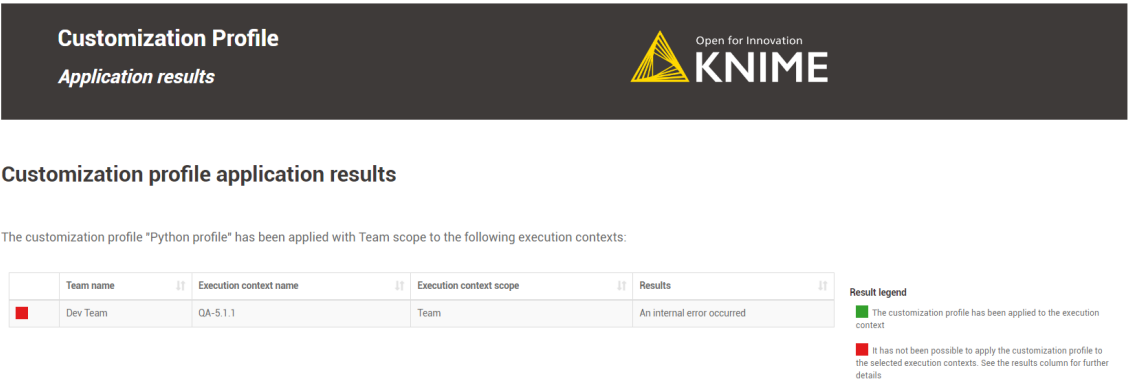


Figure 20. A view of the error.

- c. No customization profile was selected.
- d. When no execution contexts were chosen to apply the customization profile.

If you encounter an error message, click *Next* and select the *Turn back to the choose action menu* option. Repeat the operation, fixing the error.



Figure 21. Turn back to the action menu and repeat the operation.

- 7. Click *Next* to finish the operation. You can return to the *Action Selection* menu to perform additional actions or close the data application directly.

Apply via REST request

In order to apply a customization profile to all executors running in a KNIME Business Hub **execution context**, you will need to send a request to the execution context endpoint.

First you need to get the execution context ID. To do so you can use the following GET request to get a list of all the execution contexts that are owned by a specific team:

```
GET api.<base-url>/execution-contexts/account:team:<team_ID>
```



If you are a global admin you can also GET a list of all the execution contexts available on the Hub instance with the call `GET api.<base-url>/execution-contexts/`.

Now you can apply the new customization profile to the selected execution context.

You will need to obtain the `<profile_ID>` using the following GET request:

1. For global customization profiles:

```
GET api.<base-url>/execution/customization-profiles/hub:global
```

2. For team's customization profiles:

```
GET api.<base-url>/execution/customization-profiles/account:team:<team_ID>
```



Refer to the [above section](#) to find out how to get the `<team_ID>`.

Then you need to update the execution context by using the following PUT request:

```
PUT api.<base-url>/execution-contexts/<execution-context_ID>
```

You need to select **Body** > **raw** > **JSON** and add the following to the request body:

```
{
  "customizationProfiles": [
    "<profile_ID>"
  ]
}
```

This will cause a restart of all executor pods of this execution context, after which the profile will be applied.



At present, this will also terminate all jobs currently running on the executor.

Update a customization profile

It is possible to update customization profiles uploaded to KNIME Business Hub. This can be done even if the customization profile *has already been applied* to the KNIME Business executor via execution context. You can update two features:

1. The customization profile **name**.
2. The customization profile zip file can be overwritten. Refer to the [create a customization profile](#) section to learn how to make one.

Update via Data Application

Users can use the customization profile data app to **update** customization profiles previously uploaded to the KNIME Business Hub.

Depending on the user's role as a Global Admin or a Team Admin, they can access and update specific customization profiles:

- **Global Admin:** can **update** all uploaded customization profiles within the current KNIME Business Hub instance, either team or global-scoped customization profiles.
- **Team Admin:** can only update team-scoped customization profiles from their own teams.

-
1. Learn how to download the data app from Community Hub, upload and deploy it in KNIME Business Hub, and authenticate with your application password in the [Upload a customization profile](#) section.
 2. Select *Update an existing Customization Profile* in the *Action Selection* menu to update a Customization Profile. Choose the desired customization profiles from the table and click *Next*.

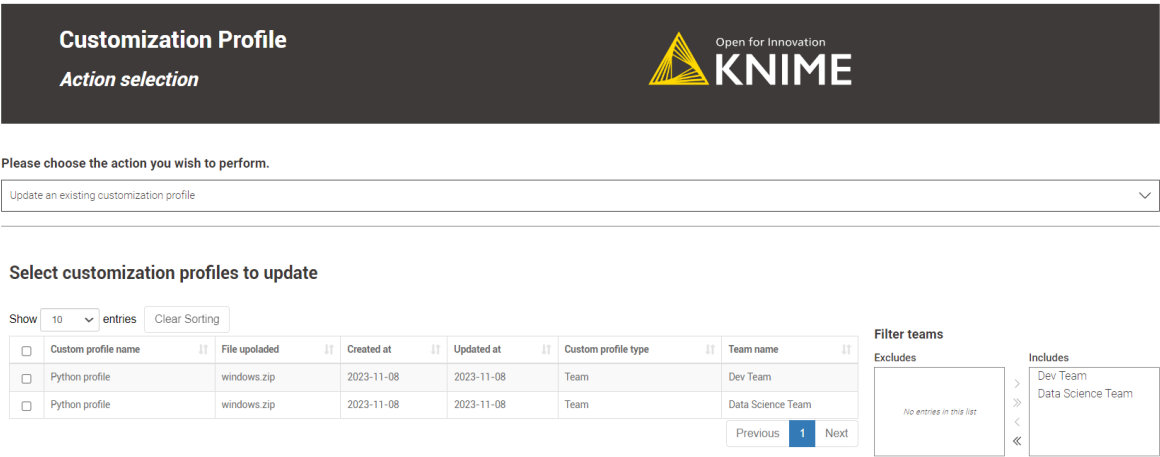


Figure 22. A team admin could select only team-scoped customization profiles.

i You can choose to update multiple customization profiles at the same time.

3. In the *Update Customization Profile* step, you can change the customization profile name by **double-clicking** the corresponding table cell and uploading a new customization profile to overwrite the existing one.

Refer to the "File uploaded" column for the currently assigned customization profile file.

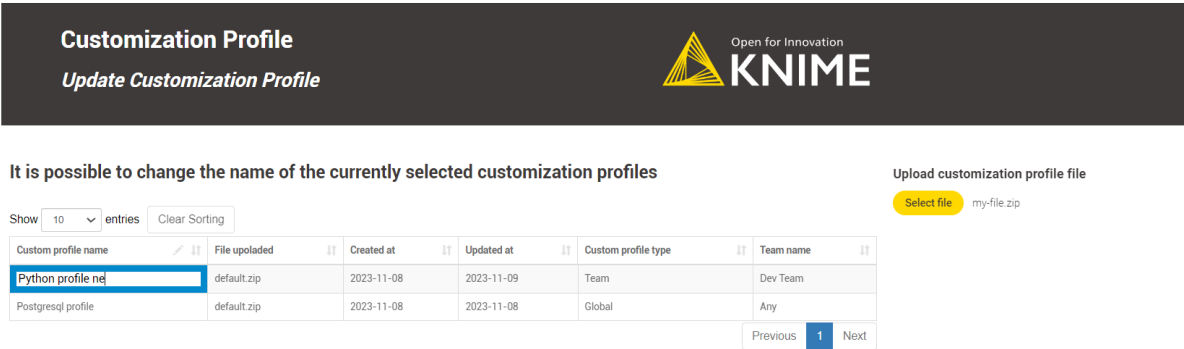


Figure 23. A Global Admin updating two customization profiles at once.


i Updating multiple customization profiles at once is possible, but it's important to note that you can only upload a single customization profile file for all the customization profiles. Although each customization profile can be renamed, they will all share the same customization profile file.

4. In the *Update Results* step, a table displays the results of the customization profile updating operation, along with a legend.

If everything goes well, the green light indicates that your selected customization profiles have been successfully updated with new names and a new customization profile file.


Customization Profile

Update results




Customization profile updated details

The updated customization profiles will be automatically applied to execution contexts or KNIME Analytics Platform clients using it.

	Customization profile name	Uploaded file name	Created at	Updated at	Results
	Python profile data team	default.zip	2023-11-08 GMT	2023-11-09 GMT	The customization profile has been successfully updated.

Result legend

 The customization profile has been successfully updated


 It has not been possible to update the customization profile. See the results column for further details

Figure 24. The outcome of updating two customization profiles.




The updated customization profiles will be automatically applied to execution contexts or KNIME Analytics Platform clients using it.

5. **Error Handling:** The data application handles two types of errors for the update option:
- a. An error will occur if the user does not upload a new customization profile file during the *Update Customization Profile* step.

Customization Profile

Update results



Customization profile updated details

An error occurred: No .zip file was uploaded. Click "Next" to return to the action selection menu.

Figure 25. A view of the error.

- b. No customization profiles were selected in the *Update Customization Profile* step.

If you encounter an error message, click *Next* and select the *Turn back to the choose action menu* option. Repeat the operation, fixing the error.



Figure 26. Turn back to the action menu and repeat the operation.

6. Click *Next* to finish the operation. You can return to the *Action Selection* menu to perform additional actions or close the data application directly.

Update via REST request

Updating a customization profile is like replacing an existing profile with a new file and name.

If you want to update a customization profile via REST request, the process is similar to the [uploading process](#). The only difference is that instead of a POST Request, you need to perform a **PUT Request** and specify the ID of the customization profile.

1. If you have not done so already, on KNIME Business Hub, create an **application password** for the user uploading the profile.
 - a. For global customization profiles, use a global admin user.
 - b. For team's customization profiles, the user can be either a global admin user or a team admin user.
2. To begin, you will need to get the list of all uploaded customization profiles on KNIME Business Hub:

You can obtain the <profile_ID> using the following GET request:

For global customization profiles:

```
GET api.<base-url>/execution/customization-profiles/hub:global
```

For team's customization profiles:

```
GET api.<base-url>/execution/customization-profiles/account:team:<team_ID>
```



Refer to the [above section](#) to learn how to get the <team_ID>.

3. **Updating** a customization profile is similar to uploading it via a REST Request. However, unlike uploading, we only need to provide the name and file of the customization profile for updating. So, we don't need to provide the scope as it remains the same.

- a. Send a PUT request to `https://api.<base-url>/execution/customization-profiles/<customization-profile-ID>` with the following set up:
- b. Authorization: select Basic Auth, using username and password from the created application password
- c. Body: select form-data as request body type and add the following entries:
 - i. Add a new key, set content to "File". Name of the key needs to be "content". The value of the key is the .zip file containing the profile, and the content type is application/zip.
 - ii. Add a new key, set content to "Text". Name of the key needs to be "metadata".
 - A. The value of the key is the same for global and team-scoped customization profiles:

```
{
  "name": "<profile_name>"
}
```

- B. Set the content type for the second key to application/json

- d. When using Postman there is no need to manually adjust the headers

4. Send the request

Detach a customization profile

It's important to know that detaching a customization profile is not the same as deleting it. When you detach a customization profile, it isn't removed from KNIME Business Hub. Instead, it just means that the customization profile won't be applied to the KNIME Analytics Platform or the KNIME Business executors.

However, the customization profile is still available in KNIME Business Hub, so it can be reused again whenever needed.

Detach a customization profile from a KNIME Business Hub executor

Detaching a customization profile applies **only** to those customization profiles applied to a KNIME Business Hub executor via execution context. Separating a customization profile from its execution context is the prerequisite step to deleting a customization profile from a KNIME Business Hub instance.

Detaching a customization profile from an execution context can also be done if the user no longer wants to apply the customization to the executors.

Detach via Data Appliation

The data application allows users to **detach** customization profiles applied to execution contexts in KNIME Business Hub.


For instructions on how to apply a customization profile to an execution context, [refer to this section](#).

The customization profiles available for detachment depend on the user type:

- **Global Admin:** can **detach** all applied customization profiles within the current KNIME Business Hub instance, either team or global-scoped customization profiles.
- **Team Admin:** can only detach team-scoped customization profiles from their own teams.
 1. Learn how to download the data app from Community Hub, upload and deploy it in KNIME Business Hub, and authenticate with your application password in the [Upload a customization profile](#) section.
 2. Select *Detach a customization profile from an execution context* in the *Action Selection* menu to detach a customization profile. Choose the desired customization profiles from the table and click *Next*.

Customization Profile

Action selection



Please choose the action you wish to perform.

Detach a customization profile from an execution context

Select customization profiles to detach from execution contexts

Show 5 entries Clear Sorting

Custom profile name	Custom profile type	File uploaded	Created at	Updated at	Team name
<input checked="" type="checkbox"/> Python profile data team	Team	default.zip	2023-11-08	2023-11-09	Dev Team
<input type="checkbox"/> PostgreSQL profile new	Global	windows.zip	2023-11-08	2023-11-09	Any
<input checked="" type="checkbox"/> Python new profile	Global	windows.zip	2023-11-08	2023-11-08	Any
<input type="checkbox"/> DB configuration	Global	default.zip	2023-11-08	2023-11-08	Any

Previous 1 Next

Filter teams

Excludes

No entries in this list

Includes

Dev Team

Any

Customization profile scope

☒ Global ☒ Team

Figure 27. A Global admin is able to detach either team or global-scoped customization profiles.


i

You can choose to detach multiple customization profiles at the same time.

3. Once you click on *Next*, a table will be displayed showing all the execution contexts where the selected customization profiles have been applied.
4. At this point, you can choose multiple execution contexts to detach the customization profile.
5. Notice that for execution contexts with `global` scope, all teams where they are in use are listed in the column *Team name*. By detaching the customization profile from such an execution context, you will detach it from all the teams where the execution is being applied.

Customization Profile

Detach customization profile



Select from which execution contexts detach the customization profiles

Execution context name	Execution context scope	Custom profile name	Custom profile type	File uploaded	Created at	Team name
<input type="checkbox"/> Common execution context	Global	Python new profile	Global	windows.zip	2023-11-08	The A team, Dev Team, knime_admins team
<input type="checkbox"/> official-nightly	Team	Python profile data team	Team	default.zip	2023-11-08	Dev Team
<input type="checkbox"/> QA-5.1.1	Team	Python profile data team	Team	default.zip	2023-11-08	Dev Team

Previous 1 Next

Warning: Detaching a profile from an execution context doesn't remove it from the Business Hub or other execution contexts using the same profile.

Filter by teams

The A team

Dev Team

knime_admins team

☒ Include disabled execution contexts

Figure 28. Here's an example of a global execution context shared between three teams.

6. After you perform the *Detach Results* step, a table will appear showing you the customization profile detach operation results.

If everything goes well, the green light indicates that your selected customization profiles have been successfully detached from the selected execution contexts.

Customization Profile							
Detach results							
Customization profile detached from the following execution contexts							
Execution context name	Execution context scope	Team name	Custom profile name	Custom profile type	File uploaded	Created at	Results
Common execution context	Global	The A team, Dev Team, knime_admins team	Python new profile	Global	windows.zip	2023-11-08	The customization profile has been detached from the execution context
QA automation	Team	QA automation team	DB configuration	Global	default.zip	2023-11-08	The customization profile has been detached from the execution context

Result legend

■ The customization profile has been detached from the execution context

■ It has not been possible to detach the customization profile from the selected execution contexts. See the results column for further details

Figure 29. The outcome of detaching two customization profiles.



Detaching a customization profile from an execution context doesn't remove it from the Business Hub or other execution contexts using the same profile.

7. **Error Handling:** The data application handles three types of errors for the detach option:
- An error will occur if no customization profiles are applied to any KNIME Business Hub execution context. In this case, you must select the option to **Apply an existing customization profile**.
 - An error is raised if the user:
 - Does not select any customization profile to detach.
 - If there are no execution contexts selected from where to detach the customization profiles.




Figure 30. A view of one of the possible errors.

If you encounter an error message, click *Next* and select the *Turn back to the choose action menu* option. Repeat the operation, fixing the error.



Figure 31. Turn back to the action menu and repeat the operation.

8. Click *Next* to finish the operation. You can return to the *Action Selection* menu to perform additional actions or close the data application directly.

Detach via REST request

You can detach a customization profile from a KNIME Business Hub execution context, which is the inverse of applying it. The steps for detaching a customization profile are similar to applying one.

To detach a customization profile from all executors running in a KNIME Business Hub **execution context**, you must send a request to the execution context endpoint, not including the customization profile ID that you want to detach.

1. If you have not done so already, on KNIME Business Hub, create an **application password** for the user uploading the profile.
 - a. For global customization profiles, use a global admin user.

- b. For team's customization profiles, the user can be either a global admin user or a team admin user.
2. First, you need to get the execution context ID. To do so, you can use the following GET request to get a list of all the execution contexts that are owned by a specific team:

```
GET api.<base-url>/execution-contexts/account:team:<team_ID>
```



If you are a Global Admin you can also GET a list of all the execution contexts available on the Hub instance with the call GET api.<base-url>/execution-contexts/.

3. Retrieve the existing customization profiles in the execution context from the above Get Request response. Look for a key in the JSON body similar to:

```
"customizationProfiles" : [ "<customization-profile-ID_1>",<customization-profile-ID_2>" ]
```

4. Now, you can detach the target customization profile from the selected execution context. To do so, you need to update the execution context by using the following PUT request:

```
PUT api.<base-url>/execution-contexts/<execution-context_ID>
```

5. To detach a customization profile, e.g. <customization-profile-ID_1>, from the target execution context, follow these steps. Select Body > raw > JSON and ensure you do not include the customization profile you wish to remove. Use the syntax of the request body shown below:

```
{
  "customizationProfiles": [
    "<customization-profile-ID_2>"
  ]
}
```

6. If the target execution context has only **one** customization profile attached, you can detach it by doing an empty request.

```
{
  "customizationProfiles": []
}
```

Delete a customization profile

Deleting a customization profile from KNIME Business is a straightforward operation, and you can do it either through a REST request or the data application. Please follow the step-by-step guide below to understand how it works.

Please note that the type of user determines which customization profiles they can delete:

- **Global Admin:** can **delete** all customization profiles within the current KNIME Business Hub instance, either team or global-scoped customization profiles.
- **Team Admin:** can only **delete** team-scoped customization profiles from their own teams.



Deleting a customization profile from KNIME Business Hub requires first **detaching** it from any execution context where it was applied.

Please refer to the [Detach a customization profile from a KNIME Business Hub executor](#) section to understand how to detach a customization profile.

Delete via Data Application

1. Learn how to download the data app from Community Hub, upload and deploy it in KNIME Business Hub, and authenticate with your application password in the [Upload a customization profile](#) section.
2. Select *Delete a customization profile from your Business Hub instance* in the *Action Selection* menu to delete a customization profile. Choose the desired customization profiles from the table and click *Next*.

Customization Profile
Action selection

Please choose the action you wish to perform.

Delete a customization profile from your Business Hub instance

Select customization profiles to delete

Show 5 entries Clear Sorting

<input checked="" type="checkbox"/>	Customization profile name	Customization profile type	File uploaded	Created at	Updated at	Team name
<input checked="" type="checkbox"/>	Python profile data team	Team	default.zip	2023-11-08	2023-11-09	Dev Team

Previous 1 Next

Filter teams

Excludes

Includes

Dev Team

No entries in this list

Figure 32. Team admins can only delete customization profiles from teams they are admin of.



You can choose to delete multiple customization profiles at the same time.

- As deleting is a delicate operation, in the next step, we offer a preview of the selected customization profiles to be deleted to ensure the user is performing the desired action.

Customization Profile
Discard Customization Profiles

Below are the custom profiles selected for deletion

Click "Next" to confirm deletion

Show entries

Customization profile name	Customization profile type	File uploaded	Created at	Updated at	Team name
Python profile data team	Team	default.zip	2023-11-08	2023-11-09	Dev Team

Previous
1
Next

Figure 33. Preview of the selected customization profile.



Clicking *Next* will delete the customization profile directly.

- The *Discard results* section displays a table with a legend indicating the success or failure of deleting a customization profile. Green means success, red means failure, and orange means you need to detach the customization profile from the execution contexts first.

Customization Profile
Discard results

Discarded customization profiles

Customization profile name	Customization profile type	File uploaded	Created at	Updated at	Results
Postgresql profile new	Global	windows.zip	2023-11-08	2023-11-09	The customization profile has been discarded

Previous
1
Next

Result legend


- The customization profile has been discarded
- It has not been possible to discard the customization profile. See the results column for further details
- It has not been possible to discard the customization profile. It is already attached to an execution context

Figure 34. The deletion of a customization profile from KNIME Business Hub was successful.


- Error Handling:** The data application handles two types of errors for the delete option:
 - If the user doesn't select any customization profile.

- b. When deleting a customization profile still attached to an execution context, reference the [Detach a customization profile from a KNIME Business Hub executor](#) section for resolution.




Customization Profile
Discard results



Discarded customization profiles

	Customization profile name	Customization profile type	File uploaded	Created at	Updated at	Results
	Python profile data team	Team	default.zip	2023-11-08	2023-11-09	The customization profile is attached to an execution context. Please return to the menu and detach it before deleting it.

Result legend

-  The customization profile has been discarded
-  It has not been possible to discard the customization profile. See the results column for further details
-  It has not been possible to discard the customization profile. It is already attached to an execution context

Previous
1
Next

Figure 35. A view of one of the possible errors.

If you encounter an error message, click *Next* and select the *Turn back to the choose action menu* option. Repeat the operation, fixing the error.

Customization Profile
Action selection



Please choose the action you wish to perform.

☒ Turn back to the choose action menu
☐ I'm ok, I want to close the data application

Figure 36. Turn back to the action menu and repeat the operation.

6. Click *Next* to finish the operation. You can return to the *Action Selection* menu to perform additional actions or close the data application directly.

Delete via REST request

Deleting a customization profile from KNIME Business Hub is possible via a REST request.

Below are the steps to accomplish this:

1. If you have not done so already, on KNIME Business Hub, create an [application password](#) for the user uploading the profile.

- a. For global customization profiles, use a global admin user.
 - b. For team's customization profiles, the user can be either a global admin user or a team admin user.
2. To start, you need to retrieve the list of the uploaded customization profiles to KNIME Business Hub:

You can obtain the <customization_profile_ID> using the following GET request:

For global customization profiles:

```
GET api.<base-url>/execution/customization-profiles/hub:global
```

For team's customization profiles:

```
GET api.<base-url>/execution/customization-profiles/account:team:<team_ID>
```



Refer to the [above section](#) to find out how to get the <team_ID>.

3. Once you have the <customization-profile-ID> that you want to delete, perform a DELETE Request.
 - a. Send a DELETE request to `https://api.<base-url>/execution/customization-profiles/<customization-profile-ID>` with the following set up:
 - i. Authorization: select Basic Auth, using username and password from the created application password

After a successful deletion, you will receive a 20* status code.

Advanced configuration

This section covers some of the configuration settings that are available for your KNIME Business Hub instance.

The following configurations are available in the KOTS Admin Console and can be changed after the installation and first minimal configuration steps are concluded successfully.

You can access the KOTS Admin Console via the **URL and password** you are provided in the output upon installation.

Configure networking

In the "Networking" section of the KOTS Admin Console you can:

- Deploy an external load balancer for traffic ingress: this feature takes effect only if your cloud provider and kubernetes distribution support automatic load balancer provisioning.
- Enable Transport Layer Security (TLS): the encryption protocol that provides communications security is highly recommended especially for KNIME Business Hub instances deployed in a production environment.



Please, be aware that if TLS is not enabled some HTTPS-only browser's features will not be available. For example, it will not be possible for a user to copy generated **application passwords**.

- Enable advanced ingress configuration: you can customize the ingress proxy behavior, for example configuring the read/send/connect timeouts.

Networking

Configuration for external load balancer and TLS.

☐ **Deploy Load Balancer**

Deploy an external load balancer. Your cloud provider and kubernetes distribution (e.g., AWS EKS or Azure AKS) must support automatic load balancer provisioning in order for this feature to take effect. See [kubernetes documentation](#) for more.

☐ **Enable TLS** Recommended

Enable Transport Layer Security (TLS). This option is highly recommended for any KNIME Business Hub deployed in a production environment. A certificate must be created for all URLs configured above, including a wildcard certificate for `hub.example.com` and `*.hub.example.com`.

☐ **Enable Advanced Ingress Configuration**

Customize the ingress proxy behavior such as read/send/connect timeouts. Changes to these settings require a restart of the `ingress-nginx-controller` pod before taking effect.

☐ **Enable Custom Certificate Authority (CA) Certificate**

Enable the addition of a custom CA certificate. Some organizations may use their own CA, and adding it here will allow certain KNIME Business Hub services to communicate with external resources that require a custom CA.

Configure TLS

If you enable the Transport Layer Security (TLS) you need to have a certificate that is valid for all the **URLs defined during the installation**. We recommend to create a wildcard certificate for `<base-url>` and `*.<base-url>`, e.g. `hub.example.com` and `*.hub.example.com`.

Check *Enable TLS* in the "Networking" section of the KOTS Admin Console.

- **Upload your own certificate:** Select *Upload your own certificate* to be able to upload the certificate files.

You will need an unencrypted private key file and a certificate file that contains the full certificate chain. In the certificate chain the server certificate needs to be the first in the PEM file, followed by the intermediate certificate(s). You usually can get a certificate from your company's IT department or Certificate Authority (CA).

Another possibility, if you have a public domain name, is to use letsencrypt to obtain a certificate.

Both certificates need to be PEM formatted as requested by the `ingress-nginx-controller` (see the relevant documentation [here](#)).

Networking

Configuration for external load balancer and TLS.

☐ Deploy Load Balancer

Deploy an external load balancer. Your cloud provider and kubernetes distribution (e.g., AWS EKS or Azure AKS) must support automatic load balancer provisioning in order for this feature to take effect. See [kubernetes documentation](#) for more.

☒ Enable TLS Recommended

Enable Transport Layer Security (TLS). This option is highly recommended for any KNIME Business Hub deployed in a production environment. A certificate must be created for all URLs configured above, including a wildcard certificate for `hub.example.com` and `*.hub.example.com`.

Certificate Authority

Select the TLS certificate authority to use for KNIME Business Hub.

☒ Upload your own certificate ☐ Existing TLS Secret ☐ AWS ACM Certificate

Private Key File

The private key file should be pem formatted.

Upload a file

 Browse files for Private Key File

Certificate File

The certificate file should be the full certificate chain in pem format.

Upload a file

 Browse files for Certificate File

☐ Enable Advanced Ingress Configuration

Customize the ingress proxy behavior such as read/send/connect timeouts. Changes to these settings require a restart of the `ingress-nginx-controller` pod before taking effect.

☐ Enable Custom Certificate Authority (CA) Certificate

Enable the addition of a custom CA certificate. Some organizations may use their own CA, and adding it here will allow certain KNIME Business Hub services to communicate with external resources that require a custom CA.

- **Existing TLS Secret:** Select *Existing TLS Secret* to specify the name of of an existing Secret of type `kubernetes.io/tls` in the `knime` namespace. It needs to have keys `tls.crt` and `tls.key`, which contain the PEM formatted private key and full chain certificate.

This option is recommended if you have an automatic process that can create and renew `kubernetes.io/tls` Secrets in the cluster, like the [cert-manager project](#).

See [ingress-nginx](#) and [kubernetes documentation on TLS secrets](#) for more details.

Networking

Configuration for external load balancer and TLS.

☐ Deploy Load Balancer

Deploy an external load balancer. Your cloud provider and kubernetes distribution (e.g., AWS EKS or Azure AKS) must support automatic load balancer provisioning in order for this feature to take effect. See [kubernetes documentation](#) for more.

☒ Enable TLS Recommended

Enable Transport Layer Security (TLS). This option is highly recommended for any KNIME Business Hub deployed in a production environment. A certificate must be created for all URLs configured above, including a wildcard certificate for `hub.example.com` and `*.hub.example.com`.

Certificate Authority

Select the TLS certificate authority to use for KNIME Business Hub.

☐ Upload your own certificate ☒ Existing TLS Secret ☐ AWS ACM Certificate

Existing TLS Secret

Specify the name of an existing Secret of type `kubernetes.io/tls` in the `knime` namespace. It needs to have keys `tls.crt` and `tls.key`. See [ingress-nginx](#) and [kubernetes documentation](#) for more. This option is recommended if you have an automatic process that can create and update `kubernetes.io/tls` Secrets in the cluster.

business-hub-crt-secret

☐ Enable Advanced Ingress Configuration

Customize the ingress proxy behavior such as read/send/connect timeouts. Changes to these settings require a restart of the `ingress-nginx-controller` pod before taking effect.

☐ Enable Custom Certificate Authority (CA) Certificate

Enable the addition of a custom CA certificate. Some organizations may use their own CA, and adding it here will allow certain KNIME Business Hub services to communicate with external resources that require a custom CA.

- Select *AWS ACM Certificate* if, instead, you have deployed an AWS Elastic Load Balancer (ELB). In this case you can use AWS Certificate Manager (ACM) and set the certificate as an annotation directly on the loadbalancer. You can find more information in AWS documentation for ACM [here](#).

Once you obtained the certificate Amazon Resource Name (ARN) in the form `arn:aws:acm:<region>:<account-id>:certificate/<certificate-id>`, insert the ARN in the corresponding field as shown in the image below.

Networking

Configuration for external load balancer and TLS.

☐ Deploy Load Balancer

Deploy an external load balancer. Your cloud provider and kubernetes distribution (e.g., AWS EKS or Azure AKS) must support automatic load balancer provisioning in order for this feature to take effect. See [kubernetes documentation](#) for more.

☒ Enable TLS Recommended

Enable Transport Layer Security (TLS). This option is highly recommended for any KNIME Business Hub deployed in a production environment. A certificate must be created for all URLs configured above, including a wildcard certificate for `hub.example.com` and `*.hub.example.com`.

Certificate Authority

Select the TLS certificate authority to use for KNIME Business Hub.

☐ Upload your own certificate ☐ Existing TLS Secret ☒ AWS ACM Certificate

AWS Certificate Manager ARN

When deploying in an existing AWS EKS cluster, and selecting the loadbalancer option above, you can attach a certificate from the AWS Certificate Manager to the loadbalancer here. Set the certificate ARN below in the format: `arn:aws:acm:<region>:<account-id>:certificate/<certificate-id>`

`arn:aws:acm:<region>:<account-id>:certificate/<certificate-id>`

☐ Enable Advanced Ingress Configuration

Customize the ingress proxy behavior such as read/send/connect timeouts. Changes to these settings require a restart of the `ingress-nginx-controller` pod before taking effect.

☐ Enable Custom Certificate Authority (CA) Certificate

Enable the addition of a custom CA certificate. Some organizations may use their own CA, and adding it here will allow certain KNIME Business Hub services to communicate with external resources that require a custom CA.

Configure Browser Security

In the "Browser Security" section of the KOTS Admin Console you can:

- Specify a custom Content Security Policy for Data App execution. It may be necessary to override the default if you are using custom JavaScript views that load external resources. The default works for all standard KNIME views. For more information about how to write the CSP statement, please refer to this [resource](#).
- Configure the X-Frame-Options header being set by webapps. This header is used to avoid click-jacking attacks, by ensuring that the sites content is not embedded into other sites. See [here](#) for more information.

Browser Security

This section contains settings for browser security.

☒ Enable Content Security Policy for Data Apps

Enabling this option allows you to set a custom Content Security Policy for Data Apps below. If disabled, no Content Security Policy header is set.

Content Security Policy for Data Apps

Specifies a custom Content Security Policy for Data App execution. It may be necessary to override the default if you are using BIRT report generators or custom JavaScript views that load external resources. The default works for all standard KNIME views. For more information about how to write the CSP statement, please refer to this [resource](#).

```
default-src 'self'; script-src 'unsafe-inline' 'unsafe-eval' 'self'; style-src 'unsafe-inline' 'self'; img-src 'self' data:; font-src 'self'
```

Default value: `default-src 'self'; script-src 'unsafe-inline' 'unsafe-eval' 'self'; style-src 'unsafe-inline' 'self'; img-src 'self' data:; font-src 'self' data;;`

X-Frame-Options Header

Sets the `X-Frame-Options` header to the selected option, or doesn't set the header if `none` is selected. This header is used to avoid click-jacking attacks, by ensuring that the sites content is not embedded into other sites. See [here](#) for more information.

☒ SAMEORIGIN ☐ DENY ☐ none

Node affinity

Node affinity makes it possible to ensure that cluster resources intended for a specific task, e.g. execution resources, run on a specific set of nodes. There are two roles that each pod is grouped into: `core` and `execution`. Pods in the `core` group consist of KNIME Business Hub control plane resources, and pods in the `execution` group relate to execution contexts.

In order to use the node affinity feature in your KNIME Hub cluster, you can apply one or both of the following labels to nodes within your cluster:

- `hub.knime.com/role=core`
- `hub.knime.com/role=execution`

To label a node, you can execute the following command (where `<node-name>` is the name of the node you want to label):

```
kubectl label node <node-name> hub.knime.com/role=core
```



For more information about labeling nodes, see the [Kubernetes documentation](#).

Pods will have to be restarted in order to be rescheduled onto labeled nodes. You can use the following example commands to restart the pods in a live cluster:

- `kubectl rollout restart deployment -n istio-system`
- `kubectl rollout restart deployment -n hub`
- `kubectl rollout restart deployment -n knime`
- `kubectl delete pods --all --namespace hub-execution`



This command will restart all execution context pods.

There are a few things to note about the behavior of this feature:

- Node affinity uses a "best effort" approach to pod scheduling.
 - If one or both of the `hub.knime.com/role` labels are applied, cluster resources will attempt to be scheduled onto the nodes based on their role.
 - If no nodes have a `hub.knime.com/role` label, pods will be scheduled onto any available node.
 - If labeled nodes reach capacity, pods will be scheduled onto any available node.

- If a labeled node is shut down, pods will be rescheduled onto other nodes in the cluster with a preference towards using nodes that have a matching label.
- Node affinity for KNIME Business Hub uses the `preferredDuringSchedulingIgnoredDuringExecution` approach (see the [Kubernetes documentation](#) for more details).
- It is possible to use only one of the labels above, e.g. labeling nodes for the `execution` role but not specifying any node labels for the `core` role.

Create a collection

It is possible to create collections on your KNIME Business Hub instance.

KNIME Collections on KNIME Hub allow upskilling users by providing selected workflows, nodes, and links about a specific, common topic.

One example of a collection can be found on KNIME Community Hub [here](#).



This is a feature of KNIME Business Hub - Enterprise edition.

In order to create a new collection page you need to be a global admin of your KNIME Business Hub instance.

The creation of a collection is possible via REST API, and a description of the different configurations can be found in your KNIME Business Hub API doc at the following URL:

```
api.<base-url>/api-doc/?service=catalog-service#/Collections
```

e.g. `api.hub.example.com/api-doc/?service=catalog-service#/Collections`.

In order to create a collection the items (i.e. workflows and nodes) that are collected need to be stored and accessible on the same KNIME Business Hub instance where collection is created.

To create the collection you will need then to build a json file with the schema that is available in the API doc in the Collections section, under the POST `/collections` request description.

The following is an example that would allow you to build a collection, similar to the one available on KNIME Community Hub [here](#).

In the first section you can for example set up a title, a description, a so-called hero, which is the banner image at the top right of the example collection page, and tags:

```
{
  "title": "Spreadsheet Automation",
  "description": "On this page you will find everything to get started with spreadsheet automation in KNIME",
  "ownerAccountId": "account:user:<global-admin-user-id>",
  "hero": {
    "title": "New to KNIME?",
    "description": "Get started with <strong>KNIME Analytics Platform</strong> to import all the examples and nodes you need for spreadsheet automation right now!",
    "actionTitle": "Download",
    "actionLink": "http://download-me-quick.999.com"
  },
  "tags": [
    "Excel",
    "XLS"
  ],
}
```

Next you can add different sections and subsections, each with a title and a description, choose a layout, and select the `itemType` such as *Space*, *Component*, *Workflow*, *Node*, *Extension*, or *Collection*. For each of these items you will need to provide the `id` under which they are registered in your Business Hub installation.

The `id` for workflows, spaces, components, and collections can be build by taking the last part of their URL, after the `~`, and adding a `*` at the beginning. For example, the following workflow on the KNIME Community Hub has URL `https://hub.knime.com/-/spaces/-/latest/~1DCip3Jbxp7BWz0f/` so its `id` would be `*1DCip3Jbxp7BWz0f`. The `id` for node and extensions instead needs to be retrieved with a REST call, for example to the `search` endpoint of your KNIME Business Hub instance.

```

"sections": [
  {
    "title": "Workflow examples",
    "description": "Some subtitle text here. Can have <strong>bold format</strong>",
    "iconType": "Workflow",
    "subsections": [
      {
        "title": "How to do basic spreadsheet tasks in KNIME",
        "description": "Some examples on how to do common things",
        "layout": "SingleColumn",
        "numberOfTeaseredItems": 2,
        "items": [
          {
            "title": "Click Here!",
            "itemType": "Link",
            "absoluteUrl": "https://knime.com"
          },
          {
            "id": "*SJW5zSkh1R3T-DB5",
            "itemType": "Space"
          },
          {
            "id": "*vpE_LTbA0n96Z0g9",
            "itemType": "Component"
          },
          {
            "id": "*MvnABULB035AQcAR",
            "itemType": "Workflow"
          },
          {
            "showDnD": true,
            "id": "*yiAvNQVn0sVwCwYo",
            "itemType": "Node"
          },
          {
            "id": "*bjR3r1yW0znPIEXS",
            "itemType": "Extension"
          },
          {
            "id": "*QY7INTkMW6iDj7uC",
            "itemType": "Collection"
          }
        ]
      }
    ]
  }
]

```

Administrator workflows

The workflows described in this section of the documentation aim to support KNIME Business Hub administrators, or heavy KNIME Business Hub users, to clean up, monitor and better administrate their Business Hub instance.

The functionalities provided are a time saver for monitoring or administrating KNIME Business Hub, eliminating the need for manual work and centralizing information from various applications.

Workflows overview

The user can access the workflows on the KNIME Community Hub in a [public space](#) owned by KNIME. Additionally, the user can find them on a dedicated [collection page](#). To use them, download the workflows from the Community Hub and upload them into an existing team space in your KNIME Business Hub installation.

Business Hub has three types of user roles (global admin, team admin, and team member). All the users with access to the “Admin Space” can run the workflows. The user’s role defines their allowed actions when running the different workflows.

The workflows can be run as data applications on-demand or directly scheduled using the Business Hub UI. First, you must deploy the workflows as a Data app or Schedule.

Below the list of workflows within the “Admin Space”, click on them to read further details:

- [Discard Failed Jobs](#)
- [List All Jobs](#)
- [Delete Old Versions](#)
- [Scheduled Workflows Kick-Off Times](#)
- [Count Workflows Running Per Day](#)
- [Workflows' Run Time](#)
- [Monitor Users' Usage](#)

Requirements and prerequisites

Requirements

- The user needs to exist and be at least a Team member (no matter the user's role) where the "Admin Space" is located
- Also, the user needs at least view access to the "Admin Space."

Prerequisites

- The user should be familiar with new concepts on the Hub. See the KNIME Business Hub User Guide.
- The user needs to create an application password specific to her account on KNIME Business Hub that most applications will use.

Discard Failed Jobs

Overview

This workflow aims to keep "clean" the KNIME Business Hub installation by discarding failed jobs from any kind of execution run by the KNIME Business Hub users.

Workflow specs

Without applying any time range, the workflow discards all failed jobs for the following execution types: **ad-hoc executions**, **triggers**, **data apps**, **schedules** or **shared deployments**.

We consider as failed jobs all those with any of the following states after execution: "Execution_Failed", "Execution_Failed_With_Content", "Execution_Canceled", or "Vanished".

The failed jobs a user can discard depend on the role of the user running the workflow:

Global admin: can discard all failed jobs in any team and space from any execution type.

Team admin: can discard all failed jobs of the teams of which it is an admin from any execution type.

Team member: can discard only self-generated failed jobs from any execution type, any team and space of which it is a member (no matter the **user's right's** on space items). It also includes deployments shared with the user from teams where the user is not a member.

Deployment configuration

This workflow can be deployed as a **data app** or **schedule** deployment.

In both cases, you can provide the following information to deploy the workflow:

1. Hub URL: The URL of your KNIME Business Hub instance, e.g. "https://my-business-hub.com".
2. Application Password ID: User-associated application password ID.
3. Application Password: User-associated application password.

If you want to know how to create an application password, follow **these steps**.

Data app


After deploying the workflow as a Data App, you can run it. To do so, follow **these instructions**.

Below are the steps for the data app:

1. **Business Hub connection:** you need to connect to the KNIME Business Hub instance through the previously generated application password.
2. **Select Job State:** it is possible to customize which types of failed jobs you want to discard. Max 4 job states should be available: "Execution_Failed", "Execution_Failed_With_Content", "Execution_Canceled", or "Vanished".
3. **Discard Results:** A table with the discard jobs results is displayed by default. There is also the possibility to see an illustrated version of the table by selecting the "Switch to charts" option.





Discard Failed Jobs

Discard Results





Discard Jobs Results

☐ Switch to charts

	Execution name	Execution type	Workflow path	Space	Results
	Scheduled job 2023-04-07 03.58.00	schedule	/Users/Continuous Deployment Setup/Public test environment/Faling deployments/failing_workflow	Public test environment	The job has been discarded
	Scheduled job 2023-04-10 03.58.00	schedule	/Users/Continuous Deployment Setup/Public test environment/Faling deployments/failing_workflow	Public test environment	The job has been discarded
	Scheduled job 2023-04-10 23.58.00	schedule	/Users/Continuous Deployment Setup/Public test environment/Faling deployments/failing_workflow	Public test environment	The job has been discarded
	Scheduled job 2023-04-07	schedule	/Users/Continuous Deployment Setup/Public test environment/Faling	Public test environment	The job has been

Result legend

 The job has been discarded

 It has not been possible to discard the job. See the results column for further details


Team filter

☒ Continuous Deployment Setup

← Back

Discard Failed Jobs

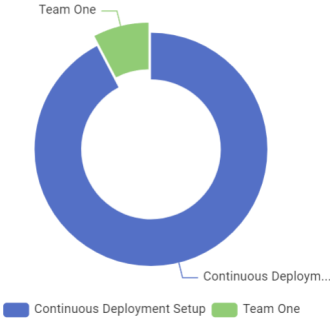
Discard Results



Discard Jobs Results

☒ Switch to charts

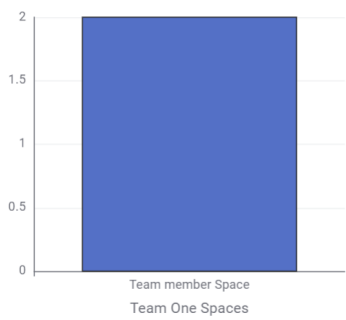
Count of Jobs by team



Team One

Continuous Deployment Setup

Count of jobs for Team One spaces



Team member Space

Team One Spaces

Team filter

☒ Team One

☒ Continuous Deployment Setup

← Back

Schedule

- 1. Define when the workflow should be executed through the schedule options. For more information, follow [this guide](#).
- 2. Ensure the Schedule Deployment is active (*Status* column).
- 3. In the Team Deployments page or the workflow page, Deployments section, you can

check the number of executions and their status.

The screenshot shows the 'Deployments' section of the KNIME Business Hub Admin Guide. It features a table with columns: Type of deployment, Name, Workflow, Version, Date of creation, Execution context, and Status. The table lists two rows of deployment jobs, both with a status of 'Execution finished'. A sidebar on the right contains links to Overview, Used extensions & nodes, Deployments, and Ad hoc executions.

Type of deployment	Name	Workflow	Version	Date of creation	Execution context	Status
Schedule	Schedule - Dis...	...rd Failed Jobs/	Testing links in...	Apr. 12, 2023, ...	4.7.1 Context	Execution finished
Schedule	Schedule - Dis...	...rd Failed Jobs/	Testing links in...	Apr. 12, 2023, ...	4.7.1 Context	Execution finished

List All Jobs

Overview

The workflow scope is to list all the jobs (excluding the failed ones) accessible to the user.

Additionally, the user can use the workflow to easily pinpoint any irregularities in the workflow's execution or sort jobs that have been in an execution state for an extended period.

Furthermore, it is possible to select the jobs and discard them.

Workflow specs

Without applying any time range, the workflow lists all jobs for the following execution types: **ad-hoc executions**, **triggers**, **data apps**, **schedules** or **shared deployments**.

We only consider not failed jobs. This means we exclude any jobs with states such as: "Execution_Failed", "Execution_Failed_With_Content", "Execution_Canceled", or "Vanished".

The job information a user could retrieve depends on the user role running the deployed workflow:

Global admin: can recover all workflow jobs in any team and space.

Team admin: can recover all workflow jobs within the team where it is an admin.

Team member: can recover workflow jobs from any team and space where is a member (no matter the **user's rights** on space items). It also includes shared deployments from teams where the user is not a member.

The jobs a user can discard depend on the role of the user running the deployed workflow:

Global admin: can discard all jobs in any team and space from any execution type.

Team admin: can discard all jobs of the teams of which it is an admin from any execution type.

Team member: can discard only self-generated jobs from any execution type, team and space of which it is a member (no matter the user's right on space items). It also includes deployments shared with the user from teams where the user is not a member.

Deployment configuration

This workflow can be deployed as a **data app**.

You can provide the following information to deploy the workflow:

1. **Hub URL:** The URL of your KNIME Business Hub instance, e.g. "https://my-business-hub.com".
2. **Application Password ID:** User-associated application password ID.
3. **Application Password:** User-associated application password.

If you want to know how to create an application password, follow **these steps**.

Data app

After deploying the workflow as a Data App, you can run it. To do so, follow the **instructions**.

Below are the steps for the data app:

1. **Business Hub connection:** you need to connect to the KNIME Business Hub instance through the previously generated application password.
2. **Explore and select Jobs:** this feature displays a table of all available jobs for the user. Each job is listed with its name, state, runtime information, and corresponding workflow deployment.

Two exceptions related to the job's deployment information:

- a. When ad-hoc executions generate jobs, the deployment name is not available.
- b. If the workflow is executed when a deployment that has generated a job is not available anymore in the KNIME Business Hub (because it has been discarded or, in the case of schedules, the deployment has ended), retrieving the deployment


information is not possible. It displays a message in the "Deployment name" column: *"This deployment is no longer available."*

On the right side, the user can find three filters:

- Run time anomalies: it detects outliers using the **Numeric Outliers** node. As a user, you can focus on "Outliers", which will help you identify jobs that take significantly longer or shorter to execute than others within the same workflow deployment.
 - Job state: It allows filtering by specific job states. The available job states will be shown based on the currently listed jobs.
 - Team: the user can filter by team.
3. **Discard Results:** A table with the discard jobs results is displayed by default. There is also the possibility to see an illustrated version of the table by selecting the "Switch to charts" option.

List All Jobs

Explore and select Jobs



Explore jobs by sorting and using the filters, then choose which ones to discard

Show 5 entries

<input type="checkbox"/>	Job name	Job state	Run time (seconds)	Workflow	Started execution at	Finished execution at	Deployment type	Deployment name
<input type="checkbox"/>	KNIME_project39 2023-05-08 12:16:39	NOT_EXECUTABLE	0.01	KNIME_project39	12:16:39 - 2023-05-08	12:16:39 - 2023-05-08	ad-hoc execution	?
<input type="checkbox"/>	Scheduled Workflows kick-off time - release version test 2023-05-10 09:29:33	INTERACTION_REQUIRED	16192.31	Scheduled Workflows kick-off time	09:29:39 - 2023-05-10	13:59:32 - 2023-05-10	data-app	Scheduled Workflows kick-off time release version test
<input type="checkbox"/>	List All Jobs 2023-05-08 14:16:04	INTERACTION_REQUIRED	198996.36	List All Jobs	14:16:10 - 2023-05-08	21:32:47 - 2023-05-10	ad-hoc execution	?
<input type="checkbox"/>	List All Jobs	INTERACTION_REQUIRED	199009.48	List All Jobs	14:15:57 -	21:32:47 -	ad-hoc	?

Filters


Run time anomalies
☒ Outliers
☒ Inliers

Job State
☐ EXECUTION_FINISHED
☒ NOT_EXECUTABLE
☒ INTERACTION_REQUIRED

Team
☒ Dev Team




← Back
Cancel
Next

List All Jobs
Discard Results





Discard Jobs Results

☐ Switch to charts

	Job name	Job state	Run time (seconds)	Workflow	Started execution at	Finished execution at
	test123 2023-05-02 09.35.43	EXECUTION_FINISHED	0.82	Housing Data App	2023-05-02	2023-05-02
	04_Example_for_Pie_Chart 2023-05-08 14.21.35	EXECUTION_FINISHED	0.03	04_Example_for_Pie_Chart	2023-05-08	2023-05-08
	04_Example_for_Pie_Chart 2023-05-08 14.25.51	EXECUTION_FINISHED	6.28	04_Example_for_Pie_Chart	2023-05-08	2023-05-08


Result legend

 The job has been discarded

 It has not been possible to discard the job. See the results column for further details

← Back

List All Jobs
Discard Results



Discard Jobs Results

☒ Switch to charts

Count of jobs by team

Workflow...

Dev Team 2 (66.67%)

Dev Team

Dev Team Workflow Building

Count of jobs for Dev Team spaces

Initial Space Daniel

Dev Team Spaces

← Back

Delete Old Versions

Overview

The workflow aims to delete old item versions that aren't involved in any deployment. The sought effect of this operation is to avoid a massive proliferation of item versions within the

Business Hub installation, impacting disk space.

Workflow specs

The workflow deletes all item versions older than the specified number of days, e.g. older than seven days.

The deletion will only be applied to the selected teams and spaces using the workflow as a data app or schedule deployment.

In cases where a rule applies to all versions in a space, the latest version will not be deleted, even if it is affected by the rule. Additionally, item versions that are involved in deployments are exempt from the rule.

Depending on your user role on the KNIME Business Hub, if you execute this workflow, you will have the following permissions:

Global admin: Can delete every item version not used in any deployment from any team on the KNIME Business Hub instance.

Team admin: Can only delete the item versions not involved in any deployment in the team where it's an admin.

Team member: Can only delete the item versions not involved in any deployment in the teams where it's a member. The Team member must have "Edit" permissions for the targeted spaces to perform the version deletion.

Deployment configuration

This workflow can be deployed as a **data app** or **schedule** deployment.

You need to provide the following information to deploy the workflow as a data app or to schedule it:

1. Hub URL: The URL of your KNIME Business Hub instance, e.g. "https://my-business-hub.com".
2. Application Password ID: User-associated application password ID.
3. Application Password: User-associated application password.
4. Team name: Use "*" to select all and "/" to choose a subset: "Team1/Team2".
5. Space name: Use "*" to select all and "/" to choose a subset: "Space1/Space2".
6. The number of days: All item versions older than that will be deleted (default is 30

days).

To create an application password, follow [these steps](#).


Data app

After deploying the workflow as a data app, you can run it. To do so, follow these [instructions](#). Below are listed the steps for the data app::

1. **Business Hub connection:** You need to connect to the KNIME Business Hub instance through the application password generated previously.
2. **Team and Space selection:** Select the teams and associated spaces from which the workflow should delete the old versions.
3. **Define version deletion rule:** Here, you can set a version deletion rule. All versions older than the given number of days will be deleted.
4. **Deletion result:** A table showing the deletion result with the version information will appear by default.

Switching to a chart view is possible using the “Switch to charts” option.

Delete Old Item Versions
Deletion results



Deleted versions

Team

Workflow Building

☐ Switch to charts





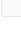
Space

☒ Migration Workflow Updated with Item Versioning
☐ Trigger Demo 2
☒ Trigger Demo

Show

5

 entries


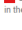
	Item path	Creation date	Version details	Results
	/Users/Workflow Building/Migration Workflow Updated with Item Versioning/KNIME Server to KNIME Hub Migration	2023-07-04 UTC	Version title: latest - Version index: 1 - Version description :	Item version deleted
	/Users/Workflow Building/Trigger Demo/Training Workflow	2023-04-03 UTC	Version title: Version created by Trigger Deployment (migrated from space version 35) - Version index: 35 - Version description :	Item version deleted
	/Users/Workflow Building/Trigger Demo/Training Workflow	2023-04-03 UTC	Version title: Version created by Trigger Deployment (migrated from space version 34) - Version index: 34 - Version description :	Item version deleted
	/Users/Workflow Building/Trigger Demo/Training Workflow	2023-04-03 UTC	Version title: Version created by Trigger Deployment - Version index: 33 - Version description :	Item version deleted
	/Users/Workflow Building/Trigger Demo/Training Workflow	2023-04-03 UTC	Version title: Version created by Trigger Deployment - Version index: 32 - Version description :	Item version deleted

Previous

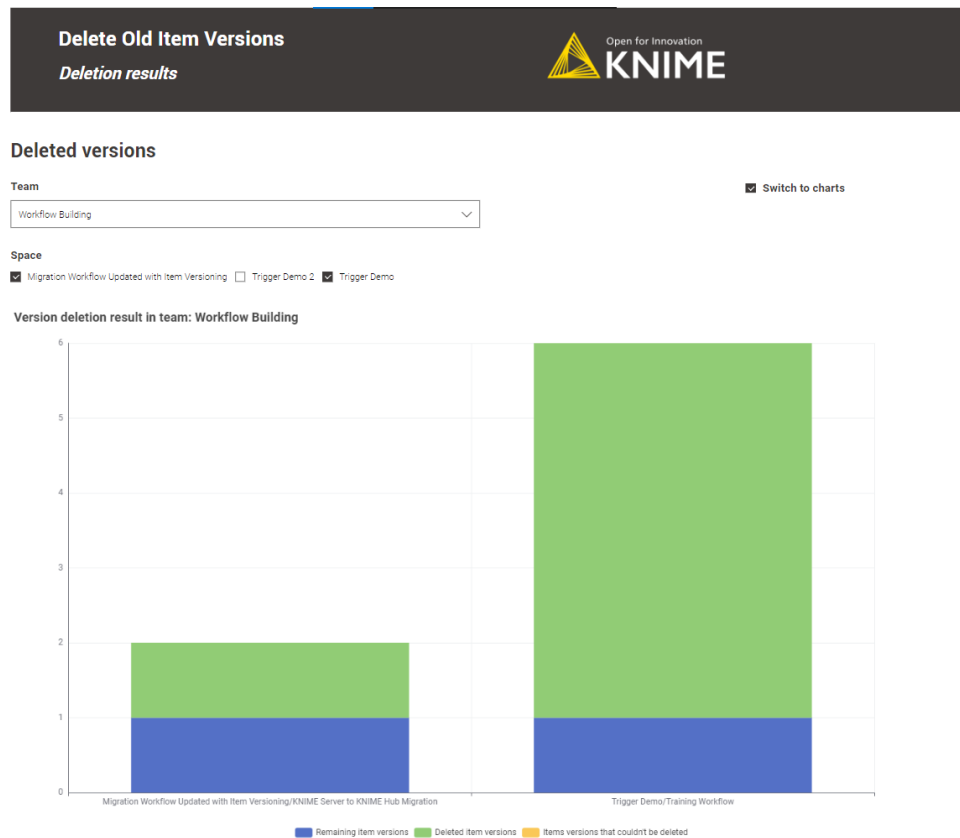
1

2

Next

Result legend
 The item version has been deleted successfully.
 Something went wrong. See the *Results* column in the table for further details.

← Back



Schedule

1. Define when the workflow should be executed through the schedule options. For more information, follow [this guide](#).
2. Ensure the Schedule Deployment is active (*Status* column).
3. In the Team Deployments page or the workflow page, Deployments section, you can check the number of executions and their status.

Scheduled Workflows Kick-Off Times

Overview

The workflow aims to prevent scheduled deployments from overlapping and, consequently, helps to avoid overloading KNIME Business Hub executors.

Running on demand as a data app offers the user a visual overview of the scheduled deployments within the given number of days by team and execution context.

Workflow specs

The workflow visualizes all the scheduled deployments by KNIME Business Hub execution context and team in the following days. Notice that the maximum number of days is set up to seven.

The workflow shows only those schedules whose deployment is active and has not been disabled via the KNIME Analytics Platform and those with a valid next execution date.

Depending on the user role, you can view the following schedules:

Global admin: monitor the scheduled workflows of any team and any KNIME Business Hub execution context.

Team admin: monitor the scheduled workflows for any KNIME Business Hub execution context in the team.

Team member: monitor the scheduled workflows for any KNIME Business Hub execution context in the team.

Deployment configuration

This workflow can be deployed as a **data app**.

You can provide the following information to deploy the workflow as a data app:

1. **Hub URL:** The URL of your KNIME Business Hub instance, e.g. "https://my-business-hub.com".
2. **Application Password ID:** User-associated application password ID.
3. **Application Password:** User-associated application password.

To create an application password, follow **these steps**.


Data app

After deploying the workflow as a data app, you can run it. To do so, follow these **instructions**. Below are the steps for the data app:

1. **Business Hub connection:** You need to connect to the KNIME Business Hub instance through the application password generated previously.
2. **Scheduled Workflows Monitoring:** By default, it visually represents the scheduled workflows in the coming days.

The number of scheduled workflows per day and hour and the number of schedules by team and execution context are represented.

Diving deep into each scheduled deployment is possible by switching to a “Table” mode. Use the “Switch to schedules details (Table)” option.



Scheduled Workflow Kickoff Times
Scheduled Workflows Monitoring

Open for Innovation
KNIME

Days forward: 7

☒ Switch to schedules details (Table)

Schedule name	Created by	Workflow path	Running each	Running times filter	Running days of week filter	Running days filter	Running months filter
ETL workflow	knime	/Users/Team One/Team member Space/failing_workflow	6 hours	[{"start": "00:00", "end": "23:59:59"}]	[MONDAY, TUESDAY, WEDNESDAY, THURSDAY, FRIDAY, SATURDAY, SUNDAY]	[1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31]	[JANUARY, FEBRUARY, MARCH, APRIL, MAY, JUNE, JULY, AUGUST, SEPTEMBER, OCTOBER, NOVEMBER, DECEMBER]
knime_user_failing_workflow	knime	/Users/Continuous Deployment Setup/Public test environment/failing_deployments/failing_workflow	2 hours	[{"start": "00:00", "end": "23:59:59"}]	[MONDAY, TUESDAY, WEDNESDAY, THURSDAY, FRIDAY, SATURDAY, SUNDAY]	[1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31]	[JANUARY, FEBRUARY, MARCH, APRIL, MAY, JUNE, JULY, AUGUST, SEPTEMBER, OCTOBER, NOVEMBER, DECEMBER]
Schedule - ETL Workflow	knime	/Users/Continuous Deployment Setup/Team admin space/failing_workflow	3 hours	[{"start": "00:00", "end": "23:59:59"}]	[MONDAY, TUESDAY, WEDNESDAY, THURSDAY, FRIDAY, SATURDAY, SUNDAY]	[1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31]	[JANUARY, FEBRUARY, MARCH, APRIL, MAY, JUNE, JULY, AUGUST, SEPTEMBER, OCTOBER, NOVEMBER, DECEMBER]
Schedule - Workflow for ETL Basics Operations	knime	/Users/Continuous Deployment Setup/Public test	2 hours	[{"start": "00:00", "end": "23:59:59"}]	[MONDAY, TUESDAY, WEDNESDAY, THURSDAY, FRIDAY, SATURDAY, SUNDAY]	[1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31]	[JANUARY, FEBRUARY, MARCH, APRIL, MAY, JUNE, JULY, AUGUST, SEPTEMBER, OCTOBER, NOVEMBER, DECEMBER]

← Back

Count Workflows Running Per Day

Overview

After a while, the KNIME Business Hub instance automatically discards already executed workflow jobs. KNIME Business Hub administrators can define this time window until a specific period, but it is not infinite.

The workflow scope is to go beyond this automatic discard timeline creating a historical file with all the workflow execution information for further consultation.

Workflow specs

The workflow saves in the root level of an existing team space a CSV file of the jobs information of any state ("Executed", "Failed", etc.) and from any execution type (Deployments and Ad-hoc executions).

The workflow job information a user could retrieve depends on the user role running the deployed workflow:

Global admin: can recover all workflow jobs in any team and space.

Team admin: can recover all workflow jobs within the team where it is an admin.

Team member: can recover workflow jobs from any team and space where is a member (no matter the **user's right's** on space items). It also includes shared deployments from teams where the user is not a member.

The team space where the CSV file can be saved depends on the user role running the deployed workflow:

Global admin: can save workflow job information in any team space.

Team admin: can save workflow job information in every team space where it is an admin.

Team member: can save workflow job information in every team space where it is a member and it has edit permissions.

The user must define the number of backward days you want to retrieve the workflow job information, e.g. last five days.

If you repeatedly target the same team space executing this workflow, the workflow will append the new information to the master file.

Deployment configuration

This workflow can be deployed as a **data app** or **schedule** deployment.

You can provide the following information to deploy the workflow as a data app or schedule it:

1. Hub URL: The URL of your KNIME Business Hub instance, e.g. "https://my-business-hub.com".
2. Application Password ID: User-associated application password ID.
3. Application Password: User-associated application password.
4. Last n days filter: Save job information of all jobs from the last n days.
5. Select the destination Team: The team in which the job information will be saved.
6. Select the destination Space: The space in which the job information will be saved.

To create an application password, follow **these steps**.

Data app

After deploying the workflow as a data app, you can run it. To do so, follow these **instructions**. Below are the steps for the data app:

1. **Business Hub connection:** You need to connect to the KNIME Business Hub instance through the previously generated application password.
2. **Set rule and file destination:** Select the last n days to save the workflow job information. And choose the destination team and space where this file will be saved.
3. **File location link and preview:** You can find the link to the destination space and a preview of the saved job information.

Workflows running per day

File location link and preview



You can find the file with the workflows running per day here: [User space](#)

Workflows running per day preview

Execution name	Execution type	Team	Hub path	Job state	Timestamp
Scheduled job 2023-04-18 07.34.00	schedule	Continuous Deployment Setup	/Users/Continuous Deployment Setup/Public test environment/Scheduled deployments/flyweight-execution-workflow	EXECUTION_FINISHED	2023-04-18
Scheduled job 2023-04-21 11.33.00	schedule	Continuous Deployment Setup	/Users/Continuous Deployment Setup/Public test environment/Scheduled deployments/middleweight-execution-workflow	EXECUTION_FINISHED	2023-04-21
Scheduled job 2023-04-17 07.34.00	schedule	Continuous Deployment Setup	/Users/Continuous Deployment Setup/Public test environment/Scheduled deployments/flyweight-execution-workflow	EXECUTION_FINISHED	2023-04-17
Scheduled job 2023-04-18 15.34.00	schedule	Continuous Deployment Setup	/Users/Continuous Deployment Setup/Public test environment/Scheduled deployments/flyweight-execution-workflow	EXECUTION_FINISHED	2023-04-18
Scheduled job 2023-04-18 21.33.00	schedule	Continuous Deployment Setup	/Users/Continuous Deployment Setup/Public test environment/Scheduled deployments/middleweight-execution-workflow	EXECUTION_FINISHED	2023-04-18

Job state filter

- ☒ EXECUTION_FINISHED
- ☐ EXECUTION_FAILED
- ☐ INTERACTION_REQUIRED

[← Back](#)

Schedule

1. Define when the workflow should be executed through the schedule options. For more information, follow [this guide](#).
It is recommended to adjust the last n days' field to the chosen schedule option, e.g. if the workflow is scheduled every seven days, set the last n days to seven to avoid duplicate information in the CSV file.
2. Ensure the Schedule Deployment is active (*Status* column).
3. In the Team Deployments page or the workflow page, Deployments section, you can check the number of executions and their status.

Workflows' Run Time

Overview

This workflow aims to list all the workflows sorted by the average execution time (the greater the first) for the last n days.

Workflow specs

We calculate the run time by solely considering the deployed schedules, as monitoring the

execution time of other deployment types is not meaningful. For example, a user could play around one hour or five minutes with a data app.

The average workflow execution time and the standard deviation are calculated for the given period.

Be aware that the number of days a workflow deployment's job is available depends on the KNIME Business Hub instance's configuration, e.g. a common value would be seven days. After that period, all jobs would be discarded automatically.

The average jobs run time, which a user can monitor, depends on the role of the user running the deployed workflow:

Global admin: can monitor the run time of all deployed schedules in any Team and Space.

Team admin: can monitor the run time of all deployed schedules within the Team where it is an admin.

Team member: can monitor the run time of all deployed schedules within the Team where it is a member.

Deployment configuration

This workflow is designed to be deployed only as a **data app**.

You can provide the following information to deploy the workflow:

1. Hub URL: The URL of your KNIME Business Hub instance, e.g. "https://my-business-hub.com".
2. Application Password ID: User-associated application password ID.
3. Application Password: User-associated application password.
4. Last n days filter: number of days from calculating the average execution time, e.g. last seven days.

If you want to know how to create an application password, follow **these steps**.

Data app

After deploying the workflow as a data app, you can run it. To do so, follow these **instructions**. Below are the steps for the data app:


1. **Business Hub connection:** You need to connect to the KNIME Business Hub instance

through the application password generated previously.

2. **Workflow performances:** A performance overview of the deployed workflows for the last n days. Diving deep into any workflow performance is possible by selecting the workflow and clicking the “workflow details” button. When diving deep into the workflow, each workflow’s average run time calculation is calculated per day (stacked area chart).

Workflows' Run Time

Workflow performances



Last n days filter

7

Filter by days

Workflow execution time

<input type="checkbox"/>	Workflow name	Mean (seconds)	Stand. dev. (seconds)
<input type="checkbox"/>	Discard Failed Jobs	3.6	0.548
<input type="checkbox"/>	flyweight-execution-workflow	0	0

Previous
1
Next

Teams' filter

☒ Continuous Deployment Setup

Workflow details

← Back

Monitor Users' Usage

Overview

The workflow aims to track KNIME Business Hub user logins to better manage the KNIME Business Hub users' quantity and instance usage.

Workflow specs

You need access to the Keycloak instance embedded in KNIME Business Hub to generate the Client Id and Secret necessary to access the KNIME Business Hub user information.

The workflow shows information about two users type: active and inactive.

Inactive users: those who have never logged into the KNIME Business Hub instance since their creation date.

Active users: those who logged into the KNIME Business Hub instance in the last n days.

It is possible to analyse the session length for actively engaged users. However, it's important to note that if a user doesn't log out after each session, it could impact the accuracy of the KPI calculation. As a result, this metric cannot be considered completely dependable.

Deployment configuration


This workflow is designed to be deployed only as a **data app**.

You can provide the following information to deploy the workflow:


1. Hub URL: The URL of your KNIME Business Hub instance, e.g. "https://my-business-hub.com".
2. Keycloak Client ID: your-ClientId from Keycloak.
3. Keycloak Client Secret: your-Client-Secret from Keycloak.

To generate a Keycloak Client ID and Client Secret:


1. Login into Keycloak Admin Console: "https://auth.my-business-hub.com > Administration Console".
2. Select your Realm, and from "Realm Settings > Events > User events settings", activate "Save events" to save KNIME Business Hub Users events. You can configure how many days to keep the user's events.



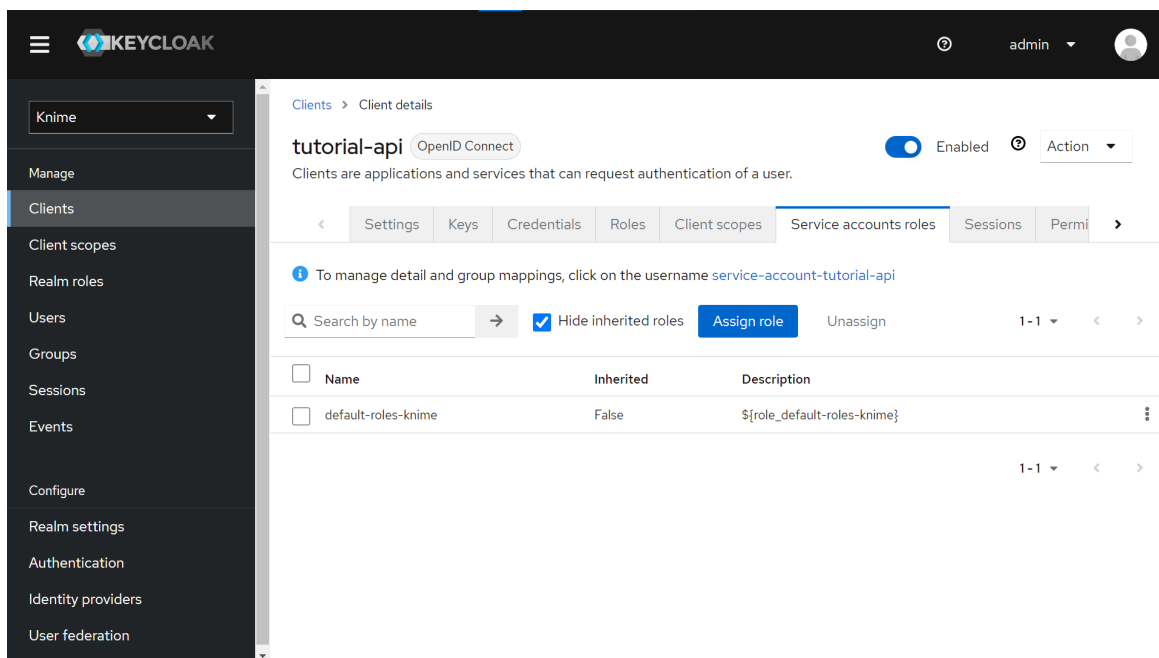
3. Create a new Client from “Clients > Create Client”. Give it a Client ID (tutorial-api in our example), a name (tutorial) and a description (lorem ipsum).



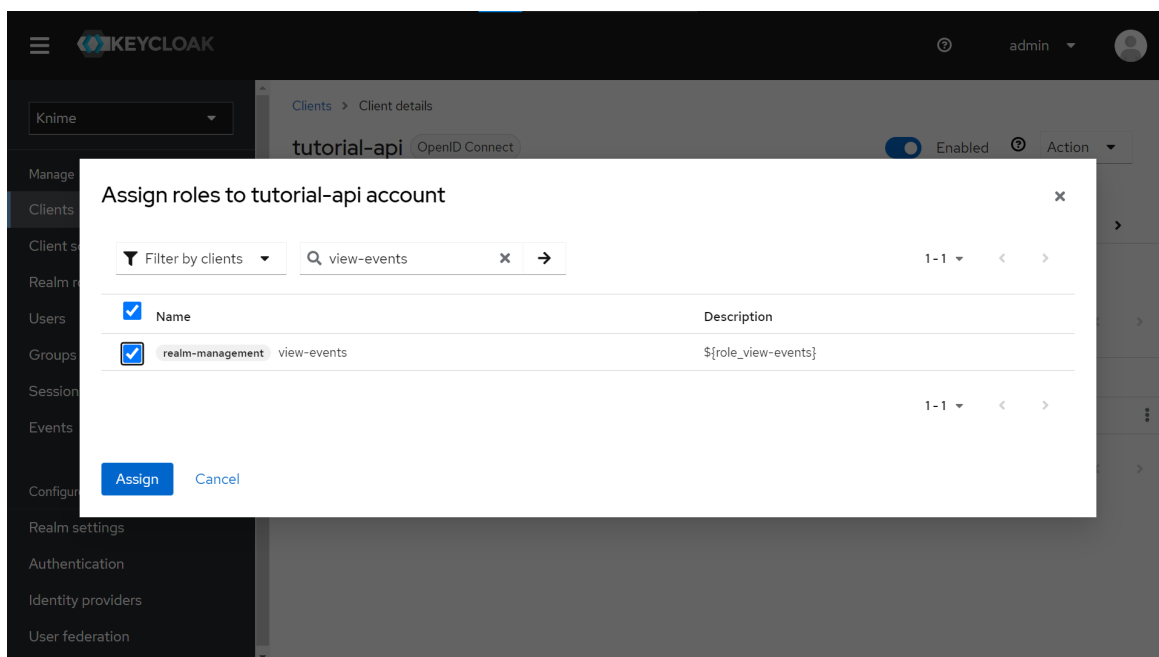
- a. In the “Capability config” section, activate “Client authentication”.
- b. Select “Service account roles” to allow you to authenticate this client to Keycloak and retrieve the access token dedicated to this client.
- c. Save your new client.



4. Click on your new Client (tutorial-api) from “Clients > Client ID (Column)”.
 - a. Go to the “Service account roles” tab > Click the “Assign role” button.

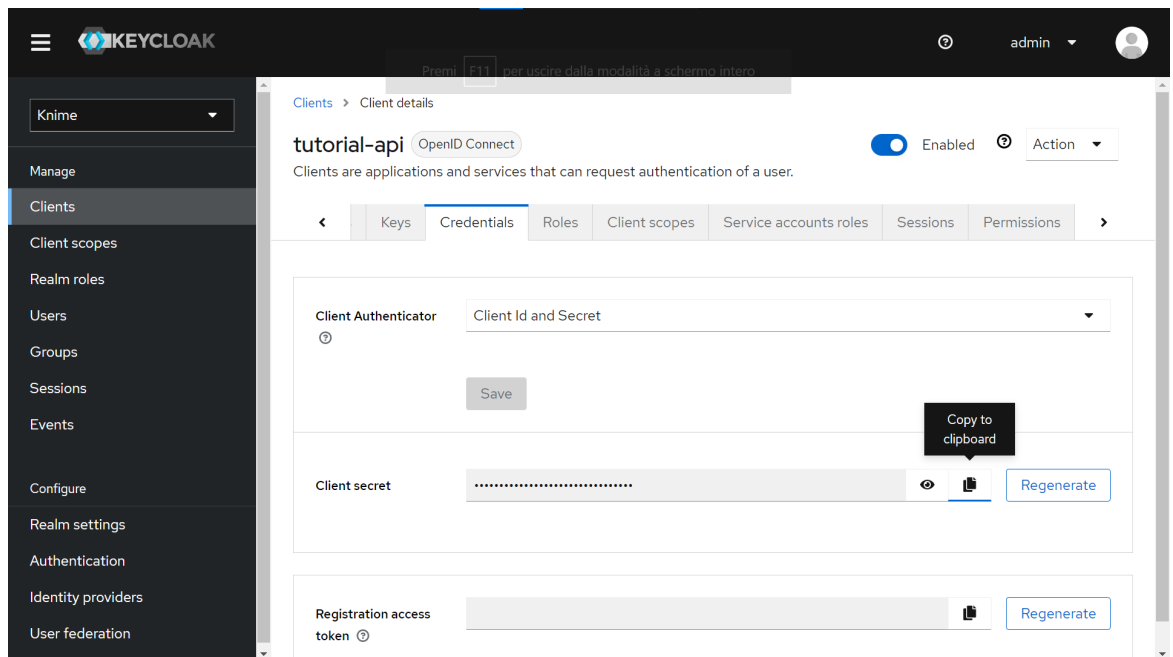


- b. Click “Filter by Roles” and select “Filter by clients”.
- c. Search for “view-events” and for “manage-users” and assign both roles to the Service account associated with our client.



5. Finally, to retrieve the Client ID and Client secret:
 - a. Go to “Clients > Client ID (Column)” and choose your client (tutorial-api in our example)
 - b. Click on the “Credentials” tab
 - c. Leave as a Client Authenticator the “Client ID and Secret” option.
 - d. Copy the Client’s Secret.

- e. Copy the Client ID from the top of the tab (“tutorial-api” in the screenshot below).



Data app

After deploying the workflow as a data app, you can run it. To do so, follow these [instructions](#). Below are the steps for the data app:


Keycloak connection: the user must connect to the Keycloak instance embedded in KNIME Business Hub through the previously generated Client ID and Secret.

User login events: Select the last n days to see KNIME Business Hub active users' login count and inactive users' information. It is possible to explore the active users' session length.

Also, the user can switch to chart mode by clicking the “Switch to charts” option.

Monitor Users Usage

User Login Events



User Usage Results

Days back

15

^

v

Confirm range

☐ Switch to charts

Active Users

<input type="checkbox"/>	Username	Logins count	Days since last access
<input type="checkbox"/>	knime	62	0
<input type="checkbox"/>	knime_admin	28	0
<input type="checkbox"/>	diego	8	0
<input type="checkbox"/>	karen	3	7
<input type="checkbox"/>	simon	3	1

Previous

1

2

Next

Inactive users - No Logins

Username	Days since creation
bernd.wiswedel	162
kevinkr	149
jim.falgout	95
test_daria	155
oole	29

Previous

1

Next

Select users in the "Active users" table and click the button to display more statistics about users' activity

Check sessions' length

← Back

Monitor Users Usage

User Login Events



User Usage Results

Days back

7

^

v

Confirm range

☒ Switch to charts

Aggregated logins count



Count

Time

Logins count

Individual logins count



Count

Time

karen knime knime_admin simon diego

Monitoring K-AI in KNIME Business Hub

A **data application** is available to help KNIME Business Hub admins monitor and govern K-AI usage.

Download the data application and upload it to your KNIME Business Hub instance and *Run* it or deploy it as a data app.



Follow the steps 1 to 3 (or 1 to 5 if you want to create a data app deployment) of the **Data apps on KNIME Hub** guide to know how to do this.

Data access

To know more about the AI history access groups, please consult the **KNIME Business Hub Installation Guide**.

Data structure

Every user on the Hub can interact with K-AI. When users open their KNIME Analytics Platform and initiate a prompt, a conversation with K-AI begins. The conversation concludes when the user closes the *Assistant* window. Each question the user asks is a prompt, and the corresponding answer is a response. This results in three distinct levels of K-AI usage data, since each user has multiple conversations consisting of several prompts and responses.

Table 2. An example of usage data for K-AI

User ID	Account created (year)	Conversation ID	User Prompt	K-AI Response	Timestamp	Tokens consumed
37373	2022	1	Hi!	Hello! How can I assist you today?	2024-01-01T00:00:00	330

User ID	Account created (year)	Conversation ID	User Prompt	K-AI Response	Timestamp	Tokens consumed
37373	2022	1	What node can I use to close a loop with the Group Loop Start node	You can use the "Loop End" node to close a loop with the "Group Loop Start" node.	2024-01-01T00:05:00	1220
37373	2022	2	How can I read an Excel file?	To read an Excel file in KNIME, you can use the "Excel Reader" node from the "KNIME Excel [...]"	2024-01-05T20:00:00	880

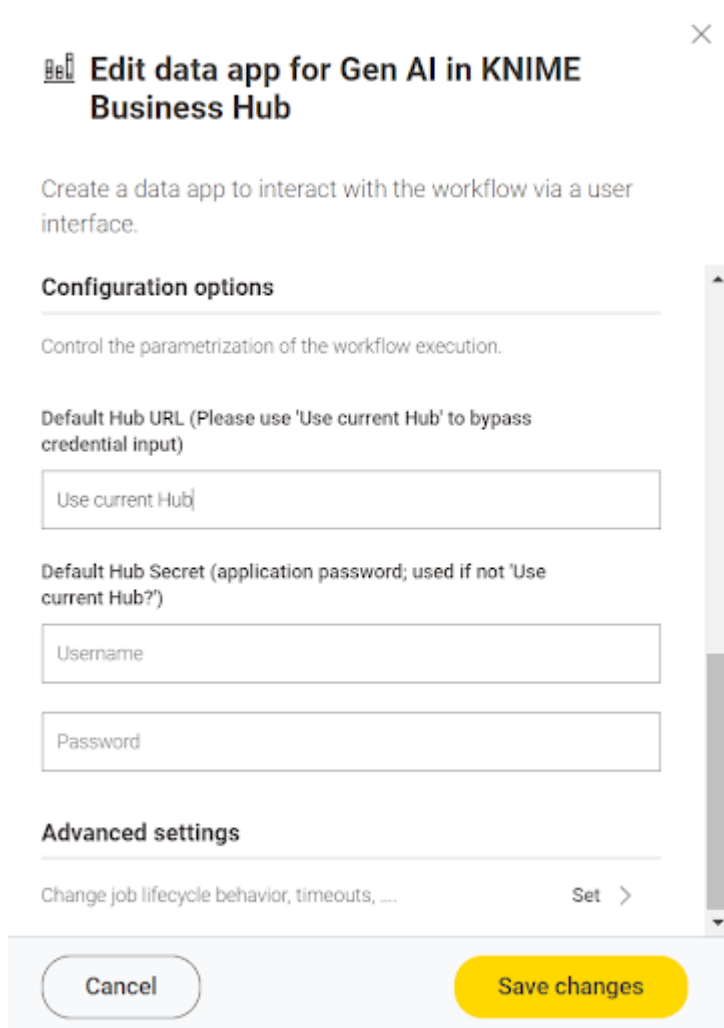
Based on the data in the table, on January 1st, 2024, a user (identified as 37373) interacted with K-AI for five minutes and made two different prompts with varying token consumption values. On January 5th, the same user issued a prompt inquiring about how to read an Excel File. This prompt was stored with a different Conversation ID, indicating that the K-AI assistant was closed between the two conversations.

Deploying the data application

To run the data application in KNIME Business Hub, it is possible to use an [Ad hoc execution](#) or run the [data application deployment](#). Please follow [this](#) link for more information about data app deployments in KNIME Business Hub.

In the data application deployment configuration, the current Hub authentication can be used if the logged user has permission for the correct rights to access the data. Otherwise, a user with permissions must provide their application password, and it is necessary to input the

Hub URL.



Edit data app for Gen AI in KNIME Business Hub

Create a data app to interact with the workflow via a user interface.

Configuration options

Control the parametrization of the workflow execution.

Default Hub URL (Please use 'Use current Hub' to bypass credential input)

Use current Hub

Default Hub Secret (application password; used if not 'Use current Hub?')

Username

Password

Advanced settings

Change job lifecycle behavior, timeouts, ... Set >

Cancel Save changes

Figure 37. Deployment using Use current Hub to bypass credential input

Data application overview

Perspective selection

The data app is divided into four different perspectives to facilitate the answer to various questions:

- **Users:** Who is using K-AI?
This perspective can answer user-related questions like “Are the new KNIME Business Hub users using K-AI?”
- **Conversations:** How are users using K-AI?
This dashboard answers questions such as how the user prompts are distributed by K-AI mode (Q&A, Workflow builder, etc.).

- **Cost:** How much does it cost to use K-AI?

The estimated cost associated with K-AI can be analyzed based on the tokens consumed.

- **All Dimensions:** This perspective combines all attributes and metrics from three previous perspectives to create a more articulated analysis.

K-AI Dashboard

Conversations | Slice and dice by conversation attributes



K-AI is an AI assistant for KNIME that supports KNIME-related questions, workflow building, and scripting nodes such as Python and Echarts. This dashboard provides four perspectives to help you gain insights into K-AI usage.

Perspective selection

☐ Users ☒ Conversations ☐ Cost ☐ All

Figure 38. Perspective selection in the K-AI Dashboard

Slice and dice

There are different slicing filters and dice attributes for every perspective. The term *Slice* refers to filtering the data by different attributes. *Dice* refers to splitting chart metrics by the same attributes. The *Slice and Dice Legend* provides a comprehensive explanation of all of them.

☐ Show slicing filters

Dice (Split by)

K-AI mode

▼ Dice and Slicing Filters Legend

The term 'Slice' refers to filtering the data by different attributes. By selecting the 'Show slicing filter' option, you access a set of filters to slice the data. This allows you to focus on specific user-, conversation-, and prompt-attributes in the current view.

The term 'Dice' refers to splitting chart metrics by the very same attributes. You select the attribute to Dice by in the Drop Down above.

Together, Slicing and Dicing by different attributes allows you to explore the data interactively.

The Dashboard handles different dimensions, each of which certain metrics are attributed to.

User Dimension: Focuses on properties of the users that interact with the AI assistant.

- **Hub Account Age Cohort:** Differentiate new and experienced users based on the year in which a user account was created. For the most recent year, the cohort is further divided into its quarters.
- **Team Memberships:** The number of teams a user is a member of.
- **Team name:** The specific team a user belongs to. Note that if a user is a member of multiple teams, the team association is grouped together.

Apply slice and dice

Figure 39. Slice and Dice Filter Legend in the K-AI Dashboard

Slicing filters

The slicing filters isolate data based on dimension attributes. For instance, from the *Conversation* perspective, you can slice the data only to analyze K-AI interactions related to Python.

The Users' slicing filters will filter users and all their conversations, while the Conversations' slicing filters will filter specific conversations with all their prompts. Cost's slicing filters only apply to the prompts.

Based on the [Table 2](#), the user with ID “37373” is filtered out along with all its conversations if you slice by “Hub account age cohort,” which equals 2024 in the *User* perspective. Conversation ID 1 is filtered out if you slice by “Conversation duration buckets” equals “10-20 min,” and finally, if you slice by “Prompt and response token buckets” equals “1000-2000 tokens,” only row 2 from the table above will be kept.

It is important to note that the slicing filters work sequentially from left to right. This means that if you select to view only K-AI interactions within Python nodes, the filtered output of the first slicing filter will be passed to the downstream filters. Please see [Figure 40](#) for a visual representation.




Figure 40. The slicing filters in the underlying workflow

Dice

Analyzing K-AI usage metrics by different attributes can help you understand how data trends are formed over time. For instance, in the *Conversations* perspective, you can observe the number of user prompts over time, split by “K-AI Mode” or “Conversation duration (minutes)” attributes.




Figure 41. Total user prompts dice by K-AI mode




Figure 42. Total user prompts dice by Conversation duration

It is possible to select the attribute to *Dice* in the drop-down menu:




Figure 43. Dice drop-down menu

Data application usage

Please note that if you modify the time range, the data application will retrieve data from the KNIME Business Hub API endpoints again. Depending on the amount of data and the time range, it may take some time to re-execute.

Instead, the granularity feature only re-executes the charts within a different time granularity (from year to day).

When switching between perspectives, e.g., from Users to Conversations, the dice and slicing

filters are reset, and the charts are re-executed.

Customization

This workflow, is a blueprint you can customize to meet your needs. For instance, you can remove or reorganize some slicing filters, choose to have only the *Cost* perspective, eliminate other perspectives, or be inspired to create your attributes.

KNIME Business Hub API documentation


Most KNIME Business Hub functionalities are also available via REST API allowing you to perform several actions.

You can access the API documentation by navigating to the following URL:

```
api.<base-url>/api-doc
```

where <base-url> is your Business Hub instance URL, e.g. `hub.example.com`.

Here you can select from the drop-down menu the service you want to use.



Support Bundles and Troubleshooting

When generating a support bundle, **no data leaves the cluster**.

If necessary, you can download the support bundle and send it to KNIME for the purpose of troubleshooting. Contact us by sending an email to support@knime.com. Under extreme circumstances, the KNIME team may forward the support bundle to the Replicated support team for additional help.

When generating a support bundle, a limited amount of information will be automatically redacted (IPv4 addresses, connection strings, etc.). You can configure additional redactions and/or manually redact information prior to sending the bundle. See the **Configuring redaction in support bundles** section for more details.

KNIME Business Hub is capable of generating support bundles in a standard format, even when the admin console isn't working. This ensures that users are able to provide all of the necessary information for KNIME to be able to identify the problem and prescribe a solution.

Generating a support bundle (GUI)

In order to help troubleshoot an installation, or to simply inspect the logs of the cluster in a user-friendly format, you will need to generate a support bundle.

Simply open the KOTS Admin Console, navigate to the **Troubleshoot** pane, and click the **Generate a support bundle** button to generate a support bundle.

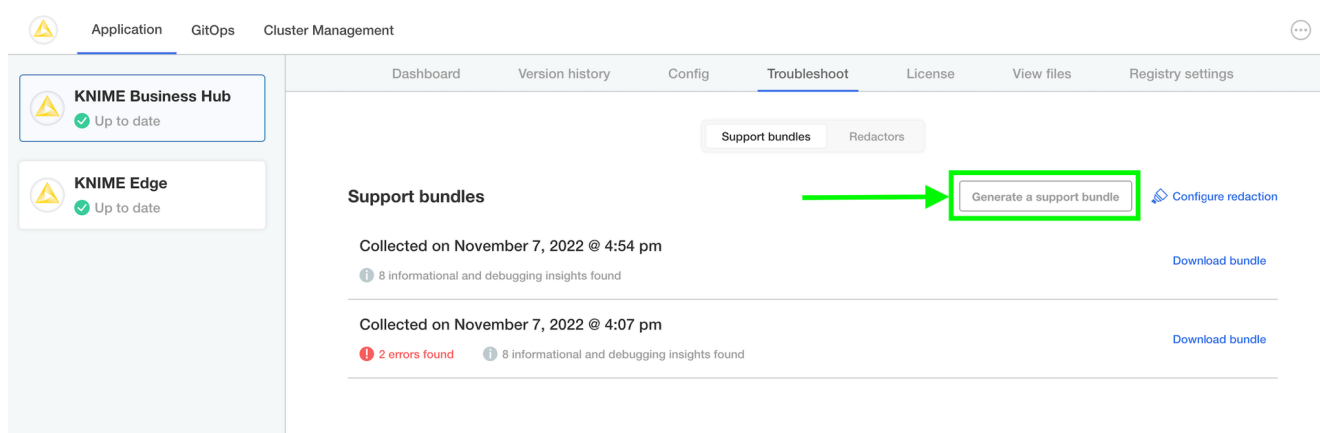


Figure 44. Generate a support bundle

All generated support bundles will display in the list above. Click the **Download bundle** button to download the bundle(s) you want to share with KNIME, and please see the **Configuring redaction in support bundles** section for information on how to redact confidential/personal information before sending.

Generating a support bundle (CLI)

See [Replicated documentation](#) for instructions on how to generate a support bundle via the Replicated CLI.

Configuring redaction in support bundles

When generating a support bundle, a limited amount of information will be automatically redacted (IPv4 addresses, connection strings, etc.) but it is not guaranteed to be a comprehensive set of redactions. You may have additional information in your logs or configuration that you do not wish to share with the KNIME engineering team.

One option is to unzip the generated .zip support bundle and manually review/redact information prior to sending the bundle to KNIME. However, there is a lot of information to review and the redaction of certain information can be automated fairly easily. The ideal option is to configure automated redactions via [Redactor](#) resources, which will automatically redact information for all future support bundles.

In order to configure automated redactors, first open the KOTS Admin Console. Navigate to the **Troubleshoot** pane and click **Configure Redaction**.




Figure 45. Configure Redaction

If you have configured your own custom redactions that you feel would be valuable to other users of KNIME Business Hub, please feel encouraged to share the configuration with KNIME so that it can be considered & potentially added to future releases.

See [this link](#) and [this link](#) for more information.

Inspecting support bundles

There are quite a number of files generated in a support bundle. Not necessarily every file

is useful for every problem. However, by collecting the same information in the same way each time, KNIME can ensure the best quality support possible for customers.

It is possible to inspect a support bundle entirely in the admin console. See below for an example screenshot.

Support bundles > November 7, 2022

Support bundle analysis

Collected on November 7, 2022 @ 4:54 pm

Download bundle

Analysis overview

File inspector

Redactor report

```
1- [
2- {
3-   "name": "required.kubernetes.version",
4-   "labels": {
5-     "desiredPosition": "1",
6-     "iconKey": "kubernetes_cluster_version",
7-     "iconUri": "https://troubleshoot.sh/images/analyzer-icons/kubernetes.svg?w=16\u0026h=16"
8-   },
9-   "insight": {
10-    "name": "required.kubernetes.version",
11-    "labels": {
12-      "iconKey": "kubernetes_cluster_version",
13-      "iconUri": "https://troubleshoot.sh/images/analyzer-icons/kubernetes.svg?w=16\u0026h=16"
14-    },
15-    "primary": "Required Kubernetes Version",
16-    "detail": "Your cluster meets the recommended and required versions of Kubernetes",
17-    "severity": "debug"
18-  },
19-   "severity": "debug",
20-   "analyzerSpec": ""
21- },
22- {
23-   "name": "container.runtime",
24-   "labels": {
25-     "desiredPosition": "1",
26-     "iconKey": "kubernetes_container_runtime",
27-     "iconUri": "https://troubleshoot.sh/images/analyzer-icons/container-runtime.svg?w=23\u0026h=16"
28-   },
29-   "insight": {
30-    "name": "container.runtime",
31-    "labels": {
32-      "iconKey": "kubernetes_container_runtime",
33-      "iconUri": "https://troubleshoot.sh/images/analyzer-icons/container-runtime.svg?w=23\u0026h=16"
```

Figure 46. Inspect a support bundle in the admin console

Here are the most important folders/files and their purposes:

Path	Purpose	Example (may have properties omitted)
./analysis.json	<ul style="list-style-type: none"> Collects the highest-level insights possible for the installation. Often times, the issue and/or resolution may be identified in this file by inspecting the <code>[] .name.insight.detail</code> property. 	<pre>[{ "name": "kotsadm.status", "insight": { "name": "kotsadm.status", "primary": "kotsadm Status", "detail": "At least 1 replica of the Admin Console API is running and ready", "severity": "debug" }, "severity": "debug", "analyzerSpec": "" }]</pre>
./logs	<ul style="list-style-type: none"> Contains logs of individual pods. Execution Context logs are stored in <code>./logs/execution-contexts</code>. 	(typical application logs)

Path	Purpose	Example (may have properties omitted)
<code>./cluster-resources</code>	<ul style="list-style-type: none">• Contains the configuration of each visible resource in the cluster.• For example, to see all pods in the cluster, navigate to the <code>./cluster-resources/pods</code> directory which contains one file per namespace in the cluster.	<pre>{ "kind": "PodList", "apiVersion": "v1", "metadata": { "resourceVersion": "1686941" }, "items": [...] }</pre>

Backup and restore with Velero Snapshots and Kotsadm

Snapshot backups and restore features are available into Replicated deployments via Velero, a tool for backing up Kubernetes cluster resources and persistent volumes.

One-time snapshots as well as automated scheduled snapshots can be managed from the *Snapshots* panel within your Kotsadm dashboard at <https://<base-url>:8800/app/knime-hub>.



Snapshot creation and restoration are disruptive processes. KNIME applications, and Replicated admin access may be unavailable during an active backup or restore operation.

Creating snapshot backups

1. First, configure storage for your backups. Navigate to the *Snapshots* tab of your Kotsadm dashboard. Click the 'Settings' button to edit backup settings where you'll be able to add a storage target.

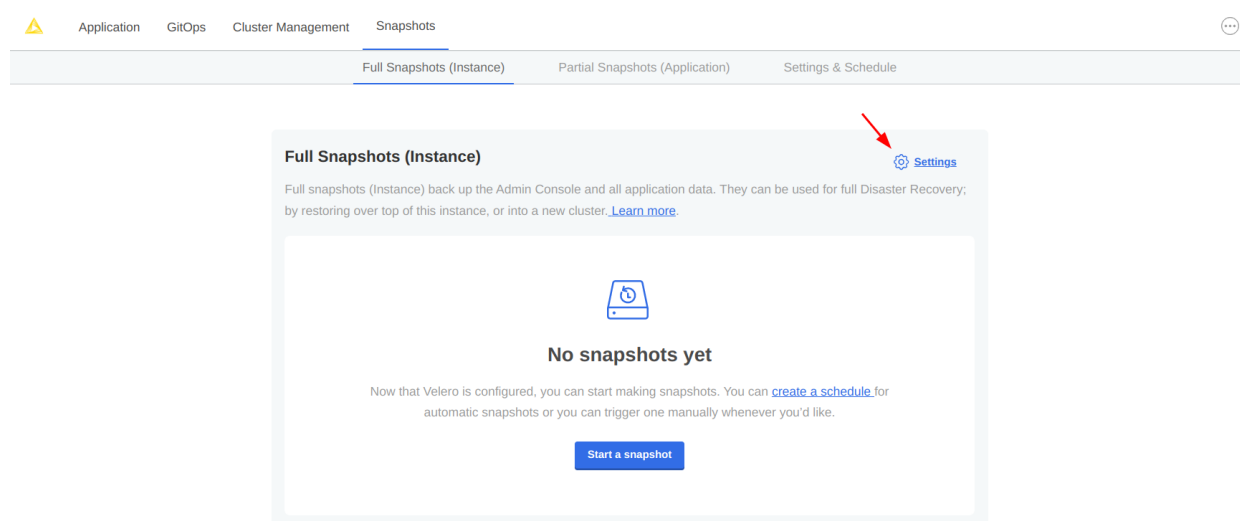


Figure 47. Snapshots tab with settings link

2. Velero supports local storage (not recommended), Amazon S3, Azure Blob Store, Google Cloud Storage, and S3 API compatible storage endpoints such as Minio. Select your preferred snapshot storage type from the 'Destination' drop-down menu, and fill in the required fields with parameters specific to your storage endpoint. Click the 'Update storage settings' button and wait for Velero to verify backup storage access.

Snapshot settings [+ Add a new destination](#)

Full (Instance) and Partial (Application) snapshots share the same Velero configuration and storage destination.

Destination

Amazon S3

Bucket **Region**

business-hub-snapshots us-east-1

Path

/snapshots/

☒ Use IAM Role

[+ Add a CA Certificate](#)

[Update storage settings](#)

All data in your snapshots will be deduplicated. To learn more about how, [check out our docs](#).

Automatic snapshots

Set up a custom schedule and retention policy for automatic snapshots of the Admin Console and all application data.

Full snapshots (Instance) **Partial snapshots (Application)**

☐ Enable automatic scheduled snapshots

Retention policy

The Admin Console can reclaim space by automatically deleting older scheduled snapshots.

Snapshots older than this will be deleted.

1 Months

[Update schedule](#)

Figure 48. Snapshots destination settings for AWS S3 storage

3. With a valid backup storage configured, you can create a Snapshot of your KNIME deployment by clicking the *Full Snapshots* tab, and then the *Start a snapshot* button. This may take a few minutes to complete.
4. Once your snapshot is complete, from the same *Full Snapshots* screen, you can click the 'Settings' button to manage snapshot retention, or configure automatic snapshots by checking the *Enable automatic scheduled snapshots* box and setting a schedule using a CRON expression.

Automatic snapshots

Set up a custom schedule and retention policy for automatic snapshots of the Admin Console and all application data.

Full snapshots (Instance)**Partial snapshots (Application)**

☒ Enable automatic scheduled snapshots

Schedule

Weekly

Cron expression

0 0 ** MON

At 12:00 AM, only on Monday

Retention policy

The Admin Console can reclaim space by automatically deleting older scheduled snapshots.

Snapshots older than this will be deleted.

1

Months

Update schedule

Figure 49. Example automatic snapshot scheduled to run at 12:00am weekly with a 1 month retention policy.

Backup Troubleshooting

Velero is installed into the embedded Kurl Kubernetes cluster with default settings and resource allocations.

As the number of objects or overall size of data to be backed up increases, it may eventually occur that the CPU and memory resources allocated for Velero processes are no longer sufficient to successfully complete the backup.

In the event that backup failures are encountered, it is recommended to increase the CPU and memory allocation **directly** to the Velero's node agent process via `kubectl`.

```
$ kubectl patch daemonset node-agent -n velero --patch \
'{"spec":{"template":{"spec":{"containers":[{"name": "node-agent", "resources":
{"limits":{"cpu": "2", "memory": "2048Mi"}, "requests": {"cpu": "1", "memory":
"512Mi"}}}]}}}}'
```

The CPU and memory resources and limit values can be adjusted as needed to find sufficient values for backup process. Typically, only the **limit** values will need to be increased.



At this time, the resource allocation override to Velero will **revert** after a Kurl upgrade has been performed. Please ensure any changes to the Velero node agent are reapplied after any Kurl cluster-level upgrades.

Restoring a snapshot

1. Navigate to the list of available snapshot restore points from your Kotsadm dashboard by browsing to *Snapshots* → *Full Snapshots*. From this screen, identify the snapshot instance you would like to use, and take note of the instance ID.

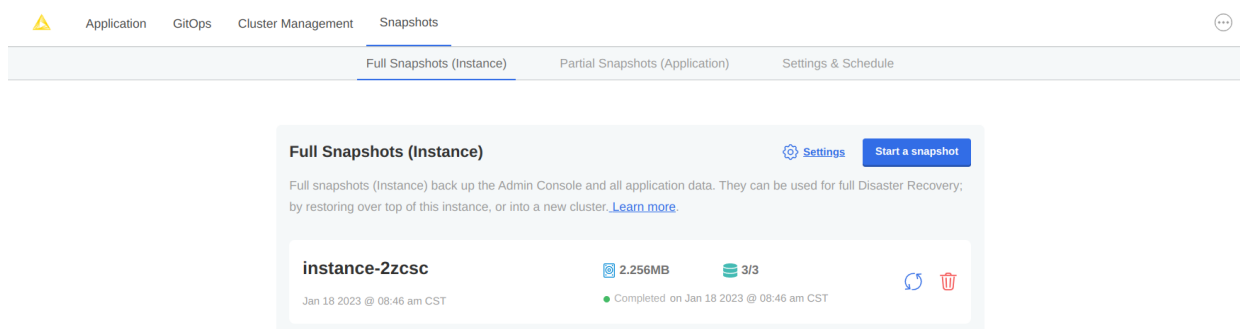


Figure 50. In this example, there is only one snapshot available and its ID is instance-2xcsc

A list of snapshots can also be retrieved by command line:

```
$ kubectl kots get backups
```

NAME	STATUS	ERRORS	WARNINGS	STARTED
instance-2zcsc	Completed	0	0	2023-01-18 14:46:26 +0000 UTC
2023-01-18 14:46:53 +0000 UTC		EXPIRES		29d

2. Now, restore the snapshot using a single CLI command.

```
$ kubectl kots restore --from-backup {Snapshot ID}
```

3. Next, redeploy the Hub instance from the KOTS Admin Console and trigger the restart of all executors by performing any change to each execution context in the Hub UI, e.g. decreasing/increasing the memory.
4. Finally, assuming the restore completed without errors, you can verify your Hub installation is functioning as expected.

Migration of versioning and compatibility

With the 1.5 release of the KNIME Business Hub we introduced the new item level versioning, migrating from space level versioning that was made available in the previous releases.

This means that users can now create, restore and delete versions of individual items as described in the [versioning section](#) of the KNIME Business Hub User Guide.

The KNIME Business Hub version 1.5 is fully compatible with KNIME Analytics Platform version 5.1, which can create versions of item to use for ad hoc execution, create deployments and update components. When upgrading to version 1.5 all items in previously versioned spaces are automatically migrated to be versioned on the single item level. However we recommend to perform a [backup](#) of your KNIME Business Hub instance.

General compatibility between KNIME Hub, KNIME Analytics Platform client and executor

- KNIME Business Hub from version 1.5.0:
 - Analytics Platform version 5.1 client is fully compatible as previously stated
 - Executor version 5.1 is fully compatible
 - All the Analytics Platform and Executor versions < 5.1 are also compatible
- KNIME Business Hub version 1.4.2:
 - Executor version 5.1 is compatible
 - New nodes for versioning of Analytics Platform version 5.1 are not usable
- KNIME Business Hub version 1.4.1:
 - Executor version 5.1 is compatible
 - New nodes for versioning of Analytics Platform version 5.1 are not usable
 - Workflow uploads do not work via Analytics Platform version 5.1

Migration from space level versioning to item level versioning

- KNIME Business Hub 1.5.0 (with item level versioning):
 - Analytics Platform version 5.1:
Space Connector node will not have the version selection functionality.
The Space Version Creator node is replaced by the Version Creator node.

- Analytics Platform version < 4.7:
Drag and drop of repository items will stop working
- Analytics Platform version < 5.1:
Space Version Creator nodes will stop working.
Space Connector nodes will stop working when the version checkbox is checked.
They will work if the version checkbox is unchecked.
- Any Executor from version 4.7:
Deployments migrated from Business Hub version 1.4.x are compatible although they might require some attention/adjustments for more complex cases.
Deployments created with Business Hub version 1.5 are fully compatible.
- KNIME Business Hub 1.4.x (with space level versioning):
 - Analytics Platform version 5.1:
Space connector nodes will not have the version selection functionality and will default to *latest*.
Version Creator node will not work.
 - Executor version 5.1 is compatible for simple deployments although some more complex deployments might require some attention/adjustments.

Changelog (KNIME Business Hub 1.9)

KNIME Business Hub 1.9.0

(released April 10, 2024)

Installation notes

Due to the Keycloak version update, Keycloak admin credentials will change during the update process.

The new credentials can be retrieved from the same secret as during the first install of Business Hub: [KNIME Business Hub Installation Guide](#)



KNIME Business Hub 1.9.0 includes the Docker executor image:
`registry.hub.knime.com/knime/knime-full:r-5.2.3-477.`

Improvements

- Expanded support for secret types, including OAuth2 and authentication mechanisms from Box, Google, and Microsoft
- New landing page that highlights relevant links for faster onboarding of users
- Elasticsearch has been updated to 8.5.1
- Keycloak has been updated to 23.0.5
- Added support for installing Business Hub on OpenShift
- Added support for Kubernetes 1.29
- Migrated UI framework to nuxt3

Changelog (KNIME Business Hub 1.8)

KNIME Business Hub 1.8.3

(released March 26, 2024)

Improvements

- Update Artemis to 2.32.0

Important Bug Fixes

- Fixed a bug where Image pull secret are not set anymore for executors in embedded registry
- Fixed license check for Azure marketplace images
- Fixed exposing of workflow paths in shared deployments

KNIME Business Hub 1.8.2

(released March 05, 2024)

Installation notes

- During the update a short downtime is expected (5-10 minutes).

Improvements

- In existing cluster installations Kubernetes Namespaces can now be created before the installation happens, in compliance with IT security standards.
- Enabled additional security context policies for multiple 3rd party tools. List of policies enabled:
 - require-node-affinity
 - disallow-capabilities
 - disallow-host-namespaces
 - disallow-host-path-volumes
 - disallow-host-ports
 - disallow-privileged-containers
 - require-non-root-groups
 - require-run-as-nonroot
 - disallow-host-process

- disallow-latest-tag
- disallow-proc-mount
- require-drop-all-capabilities
- require-read-only-root-filesystem

Important Bug Fixes

- Fixed a bug where the Velero restore process could show incomplete because of timing issues with Minio.
- Fixed a bug where editing a schedule that has credentials stored corrupted the credentials

KNIME Business Hub 1.8.1

(released February 15, 2024)

Improvements

Allow routing temp data to customer-defined volume

- Helps eliminate disk pressure in situations where large amounts of data need to be cached from jobs during execution.

Optionally set node affinity to required

- Set node affinity to be required so that users can ensure that no hub components are scheduled onto a node that does not have a particular label.

Enable Rook on Airgap HA installation

- Airgap installations now support multi-node setups with Rook.

Support external groups sharing for users in more than 200 groups

- AAD removes groups from the access token of the user is in more than 200 groups. This feature allows customers to still retrieve the groups in these cases.

Important Bug Fixes

- Fixed a bug that occurred when multiple accounts are created concurrently
- Fixed various bugs that impaired first installs on airgapped environments

Important Notes

- Some advanced settings are now hidden in the KOTS Admin Console config by default. To view them enable the *View Advanced Settings* option in the Global section at the top.

KNIME Business Hub 1.8.0

(released December 20, 2023)



KNIME Business Hub 1.8.0 includes the Docker executor image:
`registry.hub.knime.com/knime/knime-full:r-5.1.2-433`.

New features

AI Service (Business Hub Enterprise)

- Gives you control over, who the data going through the KNIME AI Assistant is shared with.
- Configured via the KOTS Admin Console
- Supports OpenAI and Azure OpenAI as LLM provider
- No data is shared with KNIME
- AP Users can connect their KNIME AI Assistant to the AI Service of a Business Hub via the KNIME AI Assistant preference page
- AI history service grants admins access to all requests made by their AP users to the KNIME AI Assistant
- Documentation: [KNIME Business Hub Installation Guide](#)

Secret Store (Business Hub Standard & Enterprise)

- Central location to manage personal and team secrets

- Sharing of team secrets with other users or teams
- No more logins stored in KNIME workflows
- Secret management is done in the KNIME Business Hub
- Secret Retriever node allows users to use secrets within KNIME workflows
- Documentation: [KNIME Business Hub User Guide](#)

Start data apps with parameters via URL

- Send customized data app links that trigger link/user-specific behavior inside the data app
- Needs a Container Input (JSON) node inside the workflow
- Documentation: [KNIME Business Hub User Guide](#)

Infrastructure updates

Minimize use of cluster roles

- There are no more cluster roles in services developed by KNIME
- Circumventing the use of cluster roles in 3rd party tools used by Business Hub will be introduced in a later release (where possible)

Backup / restore via Velero

- Stability improvements

Improvements

- Added display of CPU and RAM usage of execution contexts
- Search results indicate versioned and draft items
- Admin console now displays Business Hub version

Important bug fixes

- Date-times are now displayed in the browser's timezone

- Execution context selection list not showing all available execution contexts for deployments and ad-hoc execution, improved UX
- Fixed download button for previous workflow versions
- Fixed a bug where in some cases endless loading state could occur while editing schedules
- Fixed a bug where in some scenarios deployments could be created on disabled execution contexts
- Various other bug fixes

Changelog (KNIME Business Hub 1.7)

KNIME Business Hub 1.7.0

(released November 06, 2023)

Important installation notes

Kubernetes version 1.25-1.27 is required for Business Hub 1.7.0. Pre-flight checks in the installer make sure that the correct version is available.

Infrastructure and security updates

Support provided ingress-nginx deployments

- In existing clusters creating own ingress-nginx deployments before the installation of Hub is now enabled. Whether the Hub install process deploys the inbuilt ingress-nginx can be configured on the Kots Admin dashboard.

Security context update

- Security contexts updated for pods, containers and jobs in the hub, knime and knime-execution namespaces.

Bug Fixes

- Fixed a bug where schedules edited with KNIME Analytics Platform didn't respect job discard settings.
- Fixed a bug where saving executed job as workflows didn't work in case workflow size was larger than 50 MB.
- Fixed a bug where under certain circumstances jobs were shown duplicated in the AP explorer.
- Fixed a bug where schedules in all teams connected to a user were deactivated when the user left any team.
- Various smaller fixes.

Changelog (KNIME Business Hub 1.6)

KNIME Business Hub 1.6.0

(released September 22, 2023)

Important installation notes

Business Hub 1.6.0 comes with Kubernetes 1.25 for the embedded clusters.

We recommend updating by running the Kubernetes installer command:

```
curl -sSL https://curl.sh/knime-hub | sudo bash
```



Run the command only when updating from KNIME Business Hub version 1.5.2 or higher, or **after** you updated to version 1.6.0.

All pods will restart during the update, some downtime is expected. Creating a backup before upgrading is recommended.

New Features

Save jobs to spaces

- Jobs on KNIME Business Hub can now be saved any space where the user has permissions. Great feature to utilize for fixing a faulty job and saving the results for the future.
- Documentation: [KNIME Business Hub User Guide](#)

Shared execution contexts

- Create shared execution contexts, that allow sharing execution resources between multiple teams. Shared execution contexts are setup and maintained by the Hub admin and utilized on Team level.
- Documentation: [KNIME Business Hub Admin Guide](#)

Execution context redesign

- Improved usability of execution resource management based on customer feedback. Current job list is now available on execution context and executor level.

Auto start-stop of execution contexts

- Save money on infrastructure costs by only starting up your execution resources when they are actually used, and by shutting them down afterwards.
- Auto start-stop can be enabled for each execution context separately.

Download job logs

- Downloading job logs is now available from the Hub user interface.
- To use this feature an executor based on KNIME Analytics Platform version > 5.1 is necessary.
- Documentation: [KNIME Business Hub User Guide](#)

Improvements

Triggers listening to version creation

- It is now possible to trigger deployments when a new workflow version is created.

- Documentation: [KNIME Business Hub User Guide](#)

Deprecated nodes are not part of search results

Infrastructure updates

Embedded clusters updated

- New Kubernetes version for embedded clusters is 1.25

Kubernetes node affinity

- Node affinity available to separate execution resources from Hub core services
- Documentation: [KNIME Business Hub Admin Guide](#)

Important Bug Fixes

- Fixed a bug where two trigger deployments triggered at the same time could produce inconsistent results
- Fixed faulty URLs on the API documentation
- Fixed KNIME AP mount point registration https problem
- Fixed an issue where global proxy settings could break executor functionality
- Fixed an issue where Hub could get into a crash loop with many APs connected and large files being uploaded / downloaded at the same time

Changelog (KNIME Business Hub 1.5)

KNIME Business Hub 1.5.2

(released Aug 14, 2023)

Important installation notes

Postgres will restart during the update, short downtime is expected.

Infrastructure updates

KNIME Business Hub is now compatible with Kubernetes 1.25.

KNIME Business Hub 1.5.1

(released July 27, 2023)

Important Bugfixes

- **HUB-5628** Fixed an issue that prevented the creation of disabled schedules. Impacting KNIME Server to Hub migration.
- **HUB-5649** Fixed a bug that prevented notification service updates when a subscription for a deleted user was present.
- Fixed a configuration issue that caused the trigger-service to not start up in some scenarios.

KNIME Business Hub 1.5.0

(released July 17, 2023)

Important installation notes

All executor pods will restart during the update, some downtime is expected. Creating a backup before upgrading is recommended.

New Features

Item level versioning

- Item level versioning has been introduced with Hub 1.5.0
- Users can create / restore / delete versions of individual items
- Item versions can be utilized in Ad hoc execution, Deployments
- Using KNIME Analytics Platform 5.1 users can leverage item level versioning for Ad hoc execution, sharing component links and use new nodes like the "Version Creator"

- Space level versioning has been deprecated
- All items in previously versioned spaces are automatically migrated to be versioned on item level upon updating

Volume Mount support for Execution Contexts

- Kubernetes volumes can now be attached to Execution Contexts by editing the custom resource definition.
- This allows attaching secrets and config maps as files to Execution Contexts.

Improvements

Keycloak enforces username validation

- Username restrictions described in the [KNIME Business Hub Installation Guide](https://docs.knime.com/latest/business_hub_installation_guide/index.html#_keycloak_setup) are now enforced by KNIME Business Hub's identity provider Keycloak. Preventing problems where users could be synced into Hub but they could not login with their credentials.

Notifications redesign

- KNIME Business Hub UI notifications were redesigned to be less intrusive and to be more convenient for users.

Schedule deployment

- Initial start date of schedules can no longer be in the past.

Space permissions

- Users with viewer permission can no longer delete spaces

Infrastructure updates

Artemis update

- Artemis version has been updated to 2.29.0

Quarkus update

- Quarkus has been updated to 3.1.1 in the execution services

Execution Context security hardening

- Execution Contexts now have a more restrictive security context to help prevent containers from escalating privileges and prevent running as root user. This increases compliance with typical enterprise cluster security policies enforced with tools such as Kyverno.

Important BugFixes

- **HUB-4991** - Fixed a bug where users were not able to execute deployments shared with them, if they didn't have access to the deployed workflow.
- **HUB-5396** - Fixed the state persistor issue that could result in timeouts and various problems in workflow execution on the Hub.

Changelog (KNIME Business Hub 1.4)

KNIME Business Hub 1.4.2

Main motivation for the 1.4.2 KNIME Business Hub release is to prepare for the upcoming 5.1 KNIME Analytics Platform release. It is an optional update as 1.5.0 Business Hub is coming out with all new features before AP 5.1.

Improvements

Support for using KNIME Analytics Platform 5.1 as a workflow editor. Compatibility with uploading workflows to the Hub using KNIME AP 5.1.

Important Bugfix

- **HUB-5396** - Fixed the state persistor issue that could result in timeouts and various problems in workflow execution on the Hub.

KNIME Business Hub 1.4.1

Important Bugfixes

- **Fix Bug** disabled functionalities for interacting with components in the HUB from AP versions (4.7.0, 4.7.1, 4. 7.2, 4.7.3, 5.0.0) affected by a bug in which saving shared components might delete local workspaces
- Listing HUB spaces from affected AP versions will return the components as data objects, which don't have the potential to trigger the bug, but limits the use of components.
- Changed file extension for downloading components from .knwf to .knar.

KNIME Business Hub 1.4.0

Important installation notes

Please calculate with some downtime during which some services might not be available.

New Features

Data apps

Share data app deployments with "any signed in user"

You can now share your data apps with every Hub user in your company with one setting. Data app deployments shared this way will be available in every user's data apps portal.

Share data app deployments with external groups

Sharing data app deployments with users from externally managed groups (e.g. LDAP/AD groups) is now available. This is done by either importing LDAP groups to Keycloak, or by connecting Keycloak to an external OIDC provider, which provides groups through the access token. Usage of such external groups is for sharing of data app deployments only. The actual users who run the deployments are still managed within KNIME Business Hub.

Data Apps Portal UX enhancements

Enhanced the visual design and user experience of the data apps portal, enabling the addition of categories and descriptions to data apps. This enhancement allows data app consumers to effortlessly filter and identify relevant data apps based on their category and description.

Content Security Policy for data apps

Admins of KNIME Business Hub can now set custom Content Security Policies for data apps, to restrict which resources, from where can be loaded into the data app's user's browser. See more in: [Business Hub Admin Guide](#).

Deployments

Editing deployments

Recreating deployments is no longer needed just to change their configuration.

Admin functionality

Backup / restore via Velero

KNIME Business Hub is now complemented with the widely used, open-source backup / restore tool: Velero. You can mitigate the risk of serious data loss by setting up a backup schedule. The backup process does not require any downtime.

X-Frame-Options Header

Now admins can also select from different X-Frame-Options header options, to prevent clickjacking attacks. See more in: [Business Hub Admin Guide](#).

Security

Improved security by removing sensitive information shown in the HTML source code.

Edge

KNIME Edge integration

KNIME Edge 1.3 is now compatible with KNIME Business Hub 1.4.

Important Bugfixes

- HUB-4774: Next execution time for schedules repeating every 24 hours cannot be computed
- HUB-5093: Heap size calculation doesn't work on Ubuntu 22.04 (Cgroups V2), leading to a too low limit being set for various services

Changelog (KNIME Business Hub 1.3.0)

KNIME Business Hub 1.3.0

(released April 03, 2023)

Important installation notes

- With the release of KNIME Business Hub 1.3.0 new security features are introduced regarding executor communication. For this reason, KNIME Business Hub requires every executor to be on Analytics Platform version 4.7.1 or higher. You can find a list of the available executor images [here](#).
- With this release we changed the istio configuration of two services, to allow them to see the IP of the caller if the request is coming from within the cluster. This can cause those two services (keycloak-proxy, s3-proxy) to not start up in some scenarios, with a "too many open files" error in the istio-proxy container of those pods. This can be fixed by increasing the open file limit on the host VM, e.g. by running the following command:
`sudo sysctl fs.inotify.max_user_instances=8192.`
- The KOTS App Manager is now on version 1.96.3. This new version fixes an important bug, so it should be updated before updating to KNIME Business Hub version 1.3. To do so use the following command:
`curl -sSL https://kurl.sh/knime-hub | sudo bash.`
- If you haven't updated to KNIME Business Hub 1.2.0, yet, you need to follow update guide in the [changelog](#) for 1.2.0 before updating to 1.3.

New Features

Molecule sketcher

Use the new Molecule Widget (Labs) node to accept molecule input from your users. With this new node we made the integration of different sketchers easier and allowed to provide different sketchers also for the KNIME Analytics Platform. In this release the ketcher sketcher is included. Additional sketchers will follow in the upcoming releases.

Collections

Introduce your users to a topic using collection pages. Collection pages serve as a starting point for your users to explore examples and building blocks related to your chosen topic, facilitating their onboarding process. To start creating collections read our guide [here](#).

KNIME Analytics Platform customization profiles

Use Business Hub to distribute preference profiles to either local Analytics Platform clients or to KNIME Hub executors. This can be used to e.g. easily ship database drivers. More information about Analytics Platform customization profiles [here](#).

Trigger deployment

Trigger is a powerful new type of deployment which allows automating workflow execution based on specific conditions, such as another workflow or component being added or removed to a space. To get started with Triggers, check out our documentation [here](#).

Changelog (KNIME Business Hub 1.2.0)

KNIME Business Hub 1.2.0

(released March 13, 2023)

KNIME Business Hub 1.1.1 to 1.2.0 Upgrade

- KNIME Business Hub 1.2.0 is expected to run in a Kubernetes 1.23.x cluster. Prior versions of KNIME Business Hub ran in Kubernetes 1.21.x clusters.
- For embedded kurl cluster installations, the cluster update process is to:
 - First, check for updates in KOTS Admin Console.
 - Once the KNIME Business Hub 1.2.0 update has been identified and preflight checks have been run, trigger the deployment.
 - After the deployment process has completed, run the following command in the terminal of the host machine to upgrade the Kubernetes cluster version (among other components):
 - `curl -sSL https://kurl.sh/knime-hub | sudo bash`
 - You will have to enter y to various yes/no prompts during this upgrade process.

New Features

- New supported Kubernetes versions: 1.22, 1.23, 1.24
- Keycloak version updated to 19.0.3, see Keycloak documentation: https://www.keycloak.org/docs/19.0.3/server_admin/

- Support for airgapped environments: Installing KNIME Business Hub in environments with no internet connection. Install instructions: https://docs.knime.com/latest/business_hub_installation_guide/index.html#introduction
- Embedded docker registry for custom executor images: Embedded cluster installations using kURL now contain an embedded docker registry for storing custom executor images. Documentation: https://docs.knime.com/latest/business_hub_installation_guide/index.html#embedded-docker-registry-configuration
- Recreate strategy for Execution contexts: For upgrading execution contexts now rolling updates and recreate strategies are both available.
- Support for custom certificate authorities

Bugfixes

- Removed the unused ipv6 listener in the Business User Portal container, that could cause crashes in some environments

KNIME AG
Talacker 50
8001 Zurich, Switzerland
www.knime.com
info@knime.com