

KNIME Business Hub User Guide

KNIME AG, Zurich, Switzerland
Version 1.12 (last updated on 2024-10-11)



Table of Contents

Introduction	1
Connect to KNIME Business Hub	2
Connect to KNIME Business Hub from KNIME Analytics Platform	3
Search items in KNIME Hub	7
Nodes	8
Components	9
Workflows	10
Extensions	12
Collections	13
Drag & drop	15
Teams	17
Create a team	17
Team owned spaces	18
Manage space access	19
Manage team members	20
Customization profiles	23
Apply a customization profile to KNIME Analytics Platform client	24
Detach a customization profile from KNIME Analytics Platform client	25
Upload items to KNIME Hub	26
Move items within KNIME Hub	26
Delete items from KNIME Hub	27
Download items from KNIME Hub	27
Versioning	29
Create a version of an item	29
Show a version	29
Restore a version	30
Delete a version	30
Execution of workflows on KNIME Hub	32
Execution contexts	32
Manage shared execution contexts	39
Ad hoc execution	40
Deployments	43
Data apps	44
Service	49

Schedule.....	53
Trigger	55
Managing deployments	59
Jobs.....	62
Inspect an executed workflow.....	63
Job states	65
Data Apps Portal	66
Secrets.....	67
Known issues and compatibility.....	68

Introduction

KNIME Business Hub is a customer-managed KNIME 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.

Once you have access to a KNIME Business Hub instance available at your company, you can use KNIME Business Hub to perform a number of tasks such as:

- collaborate with your colleagues,
- test execution of workflows,
- create and share data apps, schedules, and API services
- keep track of changes with versioning.

This guide provides information on how to use KNIME Business Hub from a user perspective.

To administrate a KNIME Business Hub instance or a Team please refer instead to the [KNIME Business Hub Administration Guide](#).

Connect to KNIME Business Hub

You can use a standard web browser to connect to your KNIME Business Hub instance, once you are provided with the address (URL) and your access credentials - username and password.

After providing them you will be able to access to the home page of the KNIME Hub.

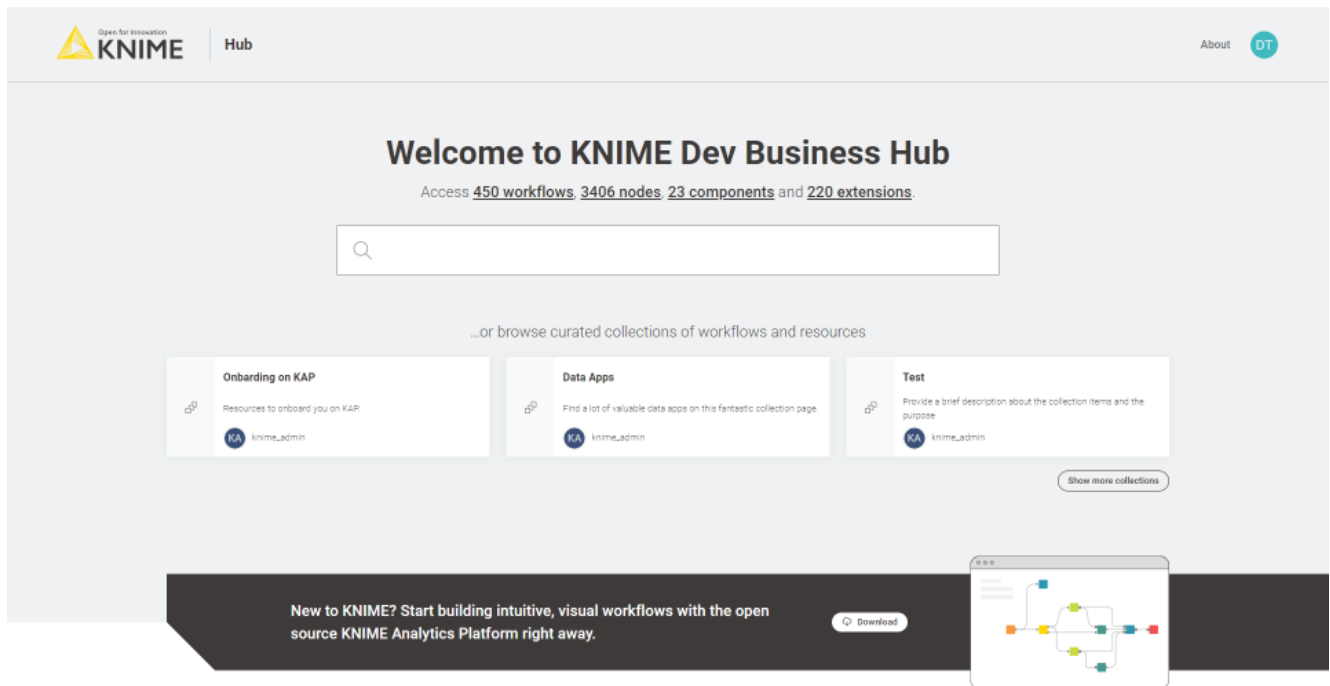


Figure 1. A KNIME Business Hub instance home page

Now you can use the search functionality by typing your search terms into the search bar. The search will show results among all the Hub items that are available to you, based on your role and the items permissions.

Under the search bar area, for KNIME Business Hub Enterprise editions, you will find collections for an easier onboarding.

i Find out more about collections [here](#).

In all the KNIME Business Hub editions instead you will be shown a button to download KNIME Analytics Platform.

i The download link can be configured by the KNIME Business Hub global admin in the [KOTS Admin Console](#).

Your user will be assigned to one or more teams so that you will be able to work on projects

together with your colleagues.

Connect to KNIME Business Hub from KNIME Analytics Platform

To connect to KNIME Business Hub from KNIME Analytics Platform you have to first add its **mount point** to the **KNIME space explorer**.

Space explorer

The space explorer on the left-hand side of KNIME Analytics Platform is where you can manage workflows, folders, components and files in a space, either local or remote on a KNIME Hub instance. A space can be:

- Your local workspace you selected at the start up of the KNIME Analytics Platform
- One of your **user's spaces** on KNIME Community Hub
- One of your **team's spaces** on KNIME Business Hub.

In the following, you will learn how to connect to the KNIME Business Hub.



Please be aware that in case you want to work with KNIME Server you need to **switch to the classic user interface**.

Setup a mount point

Click the *Preferences* button in the top-right corner to open the corresponding dialog.

In the *Preferences* window choose *KNIME* → *KNIME Explorer*. Click *New ...* to add the new Hub mount point as shown in **Figure 2**.

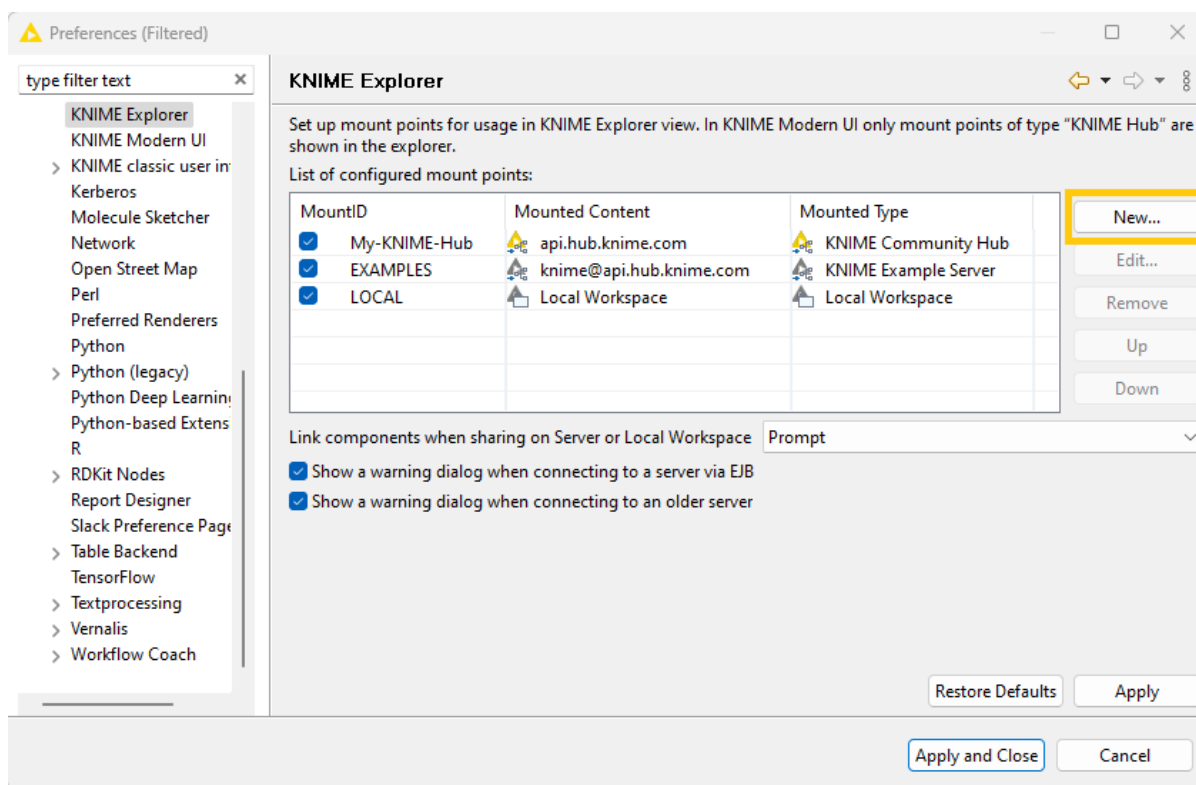


Figure 2. Add a new Hub mountpoint in the Preferences dialog

The *Select New Content* window opens as shown in Figure 3. Select *KNIME Hub* and insert the Hub address, the URL that connects the space in KNIME Analytics Platform to the KNIME Business Hub.

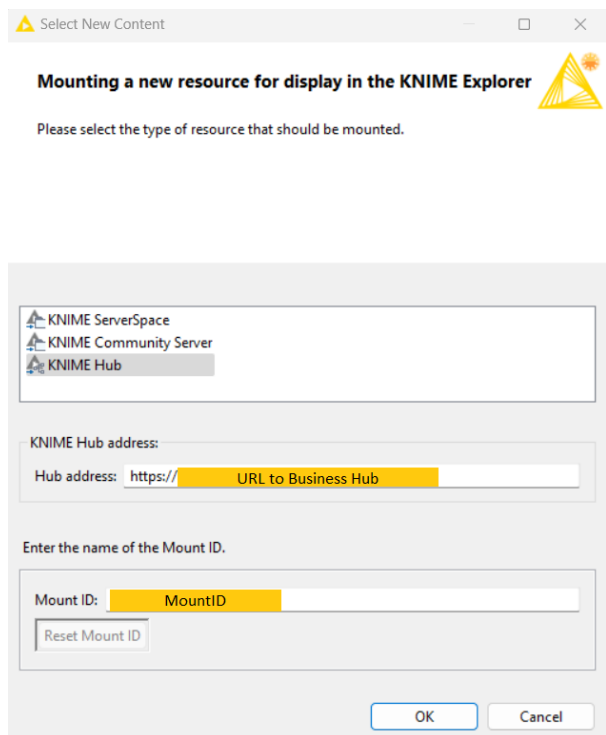


Figure 3. Select New Content window to add a new KNIME Hub mount point

In the *Mount ID* field the name of the mount point will be automatically filled. Finally click *OK*

and *Apply and Close*.

Access KNIME Hub

Via the space explorer:

On the top of the space explorer you can sign in to any of the Hub mount points and select a space. The MountID will be displayed along with the URL to the Business Hub, as shown in [Figure 4](#). Upon logging in, your private, public and team spaces will be visible.

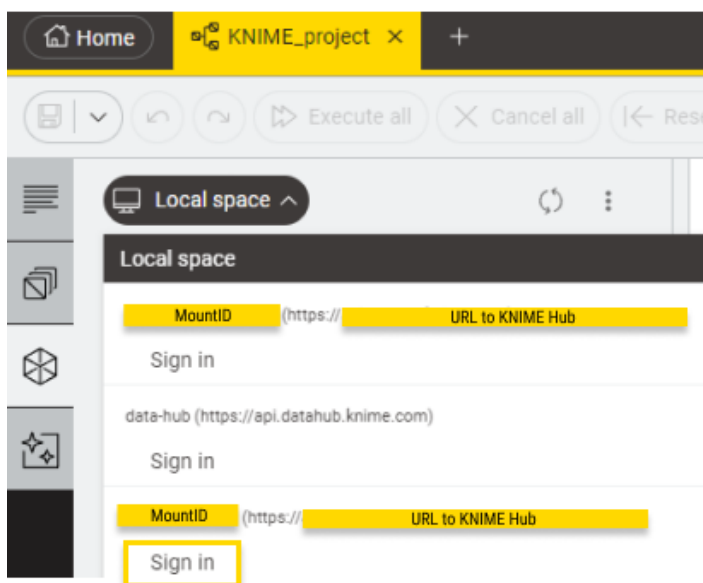




Figure 4. Connect to the Business Hub mount point in the space explorer

Here you can perform different types of operations on items:

- Delete or rename your items. To do so, right-click the desired item and select *Rename* or *Delete* from the context menu.
- Open workflows as local copy or on KNIME Hub. To do so right-click the desired item and select either *Download to local space* or *Open in Hub* from the context menu.
- Create folders to organize your items within the space. To do so, click the  icon in the upper-right corner of the space explorer and select *Create folder*.
- Create new workflows or import local workflows by clicking the  icon and selecting *Create workflow* or *Import workflow*.

Via the entry page:

Click on the *Home* tab. All your mount points are listed at the bottom of the entry page. The MountID will be displayed along with the URL to the Business Hub, as shown in [Figure 5](#). Click the yellow *Sign in* button to the right of the mount point. Now you will see all the spaces that you have access to. Click on the tile of the space you want to access. Within this space you can also import workflows, components or add data files from your local workspace to the

KNIME Hub instance or create a folder.

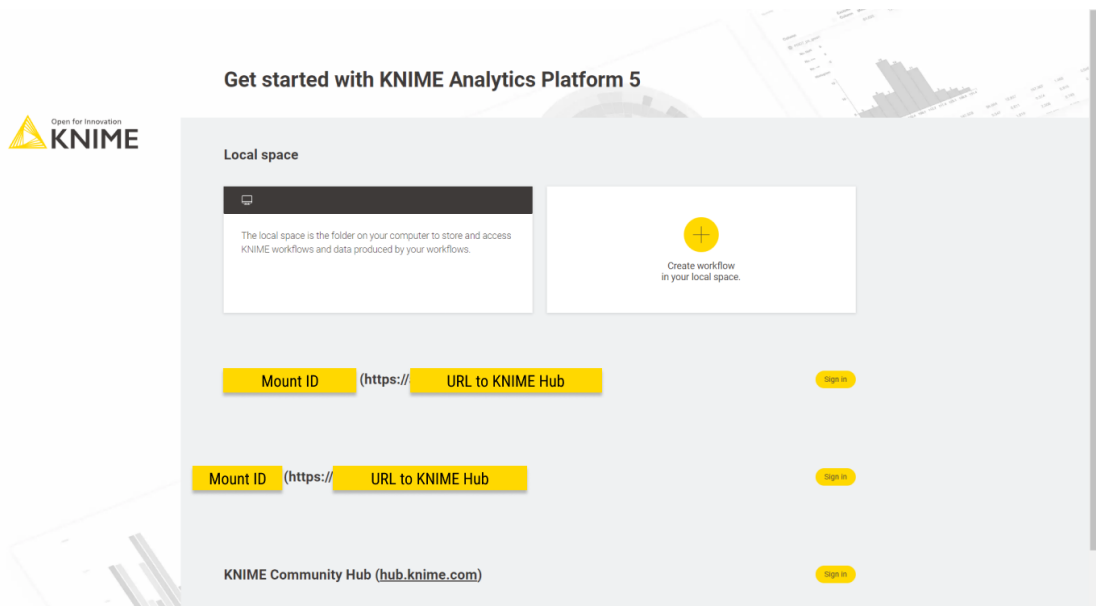


Figure 5. A KNIME Hub mount point on the entry page

Gain a more detailed insight about the [entry page](#) and the [space explorer](#) in the KNIME Analytics Platform User Guide.

Search items in KNIME Hub

Insert a keyword in the search bar to search among the following items:

- Nodes
- Components
- Workflows
- Extensions
- Collections

Press *Enter* to visualize the results.

You will see a list of all the search results in tiles, with:

- An icon indicating if they are nodes, components, workflows, extensions, or collections
- The title
- Tags, when available
- A preview of the description, when available
- The likes the item received
- The owner icon for workflows, and extensions
- The icon, for nodes and components

On top of the search results list, you can filter the results to list only nodes, components, workflows, and extensions, by clicking the respective tab, shown in [Figure 6](#).

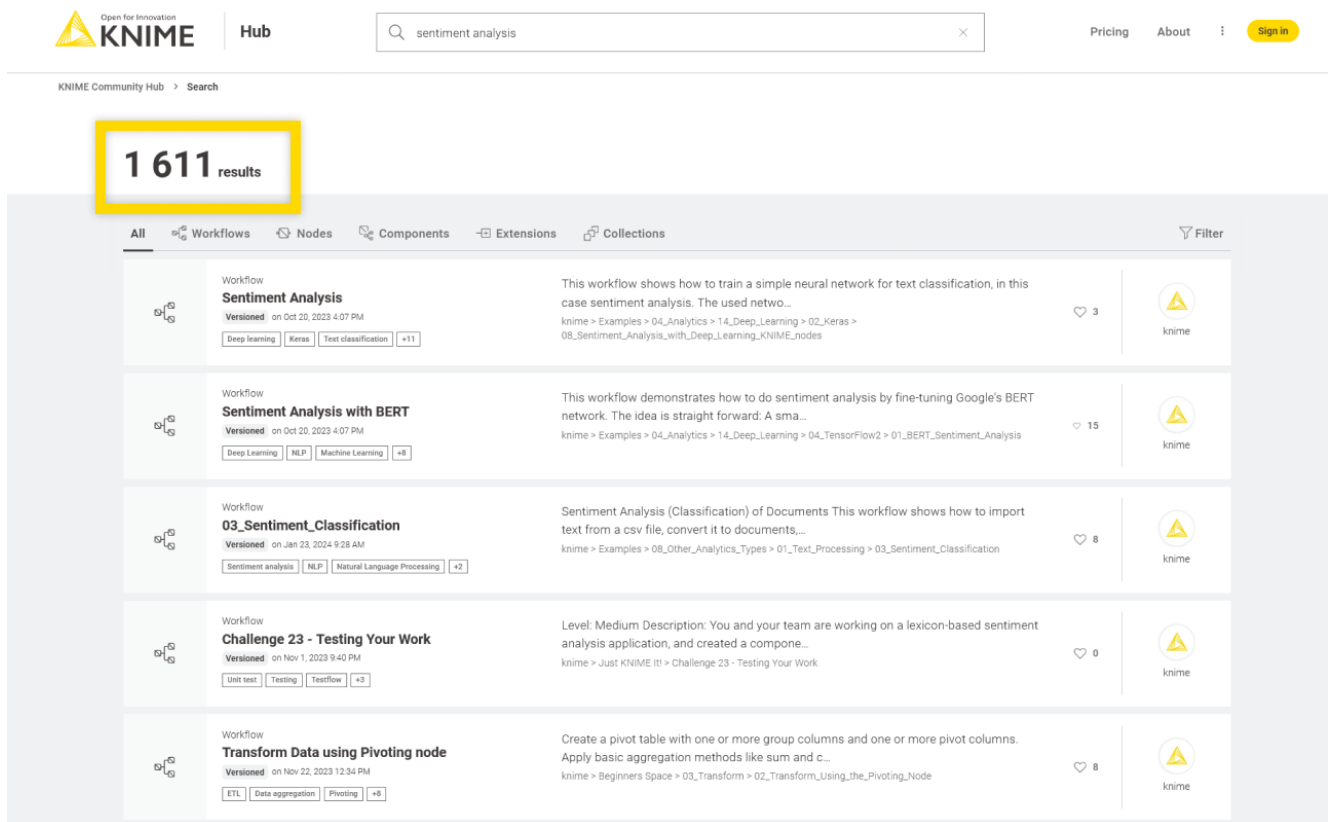


Figure 6. The KNIME Hub search results

Click on a search result to access the relative page.

Nodes

On the node page you will find:

1. The node function, name, and tags
2. The node owner
3. The node drag & drop element: You can drag & drop the element into the Workflow Editor of KNIME Analytics Platform to use the node directly to build a workflow or a component. See the [Drag & drop section](#) for more details on this feature.
4. The node likes and a *Copy link* button to copy the node short link.

For each node you can also find:

- An *Overview* with a description of the node functionality
- *Node details* where you can see information about *Input* and *Output* node ports, a description of the additional *Options*, and, when available, the *Views* that the node is able to produce

- In the section *Extension*, you can see the node extension and the extension owner
- Finally, in the section *Related workflows & nodes* you are provided with:
 - *Workflows* that are available on KNIME Hub and that contain the node
 - When available, the *Outgoing nodes*, that are the most popular nodes to follow the node when building a workflow.

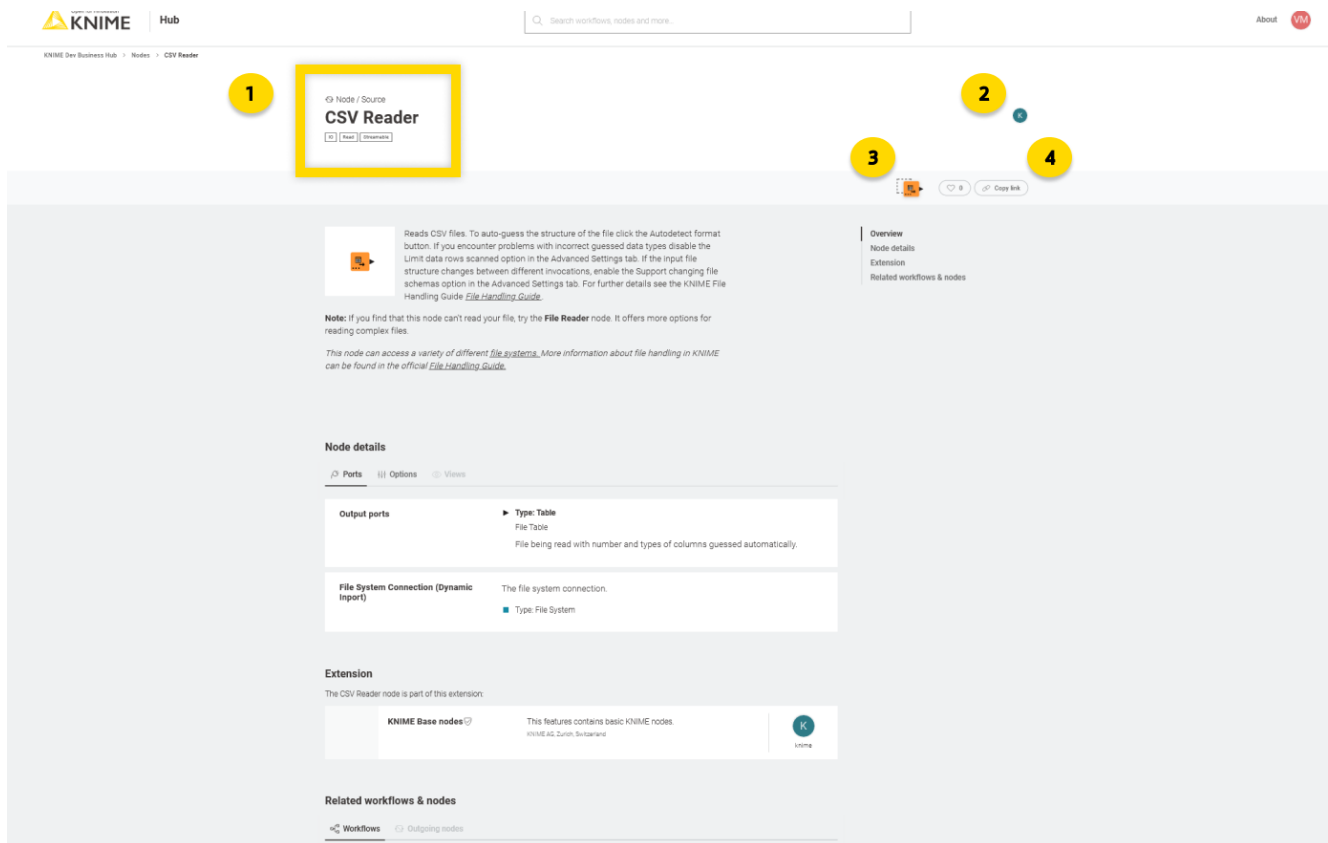


Figure 7. A node page on KNIME Hub

Components

Since **components** really are KNIME nodes that you create which bundle functionality, have their own configuration dialog and their own composite views, also in the component page you will find similar information as for nodes.

1. The component name and whether it is a draft or a version: For drafts, the time of the latest edits is indicated. For versions, the version name and date of creation is indicated. Find out more about versions in the **versioning section of this guide**.
2. The component owner
3. The component drag & drop element: You can drag & drop the element into the Workflow Editor of KNIME Analytics Platform to use the component directly to build a workflow. See the **Drag & drop section** for more details on this feature.

- The component likes, the count of downloads, and a *Copy link* button to copy the component short link.

For each component you can also find:

- An *Overview* with a description of the component functionality, when provided by the component owner
- Component details* where you can see information about *Input* and *Output* component ports, a description of the additional *Options*, and, when available, the *Views* that the component is able to produce
- In the section *Used extensions & nodes* you can see the component extensions and nodes.
- Finally, in the section *Related workflows* you are provided with *Workflows* that are available on KNIME Hub and that contain the component.

The screenshot shows the KNIME Hub interface for the 'Twitter URLs Extractor' component. The page layout includes a header with the KNIME logo and navigation links, a search bar, and a main content area. The component title 'Twitter URLs Extractor' is highlighted with a yellow box (1). Below the title, there is a description of the component's functionality, a list of used extensions and nodes, and a legal notice. The 'Copy link' button is highlighted with a yellow circle (2). The 'Like' button and download count are highlighted with a yellow circle (3). The 'Copy link' button is also highlighted with a yellow circle (4).

Figure 8. A component page on KNIME Hub

Workflows

A workflow page typically has multiple useful information about the workflow such as:

1. The workflow title and tags and whether it is a draft or a version: For drafts, the time of the latest edits is indicated. For versions, the version name and date of creation is indicated. Find out more about versions in the [versioning section of this guide](#)
2. The space owner where the workflow is contained and, if different from the space owner, the workflow developer
3. The workflow drag & drop element: You can drag & drop the element into a selected mount point in the space explorer of KNIME Analytics Platform. See the [Drag & drop section](#) for more details on this feature.
4. The workflow likes, the count of downloads, and a *Copy link* button to copy the workflow short link.

For each workflow you can also find:

- An *Overview* with a description of the workflow when provided by the workflow developer
- *External resources* links to external resources such as KNIME blog posts, KNIME documentation, or any other interesting link that the workflow developer might want to provide
- In the section *Used extensions & nodes* you can see the you can see the workflow extensions and nodes, and the KNIME Analytics Platform that has been used to create the workflow

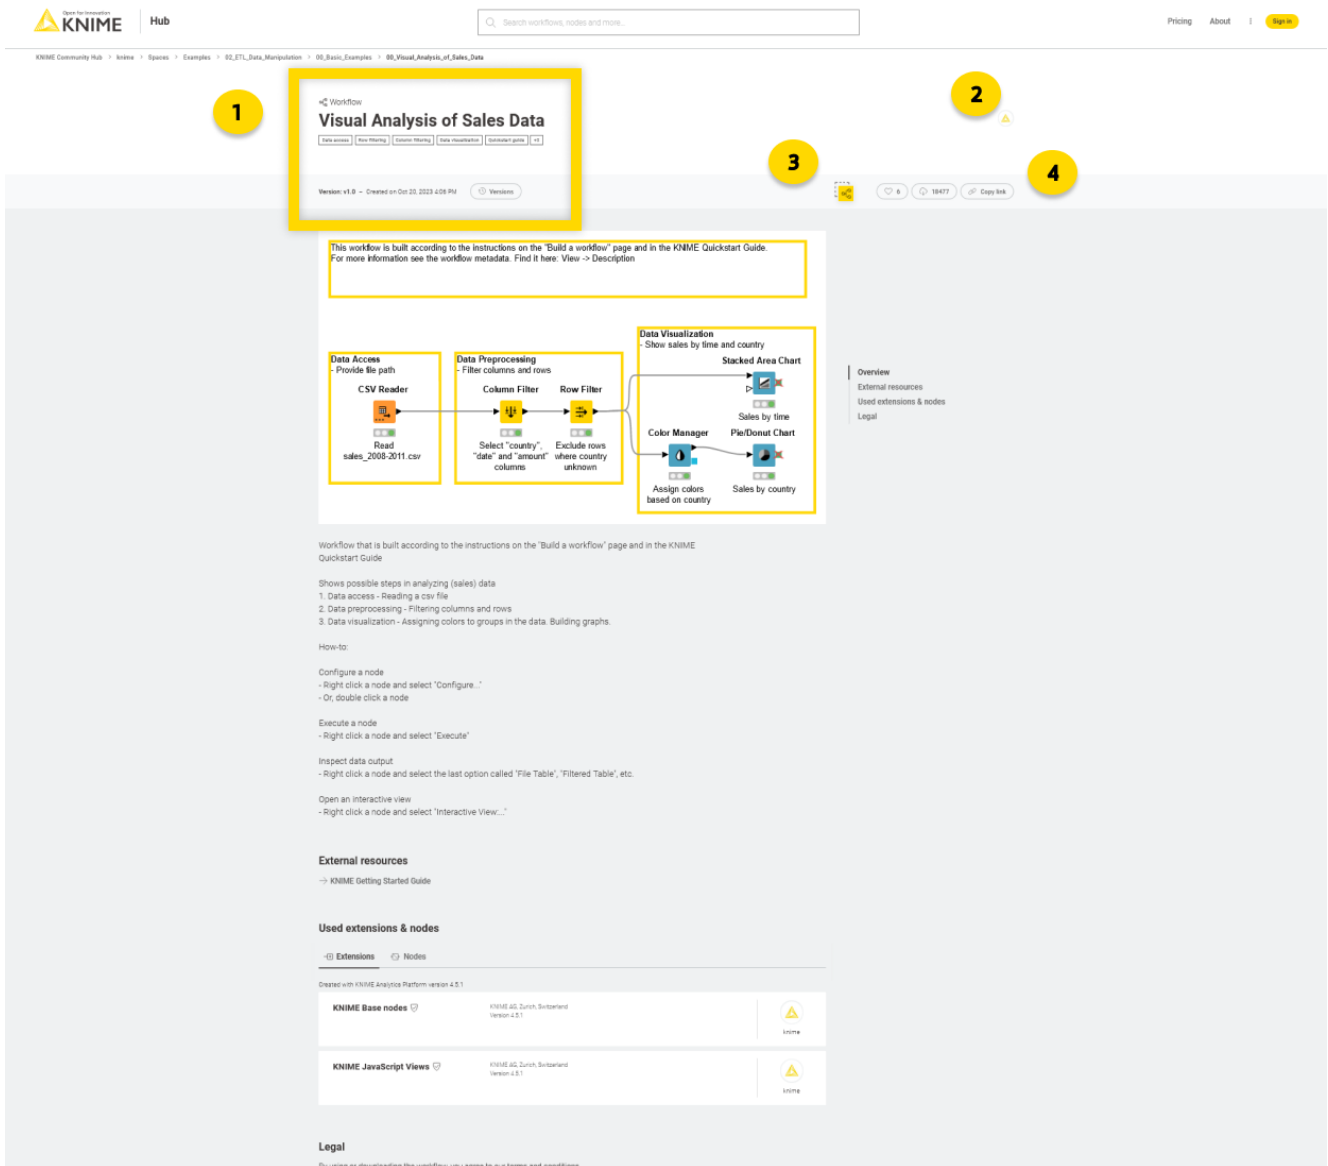


Figure 9. A workflow page on KNIME Hub

Extensions

Extensions are collection of nodes that provide additional functionality such as access to and processing of complex data types, the use of advanced algorithms, as well as the use of scripting nodes, and so on.

KNIME Extensions are developed and maintained by KNIME, allowing you to access open source projects and add their functionality to your KNIME workflows. Community Extensions instead include functionality specific to various industries and domains. Some of these community extensions are classified as Trusted Community Extensions, which have been tested for backward compatibility and compliance with the KNIME usage model and quality standards, and Experimental Community Extensions, which come directly from the labs of our community developers. Finally, also Partner Extensions are available which provide

additional capabilities offered and maintained by our partners.

All these Extensions are available on KNIME Hub.

An extension page typically has multiple useful information about the extension such as:

1. The extension name and owner, version of the extension
2. The extension drag & drop element: You can drag & drop the element into the Workflow Editor of KNIME Analytics Platform. If the extension is not already installed this will prompt an *Install Extension* window which allows you to install the extension. See the [Drag & drop section](#) for more details on this feature.
3. The extension likes, and a *Copy link* button to copy the extension short link.

For each extension you can also find:

- An *Overview* with a description of the extension when provided
- *Included nodes* with a list of all the nodes that are part of the extension and that will be available once the extension is installed
- In the section *Related workflows*, you are provided with *Workflows* that are available on KNIME Hub and that contain the nodes that are part of the extension
- Finally, in the section *Legal & update site*, you can have legal information about the copyright and the update site information, with the type of the extension, version number, and the link to the [update site](#).

Collections

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 collection, for example, contains:

- Workflow examples:

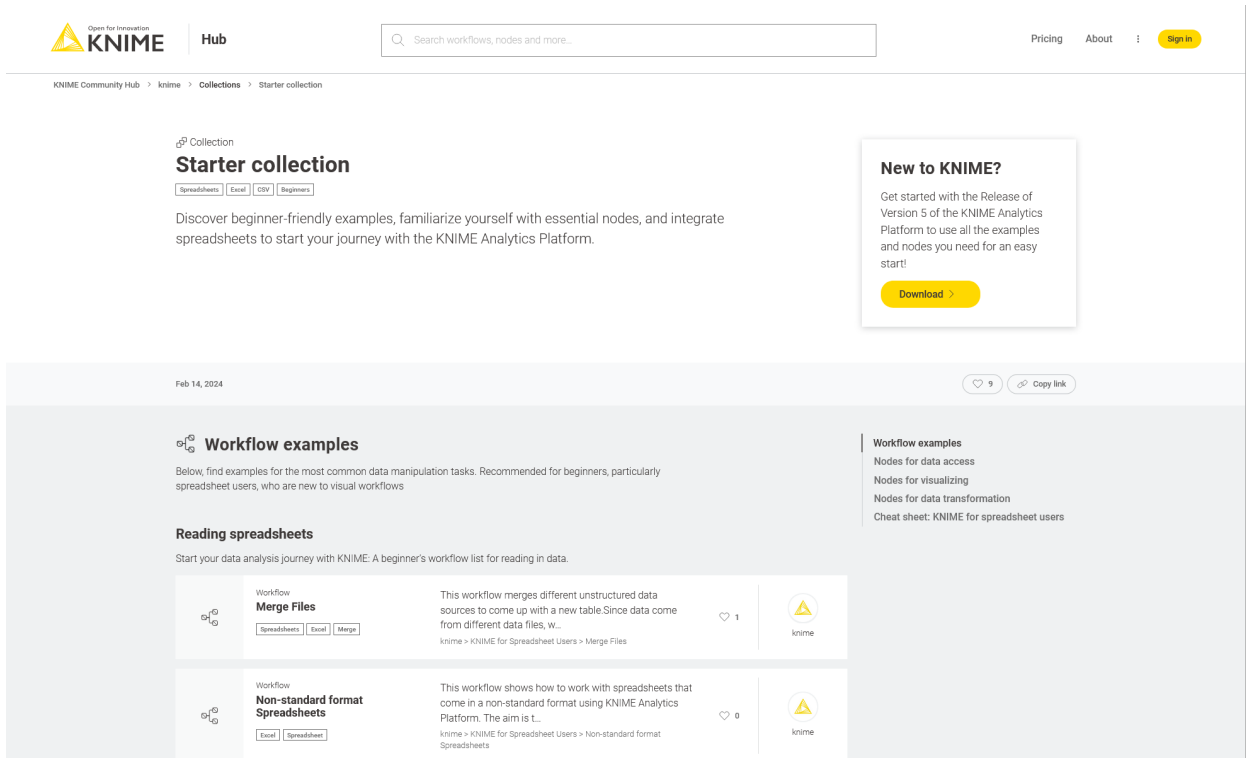


Figure 10. Example workflows in a collection on KNIME Hub

- Nodes:

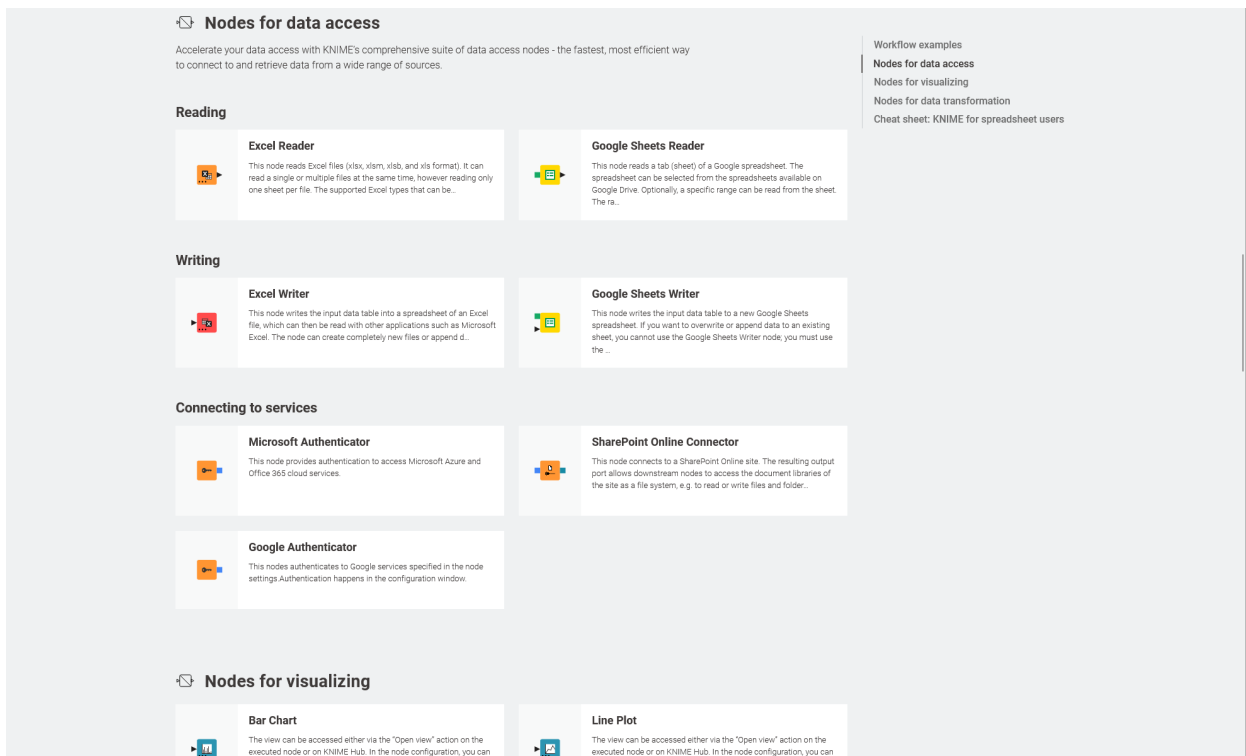


Figure 11. Nodes in a collection on KNIME Hub

- Links:

Type conversion

- String To Number**
Converts strings in a column (or a set of columns) to numbers. If the node fails to parse a string, it will generate a missing cell and append a warning message to the KNIME Console with detailed inf...
- String to Date&Time**
Parses the strings in the selected columns and converts them into Date&Time cells. The expected format can be selected from a number of commonly used formats or specified manually (see section "Type ...

Date & time handling

- Extract Date&Time Fields**
Extracts the selected fields from a Local Date, Local Time, Local Date Time or Zoned Date Time column and appends their values as corresponding integer or string columns.
- Date&Time Shift**
The node shifts a date or time with a defined duration or granularity. The user can select the Date&Time columns to shift and the shift value. The shift value can be achieved by using either a durati...

Cheat sheet: KNIME for spreadsheet users

Use this cheat sheet, designed as a companion just for KNIME Analytics Platform Version 5, as a quick reference for the most popular nodes and their purpose.

→ [Access cheat sheet](#)

Workflow examples
Nodes for data access
Nodes for visualizing
Nodes for data transformation
Cheat sheet: KNIME for spreadsheet users

KNIME
Open for Innovation
KNIME AG
Talacker 50
8001 Zurich, Switzerland

Software
Getting started
Documentation
Courses + Certification
Solutions

KNIME Hub
KNIME Forum
Blog
Events

Partner
Developers

KNIME Home
Careers
Contact us

[Download KNIME Analytics Platform](#) [Read more on KNIME Business Hub](#)

© 2024 KNIME AG. All rights reserved. Trademarks Imprint Privacy Terms & Conditions Data Processing Agreement Credits

Figure 12. Additional links in a collection on KNIME Hub

Drag & drop

You can drag & drop nodes, components, extensions, and workflows from KNIME Hub to import them into KNIME Analytics Platform and use them right away to build your own workflow, install KNIME Extensions, and execute uploaded workflow into your local installation.

- **Nodes and components:** You can drag & drop a node or a component from KNIME Hub into your Workflow Editor in KNIME Analytics Platform. If the node is part of an extension that is not yet installed in your local installation of KNIME Analytics Platform, or if the component contains nodes that are part of extensions that are not yet installed, you will be asked if you want KNIME Analytics Platform to search and install the

missing extension(s), as shown in **Figure 13**.

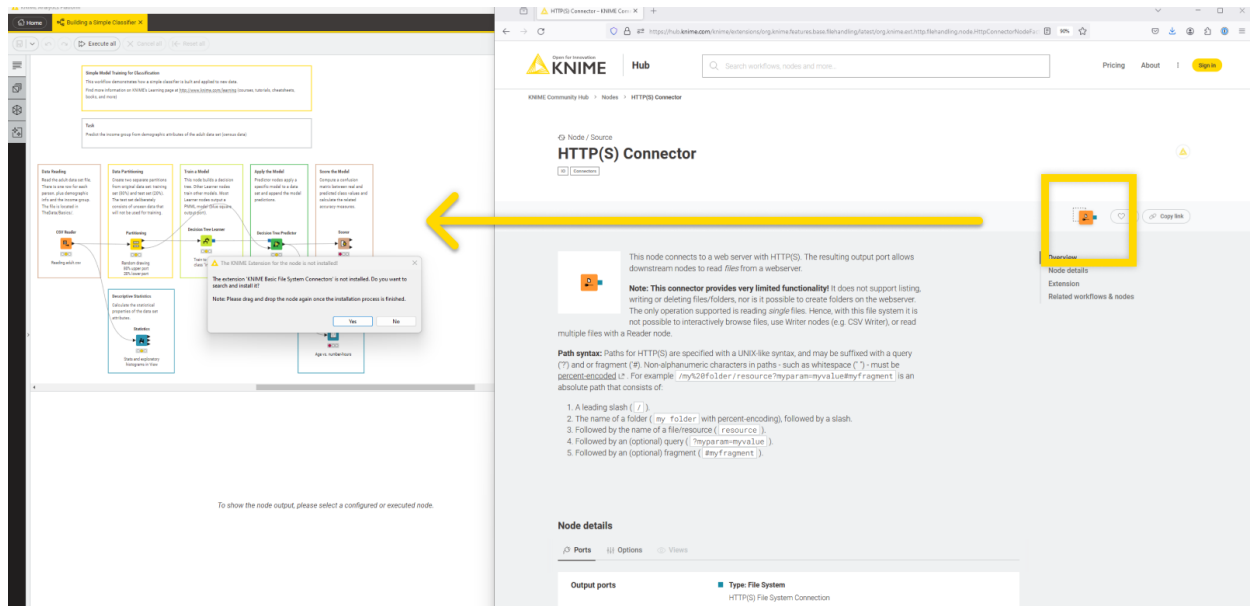


Figure 13. Drag & drop a node from KNIME Hub to KNIME Analytics Platform and install missing extension

- **Extensions:** You can drag & drop a specific extension into the Workflow Editor. KNIME Analytics Platform will search and install it. In order to be able to install the extension the extension's **update site** that is indicated at the end of the extension page in the section **Legal & update site** must be activated in KNIME Analytics Platform.

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.

The team can own public or private spaces. For more details see the section [Team owned spaces](#).

The items that are stored in a team's **public** space will be accessible to everyone with access to the Business Hub instance and will be presented as search results when searching on KNIME Hub. However, to upload and download to the public spaces of the team, you have to be a team member.

A team's **private** space, allows KNIME Hub users that are part of a team to collaborate privately on a project. Only the team members have access to the items that are stored in a team's private space.

Create a team

Teams are firstly created by the [KNIME Hub Global Administrator](#). The global admin will assign a user the role of team admin as well allocate resources to that team. The team admin will be able to start adding team members and create [execution contexts](#) that can be used by the team.

There are two types of roles a user can be assigned when part of a team:

- Administrator. A team administrator can:
 - [Add/remove users from the team](#)
 - [Promote team members to admins](#)
 - [Create execution contexts](#)
- Member. A team member can:
 - View the team page, with members list and spaces
 - [Create, modify, and delete public and private spaces](#)
- [Upload/download items to/from public and private spaces](#)
- [Delete items from public and private spaces](#)

The team creator is automatically assigned administrator role and can promote any of the team members to administrators. In order to do so please follow the instructions in the


section [Manage team members](#).

Team owned spaces

A team can own an unlimited number of both public and private spaces.

- Team owned **public** spaces: The items that are stored in a team's public space will be accessible to everyone and will be presented as search results when searching on KNIME Hub. Only team members will have upload rights to the public spaces of the team.
- Team owned **private** spaces: Only the team members instead have read access to the items that are stored in a team's private space. This will then allow KNIME Hub users that are part of a team to collaborate privately on a project.

On the team's profile page, you will find the *Create new space* tile. Click *Private space* to create a private space for your team, or *Public space* to create a public space. You can then change the name of the space, or add a description. You can change or add these also later by going to the relative space page and clicking the space name or *Add description* button to add a description for the space.

You can also change the visibility of the space from private to public and vice-versa, if you are a team admin, or delete the space. To do so, from the space page click the  icon, as shown in the image below.

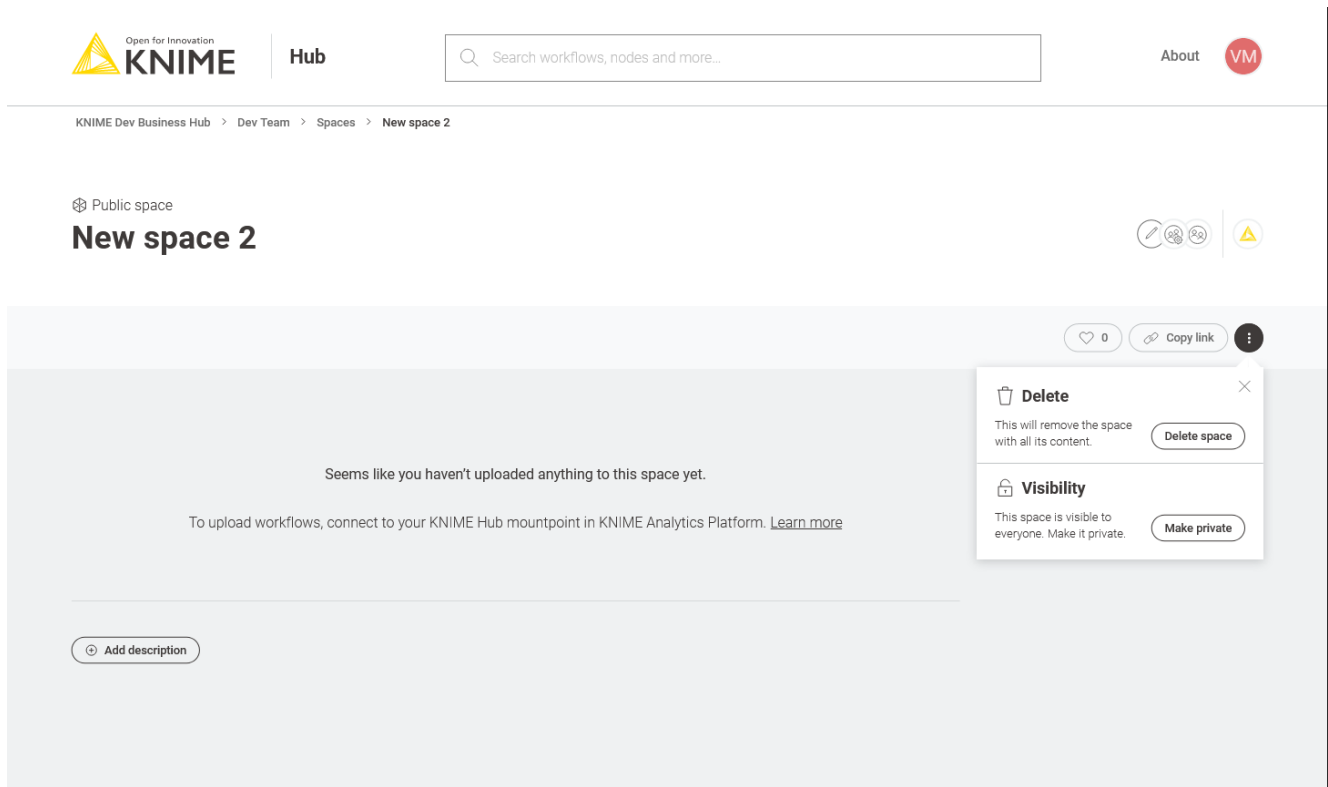



Figure 14. Change visibility or delete an existing team owned space

Manage space access

You can also manage the access to a specific space. To do so navigate to the space and click the  icon.

In the *Manage space access* side panel that opens, you can add other team members. You can change the rights they have on the items in the space - e.g. you can grant them *Edit* rights or *View* rights for this space, as shown in [Figure 15](#).

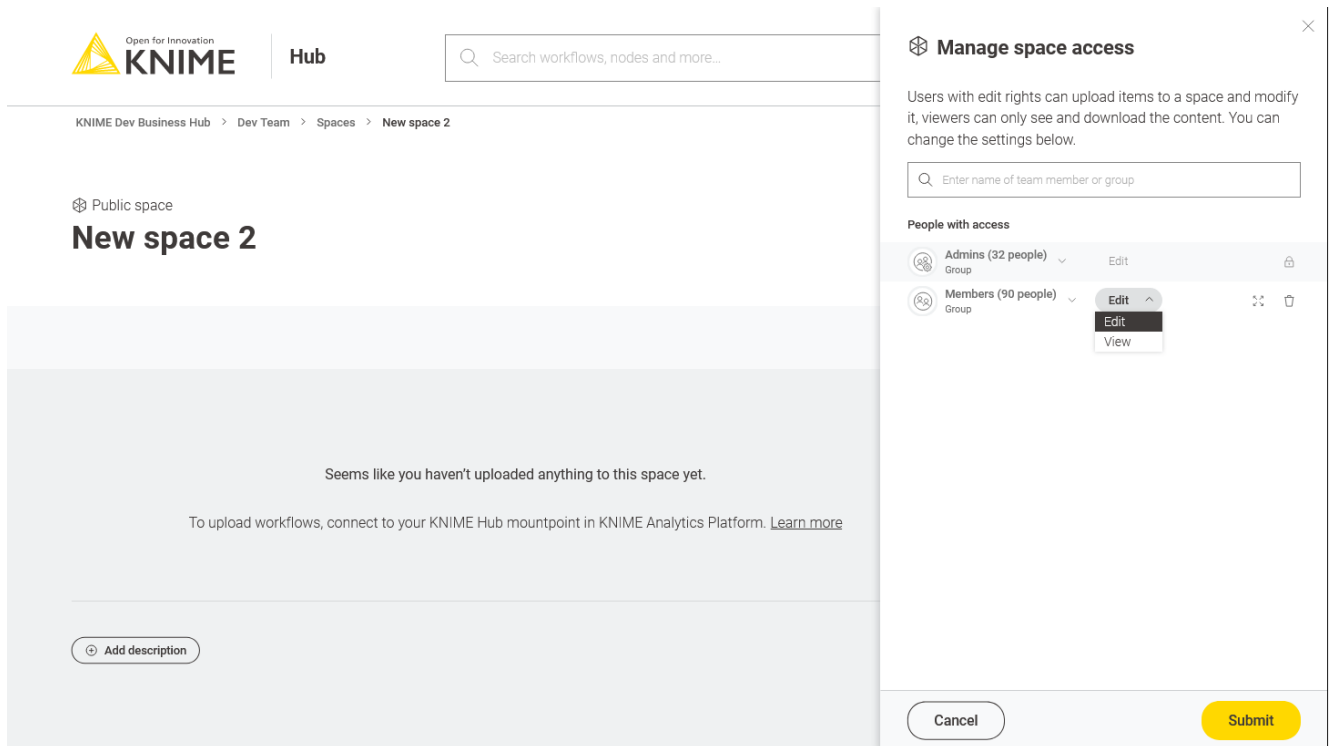


Figure 15. Manage the access to a space

Manage team members

You can manage your team by going to the team’s profile. To do so, click your profile icon in the top right corner of KNIME Hub.

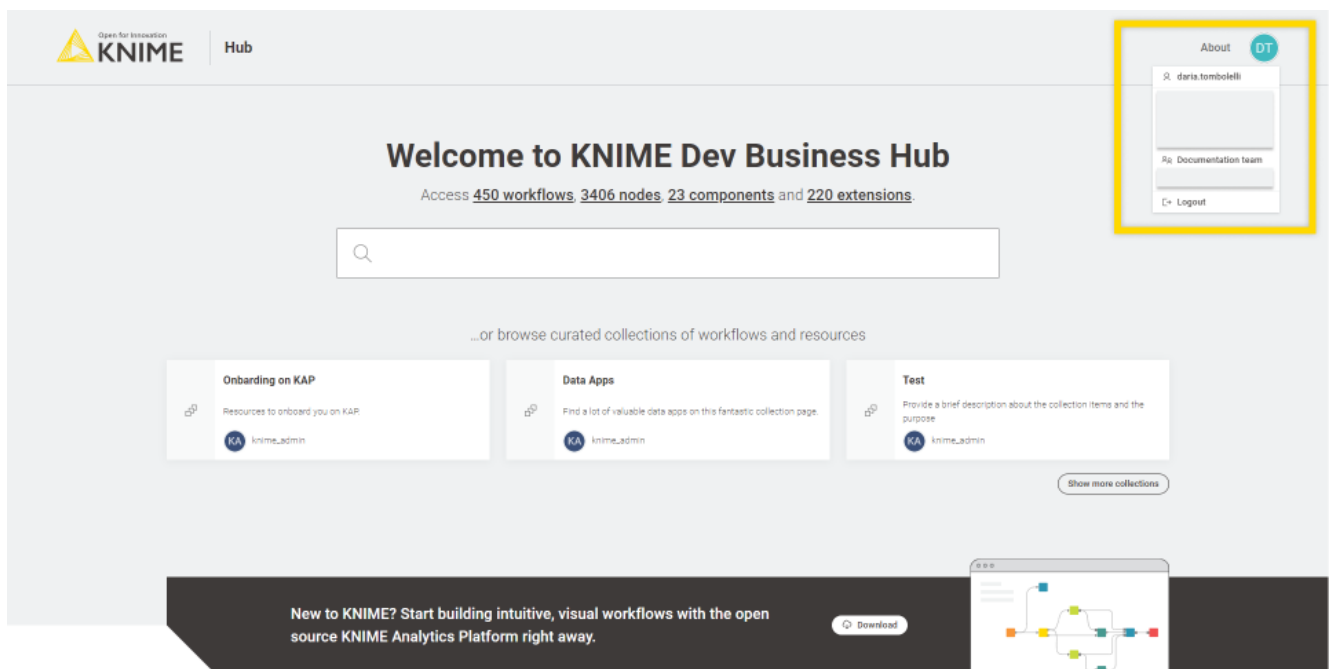


Figure 16. Select a team to go to the team’s profile page

In the dropdown menu that opens you will see your teams. Select the team you want to manage to go to the corresponding team's profile page.

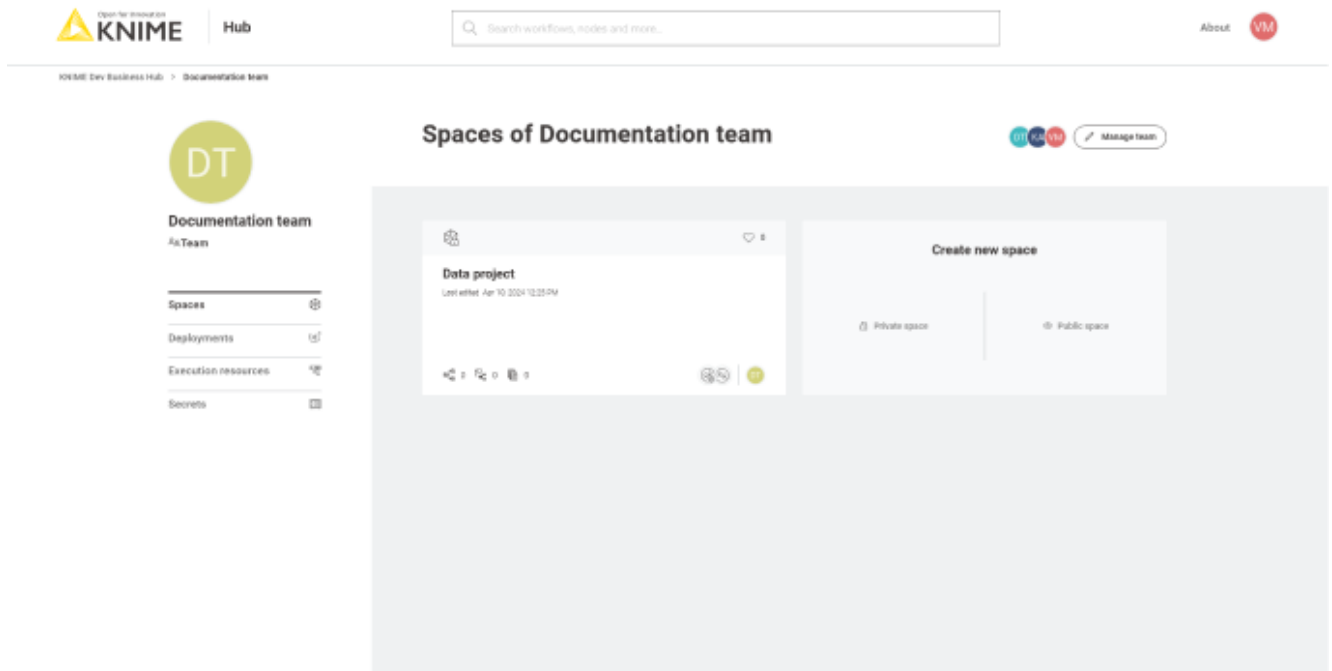


Figure 17. The team's profile page

Here, you can click the *Manage team* button to open the *Manage members* side panel, as shown in the image below.

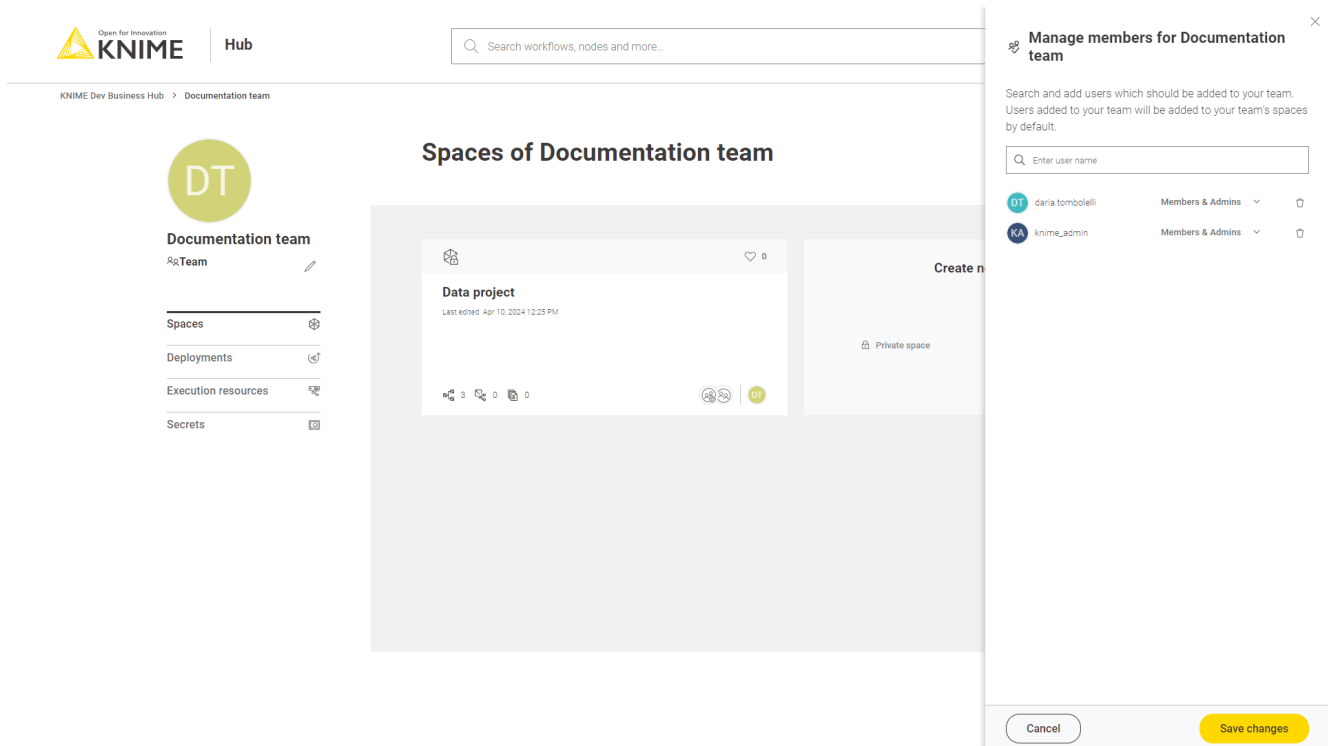


Figure 18. Manage team members

You will see here a list of the team members and their assigned role.

From here a team admin can change the role of the team members. To do so click the drop down arrow close to the name and select the roles you want to assign to each user.

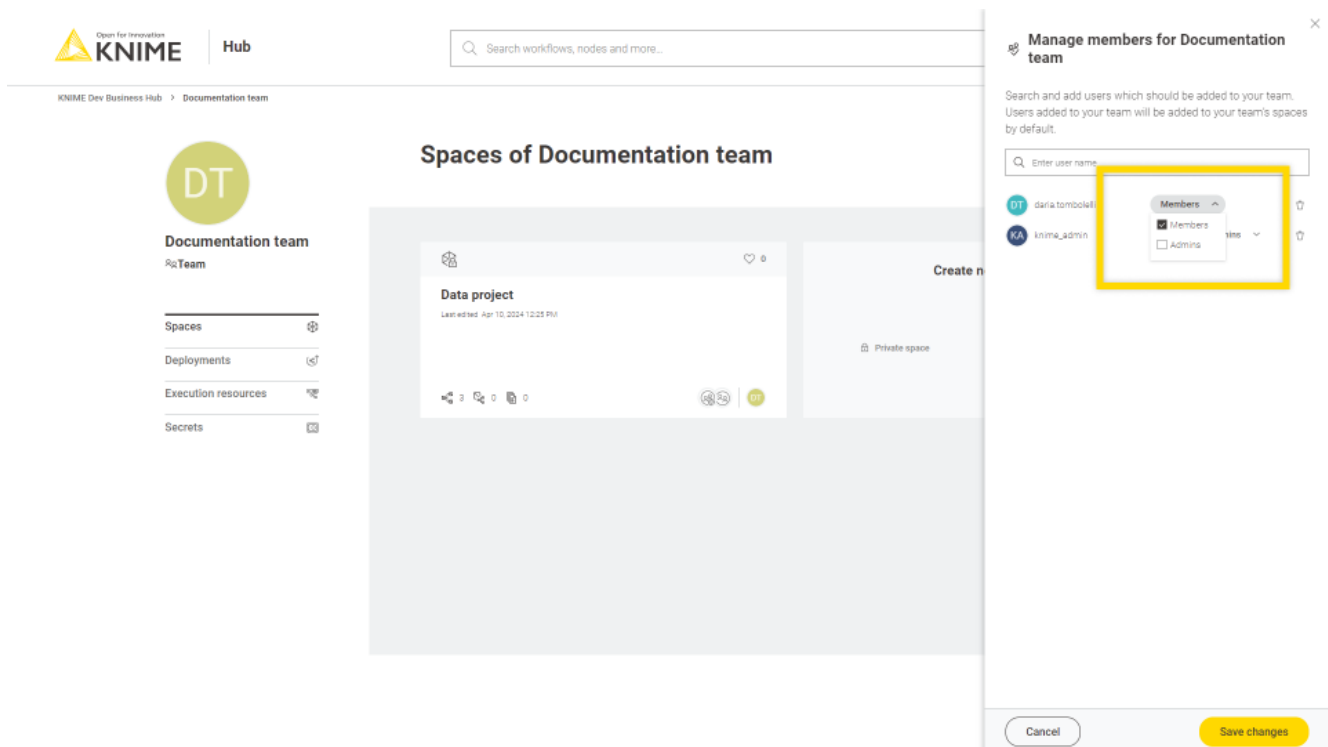


Figure 19. Manage team members roles


Then click **Save changes** button to apply your changes.

Add members to a team

To add a new member enter the user name of the users that you want to add to the team in the corresponding field in the *Manage members* panel Click *Save changes* button to apply your changes.

In case you added more users than allowed by your license you will be notified with a message. Please **remove the exceeding users** or purchase more users.

Delete members from a team

To delete a member, go to the *Manage members* panel and click the  icon for the user you want to delete. Then click the *Save changes* button to apply your changes.

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.

If you are a team admin and want to create and manage customization profiles please follow the guide in the **KNIME Business Hub Admin Guide**.

As a user of the KNIME Business Hub you can then download and use customization profiles

available for all the users or for the team members of your teams.

Apply a customization profile to KNIME Analytics Platform client

In order to apply a customization profile in a local KNIME Analytics Platform installation, you need to follow these steps:

- Go to your KNIME Analytics Platform client, and log in to the KNIME Business Hub.
- Click *Preferences* and, under *KNIME > Customization Profiles* activate the option *Enable managed customizations*.
- Select the Business Hub mount point and the available profiles will show up. You can select one or more profiles and you can also change their order. Preferences in selected profiles further down in the list will override preferences from profiles further up in the list.
- Click *Apply and Close* and you will be asked to restart the client.

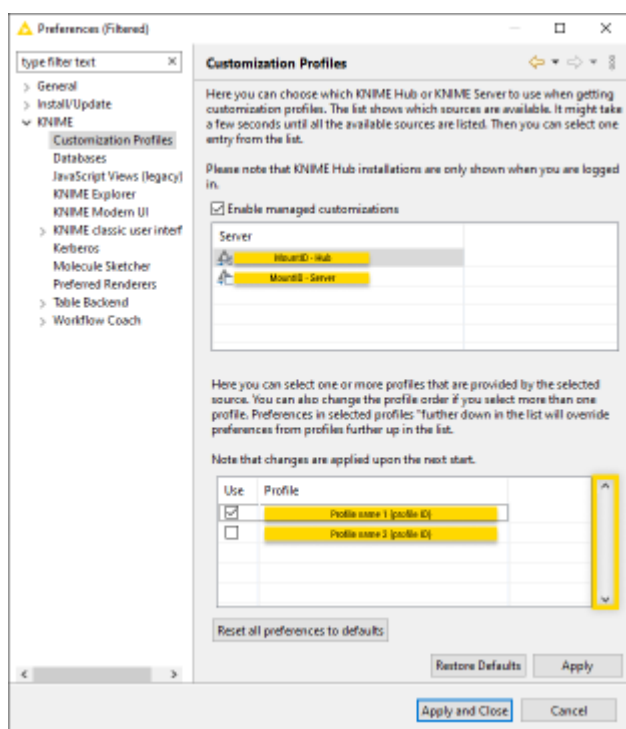


Figure 20. Customization profiles preferences window

Apply a customization profile to KNIME Analytics Platform via `knime.ini`

Another possibility to apply a customization profile in a local KNIME Analytics Platform installation, is to add it to the installation's `knime.ini` file.

If you have already changed your default settings in the KNIME Analytics Platform installation

these will not be overwritten. However, you can choose to *Restore all preferences to defaults* via the preferences page in the KNIME Analytics Platform, before the restart.

Then add the following lines to the `knime.ini`, right before the `-vmargs` line. Please note the line breaks. I.e., these need to be four individual lines in your `knime.ini`.

```
-profileLocation https://api.<base-url>/execution/customization-profiles/contents -profileList  
<profile_ID>
```



The access to the customization profile is not restricted meaning that anyone with the link can use it by adding the URL to the `knime.ini` file as explained above.

Detach a customization profile from KNIME Analytics Platform client

In order to detach a customization profile from a local KNIME Analytics Platform installation, you need to follow these steps:

- Go to your KNIME Analytics Platform client, and log in to the KNIME Business Hub.
- Click *Preferences* and, under *KNIME > Customization Profiles* deselect the customization profiles you want to detach.
- Click *Apply and Close* and you will be asked to restart the client.

Upload items to KNIME Hub

Once you have added the **KNIME Business Hub mount point** and you are connected to your KNIME Hub account from KNIME Analytics Platform you can upload the desired items to your KNIME Hub spaces.

You can upload workflows, components or files to any of your team owned spaces by right-clicking the item from the space explorer in KNIME Analytics Platform and selecting *Upload* from the context menu. A window will open where you will be able to select the location where you want to upload your workflow or component.

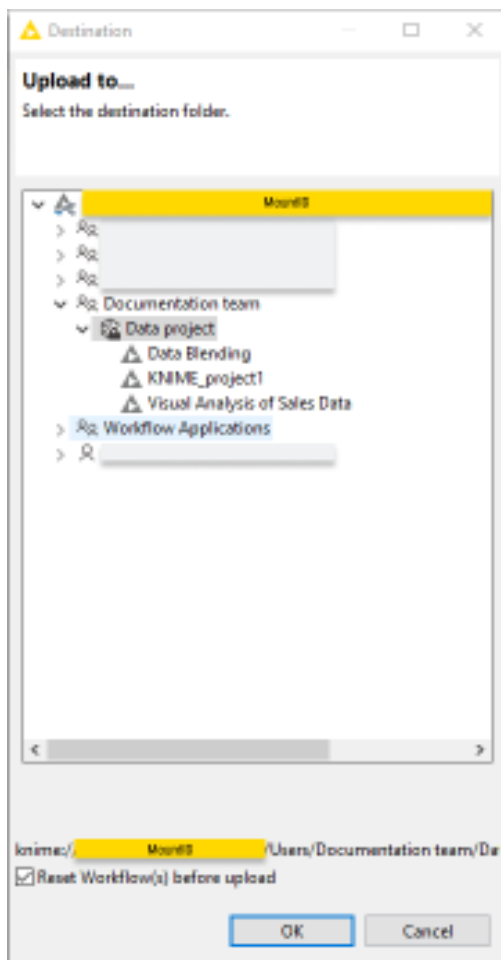


Figure 21. Upload a local item to your KNIME Hub mount point

Please notice that if the items are uploaded to a public space they will be available to everyone, hence be very careful with the data and information (e.g. credentials) you share.

Move items within KNIME Hub


You can move items that you uploaded to KNIME Hub to a new location within the space that contains the item or to a different space that you have access to. To do this you need to be

connected to the KNIME Hub mount point on KNIME Analytics Platform. You can then move the items within KNIME Hub just by dragging the item in the space explorer.

Delete items from KNIME Hub

You can also delete items that you uploaded to KNIME Hub.

To do so you can:

- Connect to KNIME Hub mount point on KNIME Analytics Platform. Right-click the item you want to delete and select *Delete* from the context menu.
- From KNIME Hub, sign in with your account and go to the item you want to delete. Click the  icon on the top right of the page and select *Delete*.

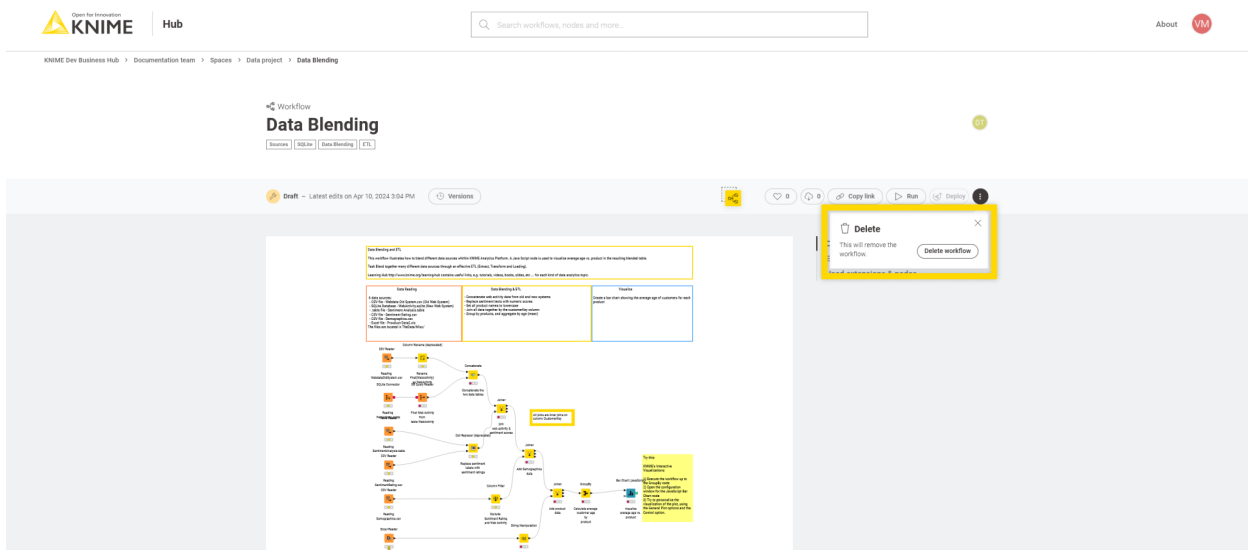




Figure 22. Delete a workflow from KNIME Hub

Download items from KNIME Hub

Of course, you can also download items from KNIME Hub to use them in your local KNIME Analytics Platform installation. You can do this from the space explorer by right-clicking the workflow from the KNIME Hub mount point and selecting *Download to local space*. Items that are downloaded like this, are directly saved to your local space, which you can access from the space explorer.

Alternatively, you can download an item from KNIME Hub in the browser. To do so, go to the specific workflow or component on KNIME Hub and click the  icon to download.



Items that are downloaded from your browser may not be saved in a folder that KNIME Analytics Platform can access directly. However, you can import the file to your KNIME Analytics Platform, by clicking the  icon in the top right corner of the space explorer and selecting either *Import workflow* or *Add files* to access items saved in your local file system.

Versioning

When **uploading items to a space on KNIME Hub** you will be able to keep track of their changes.

Your workflow or component will be saved as a draft until a version is created.

When you create versions of the items, you can then go back to a specific saved version at any point in time in the future to download the item in that specific version.

Once a version of the item is created, new changes to the item will show up as draft.

Create a version of an item

Go to the item you want to create a version of by navigating through the KNIME Business Hub instance and click *Versions*.

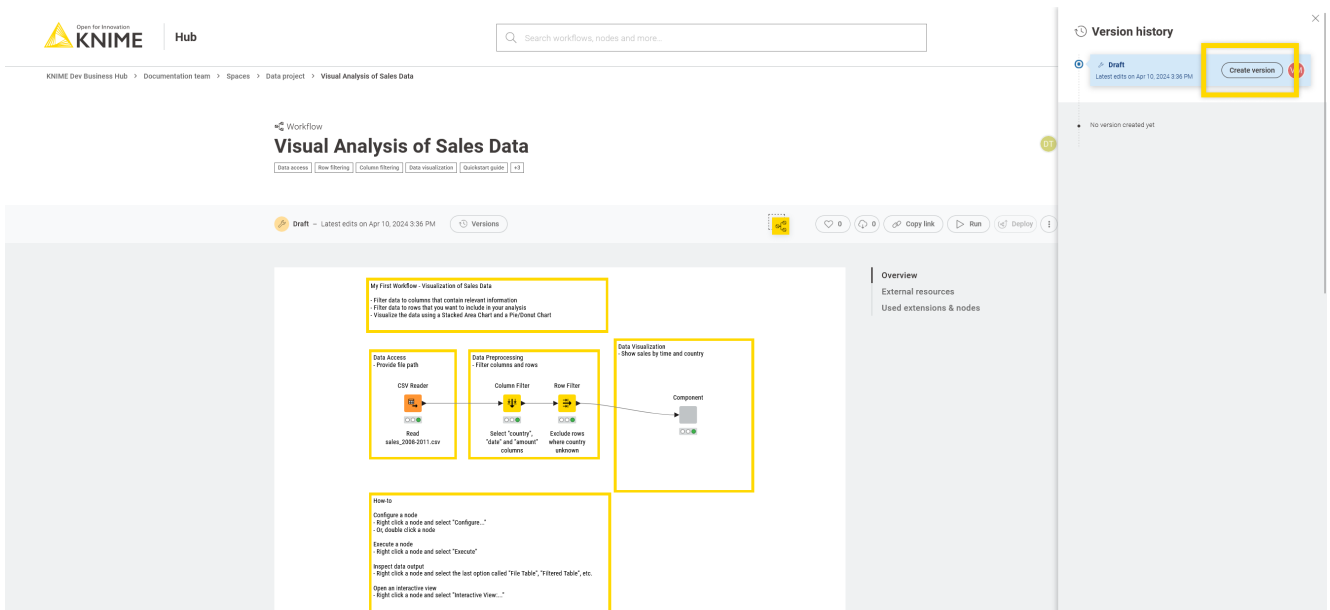



Figure 23. Versioning of a workflow

A panel on the right opens where you can see all the versions already created and all the unversioned changes of the item since the last version was created.

Click *Create* to create a new version. You can then give the version a name and add a description.

Show a version

In the *Version history* panel click the version you want to see or click the  icon and select

Show this version. You will be redirected to the item in that specific version.

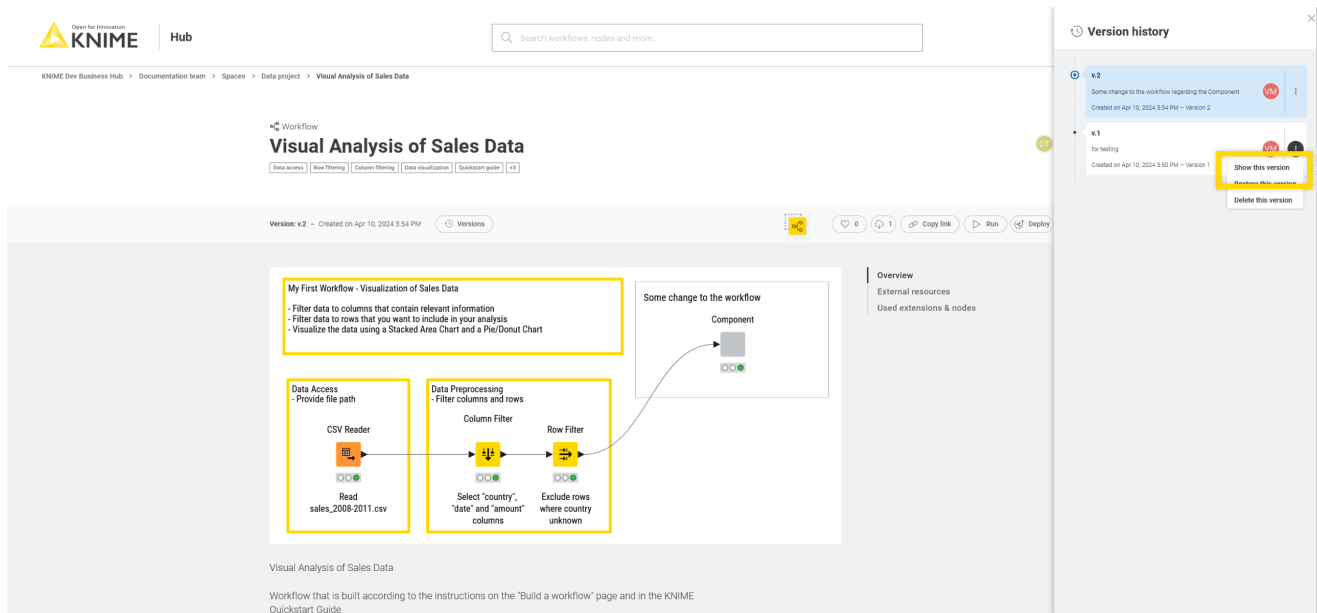



Figure 24. Show the versions of a workflow


To go back to the latest state of the item click the selected version again to unselect it.

Restore a version

To restore a version that you created click the  icon in the version tile from the item history and select *Restore this version*.

The latest state contains now the corresponding changes as unversioned edits.

Delete a version

In the *Version history* panel click the  icon for the version you want to navigate to and click *Delete this version*. Only Team administrators can delete a version.

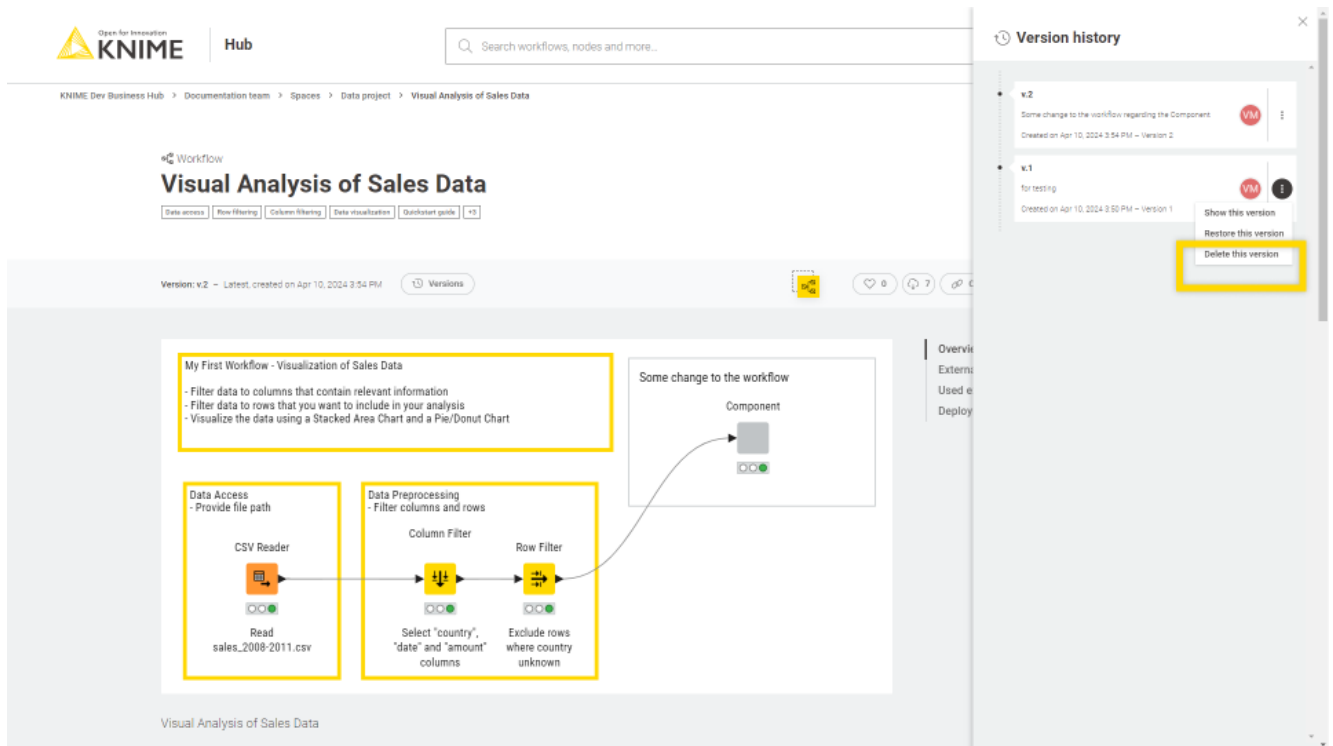


Figure 25. Delete a version

Execution of workflows on KNIME Hub

You can execute workflows on KNIME Hub. In order to do so the team you are a member of needs to have an execution context assigned to.

A team can own multiple execution contexts which are dedicated execution resources with configured execution settings.

Execution contexts

An execution context provides dedicated execution resources with configured execution settings for the execution and deployment of workflows on KNIME Hub.

They operate using a selected custom executor image, which defines custom hardware and KNIME Analytics Platform configuration for the deployments.

Execution contexts are owned and managed by the teams.

A default execution context can be assigned to a space. It is possible to create as many execution contexts as necessary. They can be used for specific deployments, for example in a development, testing, production configuration.

Execution contexts can be created by the team admin once the Docker image of the executor is available. The latter is maintained by the [global Hub admin](#). [Shared execution contexts](#) can instead be created by a Hub admin and shared with your team.




Your team can by default have a maximum of 10 execution contexts. If you need more please contact your KNIME Hub admin who can [increase the limit](#) if necessary.

Create a new execution context

The team admin can create a new execution context.

To do so go to the team profile page and click *Execution resources* in the menu on the left.

Here you will see all the available *Execution resources* of your team. You can see how many vCores are available and in use for each execution context, the version of the executor, its memory usage and CPU load.

To create a new execution context for your team click the  button.




Figure 26. An overview of the execution resources available to your team

In the side panel that opens you can define your execution context.




Figure 27. Create a new execution context

Here you can specify:

- Name: is the name of the execution context

- **Docker image:** is the name of the Docker image of the executor that you can obtain from your [KNIME Hub Global Administrator](#).

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
- registry.hub.knime.com/knime/knime-full:r-5.2.4-564
- registry.hub.knime.com/knime/knime-full:r-5.2.5-593
- registry.hub.knime.com/knime/knime-full:r-5.3.0-388
- registry.hub.knime.com/knime/knime-full:r-5.3.1-498
- knime/knime-full:r-5.3.2-564



In order to have access to the **execution context logs** for debugging purposes the execution context needs to be based on the following executor Docker images:

i

- `registry.hub.knime.com/knime/knime-full:r-5.2.5-593` or higher bugfix releases of the 5.2.x release line
- `knime/knime-full:r-5.3.2-564` or higher bugfix releases of the 5.3.x release line
- Any new major version released (e.g. 5.4) of the executor will also support this feature

i

With the release of 5.3.2 the docker image name **does not need** the `registry.hub.knime.com` prefix anymore. Find more information about KNIME Executor images provided by KNIME [here](#).

- **Blacklisted nodes:** is a list of nodes that should be blocked by the executor. You need to provide the full name of the node factory. You can get the factory name from the Hub itself by looking for the node. The last part of the URL is the node factory name. For example the Java Snippet node can be found at the following URL: <https://hub.knime.com/knime/extensions/org.knime.features.javasnapshot/latest/org.knime.base.node.jsnapshot.JavaSnapshotNodeFactory> and `org.knime.base.node.jsnapshot.JavaSnapshotNodeFactory` is its factory name. Another way to determine the factory names of the nodes you want to block is to create a workflow with all nodes that should be blacklisted. After saving the workflow you are able to access the `settings.xml` of each node under `<knime-workspace>/<workflow>/<node>/settings.xml`. The factory name can be found in the entry with key "factory".
- **Number of executors:** is the number of executors that the execution context has assigned
- **Number of vCores per executor:** is the number of vCPUs per executor
- **RAM per executor (GB):** is the amount of RAM per executor in GB

You can assign an execution context as the default for a specific space. To do so you will need to perform a PUT request to the `api.<base-url>/execution-contexts/{uuid}` where `{uuid}` is the execution context id, adding the following to the request body:



```

},
  "defaultForSpaces": [
    "*space1",
    "*space2"
  ]
}

```

where `["*space1", "*space2"]` is the list of spaces for which the execution context is the default execution context.

- **Advanced settings:** 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 *On* 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.

Advanced configuration of execution contexts

Execution contexts can be created and edited also via the Business Hub API. Here more advanced configuration are possible.

- `"defaultCpuRequirement": 0.0`,: Specifies the default CPU requirement in number of cores of jobs without a specific requirement set. The default is `0.0`.
- `"blacklistedNodes": ""`,: Add here if necessary a list of nodes that should be blocked by the executor. See the section about the [blacklisted nodes](#) for more information on the format.
- `"defaultSwapTimeout": "PT1M"`,: Specifies how long to wait for a job to be swapped to disk. If the job is not swapped within the timeout, the operation is canceled. The default is `PT1M` equal 1 minute. This timeout is only applied if no explicit timeout has been passed with the call (e.g. during execution context shutdown).
- `"perJobLogging": true`,: Enables a job trace log which records important operations on any job (loading, execution, discarding). The job trace log is enabled by default.
- `"updateLinkedComponents": false`,: Specifies whether component links in workflows should be updated right after a deployment of the workflow has been created to use the

execution context. Default is to not update component links.

- "defaultRamRequirement": 0, : Specifies the default RAM requirement of jobs without a specific requirement set. In case no unit is provided it is automatically assumed to be provided in megabytes. The default is 0MB.
- "defaultReportTimeout": "PT5M", : Specifies how long to wait for a report to be created by an executor. If the report is not created within the timeout, the operation is canceled. The default is PT5M equal 5 minutes. This timeout is only applied if no explicit timeout has been passed with the call.
- "defaultExecutionTimeout": "PT5M", : Specifies the default timeout when executing a job synchronously. If the executor does not respond within the provided time it will cancel the execution and throw an error. The default is PT5M equal 5 minutes.
- "maxJobTimeInMemory": "PT1H", : Specifies the time of inactivity before a job gets swapped out from the executor. The default is PT1H equal 1 hour. Negative numbers disable swapping.
- "startExecutionTimeout": "PT1M", : Specifies the timeout the service will wait until execution has been started by the executor. If the executor does not respond within the provided time it will cancel the execution and throw an error. The default is PT1M equal 1 minute.
- "maxLoadedJobsPerExecutor": 2147483647, : Specifies the maximum number of jobs that can be loaded at the same time into a single executor. The executor will not accept any additional jobs if this number has been reached until an existing job gets unloaded. The default is unlimited.
- "maxJobLifeTime": "PT168H", : Specifies the time of inactivity, before a job gets deleted. The default is PT168H equal 7 days. Negative numbers disable forced auto-discard.
- "maxSwapFailures": 1, : Specifies the number of times a job is attempted to be swapped if swapping fails initially.
- "maxJobInOutSize": 10, : Specifies the maximum amount of data to be passed into and out of jobs using Container Input/Output nodes. Defaults to 10 MB. If larger values are needed, the amount of memory allocated to the rest-interface service should be increased as well to avoid stability issues.
- "defaultLoadTimeout": "PT5M", : Specifies the default timeout when loading a job. If the job does not get loaded within the timeout, the operation is canceled. The default is PT5M equal 5 minutes. This timeout is only applied if no explicit timeout has been passed with the call.
- "maxJobExecutionTime": "PT24000H", : Specifies a maximum execution time for jobs. If

a job is executing longer than this value it will be canceled and eventually discarded (see `discardAfterMaxExecutionTime` option).

- `"discardAfterMaxExecutionTime": false`,: Specifies whether jobs that exceeded the maximum execution time should be canceled and discarded (`true`) or only canceled (`false`). May be used in conjunction with `maxJobExecutionTime` option. The default (`false`) is to only cancel those jobs.
- `"customizationProfiles": []`,: Specifies **customization profiles** to apply to the executors running in the execution context.
- `"saveWorkflowSummary": false`,: Specifies if the workflow summary should be stored with the job upon swapping. Default value is `false`.
- `"rejectFutureWorkflows": true`,: Specifies whether the executor should reject loading workflows that have been create with future versions of KNIME Analytics Platform client. For new installations the value is set to `true`. If no value is specified the executor will always try to load and execute any workflow by default.
- `"defaultAsyncLoadTimeout": "PT5M"`,: Specifies the default timeout when loading a job asynchronously. The default is `PT5M` equal 5 minutes.

HTML sanitization of JavaScript View nodes and Widget nodes


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 on the possible consequences of node's functionality see the following [section](#) in the KNIME Analytics Platform guide.


It is still possible to achieve the old behavior on each execution context, by turning the sanitization off for that specific execution context.

To do so you need to perform a PUT request to `api.<base-url>/execution-contexts/{uuid}` where `{uuid}` is the execution context id, adding the following to the request body:

```
{
  "operationInfo": {
    "vmArguments": [
      "-Djs.core.sanitize.clientHTML=false"
    ]
  }
}
```

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. You can also see more details about your team's execution context by clicking the corresponding tile or by clicking the  icon of the execution context and selecting *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.




The screenshot displays the 'Docs EC' execution context details page. At the top, it shows the KNIME logo and 'Hub' navigation. A search bar is present. The page title is 'Dedicated execution context Docs EC' with a 'Documentation team' indicator. A 'Running' status is shown with a green dot and text: 'Currently 1 job active. Stops automatically when not in use.' A 'Stop' button is visible. Below this, there are tabs for 'Overview' and 'Executor-ae6ec...'. The 'Usage' section contains two cards: 'CPU average' at 0.05% and 'vCore usage' showing 1 of 1 vCore token used. The 'Jobs' section features a table with 2 rows and columns for Workflow, Version, State, Created at, Start, Runtime, User, Node message, and Deployment.

Workflow	Version	State	Created at	Start	Runtime	User	Node message	Deployment
Visual Analysis of Sale...	v.2	Load error	Apr 11, 2024 4:18 PM		-	verena.mayer	-	Visual Analysis of...
Visual Analysis of Sale...	v.1	Execution finished	Apr 10, 2024 5:03 PM	Apr 10, 2024 5:04 PM	17m 6s	verena.mayer	-	

Figure 28. Additional executor information page

Edit an existing execution context


As a team admin you can also edit an existing execution context.

From the *Execution resources* overview page you can click the  icon of the execution context you want to modify and select *Edit* from the menu that opens.

You can change the parameters and configurations in the right side panel that opens.

Disable an execution context


As a team administrator you can disable an existing execution context.

You will need first to delete the jobs associated to the execution context then proceed with disabling it. From the *Execution resources* overview page you can click the  icon of the execution context you want to disable and select *Disable* from the menu that opens. You can enable the execution context again from the same menu.

Delete an execution context


As a team administrator you can delete an existing execution context.

You will need to first **disable** the execution context you want to delete. First, delete the jobs associated to the execution context then proceed with disabling it.

Now you are ready to delete the execution context. From the *Execution resources* overview page you can click the  icon of the execution context you want to delete and select *Delete* from the menu that opens.

Download an execution context's logs

You can download the log files of an execution context - this feature allows for debugging in case the executors are not working as expected.

From the *Execution resources* overview page you can click the  icon of the execution context and select *Download logs* from the menu that opens. You will download a zip file containing a log file for each executor of the execution context.

Please note that to be able to download job logs you need an executor based on the following executor Docker images:

- `registry.hub.knime.com/knime/knime-full:r-5.2.5-593` or higher bugfix releases of the 5.2.x release line
- `knime/knime-full:r-5.3.2-564` or higher bugfix releases of the 5.3.x release line

Any new major version released (e.g. 5.4) of the executor will also support this feature

Ad hoc execution

When a team owned space is provided with an execution context it is possible to execute the workflows that are uploaded to that space.

If you want to have a one time execution of a workflow you uploaded to a space, for example to test it, you can go to the workflow page on KNIME Hub and click *Run*.

The screenshot displays the KNIME Business Hub interface for a workflow titled "Visual Analysis of Sales Data". The workflow is shown in a visual editor with three main components highlighted by yellow boxes:

- My First Workflow - Visualization of Sales Data:**
 - Filter data to columns that contain relevant information
 - Filter data to rows that you want to include in your analysis
 - Visualize the data using a Stacked Area Chart and a Pie/Donut Chart
- Data Access - Provide file path:**
 - CSV Reader: Read sales_2008-2011.csv
- Data Preprocessing - Filter columns and rows:**
 - Column Filter: Select "country", "date" and "amount" columns
 - Row Filter: Exclude rows where country unknown

A side panel on the right shows "Overview" and "Used extensions & nodes". A yellow box highlights the "Run the workflow" button in the top right corner of the workflow editor.

Figure 29. Ad hoc execution of a workflow on KNIME Hub

A side panel opens where you can:

- Select the version on which the ad hoc execution will be performed
- If multiple execution contexts for the current space are available, select the execution context you want to use - otherwise the default execution context associated to that space will be used
- Enable workflow actions: You can choose if a notification via e-mail will be sent *On failure* or *On success*. You can also *Add more* actions and notify multiple e-mails on different conditions.

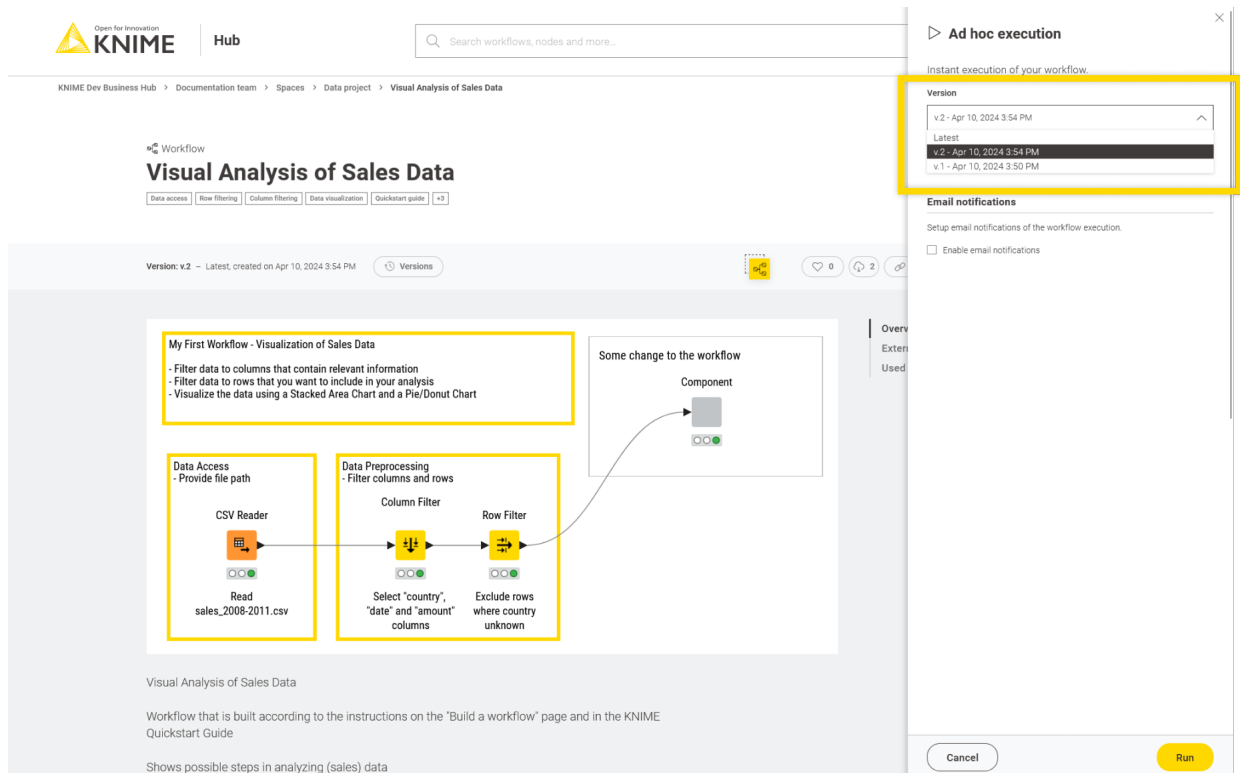


Figure 30. Ad hoc execution configuration panel

Then click *Run* and the workflow will be executed.

Finally, you can see an overview of all the ad hoc execution jobs that have been performed on a specific workflow by selecting *Ad hoc jobs* from the menu on the right in the workflow page.

Deployments

After a workflow is uploaded to KNIME Hub different types of deployments can be created.

You can create:

- **Data Apps:** **Data Apps** provide a user interface to perform scalable and shareable data operations, such as visualization, data import, export, and preview.
- **Schedules:** A workflow can be scheduled to run at specific times and perform specific actions based on the result of each execution.
- **API services:** A workflow can be deployed as a REST endpoint and therefore called by external services.
- **Triggers:** A workflow can be deployed as trigger deployment meaning that the workflow will be executed every time a specific selected event happens (e.g. a file is added to a space or a new version of a workflow is created).

To create a new deployment of a workflow that you uploaded to one of your team's spaces you first need to have created at least one version of the workflow. After you **create a version** of the workflow click *Deploy*.

The screenshot shows the KNIME Hub interface for a workflow titled "Visual Analysis of Sales Data". The workflow is in version v.2, created on Apr 10, 2024 at 3:54 PM. The workflow diagram includes nodes for CSV Reader, Column Filter, and Row Filter. A "Deploy" button is visible in the top right corner. A modal window is open, displaying four deployment options: Data app, Schedule, Service, and Trigger. Each option has a "Create" button.

My First Workflow - Visualization of Sales Data

- Filter data to columns that contain relevant information
- Filter data to rows that you want to include in your analysis
- Visualize the data using a Stacked Area Chart and a Pie/Donut Chart

Data Access - Provide file path

CSV Reader
Read sales_2008-2011.csv

Data Preprocessing - Filter columns and rows

Column Filter
Select "country", "date" and "amount" columns

Row Filter
Exclude rows where country unknown

Some change to the workflow

Component

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

Schedule
Schedule your workflow to run automatically at selected times.
Create schedule

Service
Create a service to use the workflow as an API endpoint.
Create service

Trigger
Create a trigger to execute the workflow when a specified event occurs.
Create trigger

Visual Analysis of Sales Data

Workflow that is built according to the instructions on the "Build a workflow" page and in the KNIME Quickstart Guide

Shows possible steps in analyzing (sales) data

Figure 31. Create a new deployment for a workflow on KNIME Hub

In the menu that opens you can select which kind of deployment you want to create.

Data apps

On the workflow page, click the *Deploy* button and select *Create data app* if you want to create a data app to interact with the workflow via a user interface.

In the right panel that opens you will be able to choose a *Deployment name*, select the version you want to deploy, select the execution context you want to use.

You will also be able to select a *Category* and a *Description*. Categories added here will be used to group the data apps in the [Data Apps Portal](#). Also descriptions added here will be visible in the corresponding deployment's tile in the [Data Apps Portal](#).

Email notifications

If you check the option *Enable email notifications* you can enable sending an email upon success or failure of the execution of the deployment you will create.

Add the email address and select the condition, and click *Add more* to add more actions.

The screenshot shows the KNIME Hub interface. The main area displays a workflow titled "Visual Analysis of Sales Data" with a version of v.2. The workflow consists of several nodes: "Data Access" (CSV Reader), "Data Preprocessing" (Column Filter and Row Filter), and "Data Visualization" (Stacked Area Chart and Pie/Donut Chart). The "Create data app" dialog is open on the right side, allowing the user to configure the deployment. The dialog includes fields for "Deployment name" (Visual Analysis of Sales Data), "Version" (v.2 - Apr 10, 2024 3:54 PM), "Execution context" (Docs EC), "Category", and "Description". Under "Email notifications", the "Enable email notifications" checkbox is checked, and the "Condition" is set to "Success". The "Advanced settings" section is visible at the bottom, with a "Set" button to configure job lifecycle behavior and timeouts.

Figure 32. Create a data app deployment for a workflow on KNIME Hub

Under *Advanced settings* click **Set** to change the advanced settings of the deployment. Here you can configure the options regarding:

- *Job lifecycle*: such as deciding when to discard a job, the maximum time a job will stay in memory, the job life time, or the options for timeout

- *Additional settings*: such as report timeouts, CPU and RAM requirements and so on

Notice that you can get additional information on the different fields by hovering over them with the cursor.


The values in the fields related to time need the format:


- PT, which stands for Period of Time
- the amount of time, made of a number and a letter D, H, M and S for days, hours, minutes and seconds.

For example, PT1H means a period of time of 1 hour, PT168H means a period of time of 168 hours (which is equivalent of 7 days).

Run a Data App deployment

Once you have created a Data App deployment to see the list of *Deployments* you can go on the specific workflow page.


Click the  icon on the row corresponding to the deployment you want to execute, and click *Run*.



The screenshot shows the KNIME Business Hub interface. At the top, there are tabs for 'Extensions' and 'Nodes'. Below this, there are two sections: 'KNIME Base nodes' and 'KNIME JavaScript Views', both showing 'KNIME AG, Zurich, Switzerland Version 5.3.0'. On the right side, there is a sidebar with navigation options: 'Overview', 'External resources', 'Used extensions & nodes', 'Deployments', and 'Ad hoc jobs'. The main content area is titled 'Deployments' and contains a table with the following columns: 'Type of deployment', 'Name', 'Workflow', 'Version', 'Date of creation', 'Creator', 'Execution context', and 'Status'. The table has one row with the following data: 'Data app', 'Visual Analysis of...', '/Users/Document...', 'v.2', 'Apr 10, 2024 5:22 ...', 'verena.mayer', 'Docs EC', and 'Status'. A yellow box highlights the 'Run' button in the 'Status' column of this row. Below the table, there is a section for 'Ad hoc jobs' with a table showing one job with the state 'Execution finished'.

Figure 33. Run a data app deployment

Manage access to data apps

You can manage who may access data apps deployments. In the *Deployments* section on the workflow page, click the  icon. Select *Manage access* from the menu. This opens a side panel where you can type the name of the user or **externally managed group** you want to share the deployment with.

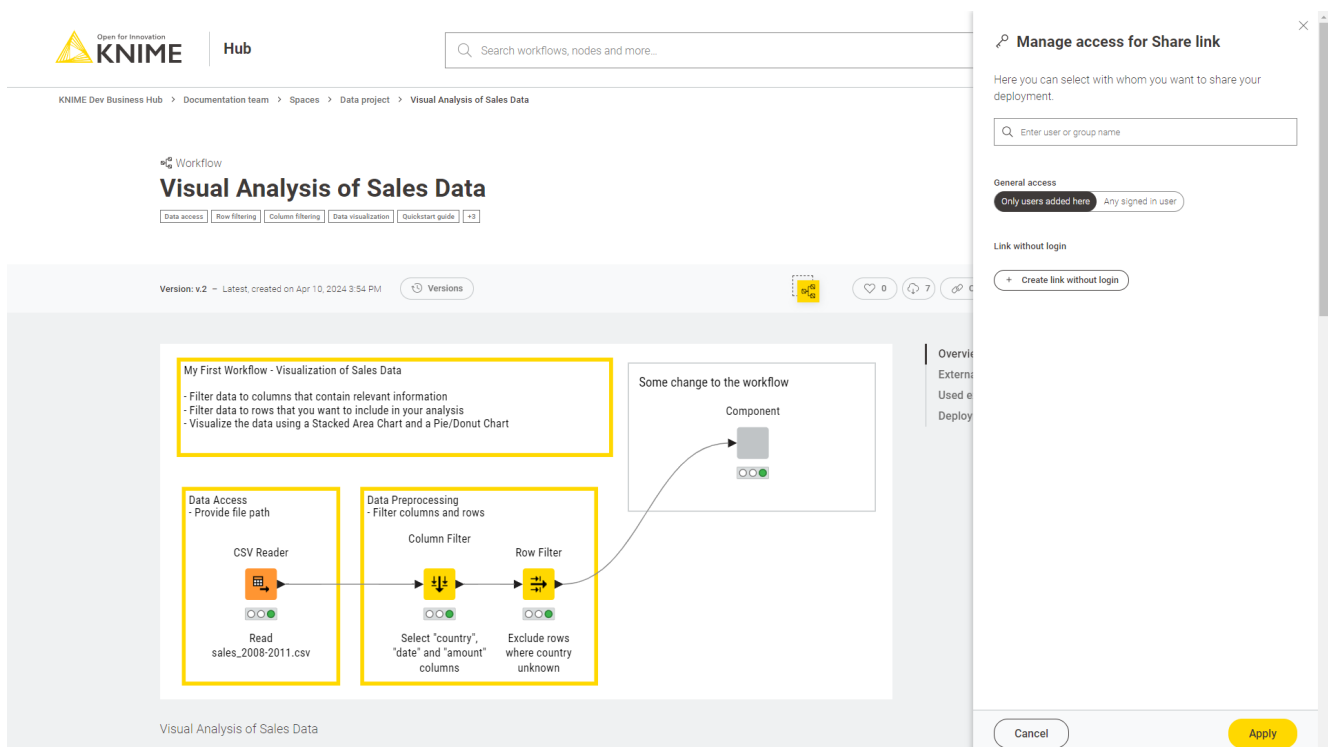
It is also possible to share the deployment with users that are not members of your team. When sharing a data app deployment, this will be available in the user profile in the *Data Apps Portal* section.

When managing the access of data apps, you can also select the option to share the deployment with all signed in users. This will allow every user who has access to the KNIME Business Hub instance to execute the deployment.



Deployments shared this way will also be available to all future users who don't have a Business Hub account yet.

To do so, select the *Any signed in user* option in the *Manage access* panel.



The screenshot shows the KNIME Business Hub interface. The main area displays a workflow titled "Visual Analysis of Sales Data" with components like "CSV Reader", "Column Filter", and "Row Filter". A side panel titled "Manage access for Share link" is open, showing options for "General access" (with "Only users added here" selected) and "Link without login" (with "Create link without login" button).

Figure 34. Manage access to a deployment of the type data app

The added user will still be available in case you change back to the default access type *Only users added here*.

Data app deployment link without login

From the *Manage* access panel you can also create links to be shared with others that do not require the user to log into the Hub instance.

Click *Create link without login* to create a link and then click *Copy link* to save the created URL in your clipboard. You can then share this link with anyone you want to give access to this specific deployment and they will be able to run the data application in the browser without having to log in.

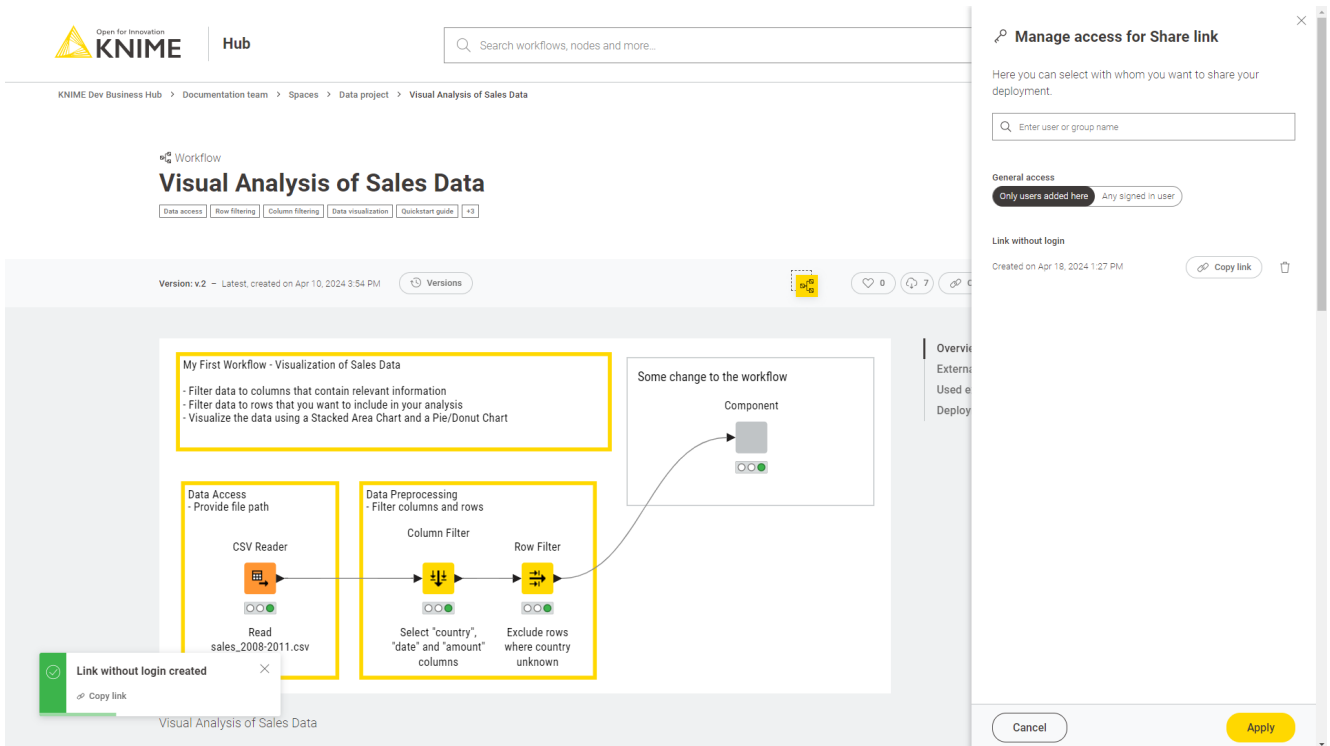


Figure 35. Create a link without login for a data app deployment

You can go back to this panel at any time and delete the link. This will prevent users who use the link from accessing the data app.

Pass data into Data Apps via URL parameters

It is possible to pass data into a KNIME Data App by using URL parameters.

The process involves configuring a Container Input (JSON) node within your KNIME workflow to receive and process the parameter passed as a plain string in the URL.


The URL needs to have the following format:

```
https://apps.<base-url>/d/<deployment-name>~<deploymentID>/run?pm:<parameter-name>=<value>
```



Note that since its name might change, the ID is enough to identify the deployment. Hence, the `<deployment-name>` is optional, can be an arbitrary string, or can be left out altogether. You will be redirected to a URL with the format given above, where the `<deployment-name>` will reflect the current name of the deployment.

The detailed steps to use this concept:

1. Create a data app deployment, go to the deployments overview and click the  icon corresponding to the data app deployment and **run** it




Figure 36. Run a data app deployment

2. Wait until the execution starts and **then** copy the URL from your browser bar
3. From the URL you copied remove the last part (i.e. everything after the last /) and add `/run?pm:<parameter-name>=<value>` to the end of the URL.



In case you want to add these parameters to a **link without login** you can use that link and add the `/run?pm:<parameter-name>=<value>` at the end of it. Anyone with this link will then be able to run the data app deployment.

4. The `<parameter-name>` is the one defined in the *Parameter Name* setting in the configuration dialog of the Container Input (JSON) node. In [Figure 37](#) the parameter name is `example-input`.

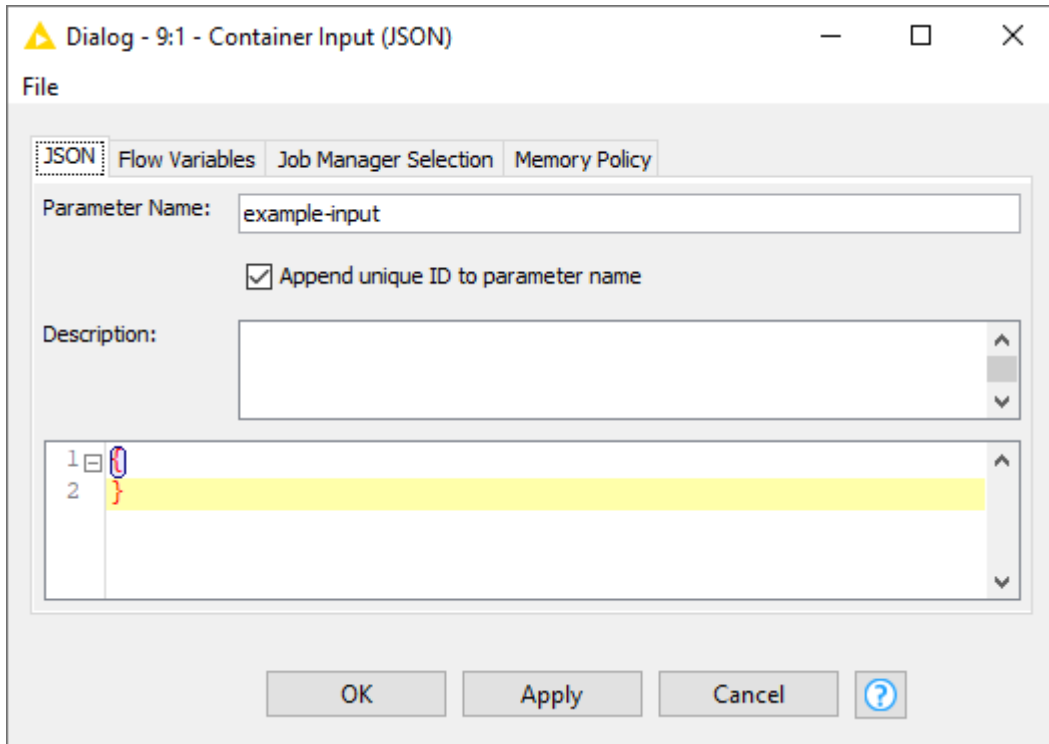


Figure 37. Container Input (JSON) node configuration dialog

The **URL encoded** <value> will be parsed into a **JSON-Value**. If you want to pass a simple string, surround it in quotes, e.g. "test string" (URL encoded: %22example%20string%22). Plain integers or null/true/false/{`"JSON-Key": "JSON-Value"`} (URL encoded: %7B%22JSON-Key%22%3A%20%22JSON-Value%22%7D) will be parsed into the corresponding JSON-Value.

Service

On the workflow page, click the *Deploy* button and select *Create service* if you want to create a service to use the workflow as an API endpoint.

In the right panel that opens you will be able to choose a *Deployment name*, select the version you want to deploy and select the execution context you want to use.

Email notifications

If you check the option *Enable email notifications* you can enable sending an email upon success or failure of the execution of the deployment you will create.

Add the email address and select the condition, and click *Add more* to add more actions.

The screenshot shows the KNIME Hub interface. At the top, there's a search bar and navigation links. The main content area displays a workflow titled "Visual Analysis of Sales Data" with a version of v.2. The workflow diagram includes a "Data Access" section with a "CSV Reader" component reading "sales_2008-2011.csv", followed by a "Data Preprocessing" section with a "Column Filter" and a "Row Filter". A callout box points to a "Component" in the workflow. On the right, a "Create service" modal is open, showing fields for "Deployment name" (Visual Analysis of Sales Data), "Version" (v.2 - Apr 10, 2024 3:54 PM), "Execution context" (Docs EC), and "Email notifications" (checked). The "Advanced settings" section is visible at the bottom of the modal, with a "Set" button.

Figure 38. Create a service deployment for a workflow on KNIME Hub

Under *Advanced settings* click **Set** to change the advanced settings of the deployment. Here you can configure the options regarding:

- *Job lifecycle*: such as deciding when to discard a job, the maximum time a job will stay in memory, the job life time, or the options for timeout
- *Additional settings* such as report timeouts, CPU and RAM requirements and so on

Notice that you can get additional information on the different fields by hovering over it with the cursor.


The values in the fields related to time need the format:

- PT, which stands for Period of Time
- the amount of time, made of a number and a letter D, H, M and S for days, hours, minutes and seconds.


For example, PT1H means a period of time of 1 hour, PT168H means a period of time of 168 hours (which is equivalent of 7 days).

Finally, return from *Advanced settings* and click *Create* at the bottom of the side panel to create your service deployment.

Open a service deployment

When a service deployment is created, a REST API endpoint is also created. To open the definition of the endpoint, navigate to the workflow page. In the *Deployments* section, click the  icon. Selecting *Open* will open the definition of the endpoint.

Manage access to service

You can manage who may access service deployments. In the *Deployments* section on the workflow page, click the  icon. Select *Manage access* from the menu. This opens a side panel where you can type the name of the user or **externally managed group** you want to share the deployment with.

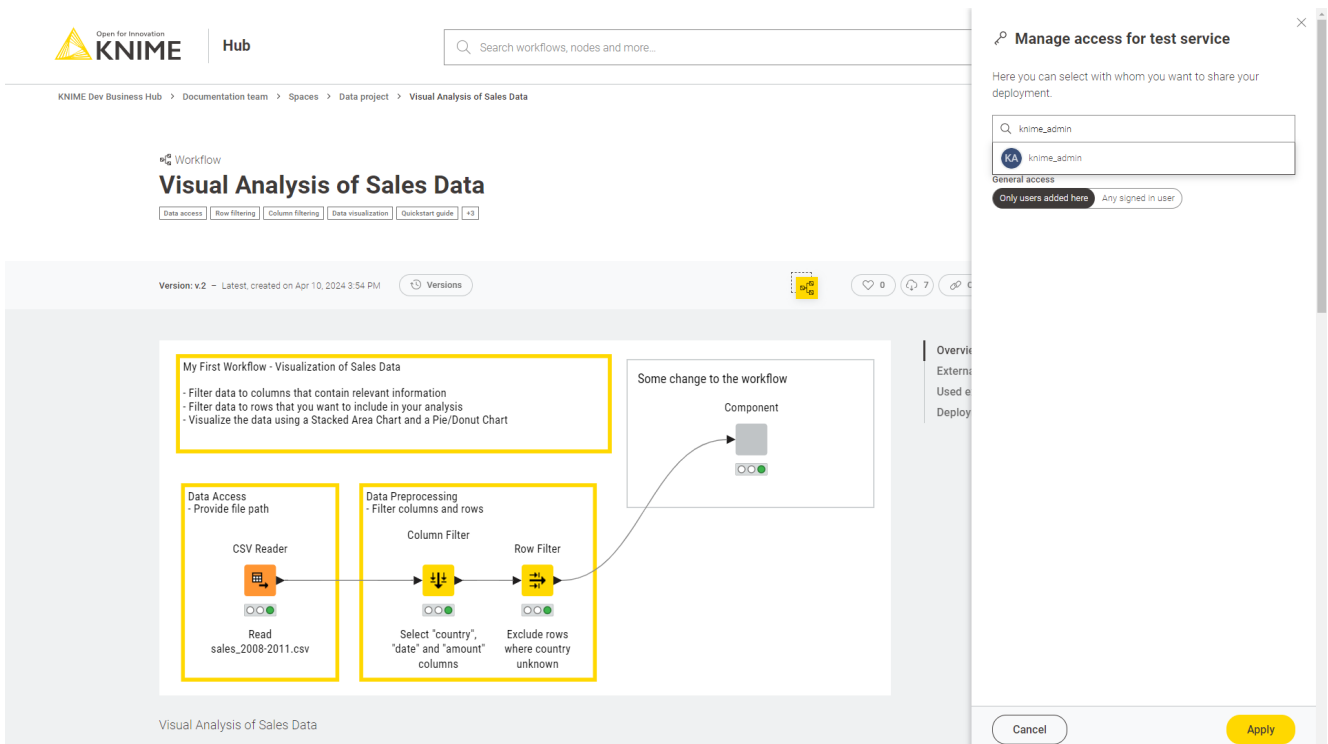
It is also possible to share the deployment with users that are not members of your team.

When managing the access of services, you can also select the option to share the deployment with all signed in users. This will allow every user who has access to the KNIME Business Hub instance to execute the deployment.



Deployments shared this way will also be available to all future users who don't have a Business Hub account yet.

To do so, select the *Any signed in user* option in the *Manage access* panel.



The screenshot displays the KNIME Business Hub interface. The main area shows a workflow titled "Visual Analysis of Sales Data" with a version of v.2. The workflow consists of several components: "Data Access" (CSV Reader), "Data Preprocessing" (Column Filter and Row Filter), and "Visualization" (Stacked Area Chart and Pie/Donut Chart). The "Manage access for test service" panel is open on the right, showing a search for "knime_admin" and a list of users. The "Any signed in user" option is selected under "General access".

Figure 39. Manage access to a deployment of the type service

The added user will still be available in case you change back to the default access type *Only users added here*.

Application passwords

Application passwords can be used to provide authentication when using the KNIME Hub REST API, for example when executing deployed REST services. To create a new application password you can go to your profile page and select *Settings* → *Application passwords* from the menu on the left.

Click *+ Create application password* and a side panel will show.

Here you can give a name in order to keep track of the application password purpose. Then click *Apply*. The ID and the password will be shown only once. You can copy them and use it as username and password in the base authentication when you want to execute a deployed REST service.



Please, be aware that the copy button will not work on HTTP, so if **TLS is not enabled**.

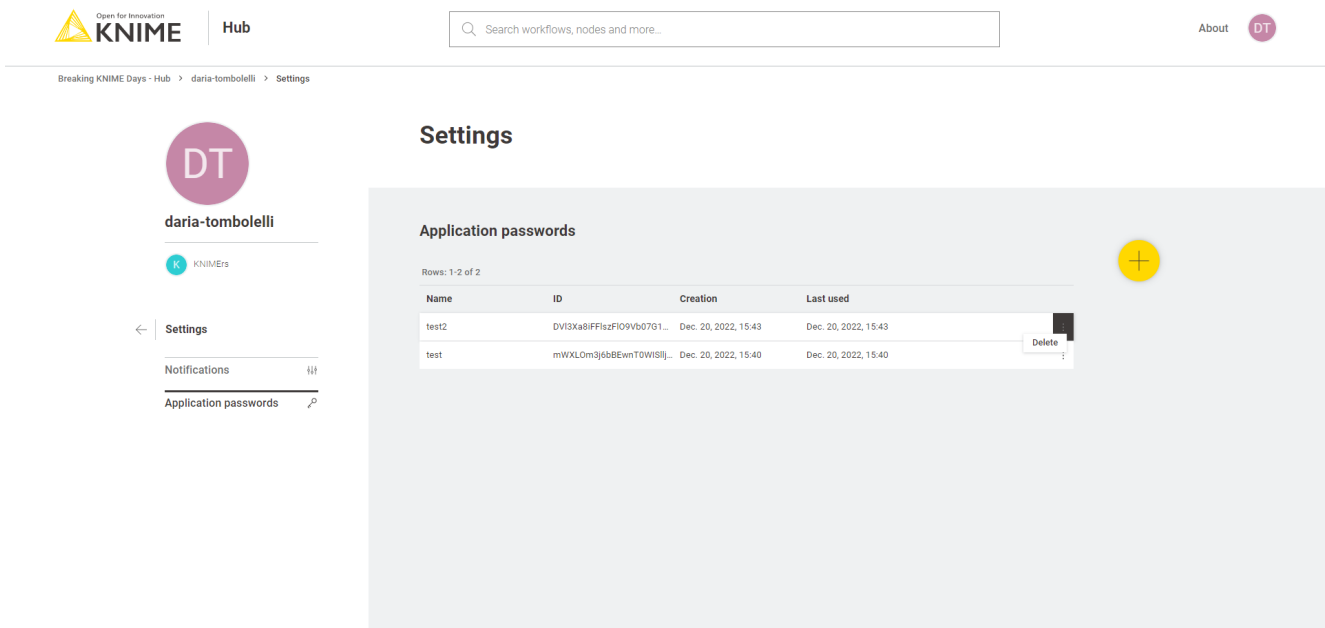



Figure 40. Application passwords

From this page you can also click the  icon and click *Delete* from the menu that opens, in order to delete a specific application password.

Schedule

On the workflow page, click the *Deploy* button and select *Create schedule* if you want to create a scheduled execution that will run your workflow automatically at selected times.

In the right panel that opens you will be able to choose a *Deployment name*, select the version you want to deploy and select the execution context you want to use.

The screenshot displays the KNIME Business Hub interface. The main area shows a workflow titled "Visual Analysis of Sales Data" with a version of v.2. The workflow diagram includes a "CSV Reader" node connected to "Data Access", followed by "Data Preprocessing" nodes: "Column Filter" and "Row Filter". A "Component" node is also visible. The right-hand panel is the "Create schedule" dialog, which includes fields for "Deployment name" (Visual Analysis of Sales Data), "Version" (v.2 - Apr 10, 2024 3:54 PM), and "Execution context" (Docs EC). The "Schedule options" section allows defining the initial execution date and time (2024-04-11 12:00:00) and includes checkboxes for "Repeat execution" and "Email notifications". The "Advanced settings" section is partially visible at the bottom.

Figure 41. Create a schedule deployment for a workflow on KNIME Hub

Schedule options

Here you can define the initial execution date and time.

Check the option *Repeat execution* if you want the execution to be repeated.

Select *Every* to set up in which interval of time the workflow will be executed, e.g. every 2 days. Select *Start times* to manually add the different times when the workflow will be executed.

If you click *Set* under *Set schedule details*, you can set up recurring executions, retries, and advanced schedule details.



If you do not check the option *Repeat execution*, you will only find the set up options for retries and the advanced schedule details.

The screenshot displays the KNIME Business Hub interface for a workflow titled "Visual Analysis of Sales Data". The workflow is shown in a visual editor with three main stages: "Data Access" (using a CSV Reader to read "sales_2008-2011.csv"), "Data Preprocessing" (using a Column Filter to select "country", "date", and "amount" columns, and a Row Filter to exclude rows where the country is unknown), and "Visualization" (using a Stacked Area Chart and a Pie/Donut Chart). The workflow is versioned (v.2) and was created on April 10, 2024, at 3:54 PM.

On the right side, the "Create schedule" configuration panel is visible. It includes a "Recurring on" section with a calendar grid for selecting days of the week (M, T, W, T, F, S, S) and a "Days" section with a calendar grid for selecting specific days of the month (1-31, Last day of the month). Below these are sections for "Months" (Jan-Dec), "Timeframes" (Add timeframe), and "Execution retries" (set to 1). At the bottom, there are checkboxes for "Reset before execution", "Skip execution if previous job is still running", and "Disable schedule".

Figure 42. Set the details for a schedule

In the section execution retries and advanced schedule details you can set up the number of execution retries, and check the following options for the advanced schedule details:

- *Reset before execution*: the workflow will be reset before each scheduled execution retries occur.
- *Skip execution if previous job is still running*: the scheduled execution will not take place if the previous scheduled execution is still running.
- *Disable schedule*: Check this option to disable the schedule. The scheduled execution will start run accordingly to the set ups when it is re-enabled again.

Email notifications

If you navigate back from *Set schedule details* to *Create schedule* and check the option *Enable email notifications*, you can enable sending an email upon success or failure of the execution of the deployment you will create.

Add the email address and select the condition, and click *Add more* to add more actions.

Advanced settings

Finally, in the section *Advanced settings* you can configure additional set ups:

- *Job lifecycle*: such as deciding in which case to discard a job, the maximum time a job will stay in memory, the job life time, or the options for timeout
- *Additional settings*: such as report timeouts, CPU and RAM requirements, check the option to update the linked components when executing the scheduled job and so on

The values in the fields related to time need the format:

- PT, which stands for Period of Time
- the amount of time, made of a number and a letter D, H, M and S for days, hours, minutes and seconds.

For example, PT1H means a period of time of 1 hour, PT168H means a period of time of 168 hours (which is equivalent of 7 days).

Change status of a schedule deployment

You can activate or deactivate a schedule after it has been created. To do so, click the toggle in the *Status* column from the list of deployments available on the respective workflow page or on the team page, as shown in [Figure 45](#), and [Figure 46](#).

Trigger

On the workflow page, click the *Deploy* button and select *Create trigger* if you want to create a trigger to execute the workflow when a specified event occurs.

In the right panel that opens you will be able to choose a *Deployment name*, the version you want to deploy, and select the execution context you want to use.

Trigger options

In the *Trigger options* section you define the events that will trigger the execution of the workflow you are deploying.

The trigger listens to a specific event and when all the conditions set here are met, the workflow is executed.

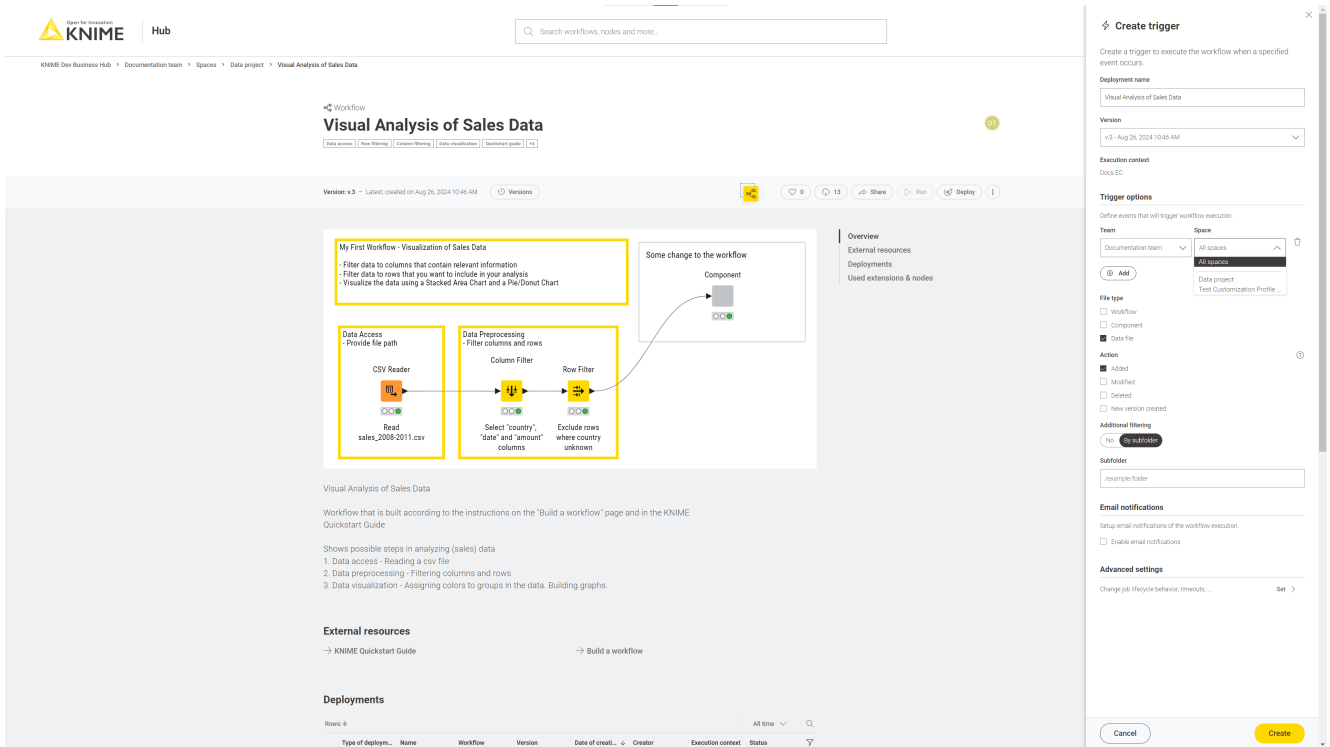


Figure 43. Create a trigger deployment for a workflow on KNIME Hub

- In the right panel that opens you will be able to choose a *Deployment name*, the version you want to deploy, and select the execution context you want to use.
- Under *Trigger options* select one or multiple spaces on the KNIME Hub in which the trigger is listening for events. The available spaces include any that are owned by teams of which you are a member.
 - **Team Members:** If you are a member of a team, you can select any spaces owned by that team for the trigger to listen in.
 - **Team Admins:** If you are an admin for a team, you have the additional option to select all spaces within your team. This selection not only includes all current spaces but also automatically applies to any future spaces created under that team.
 - **Global Admins:** As a global admin, you have the capability to select all spaces across all teams. This global trigger will monitor events across all spaces, including those in newly created teams and spaces. While this functionality is powerful for governance and similar use cases, note that it may result in a large number of job executions.
- Then select the item type between:
 - *File*
 - *Folder*
- If you select *File* as item type you can further filter by file type, choosing between

Workflow, Component or Data file. You can select one or more of these file types.

- Then select the action that is going to occur in order for the workflow to be executed. You can select multiple actions, but only one needs to happen to trigger the execution.
 - If you select *File* as item type you can select that the file is added, modified, deleted, and/or a new **version** is created.
 - If you select *Folder* as item type you can select that the folder is added and/or deleted.
- Finally, you can set up an additional filtering selecting *By subfolder*. Add the relative path (e.g. `/example/folder`) to the space the trigger deployment is listening for the event. The event will then need to happen in the space you selected, in the subfolder `/example/folder`. For example, if you selected that the trigger will execute when a *File* of the type *Data file* is added to the space *Space1* in the subfolder `/example/folder`, the workflow execution will be triggered when the data file is added to `Space1/example/folder`.



When designing workflows that you want to deploy as triggers, be cautious with file handling in Hub spaces. If the workflow changes files in the space it is monitoring, it could unintentionally trigger itself, causing an infinite loop.



It is important to note that, under special circumstances, job creation might fail (e.g. if the execution context can not accept new jobs), or the job itself fails. In these cases, the job creation or execution will not be retried and the event would not be recorded. This makes trigger deployments unsuitable for use cases that require 100% reliability.

Email notifications

If you check the option *Enable email notifications* you can enable sending an email upon success or failure of the execution of the deployment you will create.

Add the email address and select the condition, and click *Add more* to add more actions.

Advanced settings

Finally, in the section *Advanced settings* you can configure additional set ups:

- *Job lifecycle*: such as deciding in which case to discard a job, the maximum time a job will stay in memory, the job life time, or the options for timeout

- *Additional settings*: such as report timeouts, CPU and RAM requirements, check the option to update the linked components when executing the scheduled job and so on

The values in the fields related to time need the format:

- PT, which stands for *Period of Time*
- the amount of time, made of a number and a letter D, H, M and S for days, hours, minutes and seconds.

For example, PT1H means a period of time of 1 hour, PT168H means a period of time of 168 hours (which is equivalent of 7 days).

Node to get information about a trigger deployment

With KNIME Analytics Platform version 4.7.2 the Container Input (Repository Event) node is added to the [KNIME Hub Connectivity](#) extension.

If you want to receive information about the event that triggered the execution of a workflow on KNIME Hub, you can add this node to the workflow. Upload the workflow to your KNIME Hub instance and create a trigger deployment for the workflow. When the conditions set for the trigger are met and the workflow is executed the node will receive information about the event that triggered the execution and output them as flow variables. The different flow variables are listed in the node description.

The event information comprises:

- What kind of item, e.g., a workflow, component, or file;
- Where did the event happened, e.g., in a particular directory within a space;
- What type of action was performed on the item, e.g., modified, added, or deleted.

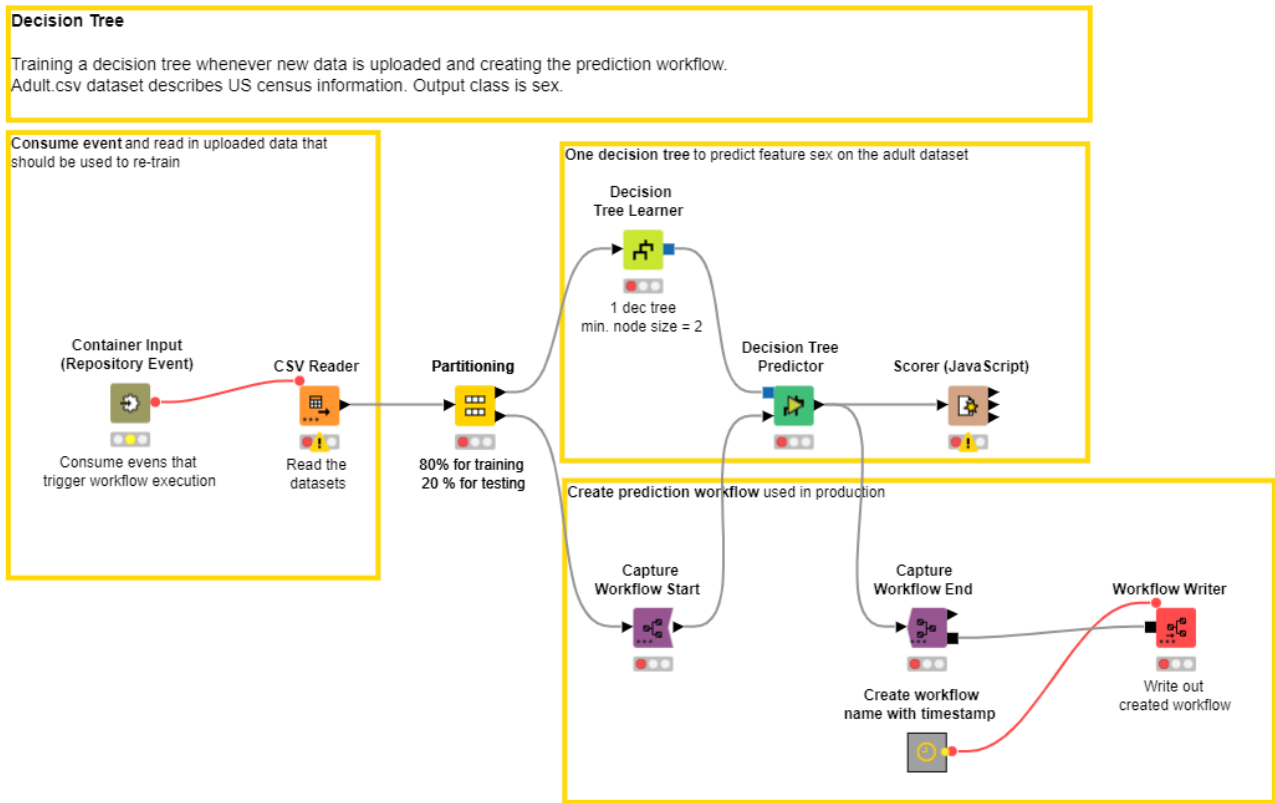


Figure 44. Example workflow using the Container Input (Repository Event) node

You can find the workflow shown in [Figure 44](#) on the [KNIME Community Hub](#).

Managing deployments

Select *Deployments* from the menu on the right in the workflow page to see a list of the deployments that have been created for a specific workflow.

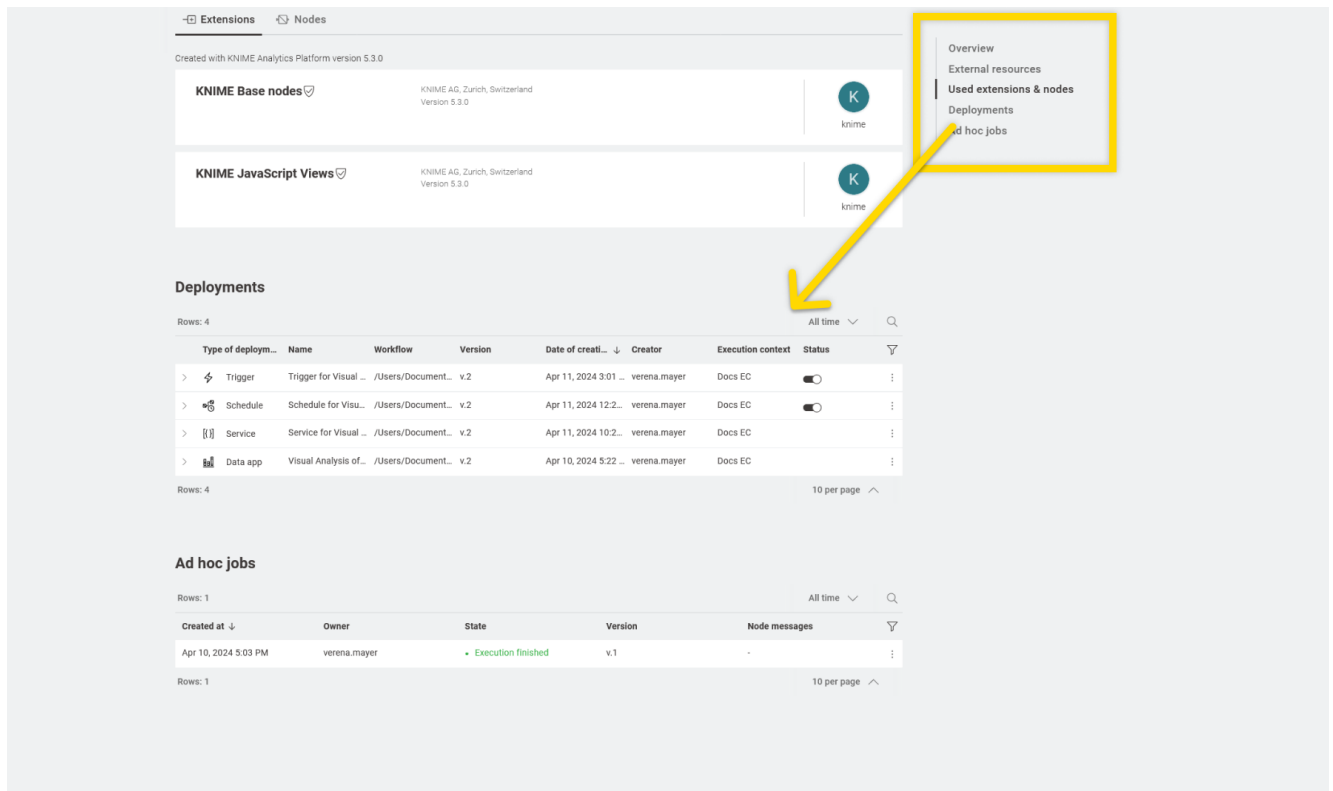


Figure 45. A list of all the deployments created for a specific workflow

To see all the deployments that were created under a specific team go to that team page and select *Deployments* from the menu on the left.

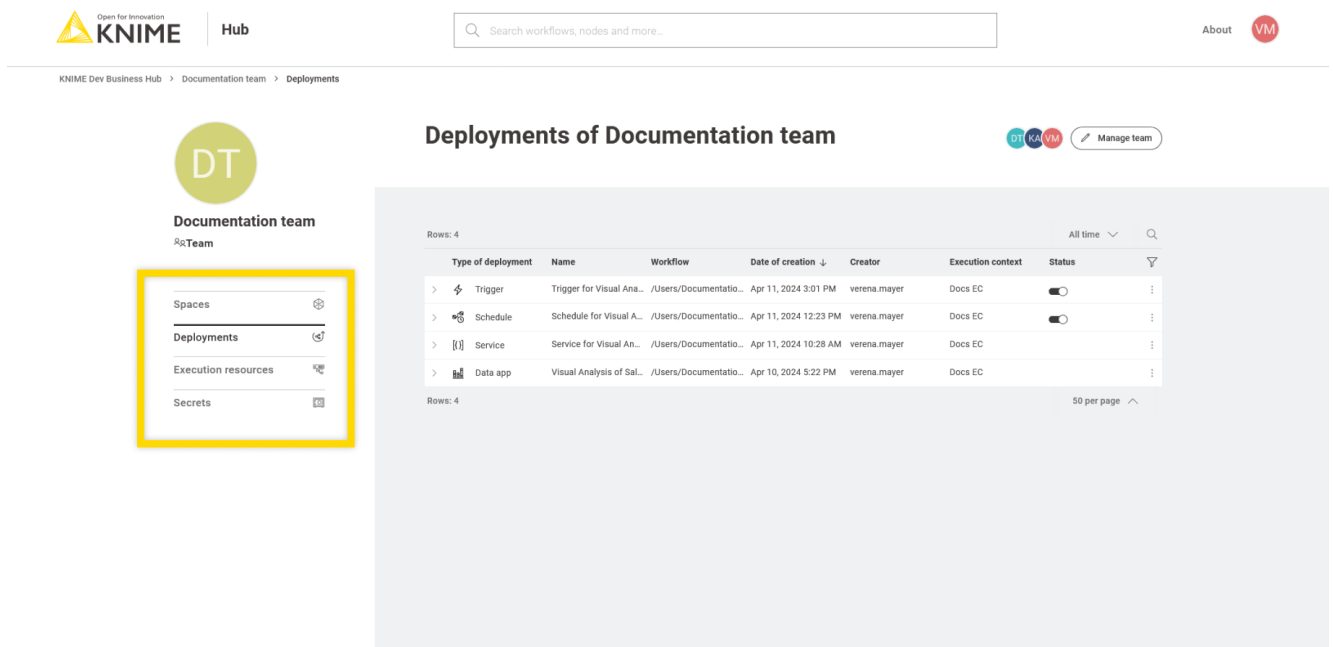



Figure 46. A list of all the deployments created by a team

By clicking the  icon at the end of the corresponding deployment row you can perform the

following actions. Notice that depending on the type of deployment some actions might not be available.


- *Run*: data apps, schedules and triggers
- *Open*: services
- *Edit*: all types of deployments
- *Manage access*: data apps and services
- *Delete*: all types of deployments


Edit a deployment

You can edit a deployment from the deployments list by clicking the  icon in the corresponding row and selecting *Edit* from the menu.

The side panel for the deployment will open and you will be able to make changes to the deployment, then click *Save changes*.

Delete a deployment

You can delete a deployment from the deployments list by clicking the  icon in the corresponding row and selecting *Delete* from the menu.



The screenshot shows the KNIME Business Hub interface. At the top, there are three sections: 'KNIME Base nodes', 'KNIME JavaScript Views', and 'KNIME Quick Forms', each with a version 4.6.0. Below these is the 'Deployments' section, which contains a table with 6 rows. The table has columns for 'Type of deployment', 'Name', 'Workflow', 'Version', 'Date of creation', 'Execution context', and 'Status'. The third row is selected, and a context menu is open over it, showing options: 'Run', 'Edit', 'Manage access', and 'Delete'. The 'Delete' option is highlighted. Below the table, there is a section for 'Ad hoc executions' which states 'No ad hoc executions available'.

Type of deployment	Name	Workflow	Version	Date of creation	Execution context	Status
Trigger	Data App: Visual Ana...	..lysis_of_Sales_Data/	trigger_test_version_2	June 22, 2023, 13:47	Initial Execution Cont...	
Schedule	Data App: Visual Ana...	..lysis_of_Sales_Data/	trigger_test_version_2	June 22, 2023, 13:46	Initial Execution Cont...	
Data app	Data App: Visual Ana...	..lysis_of_Sales_Data/	trigger_test_version_2	May 26, 2023, 12:30	Nightly-4.7.3	
Data app	test	..lysis_of_Sales_Data/	trigger_test_version_2	May 16, 2023, 15:04	Nightly-4.7.3	
Service	Service - 00_DataAp...	..lysis_of_Sales_Data/	test 7	Jan. 11, 2023, 15:38	Initial Execution Cont...	
Data app	test share 43243	..lysis_of_Sales_Data/	test v5 - data app	Dec. 16, 2022, 13:29	Initial Execution Cont...	

Figure 47. Delete a created deployment

Jobs

Every time a workflow is executed ad hoc or a deployment is executed a job is created on KNIME Business Hub.

To see the list of all the jobs that are saved in memory for the ad hoc execution of a specific workflow go to the workflow page and on the right side menu click *Ad hoc jobs*.

To see the list of all the jobs that are saved in memory for each of the deployments created for a specific workflow, go to the workflow page and on the right side menu click *Deployments*. You can expand each deployment with the > icon on the left of the deployment.

Also you can go to your team page and find a list of all deployments created within your team. Also here you can click the > icon corresponding to a specific deployment to see all its jobs.

On each job you can click the : icon on the right of the corresponding job line in the list and perform the following operations (their availability depends on the type of **deployment** or the **job state**):

- Open: For example you can open the job from a data app and look at the results in a new tab
- Save as workflow: You can save the job as a workflow in a space
- Inspect: A job viewer opens in a new tab. Here, you can investigate the status of jobs.



This functionality is available only for jobs running with executor versions 5.2.3 or later.

- Delete: You can delete the job
- Download logs: You can download the log files of a job - this feature allows for debugging in case the execution of a workflow did not work as expected. To be able to download job logs you need an executor based on KNIME Analytics Platform version > 5.1.




Your team can by default have a maximum of 10000 jobs. If you need more please contact your KNIME Hub admin who can **increase the limit** if necessary.

Inspect an executed workflow



This functionality is available only for jobs running with executor versions 5.2.3 or later.

You can inspect the status of an executed workflow, for example if the workflow execution did not succeed, or the execution is taking a long time. Click the  icon for the desired job and select *Inspect*.




Figure 48. Select *Inspect* from the menu to check the status of the execution

Here you can hover over info and error icons at the node's traffic light, to visualize the message, check where the execution stopped or the status of specific nodes. You can also inspect the data at any node's output port.




Figure 49. Inspect node messages and data at the output port of a node

You can also navigate inside components and metanodes and inspect the nodes.




Figure 50. Navigate inside components and metanodes to inspect the encapsulated nodes

Job states

Jobs exist in a variety of different states.

The possible states are:

- **Loading** - Jobs being loaded by an executor or waiting to be accepted by an Executor.
- **Executing** - Job is currently executing.
- **Execution finished** - Job has been executed successfully
- **Execution failed** - Job has been executed, but failed
- **Execution cancelled** - Job has been cancelled manually during execution
- **Interaction required** - Job is currently executing and is awaiting user input
- **Not executable** - Job contains individual, unconnected nodes
- **Discarded** - Job has been executed and discarded
- **Undefined** - This is the first state of a job, and may be seen in the case where the Executor cannot communicate with the Hub due to network issues, or the executor not having enough free CPU/RAM resources.
- **Vanished** - Job was in memory on an executor that has crashed, and is therefore lost.

Data Apps Portal

You can access the Data Apps Portal by going to the address:

```
https://apps.<base-url>
```

where `<base-url>` is the usual address of the KNIME Business Hub instance.

This page is available to every registered user. Consumers, for example, can access to this page to see all the data apps that have been shared with them, execute them at any time, interact with the workflow via a user interface, without the need to build a workflow or even know what happens under the hood.

To execute a data app click the corresponding tile. A new page will open in the browser showing the user interface of the data app.




Figure 51. The Data Apps Portal

To share a data app deployment with someone follow the instructions in the [Manage access to a deployment](#) section.

Secrets

Secrets provide a way to centrally store and manage logins to other systems. For example, a secret could be credentials to log into an external database, file system or service. Secrets are owned and managed by a user or team. User secrets are intended for managing personal logins e.g. john.smith. Team secrets on the other hand are intended for shared logins sometimes referred to as technical or service users e.g. hr_read_only, that are shared with multiple users.

For more details about how to manage secrets in the KNIME Business Hub and how to use them in your workflows go to the [KNIME Secrets User Guide](#).

Known issues and compatibility

Please be aware that the compatibility of the following nodes and functionalities with KNIME Business Hub is still under development and it will full compatibility will be available soon.

- Molecule Sketcher Widget node does not support custom widgets such as MarvinJS
- Setting job lifecycle defaults for execution contexts is only possible right now via REST API

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