

KNIME Analytics Platform User Guide

KNIME AG, Zurich, Switzerland
Version 5.3 (last updated on 2024-10-07)



Table of Contents

Introduction	1
Workspaces	2
User interface	3
Entry page	5
Workflow editor & nodes	6
Connect to KNIME Hub	15
Switch back to KNIME classic user interface	15
Space explorer	17
Building workflows	20
Node repository	21
Node description	23
Workflow description	24
KNIME AI Assistant	24
Workflow monitor	30
Node monitor	31
Help	32
Customizing the Analytics Platform	32
Reset and logging	32
Configuring KNIME Analytics Platform	33
Preferences	33
Setting up knime.ini	40
KNIME runtime options	44
KNIME tables	50
Data table	50
Column types	51
Sorting	53
Column rendering	53
Table storage	54
Shortcuts	56
General actions	56
Execution	57
Selected node actions	58
Workflow annotations	59
Workflow editor modes	60

Workflow editor actions	61
Component and metanode building	62
Zooming, panning and navigating inside the canvas	63
Panel navigation	63

Introduction

This guide covers the basics of the KNIME Analytics Platform usage, guiding you in the first steps with the platform but also providing more advanced information about the most important concepts together with indication on how to configure the platform.

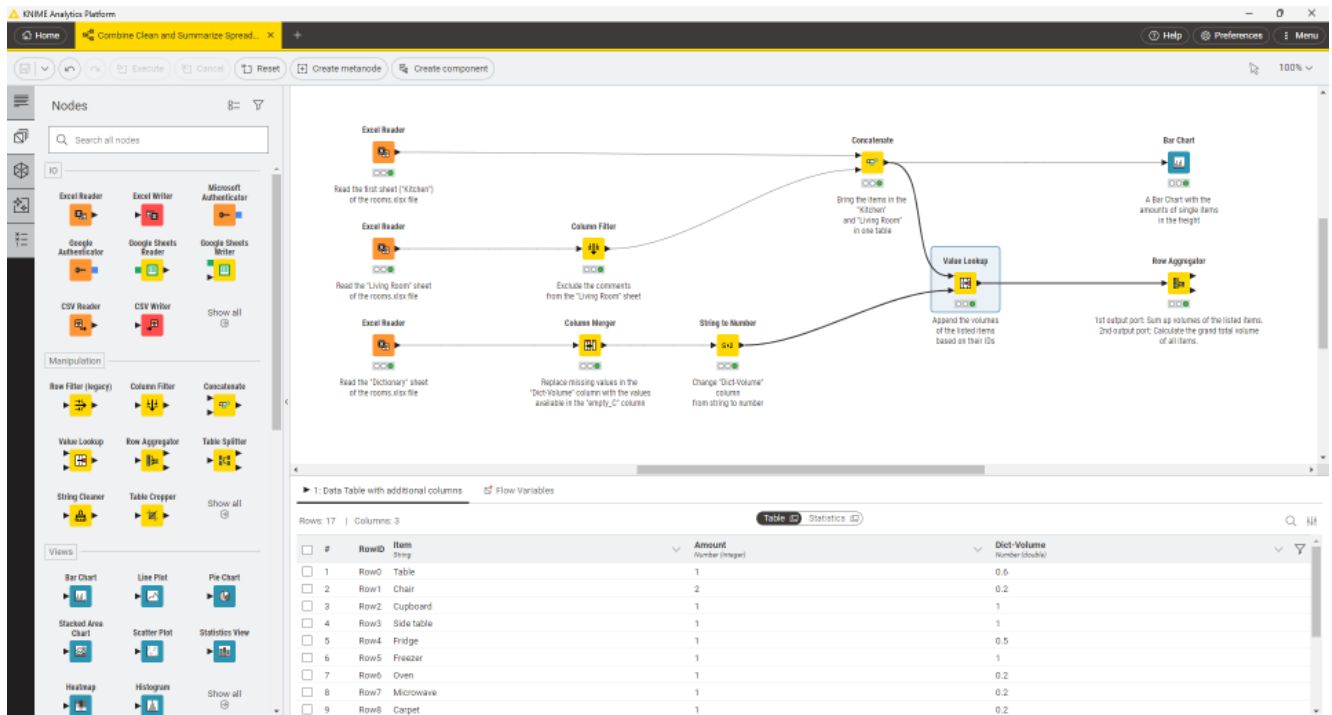


Figure 1. KNIME Analytics Platform

Workspaces

When you start KNIME Analytics Platform, the KNIME Analytics Platform launcher window appears and you are asked to define the KNIME workspace, as shown in [Figure 2](#).



The KNIME workspace is a folder on the local computer to store KNIME workflows, node settings, and data produced by the workflow.

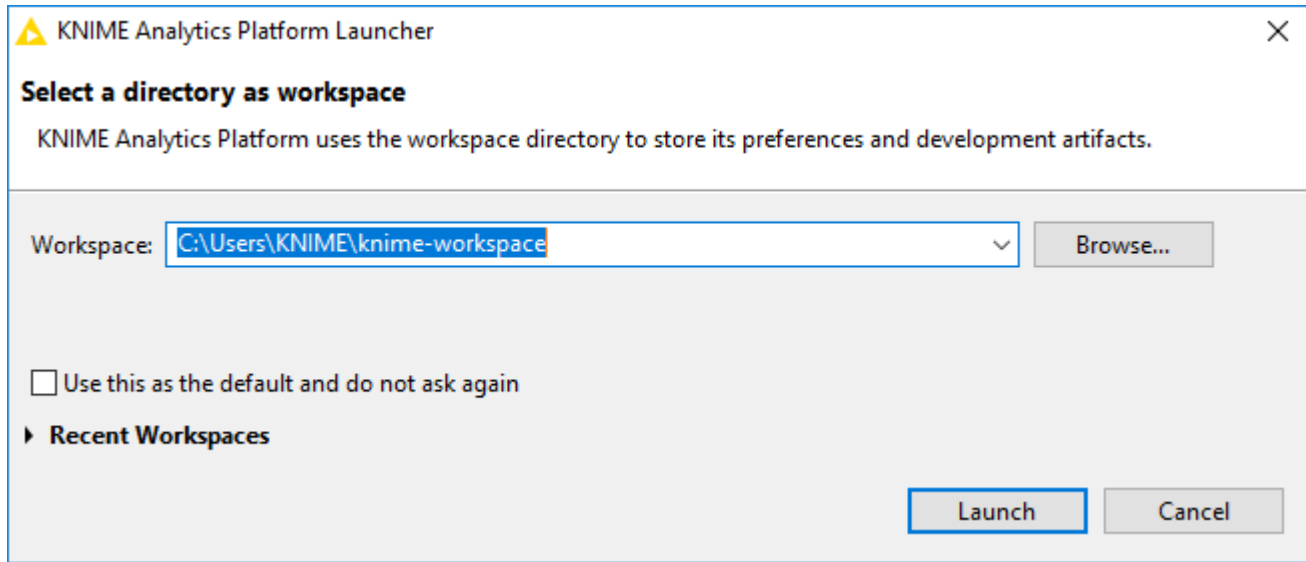


Figure 2. KNIME Analytics Platform launcher

The workflows, components and data stored in the workspace are available through the space explorer in the side panel navigation.

You can switch the workspace in a later moment under *Menu*, in the top right corner of the user interface, and select *Switch workspace*.

User interface

After selecting a workspace for the current project, click *Launch*. The KNIME Analytics Platform user interface - the KNIME Workbench - opens.

The active workflow of the KNIME Analytics Platform will be displayed after switching from an opened workflow. If you have open multiple workflows before you switch the perspective, only the active workflow and all loaded workflow tabs of the current KNIME Analytics Platform will be displayed in the KNIME Modern UI. For each workflow you will see a workflow tab after switching. After clicking the first tab (with the KNIME logo) you end up at the entry page.

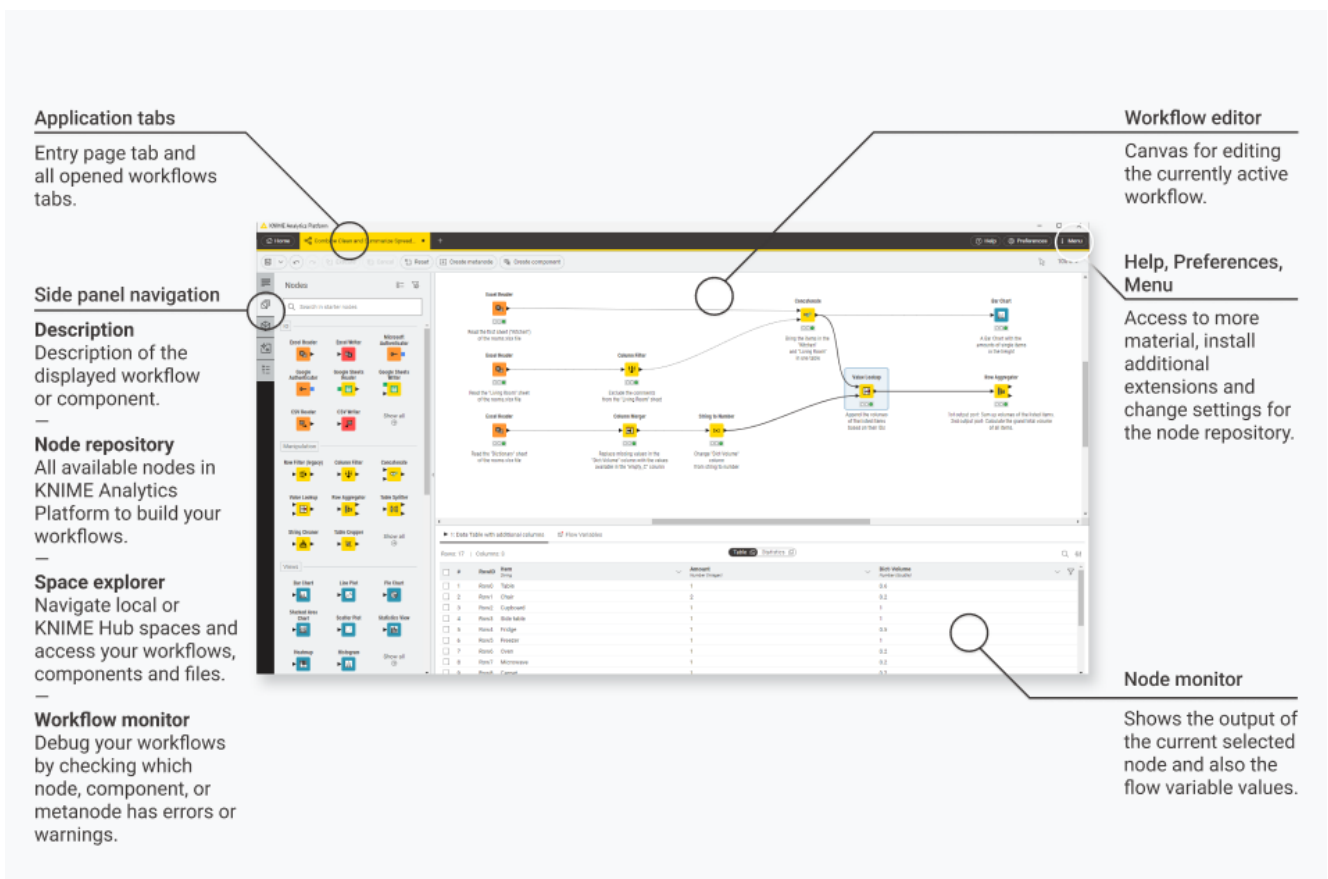


Figure 3. General user interface layout – application tabs, side panel, workflow editor and node monitor



Figure 4. User interface elements – workflow toolbar, node action bar, rename components and metanodes

In the next few sections we explain the functionality of these components of the user interface:

- [Entry page](#)
- [Workflow editor & nodes](#)
- [Connect to KNIME Hub](#)
- [Space explorer](#)
- [Node repository](#)
- [Node description](#)
- [Workflow description](#)
- [Node monitor](#)
- [Help menu](#)



You can scale the interface by going to *Menu > Interface scale*.

Entry page

The entry page is displayed by clicking the *Home* tab.




Figure 5. Entry page to create or open workflows

Here you will find:

- Recent, Local space, KNIME Community Hub. By default only the local workspace, and the link to connect to your personal KNIME Community Hub space are visible. To add a new mount point follow the instructions in the [Connect to KNIME Hub](#) section. Select:
 - *Recent* to see your recently opened workflows and components
 - *Local space* to navigate the existing workflows in your local system
 - *KNIME Community Hub* (or one of the available mount points). Click *Sign in*, provide your credentials and start navigating the available spaces.
- Three example workflows to help you get started – You can dismiss the examples by clicking the button on the top right. To restore the examples click *Help > Restore examples on home tab*.
- Create a new workflow by clicking the *+* button



Since KNIME Analytics Platform version 5.2 you can also add a **KNIME Server mount point**.

Workflow editor & nodes

The workflow editor is where workflows are assembled. Workflows are made up of individual tasks, represented by nodes.

One way to create a new workflow is to go to the space explorer, click the three dots and select *Create workflow* from the menu. Give the workflow a name and click *Create*.



Once you have a workflow open you can always create a new workflow by clicking the + in the tabs bar at the top of the user interface.

In the new empty workflow editor, create a workflow by dragging nodes from the node repository to the workflow editor, then connecting, configuring, and executing them.

Nodes

In KNIME Analytics Platform, individual tasks are represented by nodes. Nodes can perform all sorts of tasks, including reading/writing files, transforming data, training models, creating visualizations, and so on.

Facts about nodes

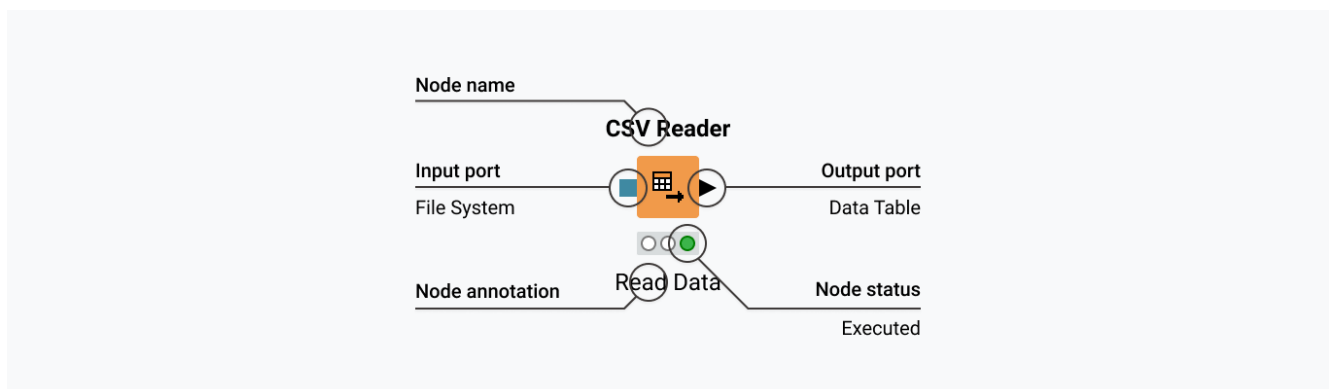


Figure 6. A node in KNIME Analytics Platform

- Each node is displayed as a colored box with input and output ports, as well as a status, as shown in **Figure 6**
- The input port(s) hold the data that the node processes, and the output port(s) hold the resulting datasets of the operation

- The data is transferred over a connection from the output port of one to the input port of another node.



For simplicity we refer to data when we refer to node input and output ports, but nodes can also have input and output ports that hold a model, a database query, or another type explained in [Node Ports](#).

A node can be in different status as shown in the [Figure 7](#).

	Not configured The node is waiting for configuration or incoming data.
	Configured The node has been configured correctly, and can be executed.
	Executed The node has been successfully executed. Results may be viewed and used in downstream nodes.
	Error The node has encountered an error during execution.

Figure 7. A node can exist in different status

Changing the status of a node

The status of a node can be changed, either configuring, executing, or resetting it.

All these options can be found:

- In the node action bar - click the different icons to configure, execute, cancel, reset and when available open the view.




Figure 8. Action bar of a node

- In the context menu of a node - open the context menu by right clicking a node.

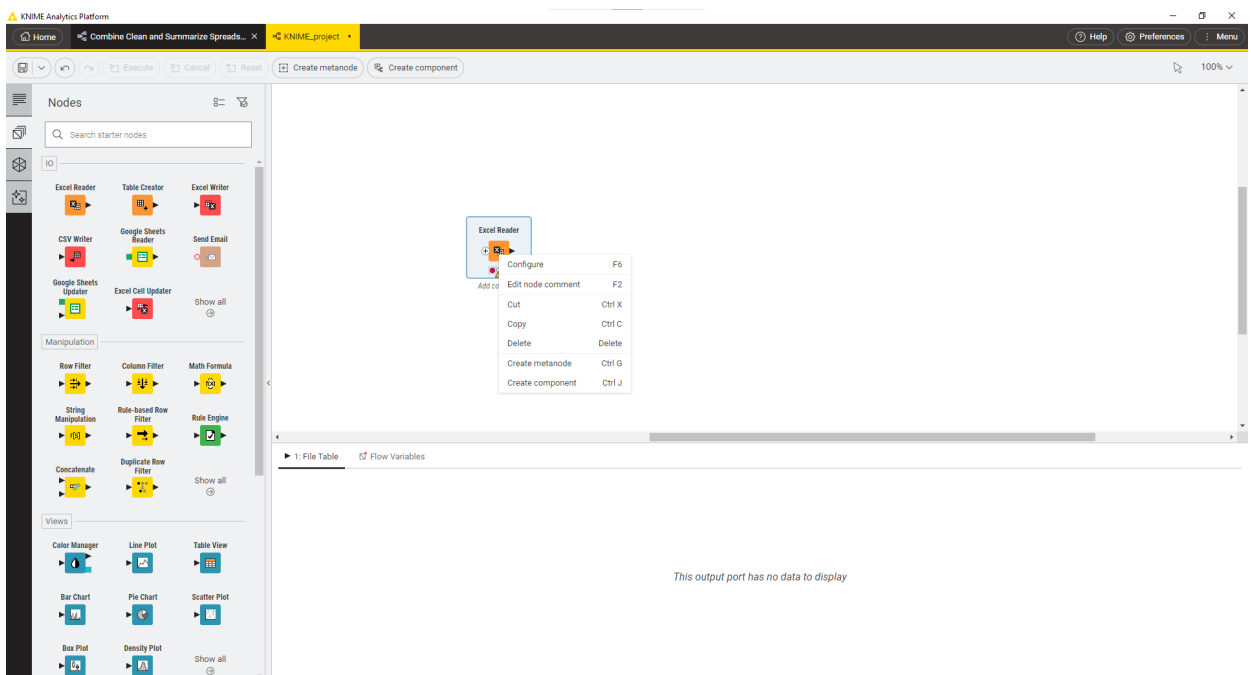


Figure 9. Context menu of a node

Identifying the node status

The traffic light below each node shows the status of the node. When a node is configured, the traffic light changes from red to yellow, i.e. from "not configured" to "configured".

When a new node is first added to the **workflow editor**, its status is "not configured" - shown

by the red traffic light below the node.

Configuring the node

The node can be configured by adjusting the settings in its configuration dialog.

Open the configuration dialog of a node by either:

- Double clicking the node
- Clicking the *Configure* button in the node action bar
- Right clicking a node and selecting *Configure* in the context menu
- Or, selecting the node and pressing F6

Executing the node

Some nodes have the status "configured" already when they are created. These nodes are executable without adjusting any of the default settings.

Execute a node by either:

- Clicking the *Execute* button in the node action bar
- Right clicking the node and selecting *Execute*
- Or, selecting the node and pressing F7

If execution is successful, the node status becomes "executed", which corresponds to a green traffic light. If the execution fails, an error sign will be shown on the traffic light, and the node settings and inputs will have to be adjusted as necessary.

Canceling execution of the node

To cancel the execution of a node click the *Cancel* button in the node action bar, or right click it and select *Cancel* or select it and press F9.

Resetting the node

To reset a node click the *Reset* button in the node action bar, or right click it and select *Reset* or select it and press F8.



Resetting a node also resets all of its subsequent nodes in the workflow. Now, the status of the node(s) turns from "executed" into "configured", the nodes' outputs are cleared.

Node ports

A node may have multiple input ports and multiple output ports. A collection of interconnected nodes, using the input ports on the left and output ports on the right, constitutes a workflow. The input ports consume the data from the output ports of the predecessor nodes, and the output ports provide data to the successor nodes in the workflow.

Besides **data tables**, input and output ports can provide other types of inputs and outputs. For each type the pair of input and output port looks different, as shown in **Figure 10**.

An output port can only be connected to an input port of the same type - data to data, model to model, and so on.

Some input ports can be empty, like the data input port of the Decision Tree View node in **Figure 10**. This means that the input is optional, and the node can be executed without the input. The mandatory inputs, shown by filled input ports, have to be provided to execute the node.



Figure 10. Common port types

A tooltip gives a short explanation of the input and output ports. If the node is executed, the dimensions of the outgoing data are shown in its data output port. A more detailed

explanation of the input and output ports is in the node description.

Adding nodes to the canvas

Currently, there are three ways of adding nodes to your canvas to build your workflow:

1. Drag and drop a node from the node repository,
2. double-click on a node inside the node repository, or
3. drop a connection into an empty area inside the workflow canvas to display the quick nodes adding panel. Up to 12 recommended nodes are displayed inside this panel. Also, you can search in the panel for all compatible nodes. Click the desired node to add it.

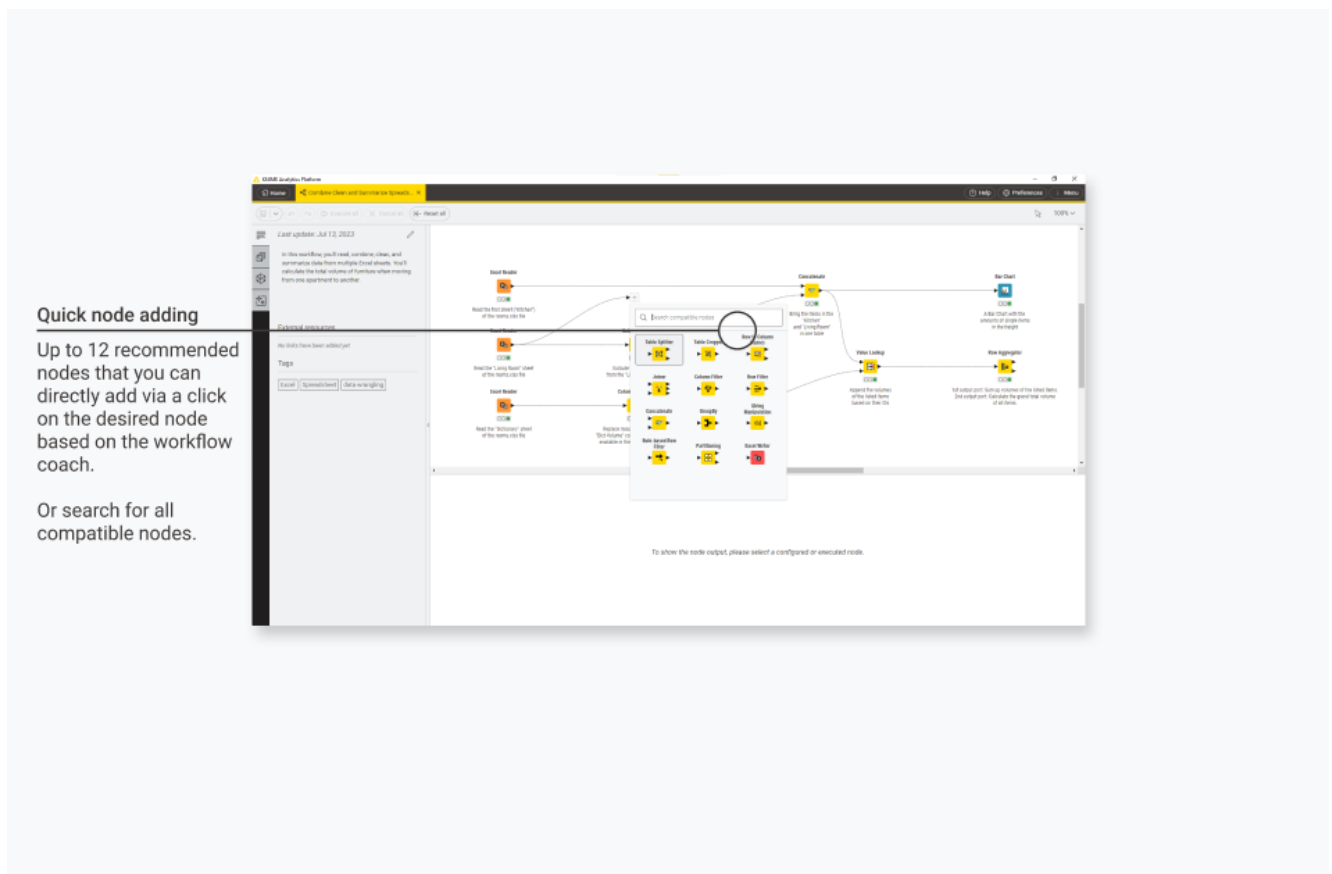


Figure 11. Quick nodes adding with recommended nodes

To use quick nodes adding you need to allow us to receive anonymous usage data. This is possible at the startup of the KNIME Analytics Platform or after switching to a new workspace by selecting Yes in the “Help improve KNIME” dialog.

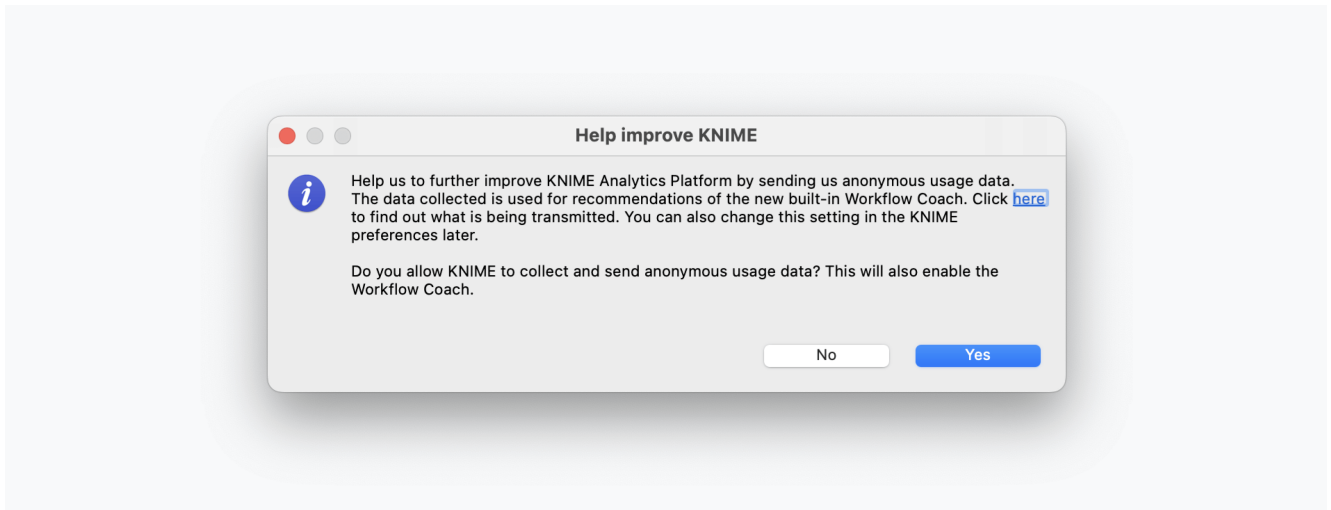


Figure 12. "Help improve KNIME" dialog

You can also activate it, via the *Open Preference* button that is displayed in the quick nodes adding panel.

[Click here](#) to find out what is being transmitted. If you don't want to do this anymore, you can deactivate it at any time in the KNIME Workflow Coach Preferences.

To open the preferences follow these steps:

1. Click *Preferences* in the top right corner of the user interface
2. Go to *KNIME* → *Workflow Coach*
3. Deactivate the setting *Node Recommendations by the Community*

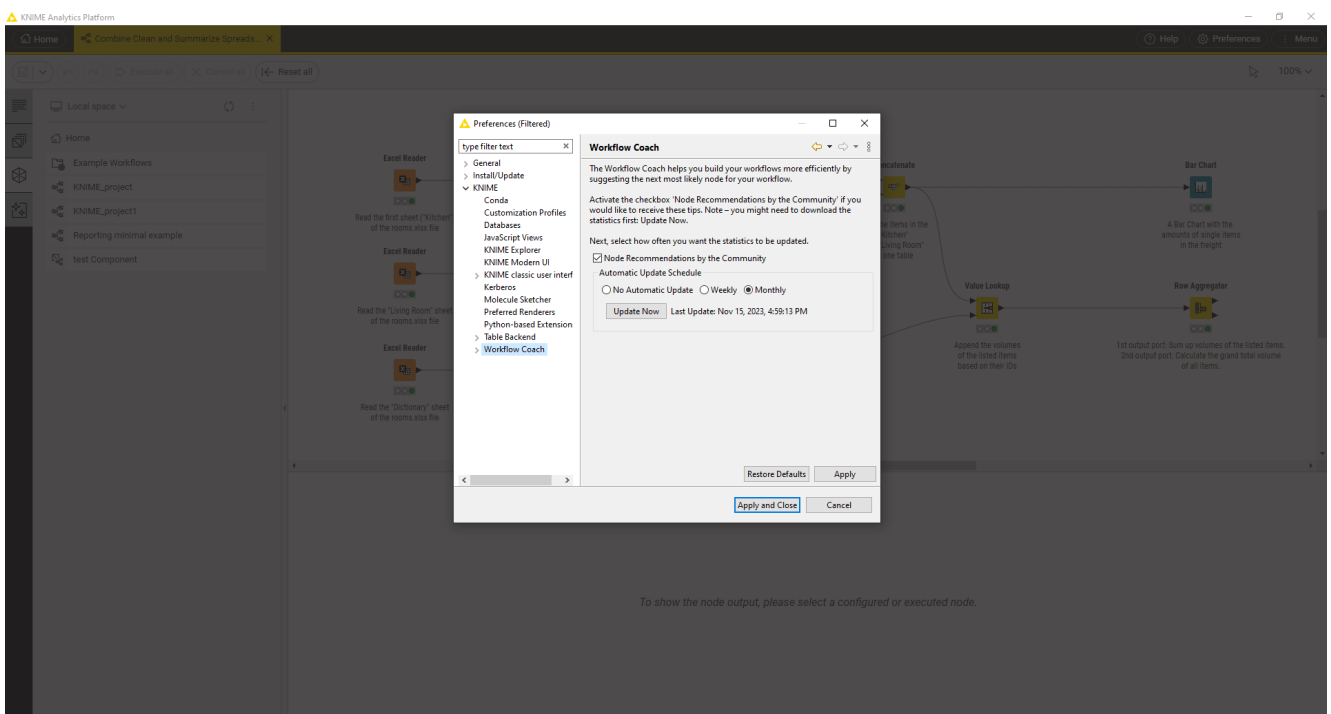


Figure 13. Workflow Coach Preferences

How to select, move, copy, and replace nodes in a workflow

Nodes can be moved into the workflow editor by dragging and dropping them. To copy nodes between workflows, select the chosen nodes, right click the selection, and select *Copy* in the menu. In the destination workflow, right click the workflow editor, and select *Paste* in the menu.

To select a node in the workflow editor, click it once, and it will be surrounded by a border. To select multiple nodes, draw a rectangle over the nodes with the mouse.

Replace a node by dragging a new node onto an existing node. Now the existing node will be covered with a colored box with an arrow and boxes inside as shown in [Figure 14](#). Releasing the mouse replaces the node.

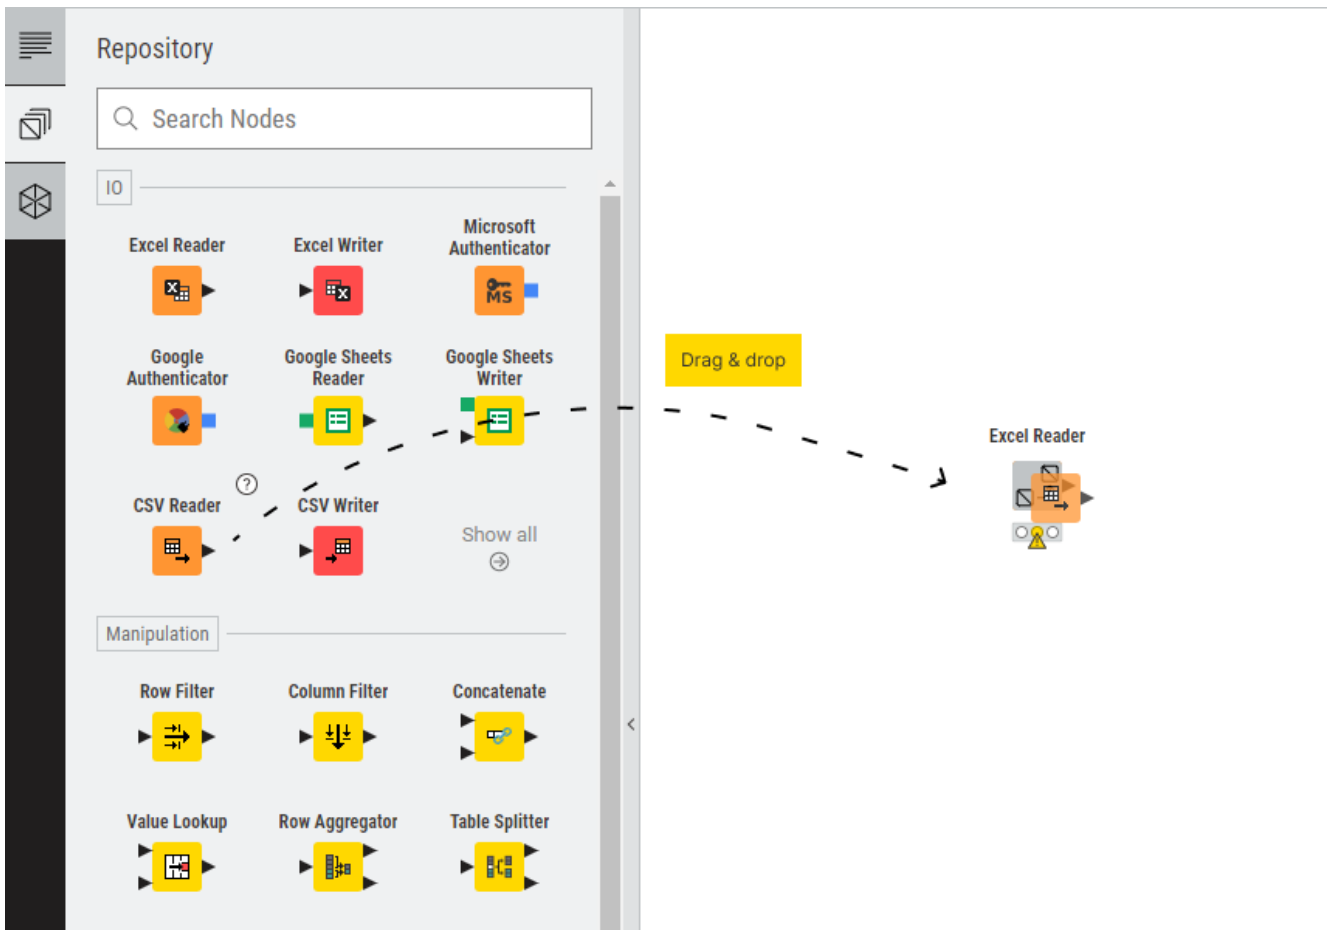


Figure 14. Replacing a node in a workflow

Comments and annotations

You have two options in the workflow editor to document a workflow:

- Node label - Add a comment to an individual node by double clicking the text field below the node and editing the text



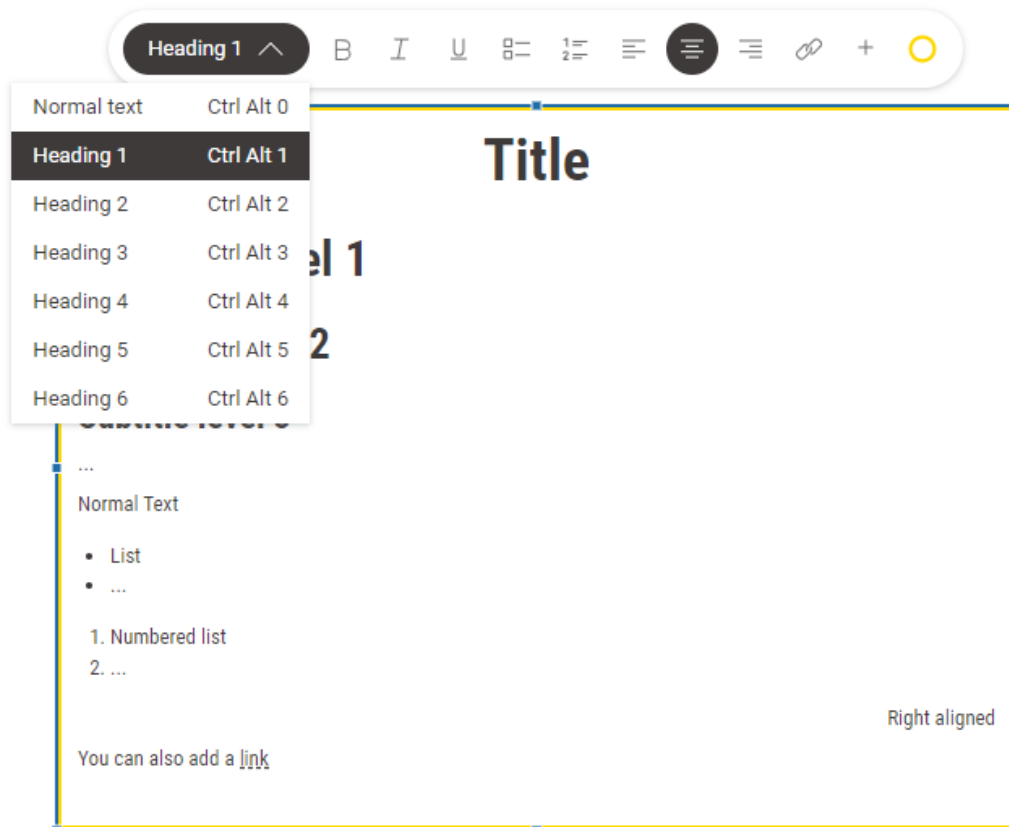


Figure 15. Writing a node comment

- Workflow annotation - Add a general comment to the workflow, right click the workflow editor and select *New workflow annotation* in the menu. Now a text box will appear in the workflow editor.



To add a new annotation you can also change the mode to annotation by clicking the  icon in the top right corner of the user interface and select *Annotation mode*, or press T to enter annotation mode.



Double click the workflow annotation to add text and format the text and change the color of the annotation outline. To change the format you can use the annotation bar or use the following syntax:

- To create a heading, add number signs (#), followed by a space, in front of a word or phrase. The number of number signs you use should correspond to the heading level (<h1> to <h6>).
- To create a bullet list, add a star sign (*) followed by a space.
- To create a numbered list, add a number followed by a point (1.), followed by a space.
- To make a text bold, italic, or underlined, select the text and press CTRL+b, CTRL+i, CTRL+u.

Finally you can click outside the annotation and click the annotation once again to move it around the canvas or to change its dimensions.

Connect to KNIME Hub

By default you can connect to your account on KNIME Community Hub from the *Home* tab.

It is possible to add a new KNIME Hub instance mount point by clicking *Preferences*, in the top right corner of the user interface.

Go to *KNIME Explorer* section and click *New...*. In the window that opens select *KNIME Hub* and add your Hub URL. Then click *Apply*.

Now the new mount point will show up in the *Home* tab.

Sign in and select the space you want to work on. The content of the space and the related operations you can do on the items are visible in the [space explorer](#).

Switch back to KNIME classic user interface

You can switch back to the classic KNIME Analytics Platform user interface under *Menu*, in the top right corner of the user interface, and select *Switch to classic user interface*.

You can switch back to KNIME Modern UI at any time by pressing the button *Open KNIME Modern UI* in the classic user interface, at the top right corner.



Really, really, *really* important disclaimer

Workflow elements such as connectors or annotations are visualized in a new way and may not look exactly like in the current KNIME Analytics Platform. Changes will therefore not look 100% the same.

Space explorer

The space explorer 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

You can switch to other spaces by:

- Going to the *Home* tab and selecting one of the available spaces. Here you can filter the space by clicking the 🔍 icon:

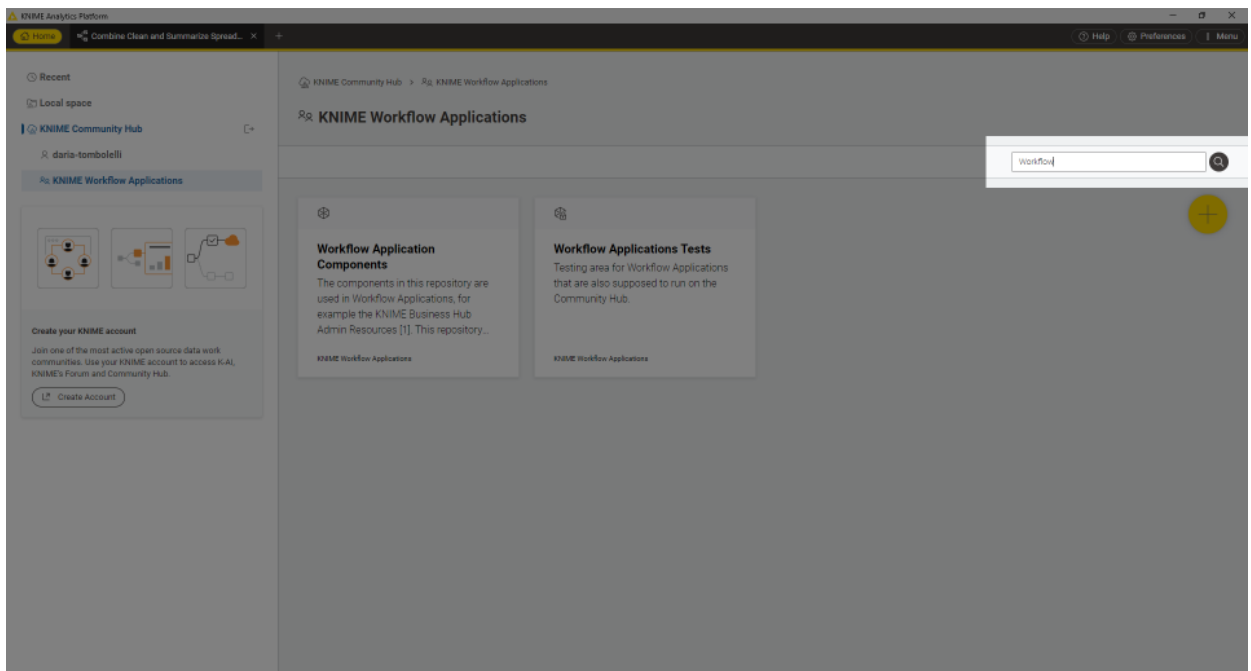


Figure 16. Filter the spaces

- On the top of the space explorer you can sign in to any of the Hub or Server mount points and select a space. You will see the spaces grouped by owner when on KNIME Hub.

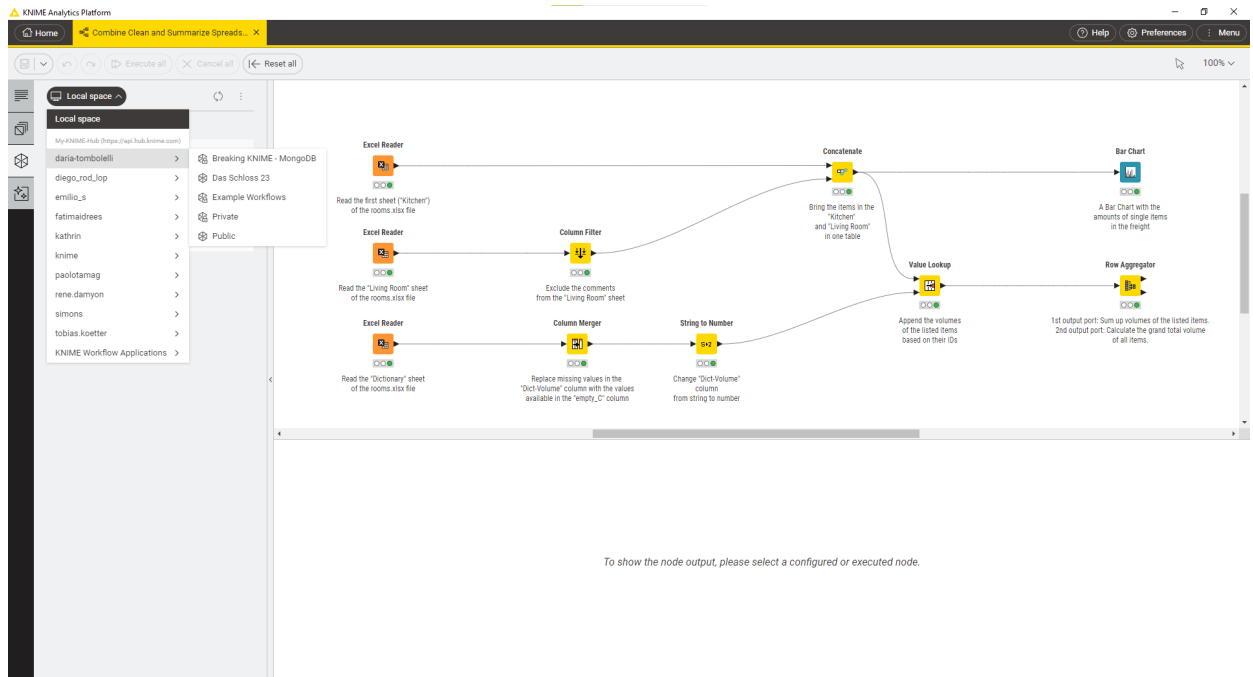


Figure 17. Select a space to explore



If you have a workflow open right-click the workflow tab at the top and select *Reveal in space explorer* to locate the workflow in the space explorer.

In the space explorer you can see:

- Workflows
- Folders
- Data files
- Components
- Metanodes

Double click a new empty workflow to open it in the workflow canvas and start adding nodes to the canvas from the node repository.



An overview on components and metanodes is available in the [KNIME Components Guide](#).

Here you can click the three dots to select one of the following actions within the current space:

- Create a new folder or a new workflow
- Import a workflow
- Add a file

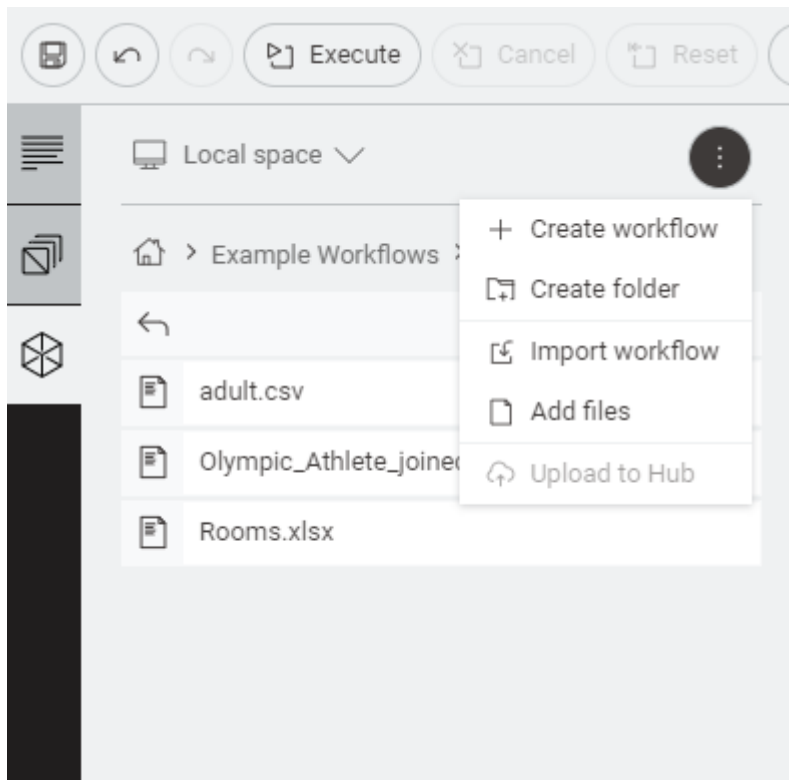


Figure 18. Space context menu

You can also drop files to the canvas. KNIME will create the appropriate file reading node automatically and preconfigure it. Finally you can drop a component to the canvas to use the component in the current workflow.

Select an item from the current space and right click on it to access the item context menu. From here you can *Rename*, *Delete*, *Export* (only available for workflows), *Upload* (if you are already connected to one of the available Hub mount points) or *Connect*.

Building workflows

When you create a new workflow, the canvas will be empty.

To build the workflow you will need to add nodes to it by dragging them from the node repository and connecting them. Alternatively you can drag an output port of a node to show the workflow coach which will suggest you the compatible nodes and directly connect them.

Once two nodes are added to the workflow editor, they can be connected by clicking the output port of the first node and release the mouse at the input port of the second node. Now, the nodes are connected. For some nodes you might have the ability to add specific ports. When hovering over these nodes you will see a + sign appearing. Click it to add a port. If the nodes supports different types of these dynamic ports a list will appear for you to scroll down to select the type of port you want to add.

You can also add a node between two nodes in a workflow. To do so drag the node from the node repository, and release it at its place in the workflow.

Node repository

Currently installed nodes are available in the node repository. You can add a node from the node repository into the workflow editor by drag and drop, as explained in the section [Building Workflows](#).

Search for a node by typing a search term in the search field on top of the node repository, as shown in [Figure 19](#).

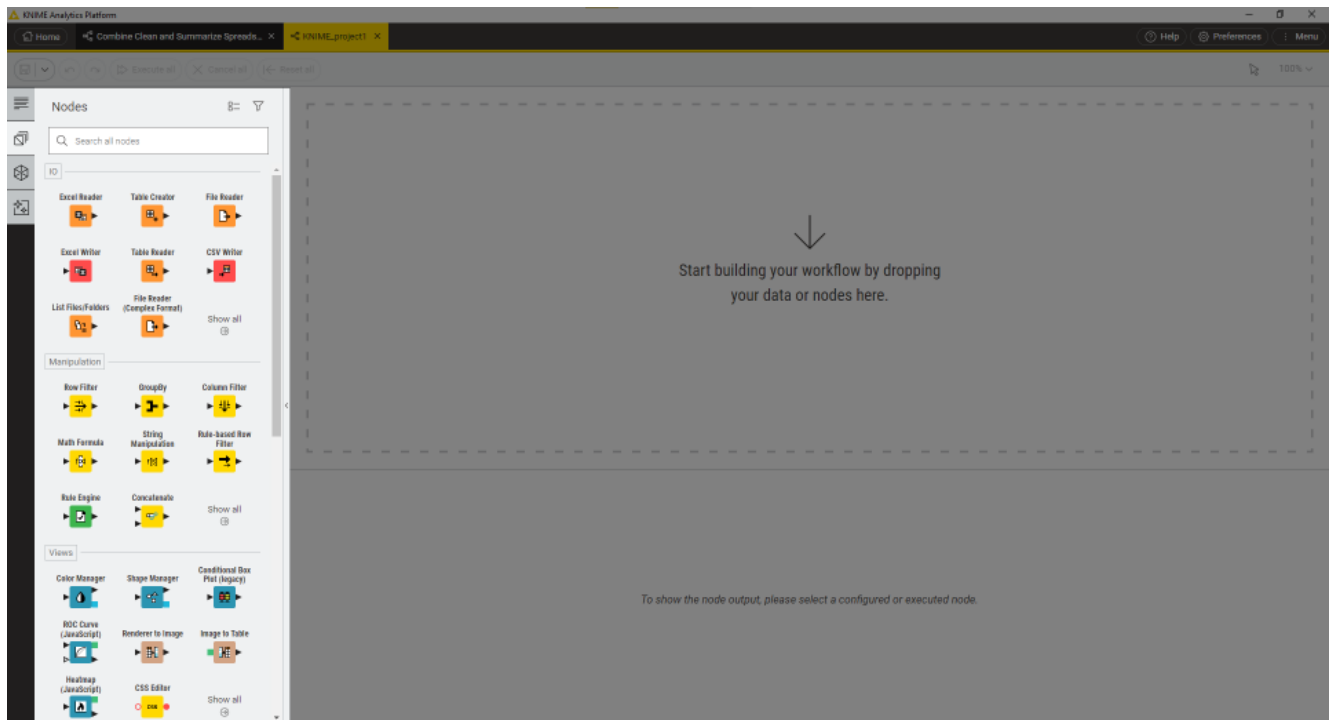



Figure 19. Node repository with two search modes

By default a specific set of nodes to help you get started with the KNIME Analytics Platform will be shown. You can expand the search results by changing the filter settings of the node repository. Click the  icon in the node repository to go to the KNIME Modern UI preferences page and change the default of the nodes included in the node repository search results.

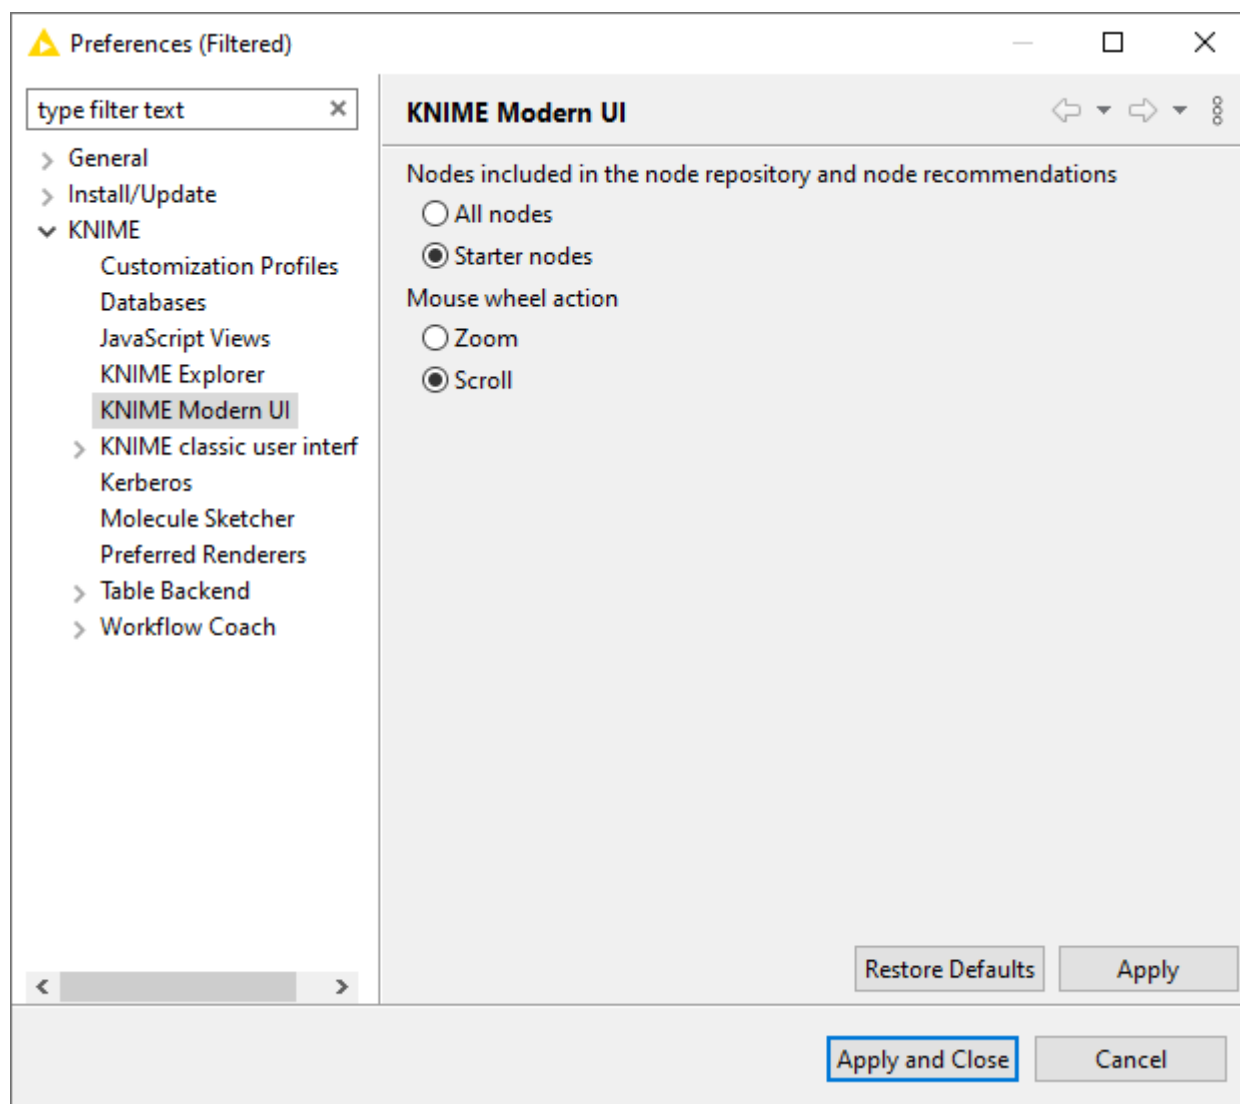




Figure 20. Change the default of the node search results

You can view your node repository as icon or list view. To switch between the two states click  (icon view) or  (list view).

Node description

You can access the node description, with information about the node function, the node configuration and the different ports available for the node in the following ways:

- Select a node you added in the canvas, go to the side panel navigation and select the first option
- Hover over a node in the node repository and click the info icon that appears. This will open the node description panel.

Workflow description

The description panel on the left of the KNIME Analytics Platform provides a description of the currently active workflow, or a selected component.

Click the pen icon to change the workflow description, add links to external resources and add tags.

Double click the workflow annotation to add text and format the text and change the color of the annotation outline. To change the format of the description you can use the formatting bar or use the following syntax:

- To create a bullet list, add a star sign (*) followed by a space.
- To create a numbered list, add a number followed by a point (1.), followed by a space.
- To make a text bold, italic, or underlined, select the text and press CTRL+b, CTRL+i, CTRL+u.



To change a component description you need to first open the component. To do so, select the component, right-click and select *Component > Open component* from the context menu.

KNIME AI Assistant

KNIME features an AI assistant designed to efficiently answer your queries about the KNIME platform and assist in constructing customized workflows, simplifying your data analysis tasks.

Installation

To install the AI assistant, first locate and open the AI assistant side panel within the KNIME interface, as shown in [Figure 21](#). Then, click the *Install AI Assistant* button and carefully follow the prompts in the installation menu to complete the setup process.

If you can not find the AI Assistant side panel, the AI Assistant might be deactivated by your administrator.

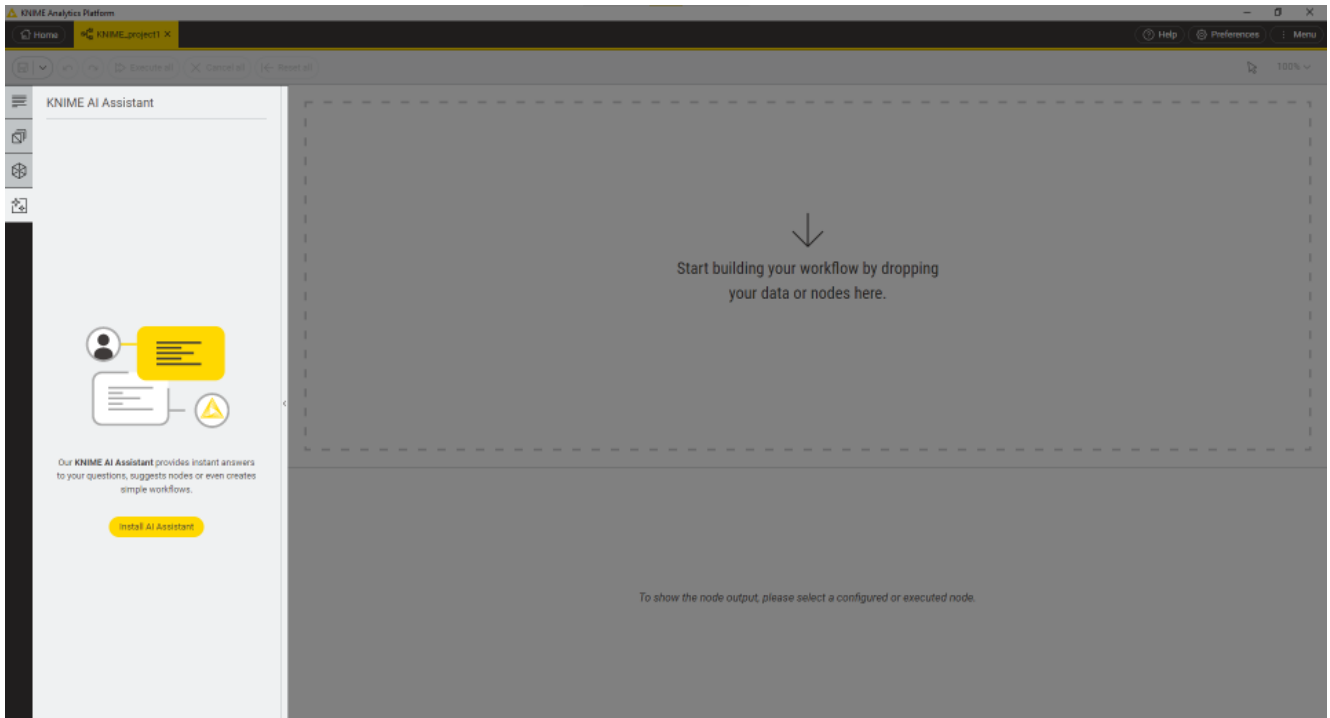


Figure 21. AI assistant side panel

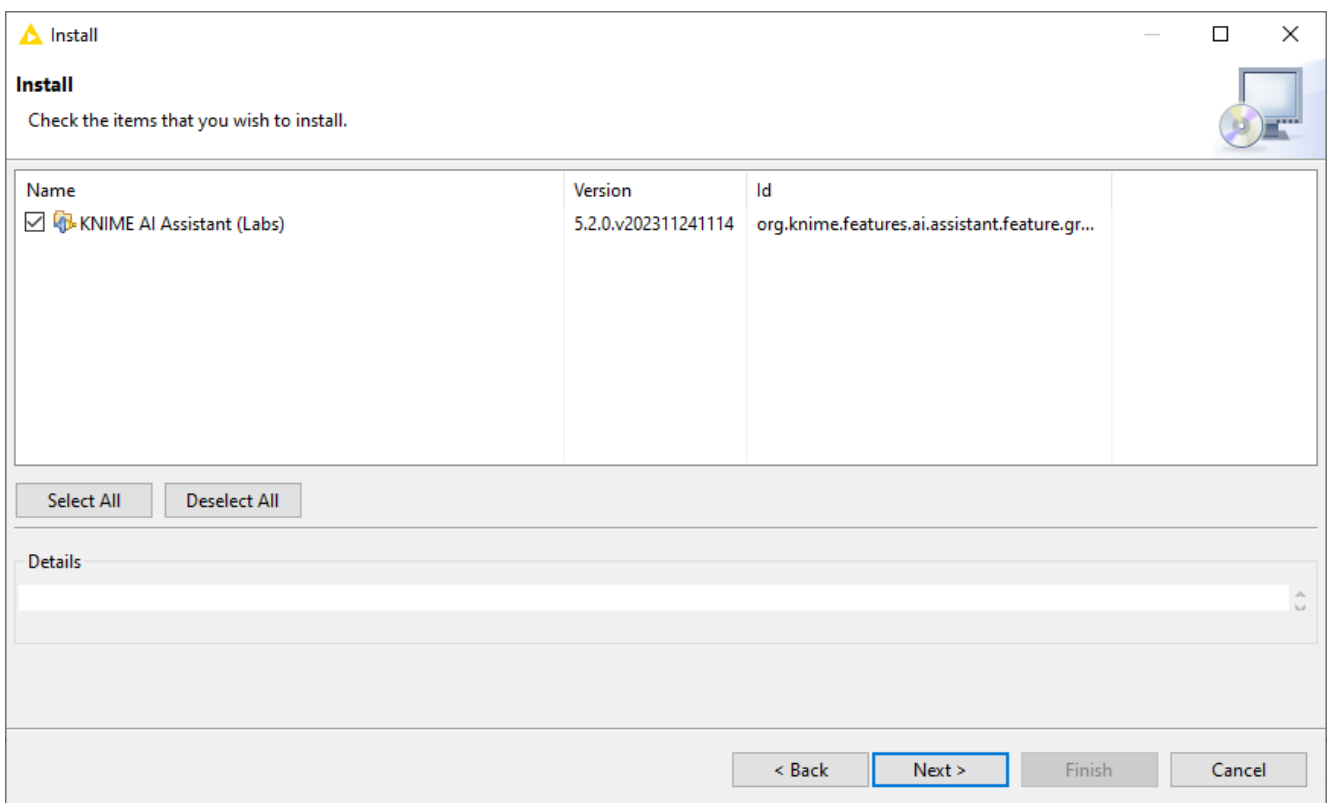


Figure 22. Installation menu

Deactivation

If you want to deactivate the AI assistant and hide it from the side panel, you can do so by adding this entry to the [knime.ini](#) file:

```
-Dorg.knime.ui.feature.ai_assistant=false
```

Usage

To access the AI assistant, please log in to KNIME Hub.

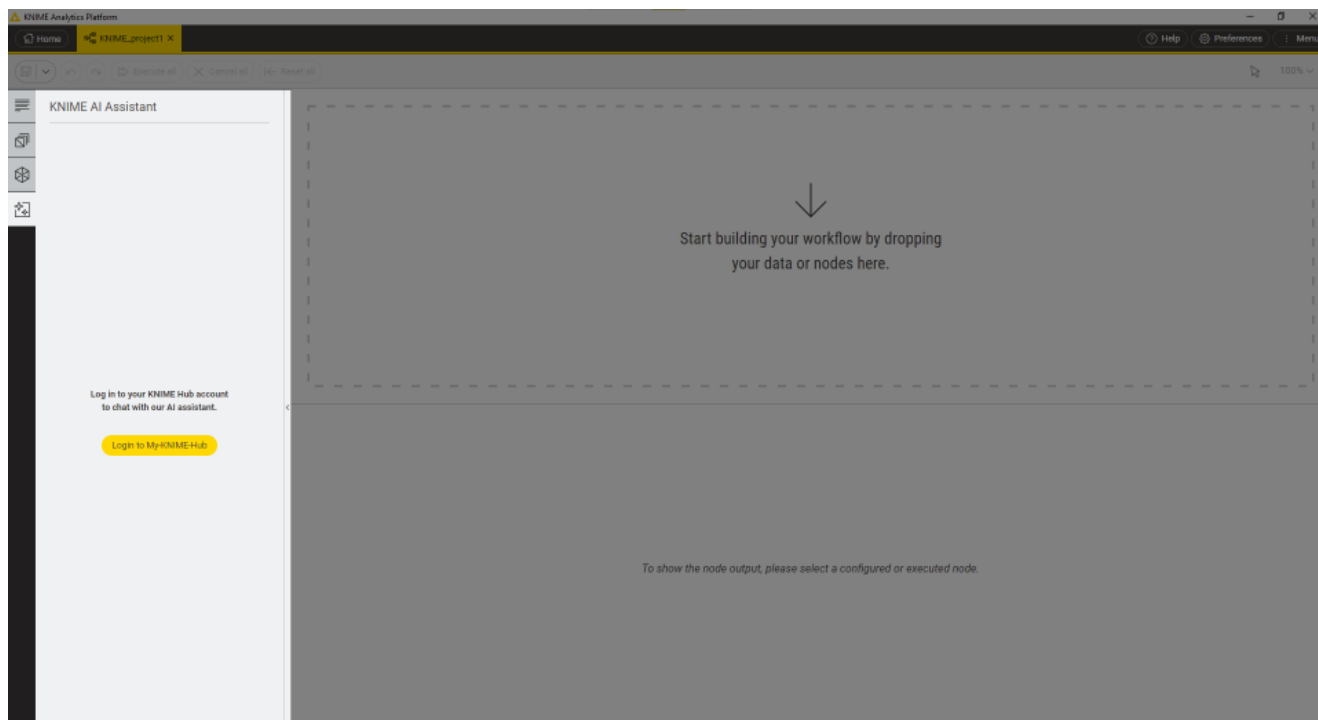


Figure 23. Login

If you have access to a KNIME Business Hub instance that is equipped with AI assistant support, you can select the specific instance to use via the AI Assistant preferences page and then log in via the AI assistant side panel.

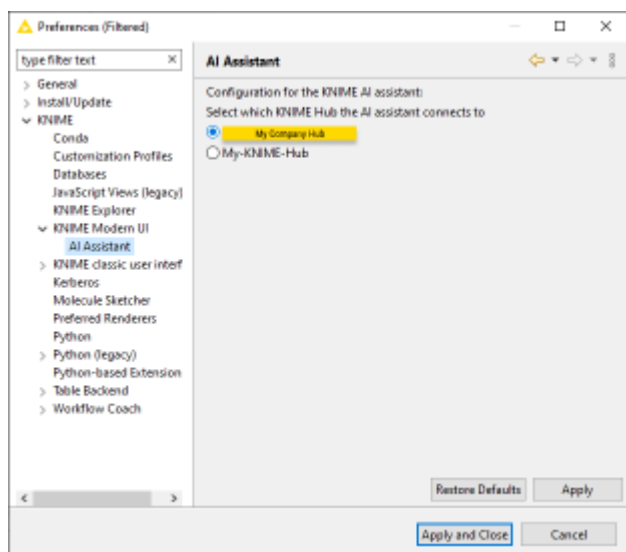


Figure 24. AI Assistant preferences page

To use the AI assistant, it is mandatory to first accept the terms outlined in the disclaimer. Please note that to deliver our services, KNIME shares data with OpenAI or Microsoft Azure. This includes all user queries, and for the **Build** mode, it includes the table specifications of selected nodes, such as column names and data types, but **not the data itself**.

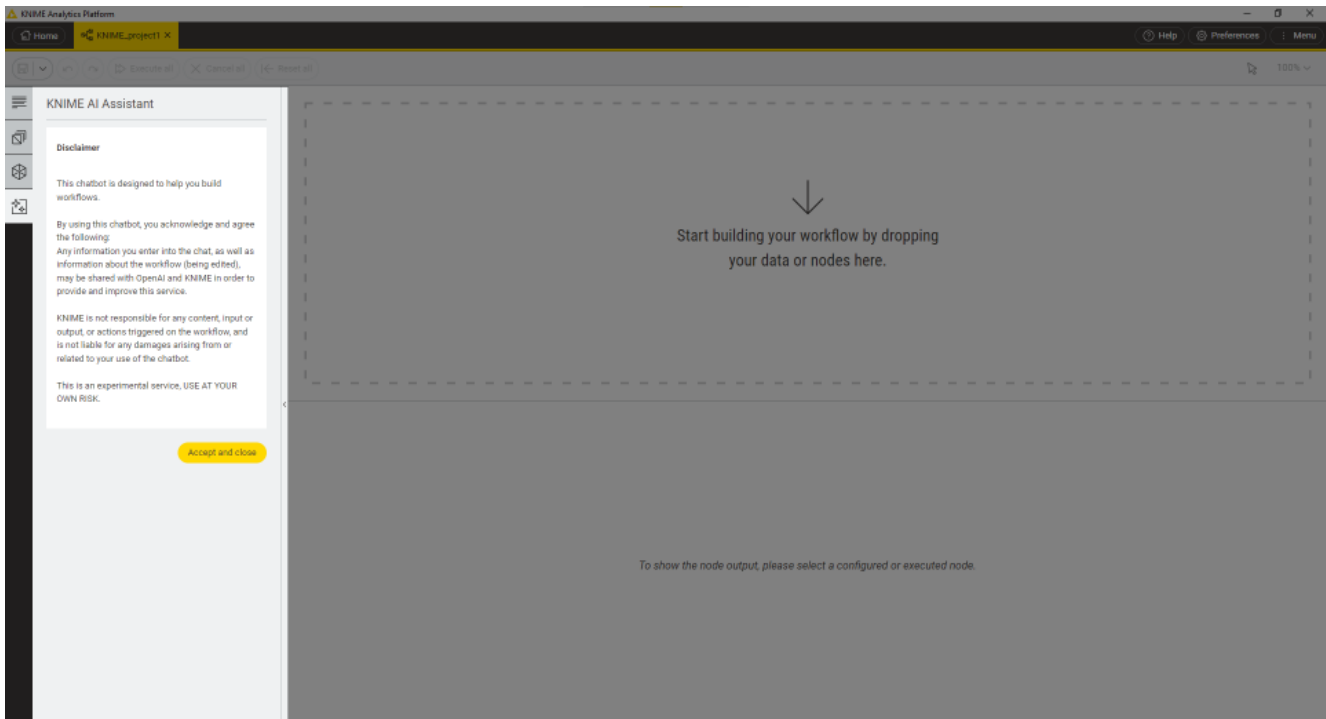


Figure 25. Disclaimer

The KNIME AI Assistant offers two modes:

- A **Q&A** mode, and
- A **Build** mode.

These can be selected using the toggle button located at the top of the side panel.

Q&A mode

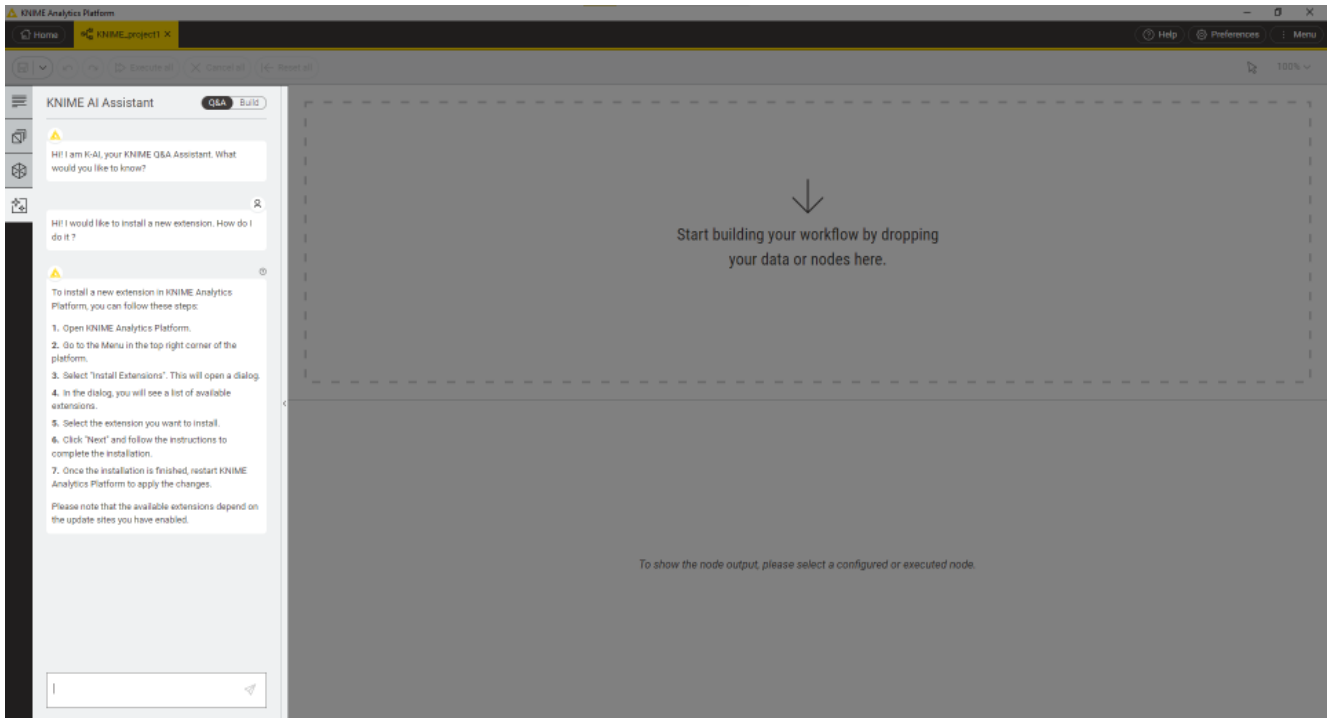


Figure 26. Q&A mode

In the Q&A mode, you can inquire about KNIME functionalities, including how to execute specific tasks, and receive informative answers.

These answers may feature recommendations for nodes effective in achieving the tasks at hand. If the suggested nodes are already installed, they can be directly dragged into the workflow. For nodes not yet installed, a link to the KNIME Hub is provided. You can then install these nodes via drag-and-drop from KNIME Hub.

By clicking the question mark located at the top of the answer, you will be provided with links to the sources that were used to generate the response.



You can leave a feedback if the answer was useful to you by hovering over the answer and clicking the thumbs up or thumbs down icons that appear.



Figure 27. Q&A mode

Build mode

The Build mode is engineered to extend workflows in response to a query. The set of nodes available to the Build mode is currently limited, but many more nodes will be added in the future. It is important to note that in Build mode, workflows cannot currently be initiated from scratch. You are required to select pre-existing nodes that already supply data, and from there, the workflow is dynamically expanded according to the your query.




Figure 28. Build mode

Workflow monitor

Access the workflow monitor tab from the side panel navigation of the user interface, shown in Figure 29. Here, you can find errors and warnings that might arise from the execution of your workflow.

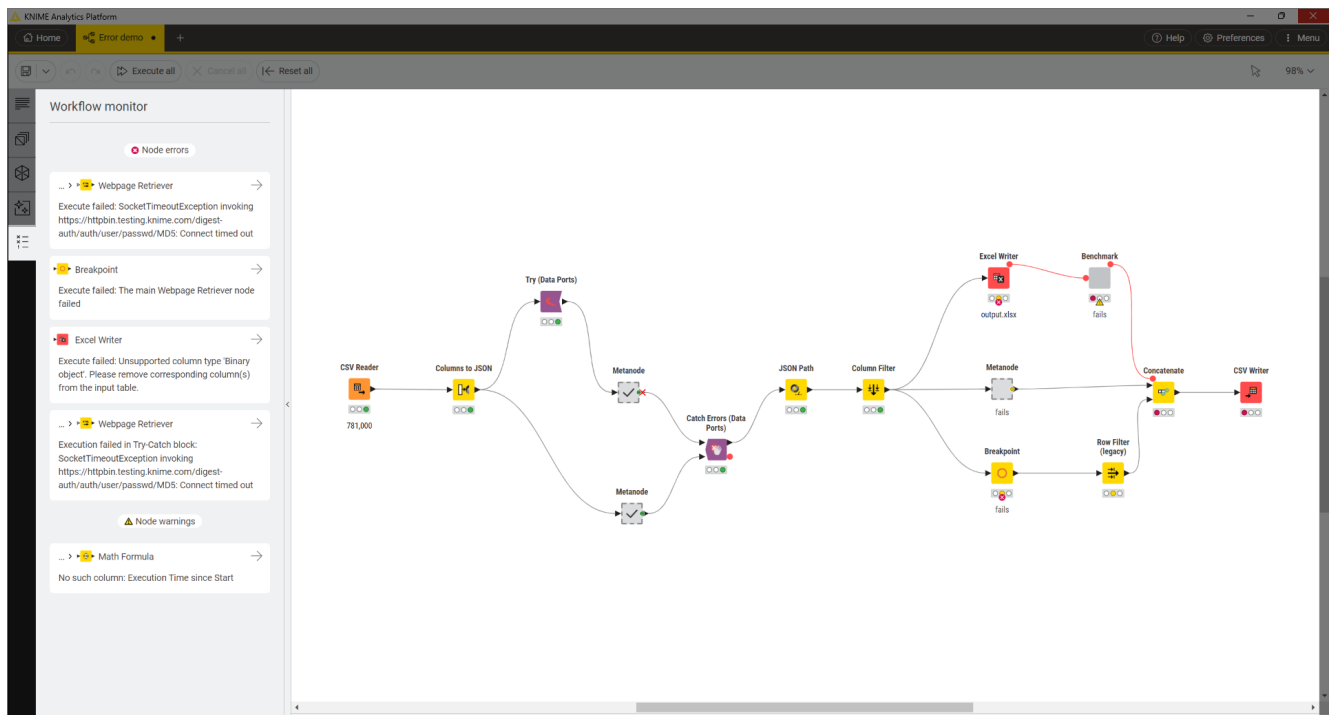


Figure 29. Workflow monitor

When a node error or a node warning occurs you can click the [16] icon to select the node that is causing the issue in the workflow.



If the node is in a component or a metanode this will automatically navigate to the level where the node that is causing the issue is present.

Node monitor

The node monitor tab is located on the bottom part of the user interface shown in [Figure 30](#). It is especially useful to inspect intermediate output tables in the workflow.

#	RowID	product	country	date	quantity	amount	card	Cust_ID
	String	String	String	String	Number (Integer)	Number (Integer)	String	String
1	Row0	prod_4	unknown	2009-12-12	1	3		Cust_3
2	Row1	prod_3	China	2009-04-10	2	160	N	Cust_2
3	Row2	prod_3	China	2009-04-10	2	160	Y	Cust_3
4	Row3	prod_3	China	2009-05-10	2	160		Cust_2
5	Row4	prod_3	USA	2009-05-20	20	1600		Cust_3
6	Row5	prod_3	Brazil	2009-06-08	15	1200		Cust_7
7	Row6	prod_1	USA	2009-07-04	2	70	Y	Cust_3
8	Row7	prod_1	USA	2009-07-14	2	70		Cust_6
9	Row8	prod_3	USA	2009-08-20	20	1600		Cust_3

Figure 30. Node monitor

Here you can choose to show the flow variables or a preview of the output data at any port of a selected node in the active workflow.

Switch to *Statistics* in order to see some basic statistics of the data.

You can also detach the table or the statistics view and open it in a new window. To do so click the icon in the respective view (*Table* or *Statistics*). This allows you to open multiple table or statistics for multiple nodes in your workflow.

Read more about the data table shown in the node monitor in the [KNIME tables](#) section.

Help

By clicking the *Help* button in the top right corner of the user interface you can see multiple useful links like:

- Keyboard shortcuts to speed up your workflow building process without relying on a computer mouse.
- Some learning resources like cheat sheets, getting started guide and documentation
- The KNIME Forum to ask the community about workflow building, tips and tricks
- About page for KNIME Analytics Platform - also with currently installed version and access to *Installation Details*
- Additional credits about open source software components

Customizing the Analytics Platform

Reset and logging

When a node is reset, the node status changes from "executed" to "configured" and the output of the node is not available anymore. When saving a workflow in an executed state, the data used in the workflow are saved as well. That is, the larger the dataset, the larger the file size. Therefore, resetting workflows before saving them is recommended in case the dataset can be accessed without any restrictions.

A reset workflow only saves the node configurations, and not any results. However, resetting a node does not undo the operation executed before. All operations done during creation, configuration, and execution of a workflow are reported in the `knime.log` file.

Click *Menu > Show KNIME log in File Explorer* in the top right corner of the user interface to see the `knime.log` location folder. The `knime.log` file has a limited size, and after reaching it the rows will be overwritten from the top.

Configuring KNIME Analytics Platform

Preferences

With the release of KNIME Analytics Platform version 5.1 the preferences have been rearranged. They can be open, from the Modern User Interface by clicking *Preferences* in the top right corner of the Analytics Platform.

Here, a list of subcategories is displayed in the dialog that opens. Each category contains a separate dialog for specific settings like database drivers, available update sites, and appearance.

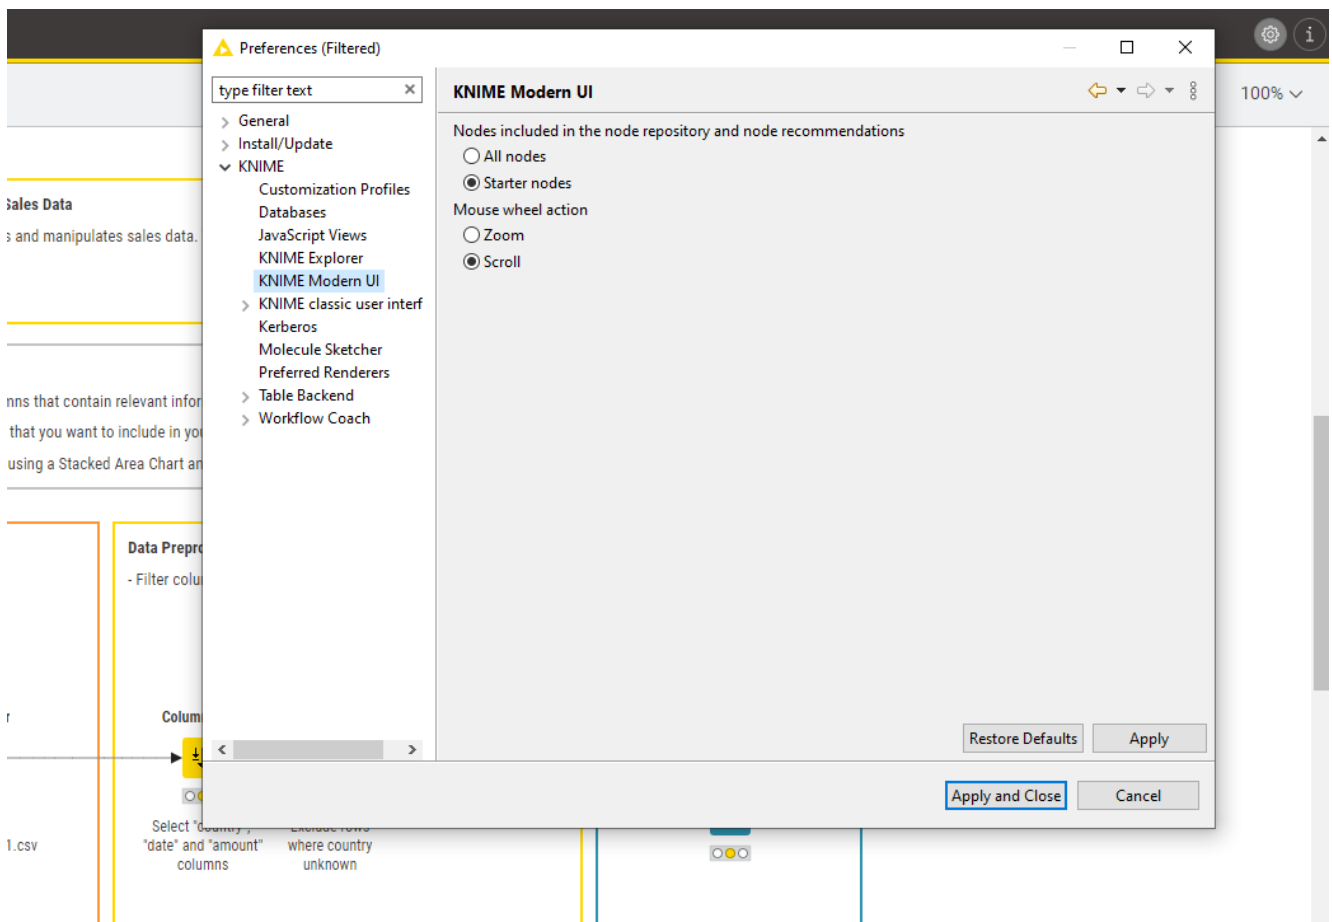


Figure 31. Open KNIME Preferences window from KNIME Analytics Platform

Network connections

Selecting *General > Network Connections* in the list of subcategories, allows you to define the networking and proxy setup of KNIME Analytics Platform.

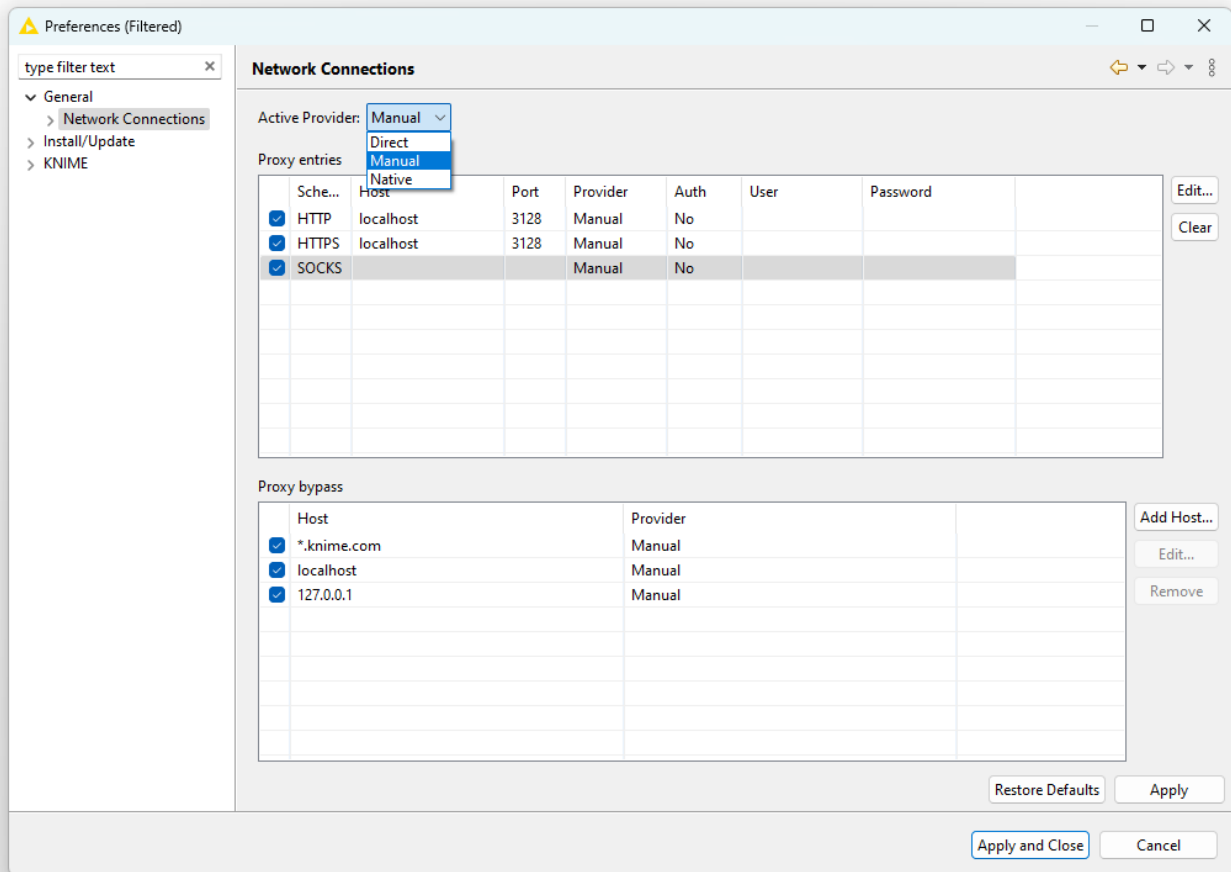


Figure 32. The Network Connections preferences page

As show in [Figure 32](#), the *Active Provider* can be set to three options:

- The *Direct* provider bypasses all proxies.
- The *Manual* provider uses the proxy configuration you see on the page.
- The *Native* provider checks for proxy settings on the OS and loads them into preferences.

Manual proxy entries are distinguished by protocol, being HTTP, HTTPS, or SOCKS. For example, the HTTPS entry only affects requests to hosts which use that protocol, such as <https://www.knime.com>.

Native proxies

Native proxies come from OS-specific, static or dynamic sources. A static proxy configuration is a hard-coded setting including the proxy host and its corresponding port. Dynamic proxy configuration is based on [proxy auto-config](#) (PAC) scripts. Proxies from this source will be labeled as *Dynamic* in the preferences.

Below you find an overview of which sources are supported in the different OSs.

Table 1. Native proxies support per OS

	Static source	Dynamic source
Windows	Proxies from the Windows system settings at <i>Network & Internet > Proxy > Manual proxy setup > Use a proxy server</i> .	PAC proxies defined at <i>Network & Internet > Proxy > Automatic proxy setup > Use setup script</i> . Note that this source does not work for all services. See below for more information.
macOS	Proxies are coming from the macOS system proxy settings.	Not supported.
Linux	First, the environment variables <code>http_proxy</code> , <code>https_proxy</code> , <code>no_proxy</code> are checked for proxies, then GNOME system settings at <i>Network > Network Proxy</i> are.	Not supported.

As mentioned above, the native proxy support on Linux is limited to GNOME systems. Additionally, loading proxies from GNOME settings requires adding this entry to the `knime.ini` file.

```
-Dorg.eclipse.core.net.enableGnome=true
```



Dynamic proxies work well for core functionality of the KNIME Analytics Platform, such as fetching updates, reloading update sites, or installing extensions. However, they are not supported for node execution in general, with some exceptions. For example, the *KNIME REST Client Extension* does support dynamic proxy sources.

Proxy authentication

KNIME supports **basic authentication** at proxies, i.e. using a username and a password. You can set the proxy credentials on the same preferences page by enabling the *Requires*

Authentication checkbox. The credentials are stored in the [Eclipse secure storage](#).

If credentials are not found, the corresponding service in the KNIME Analytics Platform will receive the HTTP response 407.

To resolve missing proxy credentials, follow these steps:

1. First, check that you correctly entered the credentials for the relevant protocol entry.
2. If you are still seeing node or log messages like `Unable to tunnel through proxy. Proxy returns "HTTP/1.1 407 Proxy Authentication Required"`, it is likely that you are missing a property in the `knime.ini` file. See [this FAQ entry](#) for more information.

Proxy exclusion

On the preferences page, you can also exclude individual hosts from using the proxy. This makes sense for local or internal hosts, for example `localhost` or `127.0.0.1`. Next to exact matching hostnames, using wildcards `*` allows you to exclude a range of hosts. It is important not to include the protocol of the URL, for example `\https://`.

Be aware that the exclusion pattern is matched to every HTTP-redirect. For example, excluding only the host `knime.com` will result in your request still using a proxy, since `knime.com` redirects to `www.knime.com` which was not excluded. For these cases, you can use the wildcard patterns, such as the pattern `*knime.com`.

KNIME

Selecting *KNIME* in the list of subcategories, allows you to define the log file log level. By default it is set to *DEBUG*. This log level helps developers to find reasons for any unexpected behavior.

Directly below, you can define the maximum number of threads for all nodes. Separate branches of the workflow are distributed to several threads to optimize the overall execution time. By default the number of threads is set to twice the number of CPUs on the running machine.

In the same dialog, you can also define the folder for temporary files.

Check the last option *Yes, help improve KNIME*. to agree to sending us anonymous usage data. This agreement activates the node recommendations in the [quick node adding panel](#).

KNIME Modern UI

In the KNIME Modern UI category you can:

- Select which nodes to include in the node repository and node recommendations
- Select which action is associated with the mouse wheel.
- Under *AI Assistant* you can also select which KNIME Hub the AI assistant connects to.

KNIME classic user interface

The *KNIME* category, contains a subcategory *KNIME classic user interface*. In this dialog, you can define the console view log level. By default it is set to "WARN", because more detailed information is only useful for diagnosis purposes.

Further below, you can select which confirmation dialogs are shown when using KNIME Analytics Platform. Choose from the following:

- Confirmation after resetting a node
- Deleting a node or connection
- Replacing a connection
- Saving and executing workflow
- Loading workflows created with a nightly build

In the same dialog, you can define what happens if an operation requires executing the previous nodes in the workflow. You have these three options:

- Execute the nodes automatically
- Always reject the node execution
- Show a dialog to execute or not


The following options allow you to define whether workflows should be saved automatically and after what time interval, also whether linked components and metanodes should be automatically updated. You can also define visual properties such as the border width of workflow annotations.

Table backend

Starting with KNIME Analytics Platform version 4.3 a new Columnar Backend is introduced, in order to optimize the use of main memory in KNIME Analytics Platform, where cell elements in a table are represented by Java objects by reviewing the underlying data representation.

The **KNIME Columnar Table Backend** extension addresses these issues by using a different underlying data layer (backed by Apache Arrow), which is based on a columnar representation.

The type of table backend used can be defined:

- At the workflow level. Open a workflow and select the Open the Description tab from the side panel navigation. Click the  icon in the top right corner of the description panel. A workflow configuration dialog will open. Here, in the tab *Table Backend* you can select the desired backend for this specific workflow from the menu.

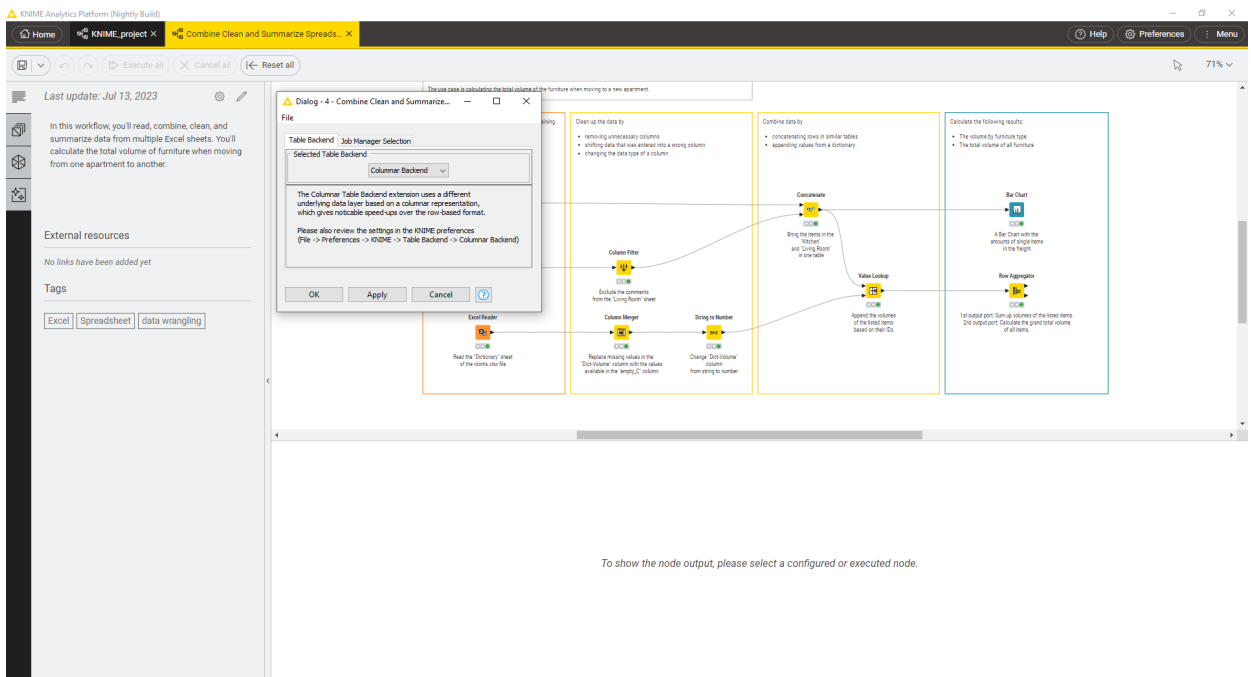


Figure 33. Configure a workflow to use Columnar Backend

- As default for all new workflows created. Open the *KNIME Preferences* and select *Table Backend* under *KNIME* in the left pane of the preferences window. Here you can select *Columnar Backend* as *Table backend for new workflows*, as shown in [Figure 34](#).




Figure 34. The Table Backend preferences page

The parameters relative to memory usage of the Columnar Backend can also be configured. Go to *File* → *Preferences* and select *Table Backend* → *Columnar Backend* under KNIME in the left pane of the preferences window, as shown in [Figure 35](#).




Figure 35. The Columnar Backend preferences page

Note that the caches of the Columnar Backend that reside in the off-heap memory region require an amount of memory in addition to whatever memory you have allotted to the heap space of your KNIME's Java Virtual Machine via the `-Xmx` parameter in the `knime.ini`. When altering the sizes of these cache via the preferences page, make sure not to exceed your system's physical memory size as otherwise you might encounter system instability or even

crashes.



For a more detailed explanation of the Columnar Backend technical background please refer to this post on [KNIME Blog](#).

High memory usage on Linux: On some Linux systems KNIME Analytics Platform can allocate more system memory than expected when using the Columnar Backend. This is caused by an unfavorable interaction between the JVM and the glibc native memory allocator. There are multiple options to circumvent this issue.

- Option 1: Reduce the number of allowed malloc areas
 1. Run KNIME Analytics platform with the environment variable `MALLOC_ARENA_MAX` set to 1.
- Option 2: Use jemalloc
 1. Install [jemalloc](#) on your OS. For ubuntu: `apt install libjemalloc2` ([link to package](#)).
 2. Find the path to jemalloc: `$ ldconfig -p | grep jemalloc`. On Ubuntu 22.04 it is `/lib/x86_64-linux-gnu/libjemalloc.so.2`.
 3. Start KNIME Analytics Platform with the environment variable `LD_PRELOAD=<path to libjemalloc.so from step 2>`.
- Option 3: Use tcmalloc
 1. Install [tcmalloc](#) on your OS. For ubuntu: `apt install google-preftools`.
 2. Find the path to tcmalloc: `$ ldconfig -p | grep tcmalloc`. On Ubuntu 22.04 it is `/lib/x86_64-linux-gnu/libtcmalloc.so.4`.
 3. Start KNIME Analytics Platform with the environment variable `LD_PRELOAD=<path to libtcmalloc.so from step 2>`.

Setting up knime.ini

When installing KNIME Analytics Platform, configuration options are set to their defaults. The configuration options, i.e. options used by KNIME Analytics Platform, range from memory settings to system properties required by some extensions.

You can change the default settings in the `knime.ini` file. The `knime.ini` file is located in the installation folder of KNIME Analytics Platform.



To locate the `knime.ini` file on macOS, open Finder and navigate to the installed Applications.

Next, right click the KNIME application, select *Show Package Contents* in the menu, and navigate to *Contents*, and open *Eclipse*.

Edit the `knime.ini` file with any plaintext editor, such as Notepad (Windows), TextEdit (macOS) or gedit (Linux).

The entry `-Xmx1024m` in the `knime.ini` file specifies how much memory KNIME Analytics Platform is allowed to use. The setting for this value will depend on how much memory is available in the running machine. We recommend setting it to approximately one half of the available memory, but this value can be modified and personalized. For example, if the computer has 16GB of memory, the entry might be set to `-Xmx8G`.

Besides the memory available, you can define many other settings in the `knime.ini` file. Find an overview of some of the most common settings in [Table 2](#) or in this [complete list](#) of the configuration options.

Table 2. Common configuration settings in `knime.ini` file

Setting	Explanation
<code>-Xmx</code> <ul style="list-style-type: none"> • default value: 1024m • example: <code>-Xmx16G</code> 	Sets the maximum amount of memory available for KNIME Analytics Platform.
<code>-Dknime.compress.io</code> <ul style="list-style-type: none"> • default value: SNAPPY • possible values: [SNAPPY GZIP NONE] • example: <code>-Dknime.compress.io=SNAPPY</code> 	Determines which compression algorithm (if any) to use when writing temporary tables to disk.
<code>-Dorg.knime.container.cellsinmemory</code> <ul style="list-style-type: none"> • default value: 5,000 • possible values: any value between 0 and 2,147,483,647 • example: <code>-Dorg.knime.container.cellsinmemory=100,000</code> 	This setting defines the size of a "small table". Small tables are attempted to be kept in memory, independent of the Table Caching strategy. By increasing the size of a small table, the number of swaps to the disk can be limited, which comes at the cost of reducing memory space available for other operations.
<code>-Dknime.layout_editor.browser</code> <ul style="list-style-type: none"> • default value version 4.7.2: swt • possible values: [cef swt] • example: <code>-Dknime.layout_editor.browser=cef</code> 	This setting defines which browser should be used to display the layout editor .

Setting	Explanation
<p><code>-Dknime.table.cache</code></p> <ul style="list-style-type: none"> • default value: LRU • possible values: [LRU SMALL] • example: <code>-Dknime.table.cache=SMALL</code> 	<p>Determines whether to attempt to cache large tables (i.e., tables that are not considered to be "small"; see setting <code>-Dorg.knime.container.cellsinmemory</code>) in memory. If set to LRU, large tables are evicted from memory in least-recently used (LRU) order or when memory becomes scarce. If set to SMALL, large tables are always flushed to disk.</p>
<p><code>-Dknime.url.timeout</code></p> <ul style="list-style-type: none"> • default value: 1,000 ms • example: <code>-Dknime.url.timeout=100</code> 	<p>When trying to connect or read data from an URL, this value defines a timeout for the request. Increase the value if a reader node fails. A too high timeout value may lead to slow websites blocking dialogs in KNIME Analytics Platform.</p>
<p><code>-Dchromium.block_all_external_requests</code></p> <ul style="list-style-type: none"> • default value: false • example: <code>-Dchromium.block_all_external_requests=true</code> 	<p>This configuration setting when set to <code>true</code> blocks all the external requests made by Chromium Embedded Framework.</p>
<p><code>-Dorg.knime.ui.feature.ai_assistant</code></p> <ul style="list-style-type: none"> • default value: true • example: <code>-Dorg.knime.ui.feature.ai_assistant=false</code> 	<p>This configuration setting when set to <code>false</code> disables the KNIME AI assistant.</p>
<p><code>-Dknime.python.cacerts</code></p> <ul style="list-style-type: none"> • default value: ENV • example: <code>-Dknime.python.cacerts=ENV</code> 	<p>This configuration setting controls which certificate authorities Python processes trust. If set to ENV, the Python process trusts the certificate authorities that are set in the Python environment. If set to AP, the Python process trusts the certificate authorities that the Analytics Platform trusts.</p>

KNIME runtime options

KNIME's runtime behavior can be configured in various ways by passing options on the command line during startup. Since KNIME is based on Eclipse, all **Eclipse** runtime options also apply to KNIME.

KNIME also adds additional options, which are described below.

Command line arguments

Listed below are the command line arguments processed by KNIME. They can either be specified permanently in the `knime.ini` in the root of the KNIME installation, or be passed to the KNIME executable. Please note that command line arguments must be specified *before* the **system properties** (see below) i.e. before the `-vmargs` parameter.

Note that headless KNIME applications, such as the batch executor, offer quite a few command line arguments. They are not described here but are printed if you call the application without any arguments.

`-checkForUpdates`

If this arguments is used, KNIME automatically checks for updates during startup. If new versions of installed features are found, the user will be prompted to install them. A restart is required after updates have been installed.

Java system properties

Listed below are the Java system properties with which KNIME's behavior can be changed. They can either be specified permanently in the `knime.ini` in the root of the KNIME installation, or be passed to the KNIME executable. Please note that system properties must be specified *after* the `-vmargs` parameter. The required format is `-DpropertyName=propValue`.

General properties

`org.knime.core.maxThreads=<number>`

Sets the maximum number of threads that KNIME is using for executing nodes. By default this number is 1.5 times the number of cores. This property overrides the value from the KNIME preferences page.

knime.tmpdir=<directory>

Sets the default directory for temporary files KNIME files (such as data files). This property overrides the value from the preferences pages and is by default the same as the `java.io.tmpdir`.

knime.synchronous.io=(true|false)

Can be used to enforce the sequential processing of rows for KNIME tables. By default, each table container processes its rows asynchronously in a number of (potentially re-used) threads. The default value is `false`. Setting this field to `true` will instruct KNIME to always handle rows sequentially and synchronously, which in some cases may be slower.

knime.async.io.cachesize=<number>

Sets the batch size for non-sequential and asynchronous handling of rows (see `knime.synchronous.io`). It specifies the amount of data rows that are handled by a single container thread. The larger the buffer, the smaller the synchronization overhead but the larger the memory requirements. This property has no effect if rows are handled sequentially. The default value is 10.

knime.domain.valuecount=<number>

The number of nominal values kept in the domain when adding rows to a table. This is only the default and may be overruled by individual node implementations. If no value is specified a default of 60 will be used.

org.knime.container.threads.total=<number>

Sets the maximum number of threads that can be used to write KNIME native output tables. By default this number equals the number of processors available to the JVM. *Note:* This value has to be greater than 0.

org.knime.container.threads.instance=<number>

Sets the maximum number of threads that can be used to write a *single* KNIME native output table. By default this number equals the number of processors available to the JVM. *Note:* This value has to be greater than 0 and cannot be larger than `org.knime.container.threads.total`.

knime.discourage.gc=(true|false)

If set to true, discourages KNIME from triggering a full stop-the-world garbage collection. Note that (a) individual nodes are allowed to disregard this setting and (b) the garbage collector may independently decide that a full stop-the-world garbage collection is warranted. Set to true by default.

org.knime.container.minspace.temp=<number>

Java property to specify the minimum free disc space in MB that needs to be available.

If less is available, no further table files & blobs will be created (resulting in an exception).

knime.columnar.chunksize=<number>

The columnar table backend horizontally divides tables into batches and vertically divides these batches into column chunks. This property controls the initial size of these chunks and thereby the number of rows per batch. A chunk is the smallest unit that must be materialized to access a single value. Changing this value can therefore impact memory footprint and overall performance. Do not change this value unless you have good reasons. The default value is 28,000.

knime.columnar.reservedmemorymb=<number>

The columnar table backend caches table data off-heap. To this end, it requires memory in addition to the JVM's heap memory, whose size is controlled via the `-Xmx` parameter. If no explicit cache sizes are set in the preferences, the default memory available for caching is computed as follows: Total physical memory minus reserved memory minus 1.25 times heap memory. The reserved memory size in this equation (in MB) can be configured via this property. The default is 4,096.

knime.columnar.verbose=(true|false)

Setting this property to true activates verbose debug logging in the columnar table backend.

-

knime.disable.rowid.duplicatecheck=(true|false)

Enables/disables row ID duplicate checks on tables. Tables in KNIME are supposed to have unique IDs, whereby the uniqueness is asserted using a duplicate checker. This property will disable this check.

Warning: This property should not be changed by the user.

knime.disable.vmlfilelock=(true|false)

Enables/disables workflow locks. As of KNIME 2.4 workflows will be locked when opened; this property will disable the locking (allowing multiple instances to have the same workflow open).

Warning: This property should not be changed by the user.

knime.database.timeout=<number>

Sets the timeout in seconds trying to establish a connection to a database. The default value is 15 seconds.

knime.database.fetchsize=<number>

Sets the fetch size for retrieving data from a database. The default value depends on the used JDBC driver.

knime.database.batch_write_size=<number>

Sets the batch write size for writing data rows into a database. The default value is 1, that is one row at a time.

knime.database.enable.concurrency=(true|false)

Used to switch on/off the database connection access (applies only for the same database connection).

Default is true, that is all database accesses are synchronized based on single connection; false means off, that is, the access is not synchronized and may lead to database errors.

knime.logfile.maxsize=<number>[mk]

Allows one to change the maximum log file size (default is 10 MB).

Values must be integer, possibly succeeded by "m" or "k" to denote that the given value is in mega or kilo byte.

knime.settings.passwords.forbidden=(true|false)

If *true*, nodes using passwords as part of their configuration (e.g. DB connection or SendEmail) will not store the password as part of the workflow on disc. Instead a null value is stored, which will cause the node's configuration to be incorrect (but valid) after the workflow is restored from disc. Default is *false*.

knime.repository.non-instant-search=(true|false)

Allows to disable the live update in the node repository search.

-

knime.macosx.dialogworkaround=(true|false)

Allows to disable the workaround for freezes when opening node dialogs under macOS.

-

knime.data.bitvector.maxDisplayBits=<number>

Sets the maximum number of bits that are display in string representations of bit vectors.

-

knime.xml.disable_external_entities=(true|false)

If set to true, all nodes that parse XML files will not read external entities defined via a DTD.

This is usually only useful when running as an executor on the server and you want prevent XXE attacks.

Plug-in dependent properties

These properties only affect some plug-ins and are only applicable if they are installed.

org.knime.cmlminblobsize=<number>[mMkK]

Allows to change the minimum size in bytes (or kilobyte or megabytes) a CML molecule must have before it is stored in a blob cell. Otherwise it is stored inline. The latter is a bit faster but needs more memory. The default is 8kB.

`org.knime.ctabminblobsize=<number>[mMkK]`

Allows to change the minimum size in bytes (or kilobyte or megabytes) a Ctab molecule must have before it is stored in a blob cell. Otherwise it is stored inline. The latter is a bit faster but needs more memory. The default is 8kB.

`org.knime.mol2minblobsize=<number>[mMkK]`

Allows to change the minimum size in bytes (or kilobyte or megabytes) a Mol2 molecule must have before it is stored in a blob cell. Otherwise it is stored inline. The latter is a bit faster but needs more memory. The default is 8kB.

`org.knime.molminblobsize=<number>[mMkK]`

Allows to change the minimum size in bytes (or kilobyte or megabytes) a Mol molecule must have before it is stored in a blob cell. Otherwise it is stored inline. The latter is a bit faster but needs more memory. The default is 8kB.

`org.knime.rxnminblobsize=<number>[mMkK]`

Allows to change the minimum size in bytes (or kilobyte or megabytes) a Rxn molecule must have before it is stored in a blob cell. Otherwise it is stored inline. The latter is a bit faster but needs more memory. The default is 8kB.

`org.knime.sdfminblobsize=<number>[mMkK]`

Allows to change the minimum size in bytes (or kilobyte or megabytes) a SDF molecule must have before it is stored in a blob cell. Otherwise it is stored inline. The latter is a bit faster but needs more memory. The default is 8kB.

KNIME tables

Data table

Very common input and output ports of **nodes** are data input ports and data output ports, which correspond to the black triangles in **Figure 36**.




Figure 36. Data input and output port

A data table is organized by columns and rows, and it contains a number of equal-length rows. Elements in each column must have the same data type.

The data table shown in **Figure 37** is produced by a CSV Reader node, which is one of the many nodes with a black triangle output port for data output. To open the table, click the node. Execute the node if it is not yet execute. The table will show up in the node monitor.

The output table has row numbers, unique RowIDs and column headers. The RowIDs are automatically created by the reader node, but they can also be defined manually. The RowIDs and the column headers can therefore be used to identify each data cell in the table. Missing values in the data are shown by a red question mark in a circle.

At the top of the node monitor you can select which output port you want to view via the tabs and the flow variable tab, which shows the available flow variables in the node output and their current values. Next row will indicate the table dimensions, meaning how many rows and how many columns there are in the table at that specific output port. Here you can also use the toggle to switch to *Statistics*. This tab shows the meta information of the table, like the columns names, columns types and some other statistics data.

► 1: File Table Flow Variables

Rows: 47 | Columns: 7 Table Statistics

#	Row...	product String	country String	date String	quantity Number (integer)	amount Number (integer)	card String	Cust_ID String
1	Row0	prod_4	unknown	2008-12-12	1	3	⊖	Cust_8
2	Row1	prod_3	China	2009-04-10	2	160	N	Cust_2
3	Row2	prod_3	China	2009-04-10	2	160	Y	Cust_5
4	Row3	prod_3	China	2009-05-10	2	160	⊖	Cust_2
5	Row4	prod_3	USA	2009-05-20	20	1600	⊖	Cust_3
6	Row5	prod_3	Brazil	2009-06-08	15	1200	⊖	Cust_7
7	Row6	prod_1	USA	2009-07-04	2	70	Y	Cust_3
8	Row7	prod_1	USA	2009-07-14	2	70	⊖	Cust_6
9	Row8	prod_3	USA	2009-08-20	20	1600	⊖	Cust_3
10	Row9	prod_2	Germany	2009-11-02	15	600	⊖	Cust_1
11	Row10	prod_2	Germany	2009-11-22	15	600	N	Cust_1
12	Row11	prod_1	Germany	2009-12-02	1	35	Y	Cust_1
13	Row12	prod_1	China	2009-12-12	1	35	Y	Cust_2
14	Row13	prod_3	USA	2010-01-03	20	1600	⊖	Cust_3
15	Row14	prod_1	Germany	2010-01-10	1	35	N	Cust_1
16	Row15	prod_3	Germany	2010-01-13	1	80	⊖	Cust_4
17	Row16	prod_2	Germany	2010-01-15	25	1000	⊖	Cust_1
18	Row17	prod_2	USA	2010-01-20	2	80	⊖	Cust_6
19	Row18	prod_2	USA	2010-02-12	6	240	Y	Cust_6
20	Row19	prod_2	USA	2010-02-22	6	240	⊖	Cust_6
21	Row20	prod_2	Brazil	2010-03-11	6	240	N	Cust_7
22	Row21	prod_3	China	2010-03-12	1	80	⊖	Cust_5

Figure 37. Data output in KNIME Analytics Platform

Column types

The basic data types in KNIME Analytics Platform are Integer, Double, and String, along with other supported data types such as Long, Boolean value, JSON, URI, Document, Date&Time, Bit vector, Image, and Blob. KNIME Analytics Platform also supports customized data types, for example, a representation of a molecule.

Switch to the *Statistics* view in an output table, to see the data types of the columns in the data table, as shown in [Figure 38](#). For numerical values, only the range of the values in the data is shown. For string values, the different values appearing in the data are shown.

► 1: File Table Flow Variables

Rows: 7 | Columns: 14 Table Statistics

Name	Type	# Missing val...	# Unique valu...	Minimum	Maximum	25% Quantile	50% Quantile ...	75% Quantile	Mean	Mean Absolu...	Standard Dev...	Sum	10 most com...
product	String	0	4	⊖	⊖	⊖	⊖	⊖	⊖	⊖	⊖	⊖	prod_1 (16; 34 ...
country	String	0	5	⊖	⊖	⊖	⊖	⊖	⊖	⊖	⊖	⊖	Germany (16; 3...
date	String	0	46	⊖	⊖	⊖	⊖	⊖	⊖	⊖	⊖	⊖	2009-04-10 (2; ...
quantity	Number (integ...	0	11	1	25	2	5	10	6.617	5.171	6.357	311	1 (11; 23.4%) ...
amount	Number (integ...	0	18	3	1,600	70	175	400	346.234	291.873	416.855	16,273	160 (6; 12.77%...
card	String	24	2	⊖	⊖	⊖	⊖	⊖	⊖	⊖	⊖	⊖	Y (15; 65.22%)...
Cust_ID	String	0	9	⊖	⊖	⊖	⊖	⊖	⊖	⊖	⊖	⊖	Cust_1 (9; 19.1...

Figure 38. Data types and data domain in "Spec" tab

The reader nodes in KNIME Analytics Platform assign a data type to each column based on their interpretation of the content. If the correct data type of a column is not recognized by the reader node, the data type can be corrected afterwards. There are nodes available to convert data types. For example: String to Number, Number to String, Double to Int, String to Date&Time, String to JSON, and String to URI.

Many of the special data types are recognized as String by the reader nodes. To convert

these String columns to their correct data types, use the Column Type Auto Cast node.

When you use the File Reader node to read a file you can convert the column types directly via the node configuration dialog. To do so go to the Transformation tab in the configuration dialog and change the type of the desired column, as shown in [Figure 39](#).




Figure 39. Change column type in File Reader node


Sorting

Rows in the table view output can be sorted by values in one column by clicking the up (ascending) and down (descending) arrow that appears hovering over the column name in the header. Note that this sorting only affects the current output view and has no effect on the node output.

To sort rows in an output table permanently, use the Sorter node. Use the Column Resorter node to reorder columns.

Column rendering

In a table view output, you can also change the way in which numeric values are displayed in a data table. For example, it is possible to display numeric values as percentages, with full precision, or replace digits by a grey scale or bars. To see these and other rendering options for a column, click the carat icon in the column header, and select the desired available renderer, as shown in [Figure 40](#). Note that these changes are temporary and have no effect on the node output.



1: File Table | Flow Variables

Rows: 47 | Columns: 7

#	Row...	product	country	date	quantity	amount	card	Cust_ID
	String	String	String	String	Number (integer)	Number (integer)	String	String
7	Row6	prod_1	USA	2009-07-04	2	70		Cust_3
8	Row7	prod_1	USA	2009-07-14	2	70		Cust_6
12	Row11	prod_1	Germany	2009-12-02	1	35		Cust_1
13	Row12	prod_1	China	2009-12-12	1	35		Cust_2
15	Row14	prod_1	Germany	2010-01-10	1	35		Cust_1
28	Row27	prod_1	USA	2010-05-17	5	175		Cust_6
30	Row29	prod_1	China	2010-06-28	10	350		Cust_5
32	Row31	prod_1	Brazil	2010-07-17	5	175		Cust_7
33	Row32	prod_1	China	2010-08-28	10	350		Cust_2
36	Row35	prod_1	China	2010-10-01	2	70		Cust_5
37	Row36	prod_1	USA	2010-10-11	2	70	Y	Cust_6
40	Row39	prod_1	USA	2011-01-10	10	350	Y	Cust_3
44	Row43	prod_1	Germany	2011-03-06	10	350	⊖	Cust_4
45	Row44	prod_1	Germany	2011-03-18	1	35	Y	Cust_4
46	Row45	prod_1	Germany	2011-03-20	11	385	N	Cust_4
47	Row46	prod_1	Brazil	2011-04-06	1	35	⊖	Cust_7
10	Row9	prod_2	Germany	2009-11-02	15	600	⊖	Cust_1
11	Row10	prod_2	Germany	2009-11-22	15	600	N	Cust_1
17	Row16	prod_2	Germany	2010-01-15	25	1000	⊖	Cust_1
18	Row17	prod_2	USA	2010-01-20	2	80	⊖	Cust_6
19	Row18	prod_2	USA	2010-02-12	6	240	Y	Cust_6
20	Row19	prod_2	USA	2010-02-22	6	240	⊖	Cust_6

Figure 40. Rendering data in table view

Table storage

When executed, many KNIME nodes generate and provide access to tabular data at their output ports. These tables might be small or large and, therefore, might fit into the main memory of the executing machine or not. Several options are available for configuring which tables to hold in memory as well as when and how to write tables to disk. These options are outlined in this section.

In-memory caching

KNIME Analytics Platform differentiates between small and large tables. Tables are considered to be small (large) when they are composed of up to (more than) 5000 cells. This threshold of 5000 cells can be adjusted via the `-Dorg.knime.container.cellsinmemory` parameter in the `knime.ini` file. KNIME Analytics Platform always attempts to hold small tables in memory, flushing them to disk only when memory becomes scarce.

In addition, KNIME Analytics Platform attempts to keep recently used large tables in memory while sufficient memory is available. However, it writes these tables asynchronously to disk in the background, such that they can be dropped from memory when they have not been accessed for some time or when memory becomes scarce. You can [configure](#) the memory consumption of a specific node to never attempt to hold its tables in memory and, instead, write them to disk on execution. This is helpful if you know that a node will generate a table that cannot be held in memory or if you want to reduce the memory footprint of a node.




Figure 41. Configuring a node's memory policy

Alternatively, by putting the line `-Dknime.table.cache=SMALL` into the `knime.ini` file, KNIME Analytics Platform can be globally configured to use a less memory-consuming, albeit much slower caching strategy. This strategy only ever keeps small tables in memory.

Disk storage

KNIME Analytics Platform compresses tables written to disk to reduce the amount of occupied disk space. By default, KNIME Analytics Platform uses the [Snappy compression](#) algorithm to compress its tables. However, you can configure KNIME Analytics Platform to use GZIP compression or no compression scheme at all via the `-Dknime.compress.io` parameter in the `knime.ini` file.

Columnar Backend

Starting with KNIME Analytics Platform version 4.3 a new Columnar Backend is introduced. This extension addresses these issues by using a different underlying data layer (backed by Apache Arrow), which is based on a columnar representation.

For information on how to set up this type of backend please refer to the [Table backend](#) section.

Shortcuts

Shortcuts in KNIME Analytics Platform allow you to speed up your workflow building process. Navigate to the *Help* button in the top right corner of the user interface. Select *Show keyboard shortcuts*. In the shortcuts window, the shortcut name is displayed on the left and the respective key sequence on the right. You can filter shortcuts according to their name and key.



The listed shortcuts are available for the KNIME Modern UI. They cannot be changed at the moment. Eclipse preferences have no impact on them.

General actions

Table 3. The supported shortcuts

Action	Mac	Windows & Linux
Close workflow	⌘ W	Ctrl + W
Create workflow	⌘ N	Ctrl + N
Save	⌘ S	Ctrl + S
Undo	⌘ Z	Ctrl + Z
Redo	⌘ ⌘ Z	Ctrl + Shift + Z
Delete	⌘ ⌘	Delete
Copy	⌘ C	Ctrl + C
Cut	⌘ X	Ctrl + X
Paste	⌘ V	Ctrl + V

Action	Mac	Windows & Linux
Select all objects	⌘ A	Ctrl + A
Deselect all objects	⌘ ⌘ A	Ctrl + Shift + A
Copy selected table cells	⌘ C	Ctrl + C
Copy selected table cells and corresponding header	⌘ ⌘ C	Ctrl + Shift + C
Close any dialog unsaved	Esc	Esc

Execution

Table 4. The supported shortcuts for execute nodes, reset and cancel node execution

Action	Mac	Windows & Linux
Configure	F6	F6
Configure flow variables	⌘ F6	Shift + F6
Execute all	⌘ F7	Shift + F7
Cancel all	⌘ F9	Shift + F9
Reset all	⌘ F8	Shift + F8
Execute	F7	F7
Open view	F10	F10
Cancel	F9	F9

Action	Mac	Windows & Linux
Reset	F8	F8
Resume loop*	⌘ F8	Ctrl + Alt + F8
Pause loop*	⌘ F7	Ctrl + Alt + F7
Step loop*	⌘ F6	Ctrl + Alt + F6
Close dialog and execute node	⌘	Ctrl + ⌘

* Find out more about loop commands in the [KNIME Analytics Platform Flow Control Guide](#).

Selected node actions

Table 5. The supported shortcuts related selected node actions

Action	Mac	Windows & Linux
Activate the n-th output port view	⌘ 1-9	Shift + 1-9
Activate flow variable view	⌘ 0	Shift + 0
Detach the n-th output port view	⌘ ⌘ 1-9	Shift + Alt + 1-9
Detach flow variable view	⌘ ⌘ 0	Shift + Alt + 0
Detach active output port view	⌘ ⌘ ⌘	Shift + Alt + ⌘
Edit node comment	F2	F2

Action	Mac	Windows & Linux
Select (next) port	⌘ P	Alt + P
Move port selection	← → ↑ ↓	← → ↑ ↓
Apply label changes and leave edit mode	⌘ ⌘	Ctrl + ⌘

Workflow annotations

Table 6. The supported shortcuts related to add workflow documentation via formatted workflow annotations

Action	Mac	Windows & Linux
Edit annotation	F2	F2
Bring to front	⌘ ⌘ PageUp	Ctrl + Shift + PageUp
Bring forward	⌘ PageUp	Ctrl + PageUp
Send backward	⌘ PageDown	Ctrl + PageDown
Send to back	⌘ ⌘ PageDown	Ctrl + Shift + PageDown
Normal text	⌘ 0	Ctrl + Alt + 0
Headline 1 - 6	⌘ ALT 1 - 6	Ctrl + Alt + 1-6
Bold	⌘ B	Ctrl + B
Italic	⌘ I	Ctrl + I
Underline	⌘ U	Ctrl + U

Action	Mac	Windows & Linux
Strikethrough	⌘ S	Ctrl + Shift + S
Ordered list	⌘ 7	Ctrl + Shift + 7
Bullet list	⌘ 8	Ctrl + Shift + 8
Add or edit link	⌘ K	Ctrl + K
Increase height	⌘ ↓	Alt + ↓
Decrease height	⌘ ↑	Alt + ↑
Increase width	⌘ ←	Alt + ←
Decrease width	⌘ →	Alt + →

Workflow editor modes

Table 7. The supported shortcuts for different modes of the workflow editor

Action	Mac	Windows & Linux
Selection mode (default)	V	V
Pan mode	P	P
Annotation mode	T	T
Leave Pan or Annotation mode	ESC	ESC

Workflow editor actions

Table 8. The supported shortcuts for workflow editor actions

Action	Mac	Windows & Linux
Quick add node	⌘ .	Ctrl + .
Connect nodes	⌘ L	Ctrl + L
Connect nodes by flow variable port	⌘ K	Ctrl + K
Disconnect nodes	⌘ ⌘ L	Ctrl + Shift + L
Disconnect nodes' flow variable ports	⌘ ⌘ K	Ctrl + Shift + K
Select node inside the quick nodes panel	← → ↑ ↓	← → ↑ ↓
Add node from quick nodes panel	⌘	⌘
Moving the selection rectangle to the next element	← → ↑ ↓	← → ↑ ↓
Select multiple elements	hold ⌘ and press ←/→/↑/↓ then select via ⌘	hold Shift and press ←/→/↑/↓ then select via ⌘
Move selected elements up	⌘ ⌘ ↑	Ctrl + Shift + ↑
Move selected elements down	⌘ ⌘ ↓	Ctrl + Shift + ↓

Action	Mac	Windows & Linux
Move selected elements right	⌘ ⌘ →	Ctrl + Shift + →
Move selected elements left	⌘ ⌘ ←	Ctrl + Shift + ←

Component and metanode building

Table 9. The supported shortcuts related to build components and metanodes

Action	Mac	Windows & Linux
Create metanode	⌘ G	Ctrl + G
Create component	⌘ J	Ctrl + J
Open component or metanode	⌘ ⌘ ⌘	Ctrl + Alt + ⌘
Open parent workflow	⌘ ⌘ ⌘ ⌘	Ctrl + Alt + Shift + ⌘
Expand metanode	⌘ ⌘ G	Ctrl + Shift + G
Expand component	⌘ ⌘ J	Ctrl + Shift + J
Rename component or metanode	⌘ F2	Shift + F2
Open layout editor	⌘ D	Ctrl + D
Open layout editor of selected component	⌘ ⌘ D	Ctrl + Shift + D

Zooming, panning and navigating inside the canvas

Table 10. The supported shortcuts related to zooming and panning

Action	Mac	Windows & Linux
Fit to screen	⌘ 2	Ctrl + 2
Fill entire screen	⌘ 1	Ctrl + 1
Zoom in	⌘	Ctrl + +
Zoom out	⌘ -	Ctrl + -
Zoom to 100%	⌘ 0	Ctrl + 0
Panning workflow canvas	<i>hold "SPACE" and drag</i>	<i>hold "SPACE" and drag</i>

Panel navigation

Table 11. The supported shortcut related to panel navigation

Action	Mac	Windows & Linux
Hide or show side panel	⌘ P	Ctrl + P

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