

# KNIME Server Preview Functionality

KNIME AG, Zurich, Switzerland  
Version 4.7 (last updated on 2018-10-24)



# Table of Contents

General introduction to KNIME Server preview functionality .....	1
Distributed executors preview installation guide .....	2
Distributed executors: Introduction .....	2
Distributed executors: Installation instructions .....	3
KNIME Server Workflow Hub usage and administration guide .....	7
Workflow Hub: Introduction .....	7
What is the Workflow Hub? .....	7
Accessing the Workflow Hub .....	7
Workflows .....	8
Workflow search .....	10
Customization .....	10
Markdown .....	11
KNIME Server: Job View usage and installation guide .....	13
KNIME Job View: Introduction .....	13
What is the Job View .....	13
Installing the Job View .....	13
Usage .....	15

# General introduction to KNIME Server preview functionality

With the release of KNIME Server 4.7 we have included three new functionalities that are available as previews. That means that the functionality is not always feature complete, or subject to change. The previews are provided to allow you to test the functionality and provide feedback that will help to shape the final product.

In case you have questions about any of the functionality in the previews please contact [support@knime.com](mailto:support@knime.com).

# Distributed executors preview installation guide

## Distributed executors: Introduction

KNIME Server 4.7 allows you to distribute execution of workflows over several executors that can sit on separate hardware resources. This allows the KNIME Server to scale much better with increasing load because it is no longer bound to a single computer. KNIME Server 4.7 implements almost all existing functionality of the RMI executors, with a couple of limitations. The table below shows you what is already available.

Feature	Available?
Workflow repository (complete functionality)	X
License server	X
Local user management (via admin pages)	X
Executing workflows via KNIME Analytics Platform	X
Executing workflows via REST	X
Scheduled execution	X
Report generation	X
Job swapping	X
Saving jobs as workflows	X
Passing inline parameters/results for jobs via REST	X
Passing files for jobs via REST	X
WebPortal execution (legacy Quick Form execution not supported)	X

Feature	Available?
WebPortal (File Upload/Download) Quickforms	
Easy installation (in AWS)	X
Dynamic executor scaling (scale-up)	X
Dynamic executor scaling (scale-down)	

All changes for enabling distributed executors are part of the standard release, only the configuration is slightly different. Nevertheless we discourage using distributed executors in production environments until the implementation is feature complete.

Installation, configuration, and operation is very similar to the single executor setup. The server communicates with the executors via a message queueing system (and HTTP(S)). We use RabbitMQ for this purpose. It can be installed on the same computer as the KNIME Server or on a different computer. In principle, the executor can also run on the same computer as the server but that is obviously only useful for testing purposes.

## Distributed executors: Installation instructions

Enabling distributed executors consists of the following steps:

- Install a new KNIME Server following the KNIME Server Administration guide.
- Shut down the server if it has been started by the installer.
- Install RabbitMQ following the instruction below.
- Adjust configuration files for the server and executor following the instructions below.
- Start the server and one or more executors.

### Installing RabbitMQ

The server talks to the executors via a message queueing system called [RabbitMQ](#). This is a standalone service that needs to be installed in addition to the KNIME Server and the executors. You can install it on the same computer as the KNIME Server or on any other computer directly reachable by both the KNIME Server and the executors.

The KNIME Server requires RabbitMQ 3.6+ which you have to install according to the [Get Started](#) documentation on their [web page](#).

Make sure RabbitMQ is running, then perform the following steps:

- Enable the RabbitMQ management plug-in by following the [online documentation](#)
- Log into the RabbitMQ Management which is available at `http://localhost:15672/` (with user `guest` and password `guest` if this is a standard installation)
- Got to the *Admin* tab and add a new user, e.g. `knime`.
- Also in the *Admin* tab add a new virtual host (select the virtual hosts section on the right), e.g. using the hostname on which the KNIME Server is running or simply `knime-server`.
- Click on the newly created virtual host, go to the *Permissions* section and set permission for the new `knime` user (all to `.*` which is the default).

## Connecting Server and executor

The KNIME Server and the executors now need to be configured to connect to the message queue.

For the KNIME Server you must specify the address of RabbitMQ instead of the path to the local executor installation in the `knime-server.config`. I.e. comment out the `com.knime.server.executor.knime_exe` option (with a hash sign) and add the option `com.knime.enterprise.executor.msgq`. The latter takes a URL to the RabbitMQ virtual host: `amqp://<user>:<password>@<rabbit-mq-host>/<virtual host>`, e.g.

```
com.knime.enterprise.executor.msgq=amqp://knime:pass4knime@rabbitmq-host/knime-server
```

Note that any special characters in the password must be URL encoded.

The same URL must also be provided to the executor as system property via the `knime.ini`:

```
-Dcom.knime.enterprise.executor.msgq=amqp://knime:pass4knime@rabbitmq-host/knime-server
```

While commands between the server and executors are exchanged via the message queue, actual data (e.g. workflows to be loaded) are exchanged via HTTP(S). Therefore, the executors must know where to reach the server. The server tries to auto-detect its own address however in certain cases this address is not reachable by the executors or – in case of https connections – the hostname doesn't match the certificate's hostname. In such cases you have to specify the correct public address in the `knime-server.config` with the option `com.knime.server.canonical-address`, e.g.

```
com.knime.server.canonical-address=https://knime-server:8443/
```

You don't have to specify the context path as this is reliably auto-detected. Now you can start the server.

Currently the executors must be started manually, the server does **not** start them. In order to start an executor (on any machine) launch the KNIME application (that has been created by the installer) with three arguments:

```
./knime -nosplash -consolelog -application  
com.knime.enterprise.slave.KNIME_REMOTE_APPLICATION
```

You can also add these arguments at the top of the `knime.ini` if the installation is only used as an executor. You can start as many executors as you like and they can run on different hosts. They will all connect to RabbitMQ (you can see them in the RabbitMQ Management in the *Connections* tab).

When you start the executor in a shell, a very simple command line interface is available to control the executor. Enter `help` at the "Executor>" prompt to get a list of available commands.

On Windows a separate window is opened for the executor process. In case there is a problem during startup (e.g. the executor cannot acquire core tokens from the server) then this window closes immediately. In this case you can add `-noexit` to the command above to keep it open and look at the log output or open at the log file which by default is `<user home>/knimeworkspace/.metadata/knime/knime.log` unless you provided a different workspace location with `-data`.

## Load limitation

If too many jobs are sent to executors this may overload them and all jobs running on that executor will suffer and potentially even fail if there aren't sufficient resources available any more (most notably memory). Therefore an executor can reject new jobs based on its current load. By default an executor will not accept new jobs any more if its memory usage is above 90% (Java heap memory, averaged over 1-minute) or the average system load is above 90% (averaged over 1-minute). These settings can be changed by two system properties in the executor's `knime.ini` file:

```
-Dcom.knime.enterprise.executor.heapUsagePercentLimit=90  
-Dcom.knime.enterprise.executor.cpuUsagePercentLimit=90
```

Moreover, if only one distributed executor is available it will currently not reject any jobs, this behavior is likely to change in subsequent releases.

# KNIME Server Workflow Hub usage and administration guide

## Workflow Hub: Introduction

The KNIME Server Workflow Hub Usage and Administration guide section covers in detail the options for the configuration and usage of the KNIME Server Workflow Hub – or Workflow Hub for short. If you are looking to install the KNIME Server you should first consult the KNIME Server Installation Quickstart Guide. For guides on connecting to the KNIME Server from the KNIME Analytics Platform, or using the KNIME WebPortal please refer to the guides: KNIME Explorer User Guide, KNIME WebPortal User Guide. Since the Workflow Hub sits on top of the KNIME Server, please consult the KNIME Server Administration guide for help setting up and configuring the KNIME Server.

## What is the Workflow Hub?

The Workflow Hub is a feature of the KNIME Server, providing users an overview over the workflows stored on the server as well as more in depth workflow information, such as a workflow image, meta information and required plugins. Additionally, the Workflow Hub allows users to give a rating, assign tags and comment on a workflow. With the integrated workflow search users can find workflows by title, author and assigned tags.

## Accessing the Workflow Hub

The Workflow Hub's path depends on the root path of the KNIME Server installation. With `<root>` being the root path as selected during the installation of the KNIME Server, the index page of the Workflow Hub can be accessed under:

```
http://server-address/knime/hub
```

The individual workflows are available under:

```
http://serveraddress/knime/hub/workflows/<workflow_path{gt}
```

where `<workflow_path>` is the URL encoded path of the workflow in the server's workflow repository, replacing forward slashes ("/") with colons (":"). The detail page of a workflow named "Data Analysis" in the workflow group "My Workflows" would be available under:

```
http://server-address/knime/hub/workflows/My%20Workflows:Data%20Analysis
```

## Workflows

Since the Workflow Hub is integrated into the KNIME Server, it shows all workflows which are also visible in the WebPortal, taking into account user and group permissions. As soon as a workflow is uploaded to the KNIME Server instance running the Workflow Hub, the workflow is available there as well. The workflow metadata, such as the description, can be edited in the KNIME Analytics Platform. The workflow description is interpreted as Markdown, a text formatting language designed to be readable in parsed and non-parsed form. Please refer to the section "Markdown" for a brief introduction.

## Workflow versioning

The Workflow Hub lists snapshots of a workflow as versions on its details page. Once a user creates a snapshot, the workflow's detail page shows a list of snapshots at the bottom of the right column, with version numbers and commit messages. Each snapshot can be accessed with an individual URL similar to the workflow details URL under:

`http://serveraddress/knime/hub/workflows/<workflow_path>/v<timestamp>`, where `<workflow_path>` is the URL encoded path of the workflow and `<timestamp>` is the creation time of the snapshot as Unix timestamp, i.e. the number of seconds that passed since 1<sup>st</sup> of January 1970. The URL

`http://serveraddress/knime/hub/workflows/My%20Workflows:Data%20Analysis/v1498651200` therefore points to a snapshot of the workflow "Data Analysis" in the workflow group "My Workflows" when the snapshot was created at the 28<sup>th</sup> of June 2017 at 12pm.

## Social features

Using the Workflow Hub, users can assign tags to workflows, rate workflows and write comments. All of those features are available on the workflow details page.

## Tags

Tags allow a flat categorization of workflows independently of the workflow group they are stored in. They can be assigned on a workflow's details page in the right column. To change the tags, click on the blue pencil symbol and enter the tags in the text box. Multiple tags are separated by comma. A list of all assigned tags is shown on the index page, providing users quick access to workflows of interest.

### Tags

Analysis PMML

Separate multiple tags by comma

## Ratings

Workflow ratings allow users to give feedback regarding the quality of a workflow on the KNIME Server. A rating can be given on a workflow's details page using the star symbols on the right side. Every user is allowed to rate a workflow only once. Subsequent ratings of the same user only change this user's rating.

### Rating

This workflow has not been rated yet

Rate this workflow:

☆☆☆☆☆

## Comments

Workflow comments are displayed below the workflow image on the workflow details page. The newest comment always appears on top of the list and new comments can be written in the text box above. Just like the workflow description, comments can be formatted in Markdown (see section "Markdown" for a brief introduction). Once a comment is submitted, it can be edited and deleted by the original author and any user with administrator privileges.

Your comment:

Comments support [Markdown](#)

## Workflow search

The workflow search allows the searching the workflow repository by title, tags and author. The results depend on the read permissions the user has for the individual workflows or their workflow groups. To search for terms in the workflow title, the query can be entered verbatim in the search field. To search for a tag, the prefix "tag:" can be prepended to the query and to search for workflows of a certain author, the prefix "author:" can be used. When searching for a tag or author that contains spaces, the tag or author name can be put in double quotes. When providing multiple queries at once, all workflows that match all of the queries are returned. The following table lists and explains some queries.

Query	Description
PMML	All workflows containing "PMML" in their name or the author's name
author:testuser	All workflows uploaded by user "testuser"
tag:analysis	All workflows having the tag "analysis"
author:"Test User" tag:"Data Analysis"	All workflows uploaded by user "Test User" and having the tag "Data Analysis"
PMML author:"Test User"	All workflows containing the text "PMML" in the title and were uploaded by user "Test User"

 sort by Title ▾

Title	Author	Rating	Created	Uploaded	Tags
<a href="#">PMML Poster</a>	<a href="#">Fritz</a>	★★★	13.01.2015 12:38:01	12.10.2017 10:00:27	<span>PMML</span>
<a href="#">What's new - Modular PMML (1)</a>	<a href="#">Frieda</a>	☆	28.01.2015 13:06:15	24.01.2017 09:39:07	
<a href="#">What's new - Modular PMML (2)</a>	<a href="#">Fred</a>	★★★	28.01.2015 13:06:15	25.07.2017 12:13:38	<span>Modular PMML</span> <span>PMML</span>
<a href="#">What's new - PMML to SQL and Java</a>	<a href="#">Franzi</a>	☆	17.11.2014 14:49:08	24.01.2017 09:39:07	

## Customization

The index page, the workflow details page and the pages containing legal information can be customized with HTML pages placed in the subdirectory "hub" in the extension directory of the KNIME Server installation, i.e. <server repository>/extensions/hub/<file>. The "hub"

directory is watched by the server process and changes to the files are usually applied within a few seconds. Deletion of the files `terms.html`, `copyright.html` and `imprint.html` removes those entries also from the "Legal" section in the page footer. When all three files are deleted, the legal section of the footer is completely hidden. When the file `terms.html` is present, a note is displayed under the workflow download button, advising the user that by downloading the workflow they agree to the terms and conditions. The HTML files are embedded in the Workflow Hub pages and have access to the website's stylesheets. The CSS framework used for styling is [Bootstrap 3](#).

The following table lists and explains the customization options. Please note that the custom content must be placed in files with the names as given in the table.

<b>Lead Text</b> – <code>lead.html</code> This content is displayed at the top of the index page and can be a general welcome text.
<b>News Text</b> – <code>news.html</code> This content is displayed as the first tile of the index page, next to the most recent workflows.
<b>Workflow Details</b> – <code>extrawfinfo.html</code> This content is displayed at the bottom of the right column of the workflow details page.
<b>Help</b> – <code>help.html</code> This content is shown on the help page.
<b>Terms and Conditions</b> – <code>terms.html</code> This content is shown on the terms and conditions page. Linked from the footer.
<b>Copyright</b> – <code>copyright.html</code> This content is shown on the copyright page. Linked from the footer.
<b>Imprint</b> – <code>imprint.html</code> This content is shown on the imprint page. Linked from the footer.
<b>Privacy Policy</b> – <code>privacy_policy.html</code> This content is shown on the privacy policy page. Linked from the footer.

## Markdown

For formatting text in the workflow description and the comments the Workflow Hub uses Markdown. This language allows users to make text cursive or bold, add headers and lists or embed images from URLs.

The following table explains some formatting options for Markdown.

Format	Example
Heading	# My Workflow
Sub-Heading	## My Workflow
Sub-Sub-Heading	### My Workflow
Italic text	<i>_italic_, *italic*</i>
Bold text	<b>__bold__, **bold**</b>
Monospace text	<code>`monospace`</code>
Horizontal rule	---
Bullet list	<ul style="list-style-type: none"><li>* Item 1</li><li>* Item 2</li><li>* Item 3</li></ul>
Numbered list	<ol style="list-style-type: none"><li>1. Item 1</li><li>2. Item 2</li><li>3. Item 3</li></ol>
Link to a website	[KNIME](http://knime.com)

Additionally, paragraphs can be separated by a blank line and line breaks can be inserted by appending two spaces to a line.

# KNIME Server: Job View usage and installation guide

## KNIME Job View: Introduction

The KNIME Job View enables users to investigate the status of jobs on the server. Whenever a workflow is executed on the KNIME Server, it is represented as a job on the server. This instance of your workflow will be executed on the KNIME Server. By viewing the job it's possible to see the current job status. That can also be helpful if there are problems with the workflow that need to be debugged.

## What is the Job View

The KNIME Job View can be used to inspect a job on the KNIME Server. With this first preview release it will enable you to see a snapshot of the workflow. You can see nodes currently executing, errors and warnings on the nodes as well as the configuration of the nodes. It is not possible to change the configuration of nodes, or to view the data.

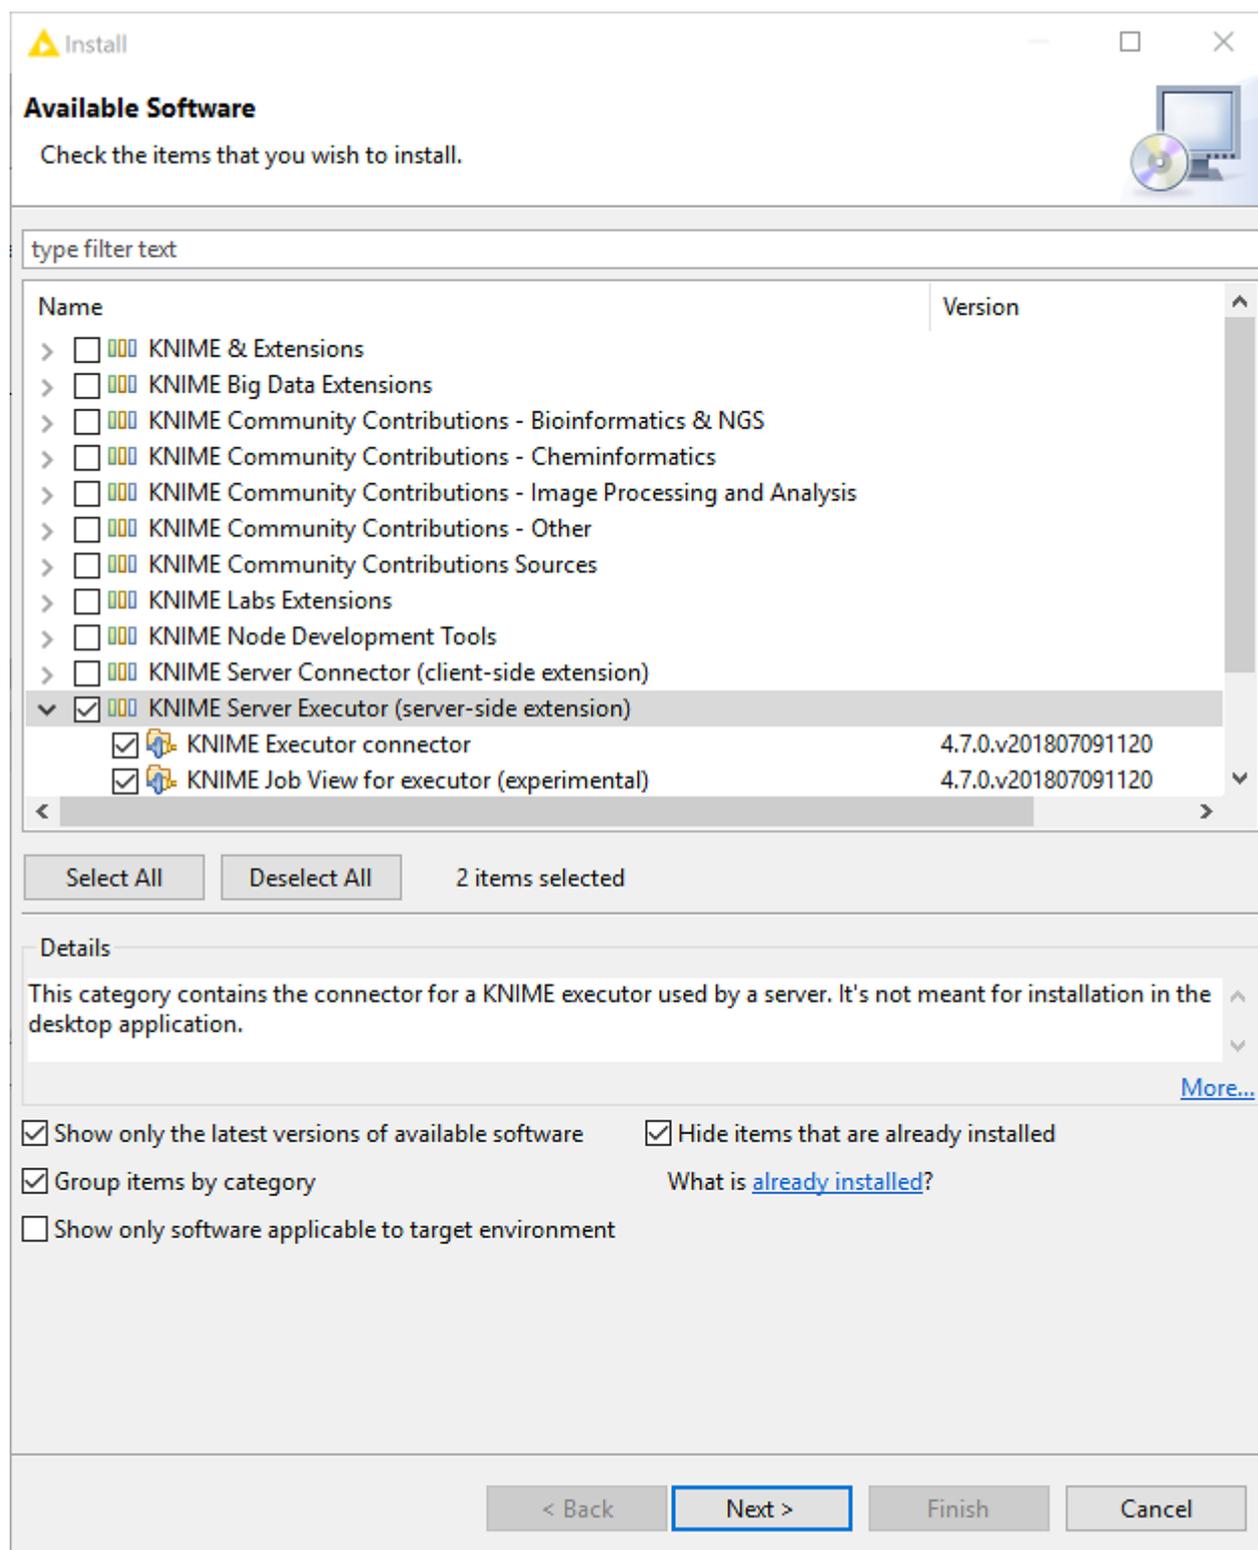
## Installing the Job View

It's required to install an extension on the KNIME Server, and then all

### Server setup

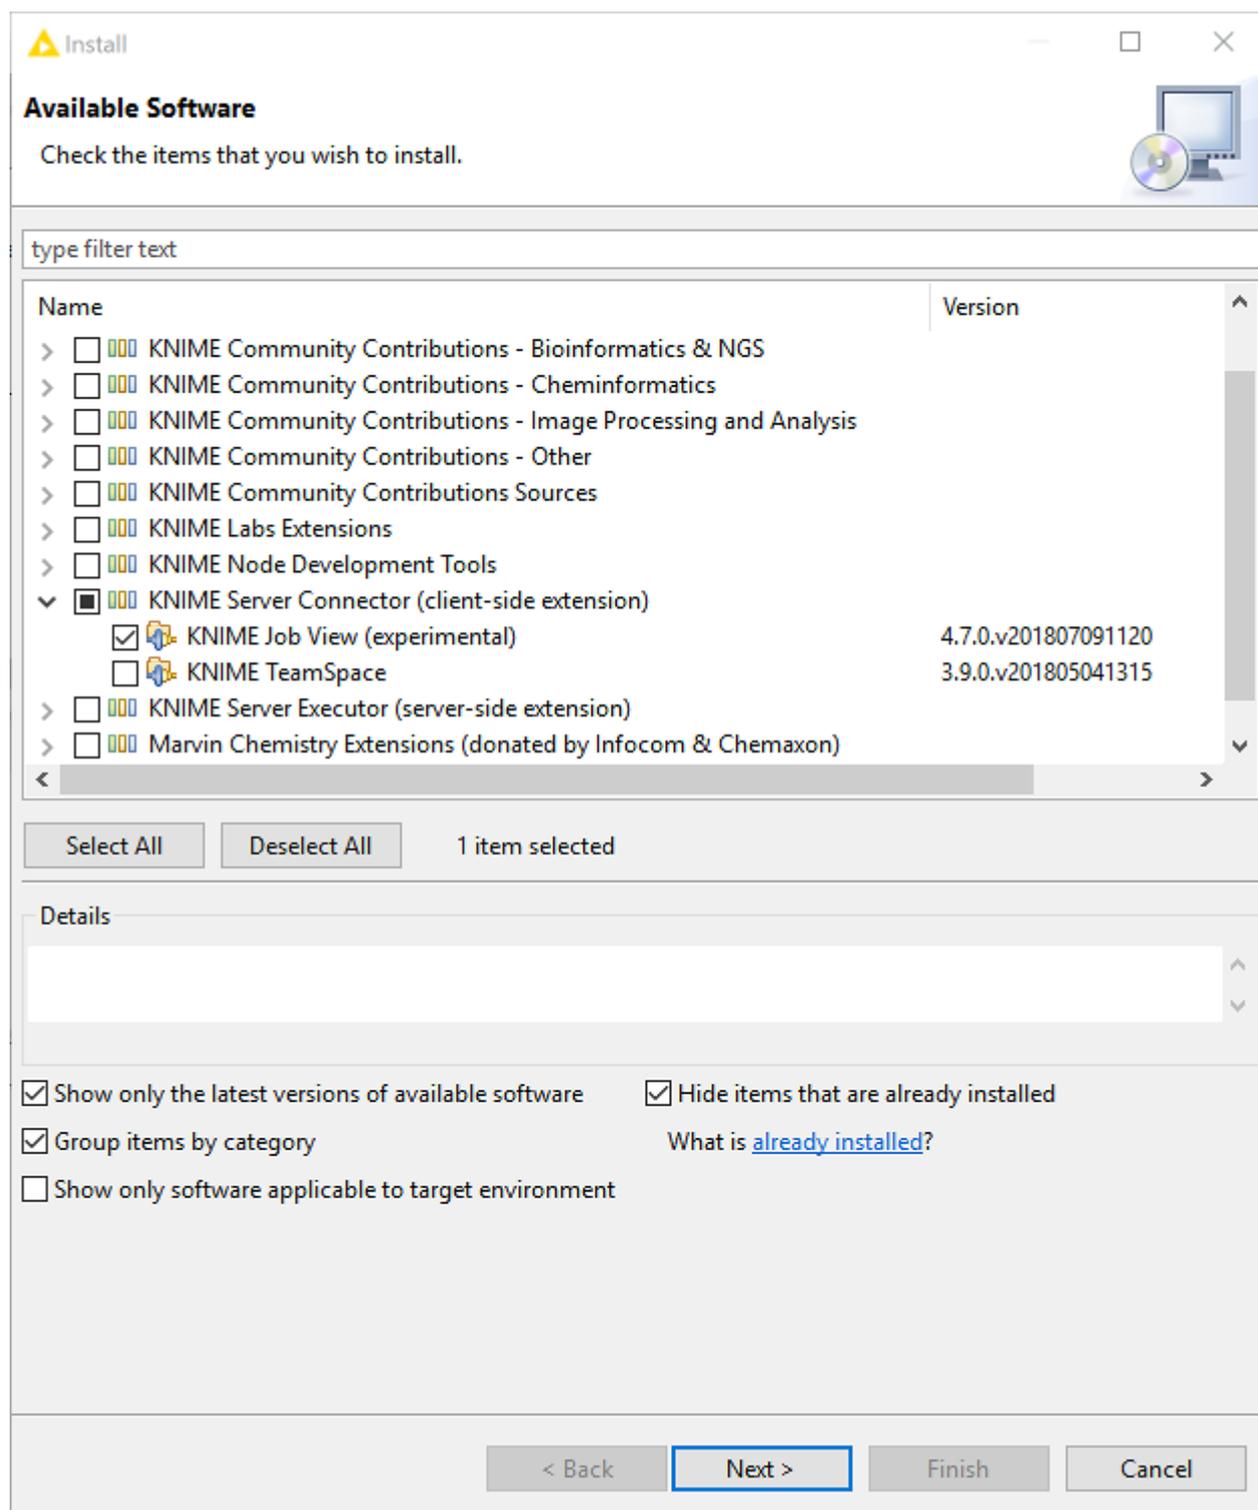
If KNIME Server is installed on Windows Server, then you may use the GUI to install the "KNIME Job View for executor (experimental)" from the "KNIME Server Executor (server-side extension)" feature. For Linux servers it is normally easier to use the command line to install the feature. Change to the KNIME Executor installation directory, and use the command:

```
./knime -application org.eclipse.equinox.p2.director -nosplash \  
-consolelog -r +http://update.knime.com/analytics-platform/3.6+ -i \  
com.knime.features.gateway.remote.feature.group -d $PWD
```



## Analytics Platform setup

The Job View feature needs to be installed in the KNIME Analytics Platform. Choose File > Install KNIME Extensions, and then select "KNIME Job View (experimental)" from the "KNIME Server Connector (client-side extension)" category.

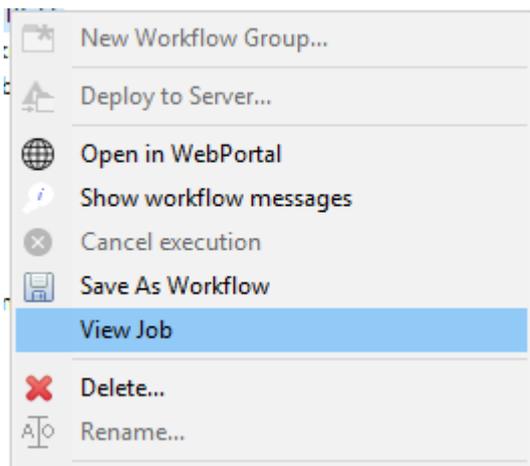


## Usage

To use the job view, you must first execute a job on your KNIME Server. This can be done via the KNIME client or by executing a workflow on the WebPortal.

This job can be visualized with a KNIME Analytics Platform using the KNIME Job View. To do so, first login to the KNIME Server instance. Select the Job and open the context menu on it

via right click. You now have the option "View Job" available in the context menu.



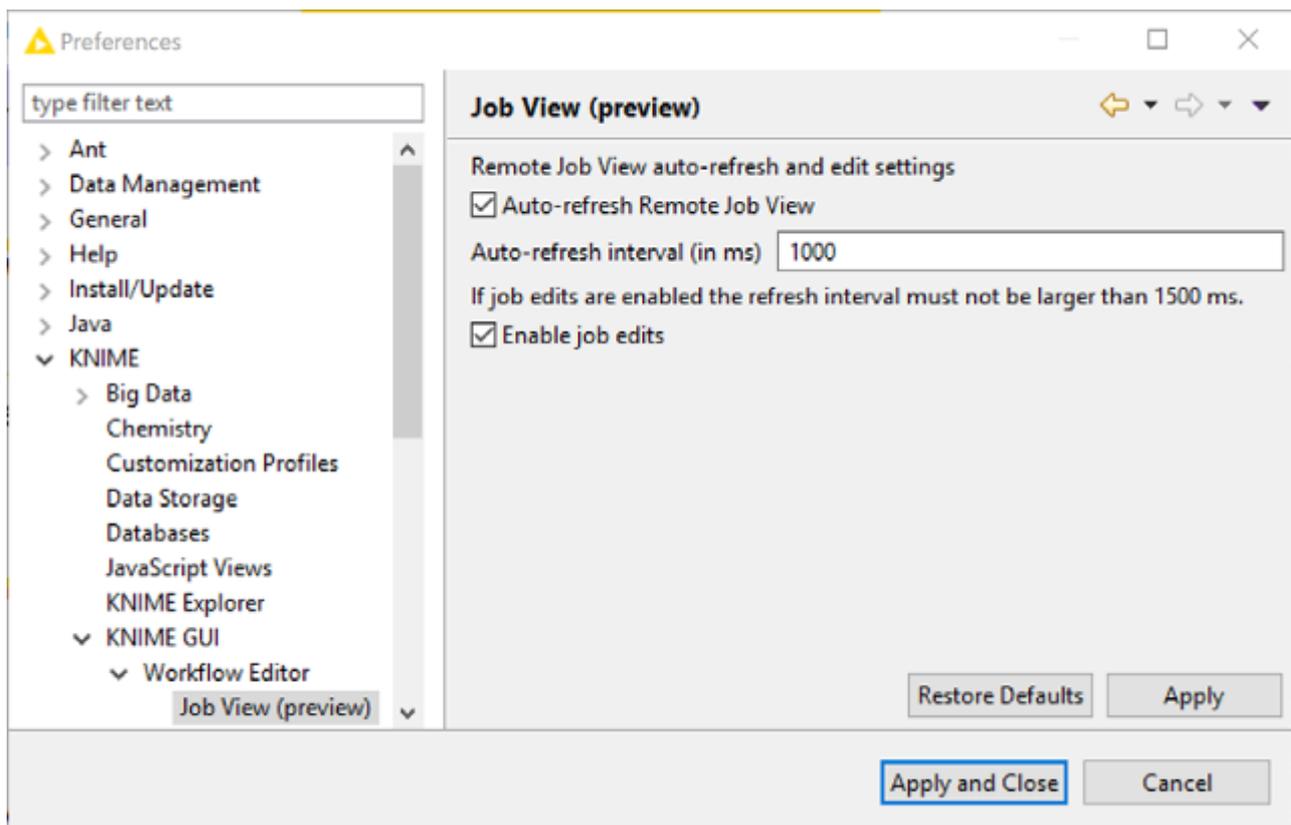
All jobs (executed from Analytics Platform, or via WebPortal) can be viewed, meaning that it's possible to see node execution progress, number of rows/columns generated, and any warning/error messages.

Jobs executed on the KNIME Server from the Analytics Platform via the 'Execute' option will in addition allow for user interaction by changing configuration parameters and resetting/executing individual nodes.

 A screenshot of the KNIME Job View interface. At the top, it says 'Welcome to KNIME Analytics Platform' and 'Job 2: Topic Detection Based on Movie Reviews'. Below that, a message states: 'This is a view on the remote job running on KNIME Server (Testserver-master). It refreshes every 1000 ms.' The main content is titled 'Topic Detection Analysis - Movie Reviews' and includes a description: 'Topic detection extracts relevant information elements from unstructured text documents and groups them to define a number of topics. This workflow illustrates how to perform a topic detection analysis on movie reviews.' The task is: 'Perform a topic detection in IMDb reviews.' Below this is a table with four columns: 'Data Reading', 'Pre-processing', 'Topic Detection', and 'Grouping'. Each column contains a description of the step. At the bottom, a workflow diagram shows four nodes: 'File Reader' (Reading), 'Document Creation and document pre-processing' (Transformation of strings to documents), 'Topic Extractor (Parallel LDA)' (4 words for, 61%), and 'GroupBy' (Conc terms for, queued).
 

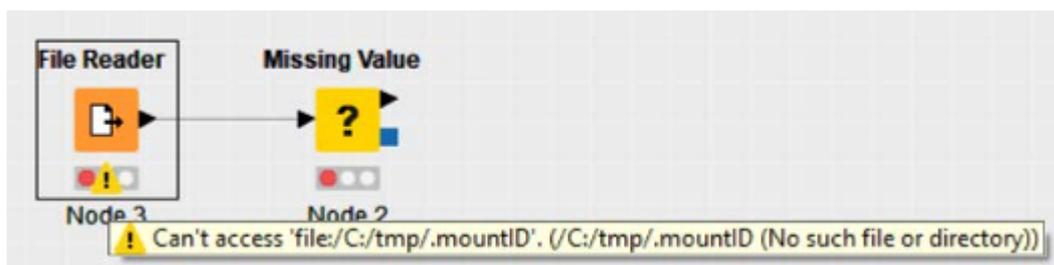
Data Reading	Pre-processing	Topic Detection	Grouping
Read IMDb reviews from a CSV file. The file is located in TheData/SocialMedia	- Classic pre-processing of documents: Punctuation Erasure, Number Filter, N Chars Filter, Stop Word Filter, Case Converter  Double-click the metanode to see the subworkflow	Build a list of topics of the pre-processed documents using the Topic-Extractor (Parallel LDA) node. Use 4 words for each topic and 8 topics.	The GroupBy node concatenates the keywords for the identified topics.

The KNIME Job View enables you to see which nodes are currently executing inside your running jobs. The view refreshes every 1 second.

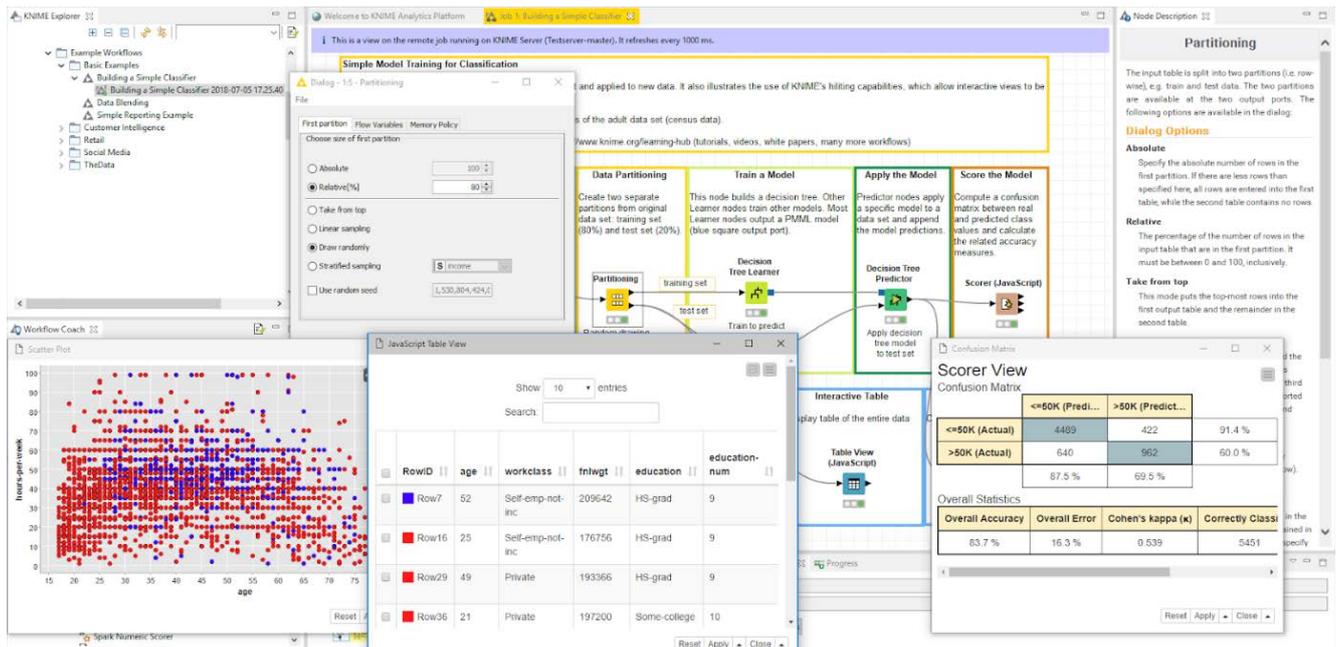


*Job View preferences allow to change the auto-refresh interval, and also optionally to uncheck the 'Enable job edits' to enforce view only mode for all workflows.*

You will be able to see which nodes are currently executing, which are already executed, and which are queued to be the next in execution. You can see errors and warning in the workflow by mouseover on the respective sign.



*Inside a failing job you can see the error and warning messages by mouse-over the respective sign*



*It's possible to view data and views by using the JavaScript Views.*

Some functionality that you have for local workflows is not available, for example adding nodes and connections to the workbench, or viewing all of the data in the node (without needing to use the Node Monitor).

With all future KNIME Releases we will continue adding functionality to the KNIME Job View. If you have any ideas or thoughts, we appreciate your input via [support@knime.com](mailto:support@knime.com).

KNIME AG  
Technoparkstrasse 1  
8005 Zurich, Switzerland  
[www.knime.com](http://www.knime.com)  
[info@knime.com](mailto:info@knime.com)