KNIME Server Release Notes and Update Guide

KNIME AG, Zurich, Switzerland

Version 4.13 (last updated on 2022-01-20)

Table of Contents

Release Notes

KNIME Server 4.13 is a feature release of the 4.x release line. All clients that have worked with KNIME Server 4.12 will continue to work with KNIME Server 4.13 without restrictions.

ℹ️ To find out which version of KNIME Server you are currently running, you can check the Administration pages on the WebPortal.

The following section is only relevant if you are using a dedicated RabbitMQ instance for communication between server and executor(s). In order to fix a bug where connections to RabbitMQ piled up when the latter was restarted some internal changes to RabbitMQ resources were necessary. This means that you have to stop RabbitMQ while upgrading from KNIME Server versions prior to 4.13.3. If you keep RabbitMQ running during the upgrade, the startup of both the server and the executor will fail.

New Features

For a list that includes the new Analytics Platform 4.4 features see here.

Highlighted new functionality is:

- Single Page Data Apps (what's new)
- Embedded Data Apps (what's new)
- Configurable number of retries on scheduled jobs (what's new)

A detailed changelog for KNIME Server 4.13 is also available.

Single Page Data Apps

KNIME Server 4.13 introduces a new feature that allows the user to build interactive, web-based applications that consist of a single page.

It is possible to re-execute parts of a Data App and therefore dynamically update parts of a page without having to leave the page with a simple "Refresh Button widget"-node.

Embedded Data Apps

Embedded Data Apps allow to access KNIME Data Apps on the WebPortal without the need to login explicitly. Therefore, it is now possible to fully embed a Data App into an existing website. Users are then able to run the workflow without providing any login credentials,

therefore providing a seamless experience.

Configurable number of retries

It is now possible to specify the number of time KNIME Server tries to re-execute a scheduled job in case the original execution run fails.

New configuration options

Various new configuration options were added to KNIME Server to improve job lifecycle management:

`com.knime.server.job.webportal.max_time_in_memory=<duration with unit, e.g. 30s, 1m>`: Allows setting the duration after which jobs started from KNIME WebPortal are swapped from memory to disk. Previously, WebPortal jobs were covered by the global setting `com.knime.server.job.max_time_in_memory`, which remains im place for all jobs not started from the WebPortal.

`com.knime.server.job.webportal.max_lifetime=<duration with unit, e.g. 30s, 1m>`: Allows setting the duration after which jobs started from KNIME WebPortal are discarded. Previously, WebPortal jobs were covered by the global setting `com.knime.server.job.max_lifetime`, which remains im place for all jobs not started from the WebPortal.

For both of the above options, the default is the same as for the existing global settings. It is recommended to only deviate from the global settings in case there are performance issues. This can for example be the case when using Embedded Data Apps, where the potentially large number of simultaneous job runs creates a need to free up memory faster.

Architecture changes

- TomEE: TomEE has been completely replaced by Tomcat. Starting with KNIME Server 4.13, TomEE based installers are not offered anymore. The same is true for upgrades. Therefore, if your current installation of KNIME Server is still based on TomEE, a fresh installation is necessary to move to KNIME Server 4.13.

- RMI: RMI connections between KNIME Server and KNIME Executors have been removed completely. It is now necessary to use either Qpid for single-executor setups or RabbitMQ in case of multiple KNIME Executors.

- Load throttling is now enabled also for single-Executor setups. This means that with

default settings the Executor will not load any new jobs if the system load is above 90%. If you want to raise the limit and continue loading jobs even if the system load is higher, you can adjust the Executor's `knime.ini`. See the section about Load throttling in KNIME Server Administration Guide for details.

Local file system access by KNIME workflows

A growing number of KNIME nodes are being revised to use a new shared framework for file access. More nodes will follow over the next releases, so that eventually the file access of all KNIME nodes uses the shared framework.

When executing on KNIME Server, a preference controls whether those nodes may access the local file system of the KNIME Server Executor or not. With the KNIME Server release 4.12, local file system access is **disallowed** by default (previously it was allowed).

To **allow** local file system access (not recommended), you can add the following line to the customization profile used by your KNIME Server Executor(s):

```
/instance/org.knime.filehandling.core/allow_local_fs_access_on_server=true
```

This preference affects all nodes that are part of the revised file handling framework. Old nodes that have not yet been ported to the new framework are not affected by this setting.

> ℹ️ Known issues: Under certain conditions, when the KNIME Server is configured for OIDC with LDAP group mapping, the server might refuse requests, if there are connection issues between the KNIME Server and LDAP. This is caused by a regression present in all currently available Tomcat releases. Once there is an updated Tomcat version that fixes this we will provide a fix as well. In case you are currently using OIDC with LDAP group mapping, please consider waiting before you upgrade your existing installation.

Introduction to the update process

This document guides you through the steps that are needed to upgrade an existing KNIME Server installation from version 4.12 to 4.13 or update by applying a bugfix to version 4.13.

You will find a complete guide to installing KNIME Server in the KNIME Server Installation Guide, and a complete description of all configuration options in the KNIME Server Administration Guide.

If you have any questions or need assistance with the update process, please contact your dedicated KNIME support specialist.

Since version 4.11, new installations of KNIME Server are based on Apache **Tomcat**, rather than Apache **TomEE** which was used in previous releases.

> ℹ️ With the release of KNIME Server 4.13.0, we discontinue support for TomEE based installations. It is not possible to upgrade a TomEE based KNIME Server to version 4.13.0 or higher. Instead, a fresh installation has to be performed.

If your current KNIME Server installation is already Tomcat based, you can follow the upgrade procedure as outlined below.

Prerequisites

Server

- You have already KNIME Server 4.12.x or KNIME Server 4.13.x running.
- A KNIME Server installation based on Tomcat application server

> ℹ️ For the sake of readability, this documentation uses `<apache-tomcat>/` to refer to the installation directory of the server application.

Client & Executor

- KNIME Analytics Platform clients with KNIME ServerSpace 4.8+ otherwise it will not be able to talk to the Server.
- KNIME Analytics Platform 4.4 as Executor with the corresponding *KNIME Executor Connector* extension version 4.13 installed on the Server side.

Backup

Even though all the data on the Server should be preserved during the update it is highly recommended to create a backup of all important data (see KNIME Server Administration Guide for backup information).

> ℹ️ With KNIME Server 4.12 the email configuration moved from the knime.xml to the `<knime-server-repository>/config/knime-server.config`, so any existing custom configuration must be moved. More information concerning the email configuration can be found in the KNIME Server Administration Guide

Upgrade/update steps

Application server upgrade/update

Start by downloading the following files and make sure that the owner of these files is your KNIME user:

- The `<knime-tomcat>.jar` file listed here

- If you are updating using Tomcat: the latest Tomcat `.war` file listed here

- The latest KNIME Server Executor file listed here

Next backup these files before starting the update process:

- The `knime.xml` file located in `<apache-tomcat>/conf/Catalina/localhost/` Please be aware that the name of the file might be different if the `.war` file had a different name (e.g. `com.knime.enterprise.server.xml` or `webportal.xml`).

  > ❗ Failure to backup this file will result in losing access to your scheduled jobs since it contains the secret key that is used to encrypt the schedules.

- The `server.xml` file located in `<apache-tomcat>/conf/`

ⓘ When upgrading a KNIME Server installation that uses RabbitMQ as message broker to **KNIME Server 4.13.3**, please make sure to shutdown RabbitMQ before starting with the upgrade. Please only start RabbitMQ after upgrading KNIME Server. This is due to some changes to how connections to RabbitMQ are handled. In case you are using a serverless message broker on e.g. AWS that can't be restarted, please delete all exchanges before proceeding with the update.

Then follow these steps:

1. Stop KNIME Server.

   a. On Windows: If you have KNIME Server installed as a service, stop KNIME Server and Executor.

      If you don't have KNIME Server installed as a service, then run the `shutdown.bat` file located in `<apache-tomcat>/bin`

   b. On Linux: If you have KNIME Server installed as a service, stop KNIME Server by running the `systemctl stop knime-server.service` command

      If you don't have KNIME Server installed as a service, then run the `shutdown.sh` file located in `<apache-tomcat>/bin`

2. Delete the existing folder `<apache-tomcat>/webapps/knime`. You may not have a folder with this name located in the `webapps` folder. If this is the case, delete the folder that matches the file name of the `.war` file located in the `webapps` folder. E.g. if you see a `webportal.war` file in the `webapps` folder, delete the `webportal` folder.

3. Move the `<filename>.war` file located in the `webapps` folder to a different folder. Then copy the new downloaded `war-4.13.x.y.war` file to the `webapps` folder and rename it to `<filename>.war`.

4. Check if the `knime.xml` file that you backed up before still exists. If not copy the backup to the previous location.

5. Copy the downloaded file `<knime-tomcat>.jar` to `<apache-tomcat>/lib`. Delete the existing old `<knime-tomcat>.jar` file that was already in that directory. Make sure to rename the new file to knime-tomcat.jar.

KNIME Executor upgrade

KNIME Server 4.13 requires KNIME Analytics Platform 4.4 or later.

Note that **upgrading** the KNIME Executor using the **standard procedures** is possible again from KNIME Analytics Platform 4.4.1 onwards. To do so please follow the instructions in the KNIME Executor upgrade from version 4.4.0 section.

Due to some larger changes instead (new Eclipse version and move to Java 11), it is not possible to upgrade to KNIME Executor 4.4.x from a version older than 4.4.0. Therefore, you need to install KNIME Executor 4.4.x version from scratch. To do so please follow the instructions in the KNIME Executor installation section.

KNIME Executor upgrade from version 4.4.0

If possible, start the executor installation in graphical mode (using the correct installation user!). If you have internet access, go to *File → Update KNIME…* and it will suggest updating to the new version. Follow the steps in the wizard. If you don't have direct internet access, you must download the zipped update sites from the commercial downloads page. Then register the ZIP files in *File → Preferences → Install/Update → Available Software Sites* and go to *File → Update KNIME…*

Under Linux, if you cannot start the graphical user interface, you can use the update-executor.sh script that is in the root of the executor installation.

Call the script on the command line and provide a list of update sites that contain the new versions of the installed extensions and all installed extension will be updated (given that an update is available):

```
./update-executor.sh https://update.knime.com/analytics-platform/4.4
```

If you get error messages when executing update-executor.sh about missing "installable units", make sure that you have provided all necessary update sites.

If you want to selectively update only certain extensions, you have to build the update command yourself. An update is performed by uninstalling (-u) and installing (-i) an extension at the same time:

```
./knime -application org.eclipse.equinox.p2.director -nosplash -consolelog -r <list-of-
update-sites> -i
<list-of-features> -u <list-of-features> -d <knime-installation-folder>
```

KNIME Executor installation

Before starting the KNIME Executor installation process create a backup of the <executor-

`folder>` renaming it to `<executor-folder>.old`.

Then proceed with the installation of KNIME Executor by downloading the KNIME Executor full build from here and extracting it. This installer includes all extensions required for running as an Executor for KNIME Server. After you successfully install the desired version of KNIME Executor, rename the new folder containing the Executor installation to the old folder name.

Alternatively you can also install KNIME Analytics Platform from scratch and install the `KNIME Executor connector` extension version 4.13 into it. In this case, you will need to run the newly installed KNIME Analytics Platform 4.4 instance as the user that owns the installation.

Please refer to the KNIME Server Installation Guide for more details about installing a KNIME Executor.

knime.ini file

In case you want to use custom settings from your old KNIME Executor's `knime.ini` file, please copy only the lines following `-vmargs` to the new `knime.ini`. The only exception to this are the four lines related to Server Managed Customizations (`-profileLocation…`), in case those are set. The remaining lines above `-vmargs` contain paths to resources that have been changed in the new release, so they are not compatible.

Also note that it is now recommended to move the lines related to Server Managed Customizations to Executor's service definition. This will help upgrading to future releases on installation where no GUI is available.

server.xml file

For fresh installations, please note that a `server.xml` file (`<apache-tomcat>/conf/server.xml`) from an older installation cannot be copied over due to some larger changes. Any custom changes applied to your existing `server.xml` need to be manually copied over to the new `server.xml`. For upgrades the existing `server.xml` continues to work, and does not need any adjustments.

Configuration of the updated KNIME Server

KNIME Server license

The KNIME Server license continues working with the new Server. If you haven't received a

license file or if it is not working correctly, please contact KNIME support or your dedicated KNIME account manager.

Enabling execution via Apache Qpid

KNIME Server 4.11 introduced an integrated message broker for communication between KNIME Server and KNIME Executor, based on Apache Qpid (https://qpid.apache.org/). This can be thought of as using an architecture that corresponds closely to distributed KNIME Executors, but runs on a single host, which is used by both KNIME Server and KNIME Executor.

Unlike distributed KNIME Executors, it is not necessary to install a separate message broker such as RabbitMQ. Instead, Apache Qpid is bundled as part of all KNIME Server installations. New installations of KNIME Server use Qpid by default. Older installations might still use RMI, which has been removed with KNIME Server 4.13.

If you are still using RMI, you need to switch your installation to Qpid. Changes are required in two locations:

Application Server

Activate queue in `knime-server.config`. The file can be found in `<knime-server-repository>/config/knime-server.config`. Set the parameter

`com.knime.enterprise.executor.embedded-broker=` to `true`

In addition, comment out the line `com.knime.server.executor.knime_exe=` by placing a # character at the start of the line.

KNIME Executor

In the KNIME Executor installation directory, add the following line to the `knime.ini` file, anywhere after the `-vmargs` line:

`-Dcom.knime.enterprise.executor.msgq=amqp://knime:20knime16@localhost/`

Alternatively, you can also add this line to the Executor's service definition.

> ℹ️ Unlike the RMI-based execution, using the Qpid message broker requires to startup KNIME Server and KNIME Executor **separately**. For starting up KNIME Server, you can still use the startup.sh/.bat file located in `<apache-tomcat>/bin`. For starting the Executor, you can use `<executor directory>/start-executor.sh` (Linux) or `<executor directory>/start-executor.bat` (Windows)

In case the Executor startup script is not present, you can start the Executor from command line by running

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

For setting up the KNIME Server service, please follow the steps outlined in the KNIME Server Installation Guide. Installing the Executor as a service follows the same procedure as described for distributed KNIME Executors.

For increased security, we recommend to run both services by different users.

Preferences file

Note that when switching from RMI to Qpid, the old `preferences.epf` file in `./workflow_repository/config` is **not** used anymore. In order to set preferences on the Executor (e.g. database drivers or Python configuration), please use a preference profile as described in the KNIME Server Admin Guide.

Server temp directory

Switching to Qpid also uses a different temp directory than RMI. On Linux, this defaults to `/tmp`, which might have size restrictions. In order to manually set a different temp directory, please add the following to the knime.ini file of the KNIME Executor, anywhere below the `-vmargs` line: `-Djava.io.tmpdir=/path/to/tmpdir`

Qpid port

In some environments, it is necessary to explicitly open a port for connections to Qpid **on localhost** (even though Server and Executor are running on the same host). By default, the port is 5672. This can be changed in the `knime-server.config` file.

Server host name in /hosts

By default, the host name of the machine is available in `/etc/hosts/`. If this is not the case, KNIME Server will not work using Qpid. This can be resolved by adding the server's host name to `/etc/hosts/`, e.g. `127.0.0.1 <hostname>`.

Compressed responses

Due to a change in Tomcat responses sent by KNIME Server 4.11 and 4.12 are not compressed even though compression is configured in the `server.xml`. In order to enable compression again, you have to add the attribute `noCompressionStrongETag=false` to all `<Connector>` elements in your `server.xml`, e.g.

```
...
<Connector noCompressionStrongETag="false" compressibleMimeType="..." compression="on"/>
...
```

If you create a new server installation using the KNIME Server Installer for 4.13.0 this attribute will be added automatically.

Troubleshooting

All schedules are lost after updating

One important step during update process is the backup of the `knime.xml` file so that it can be copied back into the KNIME Server installation folder. The `knime.xml` file contains a secret key for encrypting the communication between the server and the installation folder. If this is not restored before starting the server, the server will not be able to load jobs and schedules previously created.

Access to KNIME WebPortal version > 4.11 is not possible

Access gives error 404

If you cannot access the URL of the KNIME WebPortal version > 4.11
`https://<server-address>/knime/webportal/`
and it fails with error 404, usually indicates that during update or installation process the Context Root used to define the KNIME WebPortal URL was left empty (/). This is not possible starting with the release of KNIME WebPortal version 4.11+ for which it is necessary to **always** specify a Context Root.

Please refer to the KNIME Server Installation Guide.

Access gives error 500

If you cannot access the URL of the KNIME WebPortal version > 4.11
`https://<server-address>/knime/webportal/`
and it fails with error 500, usually indicates that the update did not work as expected.

Please contact KNIME support team via support@knime.com to get this fixed.

Update of KNIME Executor via command line is not possible

This is a known bug when updating from a KNIME Server version 4.11.1. This bug is fixed starting with KNIME Server version 4.11.2.

To solve the issue you need to follow these steps:

1. Create a copy of the file named `knime.ini-copy`

2. Remove the `-profileLocation` option from the `knime.ini` file

3. Update the KNIME Executor via command line

4. After a successful update, overwrite your `knime.ini` file with `knime.ini-copy`.

If this solution does not work, please contact KNIME support team via support@knime.com

Workflows deployed on the server cannot be executed

- If the following error message appears:

```
Unable to load workflow, it was created with a future version of KNIME
```

The KNIME Analytics Platform client with which you are connecting to the KNIME Server and with which the server was built needs to have a lower or equal version to the KNIME Executor. If this is not the case, please ask your KNIME Server Administrator to update the KNIME Server and/or KNIME Executor to the release version of the KNIME Analytics Platform client you are using locally.

- If a missing node is reported, instead, check if the node is already installed and if the version of the respective extension is the same in both the KNIME Analytics Platform client and the KNIME Executor.

- Finally, if your scripting (R or Python) snippets fail, it is possible that the respective packages are missing on KNIME Server. Please refer to the respective guides to install Python and R packages. Please also make sure the user under which you install those packages is the same user running the KNIME Executor.

Connection to KNIME Executor hosts after switching to Qpid is not possible

- Make sure the hostname of the KNIME Server is mentioned in the `/etc/hosts` file

- Check that the following ports are accessible on `localhost`. They do not need (and in fact should not) to be accessible from other systems (usually this is the case).

  - With QPid: `com.knime.enterprise.executor.embedded-broker.port` (default 5672)

  - With RabbitMQ: For KNIME Server and for the KNIME Executor that is setup on a separate machine (default 5672)

- Starting with QPid, KNIME Executor needs to be started manually. Please follow the instructions in the KNIME Server Installation Guide to do so.

- It might be necessary to change the user running the service. Please follow the instructions in the KNIME Server Installation Guide to do so.

Reading from my local file system after update to KNIME Server 4.11+ is not possible

Starting with KNIME Analytics Platform version 4.2, which is the KNIME Executor version for the KNIME Server 4.11 we are, by default, not allowing access to the local file system on the KNIME Server via the new File Util nodes. You can enable this again as described in KNIME Server Admin Guide.

Please note that this only affects nodes that are part of the revised file handling framework.

The authentication field is empty

Please update the `knime-tomcat.jar` as described in the Application server update section. This must be done before restarting the KNIME Server.

Changelog (KNIME Server 4.13)

KNIME Server 4.13.6 (released November 23, 2022)

Enhancements

- [SRV-3662] - Allow clock skew while parsing access and refresh tokens
- [SRV-3663] - Update Tomcat to 9.0.58
- [SRV-3673] - Add setting to disable 'Use KNIME Server login'

Bugfixes

- [SRV-3656] - KNIME Server only allows one workflow download at a time
- [SRV-3672] - NullPointerException when no metadata is found for snapshot
- [SRV-3675] - Fix display of jobs with colons in their name \(Windows only\)
- [SRV-3757] - CVE-2022-44748: Zip Slip vulnerability when uploading workflows

KNIME Server 4.13.5 (released January 19, 2022)

Enhancements

- [SRV-3626] - Call workflow outside of workflow instance lock
- [SRV-3399] - Add mason control for execution on an executed job with no input parameters

Bugfixes

- [SRV-3655] - Configuration dialog doesn't ignore unknown fields for representations
- [SRV-3653] - Potential deadlock when a workflow with schedules with execution retries is overwritten
- [SRV-3651] - Job is loaded even when no JobMessageDispatcher could be started
- [SRV-3649] - Typo in log message
- [SRV-3642] - Recycle Bin items cannot be deleted
- [SRV-3641] - Executor cannot acquire core tokens any more after server restart (Qpid-only)

- [SRV-3631] - File extension of input resource is not propagated to executor if it's send via message broker

- [SRV-3628] - NullPointerException in gateway client due to missing parent workflow id

- [SRV-3475] - No error message sent if an error occurs upon job loading

KNIME Server 4.13.4 (released November 8, 2021)

Enhancements

- [SRV-3615] - Update to Tomcat 9.0.54

- [SRV-3621] - Update CXF to 3.4.5

- [SRV-3616] - Automatically forward webserver root to context root in new installations

- [SRV-3613] - Server installer should restrict permission on auto-install.xml

Bugfixes

- [SRV-3612] - Directory path traversal when requesting client profiles[1]

- [WEBP-876] - Cross-Site-Scripting vulnerability in old WebPortal login[2]

KNIME Server 4.13.3 (released October 20, 2021)

Bugfixes

- [SRV-3606] - Executor doesn't send current job state upon reconnecting to message broker

- [SRV-3597] - Race condition between job loading and status update

- [SRV-3595] - Loaded jobs appear to be swapped when RabbitMQ is restarted

- [SRV-3594] - Timeout while restoring from swap may result in orphaned jobs in executor

- [SRV-3593] - Server shutdown is delayed if RabbitMQ is not reachable

- [SRV-3590] - Executor sometimes cannot send messages after RabbitMQ restart due to missing exchanges

- [SRV-3589] - KNIME Server fails during startup and does not try to reconnect to RMQ if initial connection to RMQ cannot be established

- [SRV-3588] - Each RabbitMQ restart adds 2 additional connections

- [SRV-3565] - Editing a schedule on the 31st may display wrong start/end month

KNIME Server 4.13.2 (released September 20, 2021)

Enhancements

- [SRV-3562] - Update Server Example Workflows
- [SRV-3549] - Workaround for time drifts on Windows to avoid double execution of schedules
- [SRV-3506] - Strip leading and trailing spaces from server configuration values

Bugfixes

- [SRV-3575] - Admin users logging in to the webportal with upper-case characters don't have admin access
- [SRV-3565] - Editing a schedule on the 31st may display wrong start/end month
- [SRV-3554] - Configuration page in WebPortal adds new spaces to msgq rules and names on every change
- [SRV-3452] - Jobs are not always removed from executor upon a load timeout

KNIME Server 4.13.1 (released August 25, 2021)

Enhancements

- [SRV-3533] - Update Tomcat to 9.0.52
- [WEBP-840] - Don't show passwords in cleartext in WebPortal
- [WEBP-146] - Include server version number on license page

Bugfixes

- [SRV-3560] - Potential deadlock when obtaining job status updates
- [SRV-3545] - Race condition in executor when sending executor status
- [SRV-3536] - Modifying scheduled job while its timer expired creates duplicate schedules until server restart
- [SRV-3535] - Workflow Authentication tokens don't check group permissions

- [SRV-3534] - OAuth token refresh fails with some OIDC providers due to duplicate secret
- [SRV-3530] - Executor endpoint not thread-safe
- [SRV-3528] - Changed properties of loaded job get lost after server crash
- [SRV-3524] - Inactive job not removed if post execution tasks didn't run
- [SRV-3522] - Time intervals of scheduled don't behave correctly in some cases
- [SRV-3520] - Impatient clicking on "Double-click to connect to server" freezes system
- [SRV-3519] - Scheduled job gets discarded even though there might be jobs with potential retries left
- [SRV-3510] - "Skip if previous job is still running" skips jobs if server and executor crash during execution of last job
- [SRV-3499] - Add authorization header to request to OIDC token endpoint
- [SRV-3483] - startedExecutionAt is a different timezone than finishedExecutionAt/createdAt in REST jobs endpoint
- [SRV-3333] - Disable workflow execution in installer doesn't disable Qpid
- [WEBP-841] - Workflows with execute-permission only can't be re-executed (Data Apps)
- [WEBP-834] - Refreshing page on nested folder replaces slashes in URL
- [WEBP-827] - Webportal deletes random schedule/jobs on second or subsequent pages
- [WEBP-775] - Discrepancy between Started at and Created/Finished at in job View
- [WEBP-747] - Allow line breaks in workflow description

KNIME Server 4.13.0 (released June 30, 2021)

Enhancements

- [SRV-3492] - Add parameters to the useradd command to reduce the overall setup steps required
- [SRV-3479] - Check Client AP Version against Server version when connecting to a server
- [SRV-3472] - Provide executor and AP version with the server information
- [SRV-3470] - Update to CXF 3.4.3
- [SRV-3469] - Update to Tomcat 9.0.45
- [SRV-3461] - Update embedded WebApps documentation

- [SRV-3457] - Change preferred way of providing profiles to executor

- [SRV-3447] - Add option to re-run scheduled jobs in case of failed execution

- [SRV-3446] - Add WebPortal-specific job discard setting

- [SRV-3445] - Add WebPortal-specific job swap setting

- [SRV-3427] - Make executor try reconnecting to server indefinitely the default instead of 9 times

- [SRV-3425] - Make repository path configuration read-only

- [SRV-3417] - Handle multiple incoming workflow tokens

- [SRV-3405] - Enable renaming of access tokens after first creation

- [SRV-3404] - Enable access to workflow access tokens after creation

- [SRV-3403] - Support AWS IMDSv2 for license checks

- [SRV-3395] - Support all flow variables types in remote workflow editor

- [SRV-3394] - Remove EJB support from server

- [SRV-3380] - Remove support for RMI executors

- [SRV-3369] - Enable two way SSL for RabbitMQ connections

- [SRV-3341] - Expose file size, creation and modify date for KNIME Server files

- [SRV-3332] - Allow scheduled jobs to inherit permissions from workflow

- [SRV-3327] - Support KNIME Edge compatible job result format in REST interface

- [SRV-3205] - Add a health check to the executor

- [WEBP-792] - Move access token administration to monitoring

- [WEBP-731] - Add license type or feature detection to settings.json

- [WEBP-680] - Pagebuilder re-execution UI implementation

- [WEBP-678] - WebPortal: re-execution API implementation

- [WEBP-675] - Extend WebPortal-Gateway API for partial re-execution of a page

- [WEBP-658] - Access tokens admin page

- [WEBP-657] - Share link with sidebar to create embedded links

- [WEBP-656] - New query parameter 'embed'

Bugfixes

- [SRV-3509] - Workflow Authentication tokens allowed access to other calls

- [SRV-3504] - Workflow authentication token management uses case-sensitive usernames

- [SRV-3501] - Admin username not case insensitive

- [SRV-3500] - Race condition in ObjectStore

- [SRV-3494] - update-executor.sh script uses wrong Java installation

- [SRV-3493] - Refreshing the access token yields server logout

- [SRV-3491] - Unclear reference to unknown something else in Java installation instruction

- [SRV-3488] - Response compression does not work any more due to change in Tomcat

- [SRV-3487] - Explorer sets connection timeout instead of read timeout when fetch timeout is provided

- [SRV-3485] - Hourly job check-in of server may lead to a racing condition in the executor job reporting

- [SRV-3480] - Remote workflow editor: problem de-/serializing large 'string' blob cells

- [SRV-3477] - Job with state EXECUTION_FAILED_WITH_CONTENT is not marked as finished

- [SRV-3474] - HTTP status code for several modification operations are not correct

- [SRV-3471] - Executor may reject jobs although sufficient resources are available

- [SRV-3466] - Report formats cannot be updated via Admin Portal

- [SRV-3459] - Jobs are stuck in idle during execution

- [SRV-3458] - SocketTimeoutException and Read time out when authenticating KAP with KS via LDAP or AD

- [SRV-3452] - Jobs are not always removed from executor upon a load timeout

- [SRV-3448] - Scheduled job id changes when overwriting workflow

- [SRV-3442] - Load workflow message isn't seen by every executor

- [SRV-3434] - Large reports may break server-executor communication

- [SRV-3432] - Missing descriptions in KNIME Server OpenAPI documentation

- [SRV-3397] - typo in Logging for Lifecycle Handler

- [SRV-3306] - Deleting a repository item should only require write permissions on parent not on the item itself

- [WEBP-822] - Missing SinglePageExecution causes page updates to be missed

- [WEBP-766] - Improve error handling in embedded mode

- [WEBP-757] - WebPortal backend sometimes returns 'configured' as wizard execution state

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