

Kerberos Administration Guide

KNIME AG, Zurich, Switzerland
Version 4.16 (last updated on 2022-07-26)



Table of Contents

- Overview 1
- Prerequisites..... 2
- Setting up krb5.conf 3
 - Creating krb5.conf file 3
- KNIME preferences 4
 - For KNIME Server executors 6
 - For KNIME Analytics Platform clients..... 6
- Troubleshooting 9
 - Kerberos Debug..... 9

Overview

KNIME Server executes workflows, that may try to access Kerberos-secured services such as Apache Hive™, Apache Impala™ and Apache Hadoop® HDFS™.

This guide describes how to configure KNIME Server so that it can authenticate itself against Kerberos.

To configure Kerberos on KNIME Analytics Platform, please refer to the [Kerberos User Guide](#).

Prerequisites

Setting up KNIME Server for Kerberos authentication has the following prerequisites:

- For Kerberos:
 - An existing Kerberos KDC such as MIT Kerberos or Microsoft ActiveDirectory
 - A service principal for KNIME Server. The recommended format is `knimeserver/<host>@<REALM>`, where:
 - `<host>` is the fully-qualified domain name of the machine where KNIME Server runs,
 - `<REALM>` is the Kerberos realm.
 - A keytab file for the KNIME Server service principal.
 - A Kerberos client configuration file (`krb5.conf`). Alternatively, can be created manually (see section [Setting up krb5.conf](#)).
- For KNIME Server:
 - An existing KNIME Server installation.
 - An account with administrative privileges on the machine where KNIME Server is installed. This accounts needs to be able to edit the KNIME Server configuration files and restart KNIME Server.

Setting up krb5.conf

The KNIME Server Executor might need to read the `krb5.conf` file during Kerberos authentication. In that case, a valid `krb5.conf` file needs to be obtained. In case the location of the file is unknown or the file is not available, please contact the local administrator.

Creating krb5.conf file

Alternatively, a `krb5.conf` file can be created manually. A minimal configuration file could look like this:

```
[libdefaults]
default_realm = MYCOMPANY.COM
forwardable = true

[realms]
MYCOMPANY.COM = {
    kdc = kdc.mycompany.com
    admin_server = kdc.mycompany.com
}
```

The above example declares that the Kerberos realm is called `MYCOMPANY.COM` and that the hostname of the Kerberos KDC is `kdc.mycompany.com`.

The `forwardable` flag is a prerequisite for some KNIME nodes to make use of [Kerberos constrained delegation](#) when executed on KNIME Server.

Adjust the values contained in the `krb5.conf` file as appropriate for the setup in use. Depending on the specific setup, more configuration settings may be necessary. The `krb5.conf` format is fully described as part of the [MIT Kerberos documentation](#).

KNIME preferences

KNIME Server allows to distribute **customization profiles**, which can be used to automatically distribute Kerberos preferences (including the `krb5.conf`) to all connected KNIME Server executors and KNIME Analytics Platform clients.



It is recommended to create separate customization profiles for KNIME Server executors and for KNIME Analytics Platform clients.

The Kerberos configuration preferences are stored in a preferences file (file name ends with `.epf`). The table below contains all supported Kerberos configuration options.

`/instance/org.knime.kerberos/org.knime.kerberos.conf=<VALUE>`

Specifies the Kerberos configuration options. Replace `<VALUE>` with:

- `FILE`: to use Kerberos client configuration file (`krb5.conf`).
- `DEFAULT`: to use system defaults (discouraged).
- `REALM_KDC`: to provide realm and KDC directly in the preferences file.

`/instance/org.knime.kerberos/org.knime.kerberos.conf.file=<PATH>`

Specifies the location to `krb5.conf` file. Replace `<PATH>` with the path to `krb5.conf`.

This configuration only applies if `FILE` is selected in the option above.

`/instance/org.knime.kerberos/org.knime.kerberos.kdc=<KDC_VALUE>`

Specifies KDC value. Replace `<KDC_VALUE>` with the IP or hostname of the KDC.

This configuration only applies if `REALM_KDC` is selected in the first option listed above.

`/instance/org.knime.kerberos/org.knime.kerberos.realm=<REALM_VALUE>`

Specifies Realm value. Replace `<REALM_VALUE>` with the name of the realm (the name needs to be in uppercase letters).

This configuration only applies if `REALM_KDC` is selected in the first option listed above.

```
/instance/org.knime.kerberos/org.knime.kerberos.authMethod=<VALUE>
```

Specifies the Kerberos authentication method. Replace <VALUE> with:

- KEYTAB: to use keytab and service principal. It is recommended for KNIME Server executors.
- USER_PWD: to use username and password. It is recommended for KNIME Analytics Platform clients.

```
/instance/org.knime.kerberos/org.knime.kerberos.keytabFile=<PATH_TO_KEYTAB>
```

Specifies the location to the keytab file.

This configuration only applies if KEYTAB is selected as authentication method. Keytab is recommended as the authentication method for KNIME Server executors. In this case, the keytab must not be stored in the profile folder, but needs to be present on the KNIME Server executor machine(s) so that the preferences can reference it by local path.

```
/instance/org.knime.kerberos/org.knime.kerberos.keytabPrincipal=<PRINCIPAL_VALUE>
```

Specifies the keytab service principal value.

This configuration only applies if KEYTAB is selected as authentication method.

```
/instance/org.knime.kerberos/org.knime.kerberos.showIcon=<true|false>
```

Specifies whether to show Kerberos login status bar in the lower part of KNIME Analytics Platform.

This configuration is available only for KNIME Analytics Platform clients.

```
/instance/org.knime.kerberos/org.knime.kerberos.debug=<true|false>
```

Specifies whether to enable Kerberos debug.

-

```
/instance/org.knime.kerberos/org.knime.kerberos.debugLogLevel=<LOG_LEVEL>
```

Specifies the log level if Kerberos debug is enabled.

Replace <LOG_LEVEL> with either WARN, or ERROR.

For KNIME Server executors

This section contains a step-by-step guide to create a customization profile to distribute Kerberos preferences to all KNIME Server executor(s):

1. Create a profile folder inside `<knime-server-repository>/config/client-profiles`. The name of the folder corresponds to the name of the profile. The folder should contain the preferences file and other files, such as `krb5.conf`, to be distributed.
2. Inside the profile folder, create a preferences file (file name ends with `.epf`). The content of this file depends on the Kerberos configuration and authentication method to be used. For example, a recommended Kerberos configurations for KNIME Server executors could look like the following:

```
/instance/org.knime.kerberos/org.knime.kerberos.conf=FILE
/instance/org.knime.kerberos/org.knime.kerberos.conf.file=${profile:location}/krb5.conf
/instance/org.knime.kerberos/org.knime.kerberos.authMethod=KEYTAB
/instance/org.knime.kerberos/org.knime.kerberos.keytabFile=<PATH_TO_KEYTAB>
/instance/org.knime.kerberos/org.knime.kerberos.keytabPrincipal=<PRINCIPAL_VALUE>
```



If `krb5.conf` is used, copy the `krb5.conf` file into the profile folder so that it will be distributed to all KNIME Server executors along with the preferences file.



Replace `<PATH_TO_KEYTAB>` with the path to the keytab file and `<PRINCIPAL_VALUE>` with the service principal. The keytab must not be stored in the profile folder, but needs to be present on the KNIME Server executor machine(s) so that the preferences can reference it by local path.



Please check the [table](#) above for more information on all supported Kerberos configurations.

3. KNIME Server executors need to be made aware of a customization profile so that they can request it from KNIME Server. Please consult the [respective section](#) of the [Server Admin Guide](#) for a complete reference on how to set this up.

For KNIME Analytics Platform clients

The following is the step-by-step guide to create a customization profile to distribute Kerberos preferences to all KNIME Analytics Platform clients.

1. Create a profile folder inside `<knime-server-repository>/config/client-profiles`. The name of the folder corresponds to the name of the profile. The folder should contain the preferences file and other files, such as `krb5.conf`, to be distributed.
2. Inside the profile folder, create a preferences file (file name ends with `.epf`). The content of this file depends on the Kerberos configuration and authentication method to be used. For example, a recommended Kerberos configurations for KNIME Analytics Platform clients could look like the following:

```
/instance/org.knime.kerberos/org.knime.kerberos.conf=FILE
/instance/org.knime.kerberos/org.knime.kerberos.conf.file=${profile:location}/krb5.conf
/instance/org.knime.kerberos/org.knime.kerberos.authMethod=USER_PWD
/instance/org.knime.kerberos/org.knime.kerberos.showIcon=true
```



If `krb5.conf` is used, copy the `krb5.conf` file into the profile folder so that it will be distributed to all KNIME Analytics Platform clients along with the preferences file.

Instead of writing the preferences lines manually, another possibility is to configure Kerberos preferences graphically in KNIME Analytics Platform, then export the preferences and copy the relevant lines to the profile file. Please refer to the section [Export preferences from KNIME Analytics Platform](#) for more information.

3. KNIME Analytics Platform clients need to be made aware of a customization profile so that they can request it from KNIME Server. Please consult the [respective section](#) of the [Server Admin Guide](#) for a complete reference on how to set this up. In KNIME Analytics Platform, going to *File* → *Preferences* → *KNIME* → *Customization Profiles*, opens the *Customization Profiles* page where the KNIME Server and profile to be used can be chosen. The changes will take effect after restarting KNIME Analytics Platform.

Export preferences from KNIME Analytics Platform

This section describes an alternative way to set up Kerberos preferences graphically, instead of writing the preferences lines manually (see step 2 in previous [section](#)). The Kerberos preferences can be configured graphically in the preferences page in KNIME Analytics Platform. After that the preferences can be exported and the relevant lines copied to the profile file.

1. Start KNIME Analytics Platform.
2. Set all Kerberos preferences via *File* → *Preferences* → *KNIME* → *Kerberos* and export the preferences via *File* → *Export Preferences*. Please refer to [Kerberos User Guide](#) for

more information on the Kerberos preferences page.

3. Open the exported preferences file and copy all the lines starting with `/instance/org.knime.kerberos/` into the profile file. For more information on preferences file, please check the [Preferences file](#) section of the [Server Admin Guide](#).
4. Please make sure that any paths set in the preferences are valid on the server:
 - If `krb5.conf` is used, the `krb5.conf` file needs to be copied into the profile folder so that it will be distributed to all KNIME Server executors along with the preferences file. After that, change the path to `krb5.conf` as following:

```
/instance/org.knime.kerberos/org.knime.kerberos.conf.file=${profile:location}
/krb5.conf
```

- If `keytab` is used, please make sure the path is accessible by KNIME Server.

Troubleshooting

Kerberos Debug

If Kerberos authentication fails, activating Kerberos debug logging may provide insight into why this is happening. To activate Kerberos debug logging, add the following line to the preferences file (please check [previous section](#) for more information about preferences files).

```
/instance/org.knime.kerberos/org.knime.kerberos.debug=true
```

The debug log level is set to INFO by default, but it can be changed by adding the following line:

```
/instance/org.knime.kerberos/org.knime.kerberos.debugLogLevel=<LOG_LEVEL>
```

where <LOG_LEVEL> can be replaced by either WARN, or ERROR.



Please check the [table](#) containing all supported Kerberos configuration options for more information.

Then, restart the KNIME Server Executor and run a workflow that accesses a Kerberos-secured service. The knime.log will then contain Kerberos debug messages. The knime.log can be found on the KNIME Server machine under:

```
<executor-workspace>/metadata/knime/knime.log
```

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