

Kerberos Administration Guide

KNIME AG, Zurich, Switzerland
Version 1.7 (last updated on 2023-09-21)




Table of Contents

- Overview 1
- Prerequisites..... 2
- Setting up krb5.conf 3
 - Creating krb5.conf file 3
- KNIME preferences 4
 - For KNIME Business Hub executors 6
 - For KNIME Analytics Platform clients..... 6

Overview

KNIME Business Hub 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 Business Hub 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 Business Hub 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 Business Hub. The recommended format is `knimehub/<host>@<REALM>`, where:
 - `<host>` is the fully-qualified domain name of the machine where KNIME Business Hub runs,
 - `<REALM>` is the Kerberos realm.
 - A keytab file for the KNIME Business Hub service principal.
 - A Kerberos client configuration file (`krb5.conf`). Alternatively, can be created manually (see section [Setting up krb5.conf](#)).
- For KNIME Business Hub:
 - An existing KNIME Business Hub instance.
 - An [execution context](#) that operates using a selected custom executor image. If you are working with executor version = 4.7.2, the execution context needs to be created by pointing to a specific Docker image in order to work with Kerberos. A specific Docker image is available with the following name:

`registry.hub.knime.com/knime/knime-full:4.7.2-rc-000146-1816590` If you are working with executor version > 4.7.2 you can create the execution context using the corresponding docker image for the version that you want to use. A list is available [here](#).

Follow the instructions in the [KNIME Business Hub User Guide](#) to create a new execution context and specify the above Docker image.

Setting up krb5.conf

The KNIME Business Hub execution context 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
udp_preference_limit = 1

[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 Business Hub. The `udp_preference_limit` variable value set to 1 is to ensure that when sending a message to the KDC, the library will try using TCP over UDP connection.

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



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

To create a **customization profile** you need to first create an `.epf` preferences file that will contain the Kerberos configuration preferences. 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 Business Hub 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 Business Hub executors. In this case, the keytab must not be stored in the profile folder, but needs to be present on the KNIME Business Hub 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**, **DEBUG** or **ERROR**.

For example, a recommended Kerberos configurations for KNIME Business Hub 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 Business Hub 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 Business Hub executor machine(s) so that the preferences can reference it by local path.

Now, once you created the `.epf` file you can create a `.zip` file that contains the following files:

- `preferences.epf` — created in the step above.
- `krb5.conf` — created in the previous section.
- `knimehub.keytab`

Then follow the instructions to [upload a customization profile](#) to your KNIME Business Hub instance.

For KNIME Business Hub executors

To apply the customization profile to a KNIME Hub execution context's executor, follow the instructions in the [Apply a customization profile to KNIME Hub executor](#) section.

For KNIME Analytics Platform clients

To apply the customization profile to a KNIME Analytics Platform client, follow the instructions in the [Apply a customization profile to KNIME Analytics Platform client](#) section.

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. 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.
4. Please make sure that any paths set in the preferences are valid on the Hub instance:
 - If `krb5.conf` is used, the `krb5.conf` file needs to be added together with the `.epf` file to the `.zip` file that will be uploaded to the KNIME Business Hub. 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 Business Hub.

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