

Installation Process

Jethro installation is a straightforward process, which runs as root, extracts Jethro software to `/opt/jethro`, and automatically creates and configures the Jethro OS user and Linux service.

To install Jethro:

1. Connect as root and download the Jethro RPM file to a local directory: **wget** <http://www.jethro.io/latest-rpm>
2. Run the following command to install the newly downloaded RPM file:

```
rpm -Uvh latest-rpm
```

3. Change the owner of the Jethro cache directory to the *jethro* user, which was created during the RPM installation:

```
chown jethro:jethro /mnt/jethro_local_cache
```

4. If the installation takes place on a Hadoop cluster, create a Jethro OS user for the Hadoop cluster as follows:
 - **Non-Kerberos** - Create a Hadoop OS user and a group called Jethro on the NameNode, and on any node that may serve as NameNode (such as a Standby NameNode, if configured).
If Jethro OS user and group do not exist on the NameNode, each access from Jethro to the NameNode will generate a warning and a stack trace in the NameNode logs, flooding the NameNode's log and significantly impacting its performance.
 - **Kerberos** – Create a Kerberos principal called *jethro* and generate a keytab file for it. For example, if using MIT KDC:

```
local -q "addprinc -randkey jethro"  
kadmin.local -q "ktadd -k jethro.hadoop.keytab jethro"
```

5. If using Kerberos, securely copy the newly created keytab file to the Jethro home directory and change its permissions accordingly:

```
scp kdc_owner@kdc_server:jethro.hadoop.keytab /home/jethro  
chown jethro:jethro /home/jethro/jethro.hadoop.keytab  
chmod 600 /home/jethro/jethro.hadoop.keytab
```

6. Optionally, set the password for the jethro OS user:

```
passwd jethro
```

7. Proceed with the setup as the jethro OS user:

```
su - jethro  
#To install the RPM under any user other than the default one (jethro), run:  
JETHRO_INSTALL_USER={user} JETHRO_INSTALL_GROUP={group} rpm -i {rpm file}
```



If you intend to use a non-default user, namely: any user other than '*Jethro*', before starting the installation process ensure that the requested user already exists and that you specify the installation group.

If the rpm file is already installed, you must specify the same user that was used for installing the rpm.

If a new user is going to be used for an installation that overrides an existing installation, the rpm package should be removed first (existing instances will need to be attached to directories or storage mount points, and you have to manually change the cache owner).