

Configuring Kerberos

Here you can find information how to configure AtScale to use Kerberos for authentication and to communicate with a Kerberos-secured data warehouse.



Note: There are cases where you can use Kerberos Credential Cache instead of Keytab file. For more information, see [Using Kerberos Credential Cache](#).

Before You Begin

1. Install and configure the Kerberos client package or packages for your environment on the AtScale host(s).
2. Collect the following from your System Administrator:

Information	Description
Kerberos principal under which AtScale runs	<p><code>atscale/instanceIdentifier@REALM</code></p> <p>In Clustered AtScale, this principal must be the same on every AtScale Application Host. Set the instance identifier to the fully qualified domain name (FQDN) users will use to access AtScale via the external load balancer.</p> <p>AtScale only tests with a 3-part Kerberos principal. Principal format is environmentally dependent; for maximum supportability, AtScale recommends using a 3-part principal.</p>
Kerberos keytab file	<p>Issue a keytab for a user principal with the principal instance identifier set to the fully qualified domain name (FQDN) users will use to access AtScale via the external load balancer. If running Clustered AtScale, use this keytab on all AtScale hosts.</p>
LDAP User Kerberos Principal Attribute	<p>If you want AtScale authentication to use Kerberos. See Connecting to an LDAP Server or Microsoft Active Directory for more details.</p>
LDAP User Unique ID Attribute	<p>Needed if using Delegated Authorization or Impersonation to communicate with the Data warehouse. See Connecting to an LDAP Server or Microsoft Active Directory for more details.</p>

Procedure

1. Copy the keytab file to the AtScale host. It is recommended but not required to put the keytab file in `/opt/atScale/conf`.
2. Make the file readable by the `ATSCALE_USER` (by default `atscale`)
3. Confirm that the `klist` command succeeds. For example:

```
klist -kt /opt/atScale/conf/atScaler-ad.keytab
Keytab name: FILE:/opt/atScale/conf/atScaler-ad.keytab
KVNO Timestamp          Principal
-----
 2 01/05/2019 00:26:31 atScaler/atScale-ha-node-lb.docker.infra.atScale.com@CORPTEST.INFRA.ATSCALE.COM
```

4. Open `/opt/atScale/conf/atScale.yaml`, and edit the following kerberos properties. See Figure 1 for an example.
 1. Set `enabled` to true
 2. Set `keytab` to the path to the keytab file
 3. Set `principal` to the Kerberos principal string assigned to AtScale by your system administrator

Figure 1. Example kerberos configuration in `atScale.yaml`

```
kerberos:
  enabled: true
  keytab: "/path/to/atScale.keytab"
  principal: "atScale/instanceIdentifier@REALM"
```

5. Execute `configurator.sh` with the `--apply` option to apply the new configuration.

```
su - atScale
cd /opt/atScale/versions/<package_version>
./bin/configurator.sh --apply
```

Next Steps

1. If configuring Kerberos with MapR:

1. The Hadoop administrator must generate a MapR ticket with this command:

```
maprlogin kerberos
```

The command output includes the location of the generated ticket, for example: `/tmp/maprticket_0`. Make sure the expiration is set to the desired duration with the optional parameter `-duration"`. For more information, see [maprlogin](#).

2. The AtScale administrator must copy the ticket file generated in the previous step to a location on the AtScale host (or hosts in HA mode).
3. The AtScale system administrator must then add this environment variable to the AtScale hosts' system profile. If running an AtScale cluster set this environment variable on every host in the cluster:

```
export MAPR_TICKETFILE_LOCATION=<ticket file location on disk>
```

For example: `export MAPR_TICKETFILE_LOCATION="/tmp/maprticket_0"`

2. If running an AtScale cluster, repeat the steps from the Procedure section above on every application host in the cluster.
3. Configure the Kerberos-specific Directory Service properties. Go to **SECURITY > DIRECTORY, SETUP**. Set the following properties. See [Connecting to an LDAP Server or Microsoft Active Directory](#) for more details.
 1. User Kerberos Principal Attribute - Required if you want AtScale authentication to use Kerberos.
 2. User Unique ID Attribute - Required if using delegated authorization or impersonation to communicate with the data warehouse.
4. Set the Kerberos principal on each desired data warehouse screen. See [Adding Hadoop Data Warehouses](#) for more details.
5. Enable Kerberos-specific engine settings. Go to **Settings > Engine** and enable the following properties:
 1. **THRIFTY.SASL.KERBEROS.ENABLED** - Required
6. Restart the AtScale Engine service. If running an AtScale Cluster, restart the engine service on each application host.

More Information

- ▲ If you wish to use Tableau Server Impersonation (aka "Tableau single sign-on"), see [Configure Tableau Server Impersonation](#).
- ▲ If you wish to map LDAP groups to Hadoop service accounts, see [Setting Up Impersonation of Hadoop Accounts by Directory Groups](#)