

Installation

Table of contents

| | |
|---------------------------------|---|
| 1 Introduction | 2 |
| 2 Procedure..... | 2 |
| 3 Server Commands..... | 2 |
| 4 Requirements..... | 2 |
| 5 Hadoop Distributed Cache..... | 3 |
| 6 Permissions..... | 4 |
| 7 Secure Cluster..... | 5 |

1 Introduction

Templeton is deep in the middle of development and does not yet have a smooth install procedure. It is also designed to connect together services that are not normally connected and therefore has a complex configuration. As such, this version of Templeton should only be installed by expert developers.

2 Procedure

1. Ensure that the [required related installations](#) are in place, and place required files into the [Hadoop distributed cache](#).
2. Download and unpack the Templeton distribution.
3. Set the `TEMPLETON_HOME` environment variable to the base of the Templeton installation. This is used to find the Templeton configuration.
4. Review the [Templeton configuration](#) and update or create `templeton-site.xml` as required. Ensure that site specific component installation locations are accurate, especially the Hadoop configuration path. Configuration variables that use a filesystem path try to have reasonable defaults, but it's always safe to specify a full and complete path.
5. Verify that HCatalog is installed and that the `hcat` executable is in the `PATH`.
6. Build Templeton using the command `ant jar` from the top level Templeton directory.
7. Start the Templeton server with the command `bin/templeton_server.sh start`.
8. Check that your local install works. Assuming that Tomcat is running on port 8080, the following command would give output similar to that shown.

```
% curl -i http://localhost:50111/templeton/v1/status
HTTP/1.1 200 OK
Content-Type: application/json
Transfer-Encoding: chunked
Server: Jetty(7.6.0.v20120127)

{"status":"ok","version":"v1"}
%
```

3 Server Commands

- **Start the server:** `bin/templeton_server.sh start`
- **Stop the server:** `bin/templeton_server.sh stop`
- **End-to-end build, run, test:** `ant e2e`

4 Requirements

- [Ant](#), version 1.8 or higher

- [Hadoop](#), version 0.20.205.0
- [ZooKeeper](#) is required if you are using the ZooKeeper storage class. (Be sure to review and update the ZooKeeper related [Templeton configuration](#).)
- [HCatalog](#). Version 0.4.0 or higher. The hcat executable must be both in the PATH and properly configured in the [Templeton configuration](#).
- Permissions must be given to the user running Templeton server. (see below)
- If running a secure cluster, Kerberos keys and principals must be created. (see below)
- [Hadoop Distributed Cache](#). To use the [Hive](#), [Pig](#), or [hadoop/streaming](#) resources, see instructions below for placing the required files in the Hadoop Distributed Cache.

5 Hadoop Distributed Cache

Templeton requires some files be accessible on the [Hadoop distributed cache](#). For example, to avoid the installation of Pig and Hive everywhere on the cluster, Templeton gathers a version of Pig or Hive from the Hadoop distributed cache whenever those resources are invoked. After placing the following components into HDFS please update the site configuration as required for each.

- **Hive:** [Download](#) the HCatalog tar.gz file and place it in HDFS. (If you need a version that is not yet released, you may need to build it yourself following the HCatalog instructions.)

```
hadoop fs -put /tmp/hcatalog-0.4.0.tar.gz /user/templeton/hcatalog-0.4.0.tar.gz
hadoop fs -chmod 444 /user/templeton/hcatalog-0.4.0.tar.gz
```

- **Pig:** [Download](#) the Pig tar.gz file and place it into HDFS. For example:

```
hadoop fs -put /tmp/pig-0.9.2.tar.gz /user/templeton/pig-0.9.2.tar.gz
hadoop fs -chmod 444 /user/templeton/pig-0.9.2.tar.gz
```

- **Hadoop Streaming:** Place `hadoop-streaming.jar` into HDFS. For example, use the following command, substituting your path to the jar for the one below.

```
hadoop fs -put $HADOOP_PREFIX/hadoop-0.20.205.0/contrib/streaming/hadoop-
streaming-0.20.205.0.jar \
/user/templeton/hadoop-streaming.jar
hadoop fs -chmod 444 /user/templeton/hadoop-streaming.jar
```

- **Override Jars:** Place override jars required (if any) into HDFS. *Note:* As of this writing, all released versions of Hadoop require a patch to properly run Templeton. This patch is distributed with Templeton (located at `templeton/src/hadoop_temp_fix/ugi.jar`) and should be placed into HDFS, as reflected in the current default configuration.

```
hadoop fs -put ugi.jar /user/templeton/ugi.jar
```

```
hadoop fs -chmod 444 /user/templeton/ugi.jar
```

The location of these files in the cache, and the location of the installations inside the archives, can be specified using the following Templeton configuration variables. (See the [Configuration](#) documentation for more information on changing Templeton configuration parameters.)

| Name | Default | Description |
|--------------------------------|---|---|
| templeton.pig.archive | hdfs:///user/templeton/pig-0.9.2.tar.gz | The path to the Pig archive. |
| templeton.pig.path | pig-0.9.2.tar.gz/pig-0.9.2/bin/pig | The path to the Pig executable. |
| templeton.hive.archive | hdfs:///user/templeton/hcatalog-0.3.0.tar.gz | The path to the Hive archive. |
| templeton.hive.path | hcatalog-0.3.0.tar.gz/hcatalog-0.3.0/bin/hive | The path to the Hive executable. |
| templeton.streaming.jar | hdfs:///user/templeton/hadoop-streaming.jar | The path to the Hadoop streaming jar file. |
| templeton.override.jars | hdfs:///user/templeton/ugi.jar | Jars to add to the HADOOP_CLASSPATH for all Map Reduce jobs. These jars must exist on HDFS. |

6 Permissions

Permission must be given for the user running the templeton executable to run jobs for other users. That is, the templeton server will impersonate users on the Hadoop cluster.

Create (or assign) a Unix user who will run the Templeton server. Call this USER. See the Secure Cluster section below for choosing a user on a Kerberos cluster.

Modify the Hadoop core-site.xml file and set these properties:

| Variable | Value |
|------------------------------|--|
| hadoop.proxyuser.USER.groups | A comma separated list of the Unix groups whose users will be impersonated. |
| hadoop.proxyuser.USER.hosts | A comma separated list of the hosts that will run the hcat and JobTracker servers. |

7 Secure Cluster

To run Templeton on a secure cluster follow the Permissions instructions above but create a Kerberos principal for the Templeton server with the name `USER/host@realm`

Also, set the templeton configuration variables `templeton.kerberos.principal` and `templeton.kerberos.keytab`