Configuration

Table of contents

The configuration for Templeton merges the normal Hadoop configuration with the Templeton specific variables. Because Templeton is designed to connect services that are not normally connected, the configuration is more complex than might be desirable.

The Templeton specific configuration is split into two layers:

1. **templeton-default.xml** - All the configuration variables that Templeton needs. This file sets the defaults that ship with Templeton and should only be changed by Templeton developers. Do not copy this file and/or change it to maintain local installation settings. Because templeton-default.xml is present in the Templeton war file, editing a local copy of it will not change the configuration.
2. **templeton-site.xml** - The (possibly empty) configuration file in which the system administrator can set variables for their Hadoop cluster. Create this file and maintain entries in it for configuration variables that require you to override default values based on your local installation.

The configuration files are loaded in this order with later files overriding earlier ones.

Note: the Templeton server will require restart after any change to the configuration.

To find the configuration files, Templeton first attempts to load a file from the `CLASSPATH` and then looks in the directory specified in the `TEMPLETON_HOME` environment variable.

Configuration files may access the special environment variable `env` for all environment variables. For example, the pig executable could be specified using:

```
${env.PIG_HOME}/bin/pig
```

Configuration variables that use a filesystem path try to have reasonable defaults. However, it's always safe to specify the full and complete path if there is any uncertainty.

Note: The location of the log files created by Templeton and some other properties of the logging system are set in the templeton-log4j.properties file.

1 Variables

Name	Default	Description
templeton.port	`50111`	The HTTP port for the main server.
templeton.hadoop.config.dir	`$(env.HADOOP_CONFIG_DIR)`	The path to the Hadoop configuration.
templeton.jar	`${env.TEMPLETON_HOME}/templeton/templeton-0.1.0-dev.jar`	The path to the Templeton jar file.

Name	Default	Description
templeton.libjars	`${env.TEMPLETON_HOME}/lib/zookeeper-3.3.4.jar`	Jars to add to the classpath.
templeton.override.jars	`hdfs:///user/templeton/ugi.jar`	Jars to add to the HADOOP_CLASSPATH for all Map Reduce jobs. These jars must exist on HDFS.
templeton.override.enabled	`true`	Enable the override path in templeton.override.jars
templeton.streaming.jar	`hdfs:///user/templeton/hadoop-streaming.jar`	The hdfs path to the Hadoop streaming jar file.
templeton.hadoop	`${env.HADOOP_PREFIX}/bin/hadoop`	The path to the Hadoop executable.
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.hcat	`${env.HCAT_PREFIX}/bin/hcat`	The path to the Hcatalog executable.
templeton.hive.archive	`hdfs:///user/templeton/hcatalog-0.5.0.tar.gz`	The path to the Hive archive.
templeton.hive.path	`hcatalog-0.5.0.tar.gz/hcatalog-0.5.0/bin/hive`	The path to the Hive executable.
templeton.hive.properties	`hive.metastore.local=fal hive.metastore.uris=thri localhost:9933, hive.metastore.sasl.enab`	Properties to set when running hive.
templeton.exec.encoding	`UTF-8`	The encoding of the stdout and stderr data.
templeton.exec.timeout	`10000`	How long in milliseconds a program is allowed to run on the Templeton box.
templeton.exec.max-procs	`16`	The maximum number of processes allowed to run at once.
templeton.exec.max-output-bytes	`1048576`	The maximum number of bytes from stdout or stderr stored in ram.

Name	Default	Description
templeton.controller.mr.child.opts	`-server -Xmx256m -Djava.net.preferIPv4Stac`	Java options to be passed to templeton controller map task.
templeton.exec.envs	`HADOOP_PREFIX,HADOOP_HOM`	The environment variables passed through to exec.
templeton.zookeeper.hosts	`127.0.0.1:2181`	ZooKeeper servers, as comma separated host:port pairs
templeton.zookeeper.session-timeout	`30000`	ZooKeeper session timeout in milliseconds
templeton.callback.retry.interval	`10000`	How long to wait between callback retry attempts in milliseconds
templeton.callback.retry.attempts	`5`	How many times to retry the callback
templeton.storage.class	`org.apache.hcatalog.temp`	The class to use as storage
templeton.storage.root	`/templeton-hadoop`	The path to the directory to use for storage
templeton.hdfs.cleanup.interval	`43200000`	The maximum delay between a thread's cleanup checks
templeton.hdfs.cleanup.maxage	`604800000`	The maximum age of a templeton job
templeton.zookeeper.cleanup.inter	`43200000`	The maximum delay between a thread's cleanup checks
templeton.zookeeper.cleanup.max	`604800000`	The maximum age of a templeton job
templeton.kerberos.secret	A random value	The secret used to sign the HTTP cookie value. The default value is a random value. Unless multiple Templeton instances need to share the secret the random value is adequate.
templeton.kerberos.principal	None	The Kerberos principal to used by the server. As stated by the Kerberos SPNEGO specification, it should be `USER/`

Name	Default	Description
		`${HOSTNAME}@{REALM}`. It does not have a default value.
templeton.kerberos.keytab	None	The keytab file containing the credentials for the Kerberos principal.