

GettingStartedMesos

- [Getting Started with Hama on Mesos](#)
 - [Requirements](#)
 - [Building](#)
 - [Configuration](#)
 - [Starting The BSPMaster](#)
 - [Testing](#)

Getting Started with Hama on Mesos

Requirements

In order to run Hama on Mesos it is required that Mesos already be installed on the cluster. Instructions to set up Mesos may be found at the project website: <http://mesos.apache.org/>.

Building

Be sure to build Hama with your version of Mesos specified:

```
mvn clean install -Phadoop1 -Dhadoop.version=1.x.x -Dmesos.version=0.20.0
```

The groom servers will be set up by Mesos during execution of a job. In order for Mesos to do this, you will need to upload a built version of Hama to a place where the Mesos slaves can find it, such as the HDFS:

```
hadoop fs -put hama-0.7.0-SNAPSHOT.tar.gz /hama.tar.gz
```

Configuration

There are several Mesos related properties that must be set in hama-site.xml:

Property	Recommended Value	Description
bsp.master.TaskWorkerManager.class	org.apache.hama.bsp.MesosScheduler	Instructs the scheduler to use Mesos to execute tasks of each job
hama.mesos.executor.uri	hdfs://hdfs.name.node:port/hama.tar.gz	This is the URI of the Hama distribution. Upload this yourself.
bsp.tasks.maximum.total	10	This is an override for the total maximum tasks that may be run. The default behavior is to determine a value based on the available groom servers. However, if using Mesos, the groom servers are not yet allocated. So, a value indicating the maximum number of slots available in the cluster is needed.
hama.mesos.master		This is the address of the Mesos master instance. If you're using Zookeeper for master election, use the Zookeeper address here (i.e., zk://zk.apache.org:2181/hadoop/mesos).
bsp.child.java.opts	-Xmx1024m	Java opts for the groom server child processes.

Hama requires one cpu and memory defined by `bsp.child.java.opts` for each slot. This means that a cluster with `bsp.tasks.maximum.total` set to 2 and `bsp.child.java.opts` set to `-Xmx1024m` will need at least 2 cpus and 2048m of memory.

Manually distributing the configuration is not necessary. Hama and Mesos will distribute the configuration and provide overrides where necessary.

There are also several other properties that will need to be considered when setting up hama for the first time:

Property	Default Value	Description
hama.tmp.dir	/tmp/hama-\${user.name}	Temporary directory on the local filesystem.
hama.zookeeper.quorum		Comma separated list of servers in the Zookeeper Quorum
hama.zookeeper.property.clientPort	2181	The port to which the zookeeper clients connect

For a working example of the hama-site.xml and to avoid some configuration mistakes you can check this [blog post for installing Hama on Mesos](#)

Starting The BSPMaster

With Hama on Mesos you only need to set up and start the BSPMaster. After setting the configuration the bsp master may be started:

```
% $HAMA_HOME/bin/hama-daemon.sh start bspmaster
```

Testing

In order to test the setup, run an example from Hama:

```
$HAMA_HOME/bin/hama jar hama-examples-x.x.x.jar gen fastgen -v 100 -e 10 -o randomgraph -t 2 -of json
```

Then verify that Mesos tasks are being created run. Do this by navigating to the Mesos status page and clicking on the id of the Hama framework. There should be tasks listed in either the active or completed tasks section.

Please email any questions to user@hama.apache.org