

Bigtop 1.1.0 Release

Release BOM

```
bigtop 1.1.0 stack includes the following components
  bigtop_groovy      2.4.4
  bigtop_jsvc       1.0.15
  bigtop_tomcat     6.0.36
  bigtop_utils      1.1.0
  crunch            0.12.0
  datafu            1.0.0
  flume             1.6.0
  giraph            1.1.0
  hadoop            2.7.1
  hama              0.7.0
  hbase             0.98.12
  hive              1.2.1
  hue               3.9.0
  ignite_hadoop    1.5.0.final
  kafka             0.8.1.1
  kite              1.1.0
  mahout            0.11.1
  oozie             4.2.0
  phoenix          4.6.0-HBase-0.98
  pig               0.15.0
  solr              4.9.0
  spark            1.5.1
  sqoop             1.4.5
  sqoop2           1.99.4
  tachyon          0.6.0
  tez               0.6.2
  ycsb              0.4.0
  zeppelin         0.5.6
  zookeeper        3.4.6
```

Release note

On behalf of the Apache Bigtop team, I'd love to announce the general availability of the Bigtop 1.1.0 release.

A few highlights of this release include:

- five different operating systems are supported
- a brand new ppc64le architecture was added to the support matrix
- most stable Hadoop 2.7.1 is used
- new data notebooks project Zeppelin (incubating)
- addition of Apache Hama
- toolchain is updated to the latest versions of the development tools
- standard docker containers are available for builds, testing, and alike
- on-going improvements in the unattended cluster deployment and testing
- improvements in the smoke-tests coverage and UX
- and many upgrades to the latest versions of the ecosystem projects (Crunch, Flume, Ignite, Hive, Mahout, Oozie, Phoenix, Pig, Spark, and many others)

With Bigtop 1.1 the community continues on the path of delivering the most advanced, openly-developer data analytics stack to date. This is a true bigtop project, bringing together virtually all of the Apache Software Foundation data processing projects. This release of the Apache Bigdata platform include close to 260 bug fixes and new features.

Deploying Bigtop is easy: grab the repo/list file for your favorite Linux distribution:

<http://www.apache.org/dist/bigtop/bigtop-1.1.0/repos/>

and you'll be running your very own bigdata cluster in no time!

Alternatively, you can use Bigtop provisioner to spin-up a virtual cluster with a single command as explained in

<https://cwiki.apache.org/confluence/display/BIGTOP/Bigtop+Provisioner+User+Guide>

I want to emphasize the great work done by the project contributors, and thank everyone in this and other communities who made this release possible!

Thanks,
Cos (Bigtop 1.1.0 RM)