



Architecting the Future of Big Data

Hortonworks Technical Preview for Stinger Phase 3

Released: 12/17/2013

Document Version 1.0.4





Architecting the Future of Big Data

Welcome to Hortonworks Inc, technical preview for Stinger Phase 3. The Technical Preview provides early access to upcoming features in the Hortonworks product, letting you test and review during the development process. These features are considered under development.

Although your feedback is greatly appreciated these features are not intended for use in your production systems and not considered Supported by Hortonworks.

Have fun and please send feedback to us on the Community forums:

<http://hortonworks.com/community/forums/forum/hive/>

Stinger Phase 3 Technical Preview Introduction	4
System Requirements	4
Supported Operating Systems	4
Installing Stinger Phase 3 Preview	4
Installation Step 1: Download the Stinger Quickstart Bundle	4
Installation Step 2: Install and Configure Tez.....	5
Installation Step 3: Install and Configure Hive.....	5
Optional Installation Step 4: Update Resource Manager	7
Optional Installation Step 5: Enable Large Scale Statistics	8
Optional: Run HiveServer2	8
Running Hive and Ensuring A Correct Install	9
Known Issues and Limitations	9
Further Reading.....	9
Appendix A: Parameters Reference	11

Stinger Phase 3 Technical Preview Introduction

Stinger Phase 3 combines Apache Hive and Apache Tez to provide a SQL engine capable of interactive query natively within Hadoop.

System Requirements

Stinger Phase 3 Technical Preview is optimized to work with Hortonworks HDP 2.0. This includes a full physical cluster created through Apache Ambari or a Hortonworks HDP 2.0 Sandbox which can be used for small-scale experimentation.

If you are a Hadoop expert, it may be possible to use Stinger Phase 3 Technical Preview with Apache Hadoop 2.2 based clusters, but the installation steps in this document may need to be adjusted.

Supported Operating Systems

Stinger Phase 3 Technical Preview works on clusters running the following operating systems:

- 64-bit Red Hat Enterprise Linux (RHEL) 6
- 64-bit CentOS 6

Installing Stinger Phase 3 Preview

Installing Stinger Phase 3 Preview can be done in as little as 2 steps and as many as 4. The mandatory steps do not require any disruption to the Hadoop cluster and can be done by unprivileged users. The optional steps require cluster disruptions and must be done by a privileged user. The optional steps offer performance benefits, which may be substantial. Specific details are given below.

Installing Stinger Phase 3 Preview: Overview		
Step	Description	Mandatory?
Step 1	Download Quickstart Bundle	Yes
Step 2	Install and configure Tez	Yes
Step 3	Install and configure Hive	Yes
Step 4	Update ResourceManager	No. Perform this step to further reduce query launch latency.
Step 5	Enable Statistics	No. Perform this step to improve Hive's query planner through improved table level statistics.

Installation Step 1: Download the Stinger Quickstart Bundle

Before you begin, log in to one of the nodes of your cluster as root and download the Stinger Quickstart Bundle. The Stinger Quickstart Bundle provides sample configurations that will help you deploy and tune Stinger.

```
wget
https://github.com/cartershanklin/StingerQuickstart/blob/master/Stinger-Preview-Quickstart.tgz?raw=true
tar -xzf Stinger-Preview-Quickstart.tgz
cd Stinger-Preview-Quickstart
```

Installation Step 2: Install and Configure Tez

In this step we deploy Tez 0.2 to the Hadoop cluster.

1. Download the Stinger Phase 3 Preview Tez tarball.

```
wget http://public-repo-1.hortonworks.com/HDP-2.1.0.0/repos/centos6/tars/tez-0.2.0.2.1.0.0-92.tar.gz
```
2. Extract the Stinger Phase 3 Preview Tez tarball.

```
tar -C /opt -xzf tez-0.2.0.2.1.0.0-92.tar.gz
```
3. Create a directory on HDFS for Tez.

```
hadoop fs -mkdir -p /apps/tez
hadoop fs -chmod 755 /apps/tez
```
4. Deploy Tez files to HDFS:

```
hadoop fs -copyFromLocal /opt/tez-0.2.0.2.1.0.0-92/* /apps/tez/
```
5. Deploy a sample Tez configuration from the Quickstart Bundle:
Option 1: If you are using a physical Hadoop cluster:

```
cp configs/tez-site.xml.physical \
/etc/hadoop/conf/tez-site.xml
```


Option 2: If you are in a Sandbox or small cluster:

```
cp configs/tez-site.xml.sandbox \
/etc/hadoop/conf/tez-site.xml
```

Installation Step 3: Install and Configure Hive

In this step we deploy a Hive 0.13 preview to the Hadoop cluster. This version of Hive is designed to work with Tez 0.2.

1. Download the Stinger Phase 3 Preview Hive tarball.

```
wget http://public-repo-1.hortonworks.com/HDP-2.1.0.0/repos/centos6/tars/apache-hive-0.13.0.2.1.0.0-92-bin.tar.gz
```

2. Extract the Stinger Phase 3 Preview Hive tarball.

```
tar -C /opt -xzf apache-hive-0.13.0.2.1.0.0-92-bin.tar.gz
```
3. Create a directory on HDFS for Hive with appropriate permissions.

```
hadoop fs -mkdir -p /user/hive  
hadoop fs -chmod 755 /user/hive
```
4. Deploy Hive files to HDFS:

```
hadoop fs -put \  
/opt/apache-hive-0.13.0.2.1.0.0-92-bin/lib/hive-exec-  
*.jar \  
/user/hive/hive-exec-0.13.0-SNAPSHOT.jar
```
5. Configure your environment:
Any time you want to run Hive or start HiveServer2 you should first run these commands:

```
export HIVE_HOME=/opt/apache-hive-0.13.0.2.1.0.0-92-  
bin  
export HIVE_CONF_DIR=$HIVE_HOME/conf  
export PATH=$HIVE_HOME/bin:$PATH  
export HADOOP_CLASSPATH=/opt/tez-0.2.0.2.1.0.0-  
92/*:/opt/tez-0.2.0.2.1.0.0-92/lib/*  
export HADOOP_USER_CLASSPATH_FIRST=true
```
6. Deploy Hive configuration:
To use Hive on Tez, certain configuration options must be set in hive-site.xml. This file must also contain information including login and password to your metastore.

Option 1: Copy an existing hive-site.xml.

Use this option if you don't intend to use HiveServer2 and you are willing to launch hive with extra parameters. Your existing hive-site.xml will contain all necessary URIs and logins and can be used with the new version of Hive. Additional settings will need to be given on the command line to ensure Tez is used.

Example:

```
cp /etc/hive/conf.server/hive-site.xml \  
/opt/apache-hive-0.13.0.2.1.0.0-92-bin/conf/hive-  
site.xml
```

You will then need to start Hive with an additional parameter that contains the Stinger settings. The file stinger.settings is provided in the Stinger Quickstart Bundle.

```
hive -i stinger.settings -hiveconf  
hive.optimize.tez=true
```

Option 2: Add Stinger-specific settings to an existing hive-site.xml.

If you plan to use HiveServer2 it is best to place Stinger settings directly in your hive-site.xml. The mergeHadoopConfig utility takes an existing hive-site.xml and merges in other properties needed to run Hive on Tez.

Example:

```
./mergeHadoopConfig \  
/etc/hive/conf.server/hive-site.xml \  
configs/stingerProperties.xml > \  
/opt/apache-hive-0.13.0.2.1.0.0-92-bin/conf/hive-  
site.xml
```

Option 3: If you are an expert Hive user:

If you are an expert, you can start with one of the provided hive-site.xml templates in the Stinger Quickstart Bundle and adjust to your environment. You will need to configure your metastore manually.

Option A: If you are on a physical Hadoop cluster:

Copy configs/hive-site.xml.physical to /opt/apache-hive-0.13.0.2.1.0.0-92-bin/conf/hive-site.xml.

Option B: If you are in a Sandbox or small cluster:

Copy configs/hive-site.xml.sandbox to /opt/apache-hive-0.13.0.2.1.0.0-92-bin/conf/hive-site.xml.

Appendix A contains more information on additional tuning parameters.

Optional Installation Step 4: Update Resource Manager

For substantially faster query processing the Resource Manager needs to be updated and restarted. This step is strongly suggested but not completely mandatory. It is possible to try Stinger Preview using only Steps 1 and 2 (for example on a shared cluster where you cannot restart the RM), but query startup times will be slower if you don't upgrade the RM.

Note: This step assumes you are managing your cluster with Apache Ambari.

1. Log into the Resource Manager node as root.
2. Add the Stinger Phase 3 Preview RPM repository to the machine.
wget <http://public-repo-1.hortonworks.com/HDP-2.1.0.0/repos/centos6/hdp.repo> -O /etc/yum.repos.d/stinger.repo

3. Upgrade the Resource Manager via YUM.
`yum upgrade hadoop-yarn-resourcemanager`
4. Set the following variable in YARN within Ambari:
`yarn.scheduler.capacity.schedule-asynchronously = true`
5. Restart YARN and the Resource Manager in Ambari.

Optional Installation Step 5: Enable Large Scale Statistics

Statistics allow Hive to create more intelligent plans to process queries more efficiently. In this version of Hive, statistics are computed using counters and require one counter per partition of data.

The number of counters can be set on a per-job basis but is subject to a globally defined limit.

If you plan to import or analyze a large number of partitions while you test Stinger, we recommend you increase the global counter limit. This step is not necessary if you plan to use a small number of partitions, or if you load or analyze 10 or fewer partitions at a time.

This change requires an update to `mapred-site.xml` on all nodes and requires a History Server restart. Because multiple files need to be modified it is best to do this via Ambari.

Set the following values in MapReduce Advanced Settings in Ambari:

```
mapreduce.job.counters.max = 10000
mapreduce.job.counters.group.name.max = 10000
mapreduce.job.counters.counter.name.max = 10000
mapreduce.job.counters.groups.max = 10000
```

Restart MapReduce in Ambari after making these changes.

Optional: Run HiveServer2

If you want to use HiveServer2, these 6 commands can be used to start it. Note that Tez settings must be specified in `hive-site.xml` before starting.

```
export HIVE_HOME=/opt/apache-hive-0.13.0.2.1.0.0-92-bin
export HIVE_CONF_DIR=$HIVE_HOME/conf
export PATH=$HIVE_HOME/bin:$PATH
export HADOOP_CLASSPATH=/opt/tez-0.2.0.2.1.0.0-92/*:/opt/tez-0.2.0.2.1.0.0-92/lib/*
export HADOOP_USER_CLASSPATH_FIRST=true
nohup /opt/apache-hive-0.13.0.2.1.0.0-92-bin/bin/hiveserver2 > hive-server2.log &
```


Running Hive and Ensuring A Correct Install

After setup is complete, Stinger Preview can be used either via the `hive` command line, or, if you set it up, through HiveServer2 using `beeline` or an analytics tool of your choice.

Launch Hive using:

```
hive -hiveconf hive.optimize.tez=true
```

The additional parameter causes Hive to pre-warm containers for lower query latency.

To ensure the proper configurations have been made, in the `hive` command line, issue the command `set hive.vectorized.execution.enabled;` you should see the following output:

```
hive> set hive.vectorized.execution.enabled;
hive.vectorized.execution.enabled=true
hive> █
```

If you do not see a this property set to true, ensure you have set the `HIVE_HOME` environment variable as in Step 3 before starting `hive`.

Known Issues and Limitations

Stinger Phase 3 Beta has the following known limitations:

- Stinger Phase 3 Technical Preview does not support the SQL UNION operator.
- Stinger Phase 3 Technical Preview does not support the sort-merge-bucket join.
- Stinger Phase 3 Technical Preview is not compatible with secure (kerberized) clusters.
- Stinger Phase 3 Technical Preview may return incorrect query results if you run multiple datanodes on a single host or VM. This is not a common deployment scenario.

If you have questions or encounter other issues not listed above, visit the forums for additional help at: <http://hortonworks.com/community/forums/forum/hive/>

Further Reading

[Apache Tez](#)

[Apache Hive](#)

Appendix A: Parameters Reference

For advanced tuning, this section tells you all you need to know about Stinger Preview settings.

These parameters should always be set when using Stinger.

Stinger Parameters You Should Set	
Name:	hive.optimize.tez
Description:	Causes Hive to generate Tez plans instead of Map/Reduce plans. Can be set on a query-by-query basis.
Suggested Value:	true
Name:	hive.vectorized.execution.enabled
Description:	Causes Hive to use the new vectorized execution engine.
Suggested Value:	true
Name:	hive.compute.query.using.stats
Description:	When set to true, Hive will compute certain queries such as count(*) using only metadata.
Suggested Value:	true
Name:	mapreduce.map.output.compress
Description:	Compress map output to save network and disk utilization.
Suggested Value:	true
Name:	mapreduce.map.output.compress.codec
Description:	Codec to use when compressing map output.
Suggested Value:	org.apache.hadoop.io.compress.SnappyCodec
Note:	You may need to install Snappy on your cluster if not using HDP 2.0
Name:	tez.runtime.intermediate-output.should-compress
Description:	Compress Tez intermediate output.
Suggested Value:	true
Name:	tez.runtime.intermediate-output.compress.codec
Description:	Codec to use when compressing Tez intermediate output.
Suggested Value:	org.apache.hadoop.io.compress.SnappyCodec
Note:	You may need to install Snappy on your cluster if not using HDP 2.0
Name:	tez.runtime.intermdiate-input.is-compressed
Description:	Instructs Tez that its intermediate outputs are compressed. Note that “intermediate” is misspelled in this property, which will be corrected in the future.
Suggested Value:	true
Name:	tez.runtime.intermediate-input.compress.codec
Description:	The codec used to compress intermediate output.
Suggested Value:	org.apache.hadoop.io.compress.SnappyCodec
Note:	You may need to install Snappy on your cluster if not using HDP 2.0
Name:	hive.input.format
Description:	Setting the right format ensures Hive can eliminate empty splits.
Suggested Value:	org.apache.hadoop.hive.ql.io.HiveInputFormat

Name:	hive.orc.splits.include.file.footer
Description:	Facilitates faster split calculation.
Suggested Value:	true
Name:	hive.stats.autogather
Description:	Automatically gather key table statistics.
Suggested Value:	true

These parameters can be tuned to best suit your needs.

Stinger Parameters You Can Tune	
Name:	mapreduce.map.memory.mb
Description:	Maximum amount of memory YARN will allow map tasks to be before killing them. This value should be somewhat larger than the max heap size (-Xmx) to allow for garbage collection overhead.
Suggested Value:	4096 for physical clusters, 1024 for Sandboxes.
Name:	mapreduce.reduce.memory.mb
Description:	Maximum amount of memory YARN will allow reduce tasks to be before killing them. This value should be somewhat larger than the max heap size (-Xmx) to allow for garbage collection overhead.
Suggested Value:	4096 for physical clusters, 1024 for Sandboxes.
Name:	mapred.map.child.java.opts
Description:	The actual map task command-line options. Important: Maps and reduces should use the same size to enable container re-use.
Suggested Value:	-server -Xmx3584m -Djava.net.preferIPv4Stack=true (physical clusters) -server -Xmx768m -Djava.net.preferIPv4Stack=true (Sandboxes)
Name:	mapred.reduce.child.java.opts
Description:	The actual reduce task command-line options. Important: Maps and reduces should use the same size to enable container re-use.
Suggested Value:	-server -Xmx3584m -Djava.net.preferIPv4Stack=true (physical clusters) -server -Xmx768m -Djava.net.preferIPv4Stack=true (Sandboxes)
Name:	mapreduce.reduce.shuffle.parallelcopies
Description:	Number of parallel reads a reducer is allowed to perform. A larger number tends to work better when a job has many maps. Hive on Tez tends to produce a large number of mappers so setting this above the Hadoop default of 5 is recommended. If your job produces large intermediate output files, a smaller number may work better.
Suggested Value:	30
Name:	hive.auto.convert.join.noconditionaltask.size
Description:	Hive will attempt to convert a join to a map join if the sum of sizes of all tables minus the size of the largest table is smaller than this amount. Note that in this release sizes on disk are used for this calculation, even if the files are compressed. Running joins as map joins is generally desirable so this value should be large. Having a larger heap size (mapred.map.child.java.opts) allows this value to be set higher.
Suggested Value:	64000000
Name:	tez.session.pre-warm.enabled
Description:	Instruct Tez to pre-warm containers for faster query launch. Set to true to enable this

feature. Disabled by default.	
Name:	tez.session.pre-warm.num.containers
Description:	Number of containers to pre-warm. Choose a reasonable amount relative to number of containers your cluster has. Suggested to use 10-30% of available containers.
Name:	tez.am.grouping.split-waves
Description:	The number of waves controls the number of mappers that will be launched to process a query. The number of mappers used is relative to the number of available slots in the cluster. A smaller value will cause more mappers to be launched. Larger values cause greater staggering and can be useful for dealing with stragglers.
Suggested Value:	1.7
Name:	tez.am.grouping.min-size
Description:	The smallest unit of data Tez will assign a map to work on.
Suggested Value:	16777216



About Hortonworks

Hortonworks is a leading commercial vendor of Apache Hadoop, the preeminent open source platform for storing, managing and analyzing big data. Hortonworks Data Platform provides an open and stable foundation for enterprises and a growing ecosystem to build and deploy big data solutions. Hortonworks is the trusted source for information on Hadoop, and together with the Apache community, Hortonworks is making Hadoop easier to install, manage and use. Hortonworks provides technical support, training & certification programs for enterprises, systems integrators & technology vendors.



3460 W. Bayshore Rd.
Palo Alto, CA 94303 USA

US: 1.855.846.7866
International: 1.408.916.4121
www.hortonworks.com