



# ADH - Advertising Distribution Hadoop

---

Apache Bigtop在AdMaster的应用 向磊

# Who am I?



- 向磊

AdMaster - 架构师

- 简介

2012 easyhadoop, phpHiveAdmin开源项目

2013 hadoop on ARM

# Who we are?



- 基础架构
- 技术研发
- AdMaster
  - 数字营销
  - 客户资源管理
  - 跨平台受众分析

## Why doing this?



- 面向广告行业的特性，在Hadoop生态系统中开发自定义Patch，特别针对广告数据特性进行时间复杂度优化，及作业监控、调度、机架感知相关。
- 其他一些bugfix或feature
- 为自动化运维提供RPM和DEB安装方式。



# A D H

优化合并过程，使采集数据直接合并成客户所需数据，提高处理速度

内置广告行业算法，不需编写MR即可计算PV, UV等各种维度数据

优化HBase查询，专为社会化数据定制，提高处理性能

集成数据任务调度系统，可以根据业务需求自动调整计算资源

集成Storm，优化Storm传输，减小数据延迟，实时提供数据计算结果

集成Spark，优化迭代工作负载，优化RDD序列化，提高性能和存储效率

# Screen shot from ADH



<b>Started:</b>	Thu Mar 09 15:02:29 CST 2017
<b>Version:</b>	2.6.1-adh1.0.0, rel1581abbb6ab62b0a41b7ce6141d7280bf0c53da
<b>Compiled:</b>	2017-02-24T10:19Z by xianglei from Unknown
<b>Cluster ID:</b>	CID-d56b51b3-0af0-4d66-93af-f7a4428ddddd
<b>Block Pool ID:</b>	BP-252487990-127.0.0.1-1413790391317

Admin State	Capacity	Used	Non DFS Used	Remaining	Blocks	Block pool used	Failed Volumes	Version	Rack
In Service	43.41 TB	31.19 TB	2.85 GB	12.22 TB	1112	31.19 TB (71.84%)	0	2.6.1-adh1.0.0	/K11
In Service	43.41 TB	28.67 TB	1.89 GB	14.74 TB	985	28.67 TB (66.05%)	0	2.6.1-adh1.0.0	/K12
In Service	43.41 TB	31.29 TB	649.33 MB	12.13 TB	1104	31.29 TB (72.06%)	0	2.6.1-adh1.0.0	/K10
In Service	42.74 TB	31.18 TB	76.57 GB	11.48 TB	1589883	31.18 TB (72.96%)	0	2.6.1-adh1.0.0	/J5

VCoers Used	VCoers Avail	Version
5	95	2.6.1-adh1.0.0
30	70	2.6.1-adh1.0.0
10	90	2.6.1-adh1.0.0
6	94	2.6.1-adh1.0.0
1	99	2.6.1-adh1.0.0
1	99	2.6.1-adh1.0.0
12	88	2.6.1-adh1.0.0
4	96	2.6.1-adh1.0.0
5	95	2.6.1-adh1.0.0
11	89	2.6.1-adh1.0.0
9	91	2.6.1-adh1.0.0
8	92	2.6.1-adh1.0.0



- 对Hadoop及其周边生态进行rpm和deb封装
- 解决各生态之间对hadoop与其他组件的兼容性问题
- 我的工作，是对patch后的Hadoop源码及生态组件进行重新封装发行

# You should know before start with bigtop



- Redhat Package Manager / Debian打包知识, spec脚本语法
- Shell脚本
- Gradle配置相关
- Maven / Ant编译相关
- 自建HTTP服务器, ivy等
- Puppet / Ruby相关知识
- Over the wall (Testament)

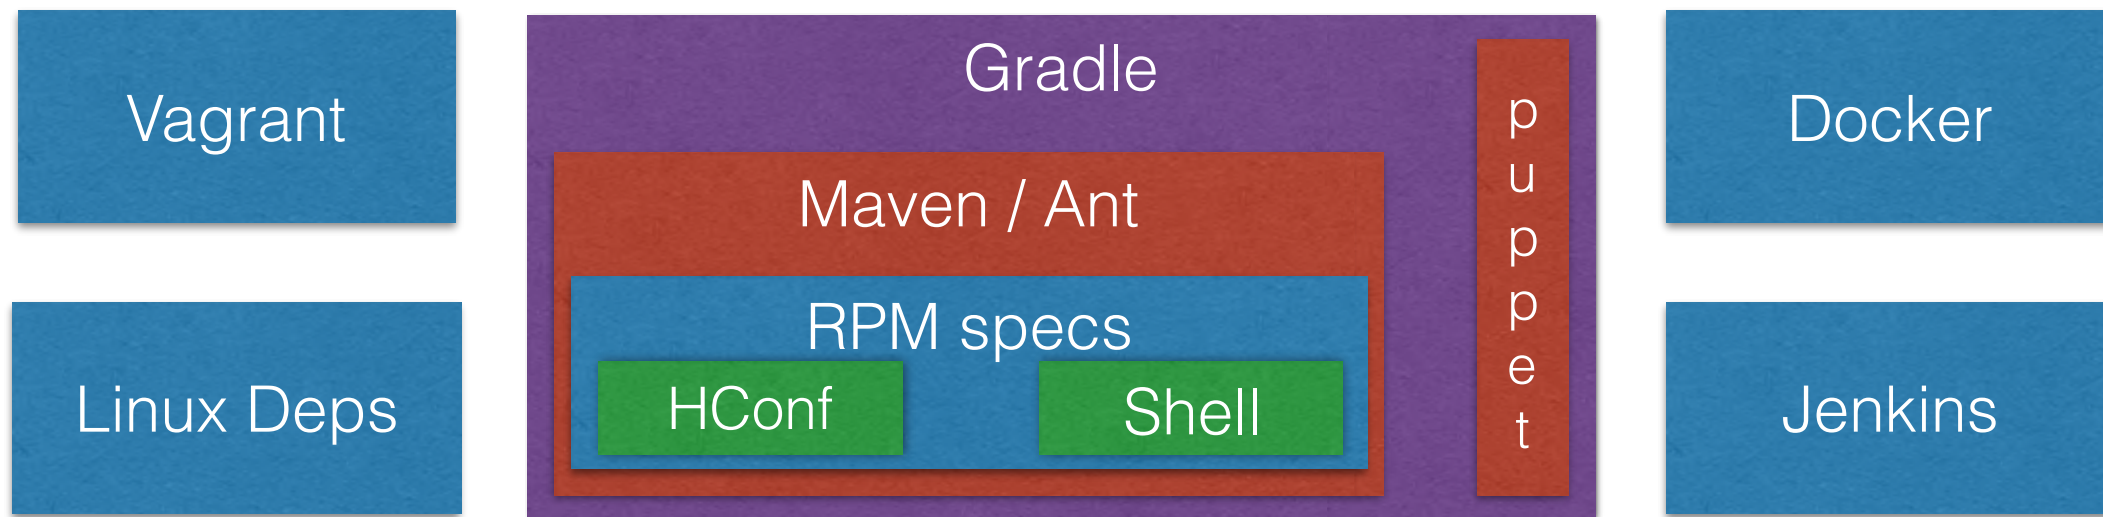


# Other knowledges



- Jenkins持续集成与交付工具，自动化编译发布
- RH/Debian创建仓库
- Linux Patch
- Docker
- make, cmake
- dependencies versions

# 1.1.0 Infrastructure



Bigtop版本演进:

before 1.0: 编译采用make, 部署采用puppet

after 1.0: 编译采用gradle, 解决依赖采用puppet

# Hadoop eco compile dependencies



- ruby, puppet, ant, maven, gradle, forrest, scala, groovy, jenkins, jdk for compile environment
- protobuf, thrift, openldap, fuse, yaml, asciidoc, xmlto...
- vagrant for deployment
- linux busybox, boostlib
- Deps version is so important

# Bigtop vs. CDH package



	Apache Bigtop	CDH-package	Solve
compile	Gradle	Make	fusion
hadoop/ hbase/hive/ pig	N/A	0.20-mr/ha	cdh spec用于Apache bigtop, 代码融合, 自行编写安装脚本
avro/sentry/ parquet/ format/llama/ impala	N/A	YES	自行编写bigtop.bom及安装脚本
pom/build	apache	cdh	adh

# diff Bigtop.spec ADH.spec



```
24 %define hadoop_name hadoop
25 %define etc_hadoop /etc/%{name}
26 %define etc_yarn /etc/yarn
27 %define etc_https /etc/%{name}-https
28 %define config_hadoop %{etc_hadoop}/conf
29 %define config_yarn %{etc_yarn}/conf
30 %define config_https %{etc_https}/conf
31 %define tomcat_deployment_https %{etc_https}/tomcat-conf
32 %define lib_hadoop_dirname /usr/lib
33 %define lib_hadoop %{lib_hadoop_dirname}/%{name}
34 %define lib_https %{lib_hadoop_dirname}/%{name}-https
35 %define lib_hdfs %{lib_hadoop_dirname}/%{name}-hdfs
36 %define lib_yarn %{lib_hadoop_dirname}/%{name}-yarn
37 %define lib_mapreduce %{lib_hadoop_dirname}/%{name}-mapreduce
38 %define log_hadoop_dirname /var/log
39 %define log_hadoop %{log_hadoop_dirname}/%{name}
40 %define log_yarn %{log_hadoop_dirname}/%{name}-yarn
41 %define log_hdfs %{log_hadoop_dirname}/%{name}-hdfs
42 %define log_https %{log_hadoop_dirname}/%{name}-https
43 %define log_mapreduce %{log_hadoop_dirname}/%{name}-mapreduce
44 %define run_hadoop_dirname /var/run
45 %define run_hadoop %{run_hadoop_dirname}/hadoop
46 %define run_yarn %{run_hadoop_dirname}/%{name}-yarn
47 %define run_hdfs %{run_hadoop_dirname}/%{name}-hdfs
48 %define run_https %{run_hadoop_dirname}/%{name}-https
49 %define run_mapreduce %{run_hadoop_dirname}/%{name}-mapreduce
50 %define state_hadoop_dirname /var/lib
51 %define state_hadoop %{state_hadoop_dirname}/hadoop
52 %define state_yarn %{state_hadoop_dirname}/%{name}-yarn
53 %define state_hdfs %{state_hadoop_dirname}/%{name}-hdfs
54 %define state_mapreduce %{state_hadoop_dirname}/%{name}-mapreduce
55 %define state_https %{state_hadoop_dirname}/%{name}-https
56 %define bin_hadoop %{_bindir}
57 %define man_hadoop %{_mandir}
58 %define doc_hadoop %{_docdir}/%{name}-%{hadoop_version}
59 %define https_services https
60 %define mapreduce_services mapreduce-historyserver
61 %define hdfs_services hdfs-namenode hdfs-secondarynamenode hdfs-datanode hdfs-zkfc hdfs-journalnode
62 %define yarn_services yarn-resourcemanager yarn-nodemanager yarn-proxyserver yarn-timelineserver
63 %define hadoop_services %{hdfs_services} %{mapreduce_services} %{yarn_services} %{https_services}
64 # Hadoop outputs built binaries into %{hadoop_build}
65 %define hadoop_build_path build
66 %define static_images_dir src/webapps/static/images
67 %define libexecdir /usr/lib
```

← Apache

```
31 %define hadoop_name hadoop
32 %define etc_hadoop /etc/%{name}
33 %define etc_yarn /etc/yarn
34 %define etc_https /etc/%{name}-https
35 %define etc_kms /etc/%{name}-kms
36 %define config_hadoop %{etc_hadoop}/conf
37 %define config_yarn %{etc_yarn}/conf
38 %define config_https %{etc_https}/conf
39 %define config_kms %{etc_kms}/conf
40 %define tomcat_deployment_https %{etc_https}/tomcat-conf
41 %define tomcat_deployment_kms %{etc_kms}/tomcat-conf
42 %define lib_hadoop_dirname /usr/lib
43 %define lib_hadoop %{lib_hadoop_dirname}/%{name}
44 %define lib_https %{lib_hadoop_dirname}/%{name}-https
45 %define lib_kms %{lib_hadoop_dirname}/%{name}-kms
46 %define lib_hdfs %{lib_hadoop_dirname}/%{name}-hdfs
47 %define lib_yarn %{lib_hadoop_dirname}/%{name}-yarn
48 %define lib_mapreduce %{lib_hadoop_dirname}/%{name}-mapreduce
49 %define lib_mapreduce_mr1 %{lib_hadoop_dirname}/%{name}-0.20-mapreduce
50 %define log_hadoop_dirname /var/log
51 %define log_hadoop %{log_hadoop_dirname}/%{name}
52 %define log_yarn %{log_hadoop_dirname}/%{name}-yarn
53 %define log_hdfs %{log_hadoop_dirname}/%{name}-hdfs
54 %define log_https %{log_hadoop_dirname}/%{name}-https
55 %define log_kms %{log_hadoop_dirname}/%{name}-kms
56 %define log_mapreduce %{log_hadoop_dirname}/%{name}-mapreduce
57 %define run_hadoop_dirname /var/run
58 %define run_hadoop %{run_hadoop_dirname}/hadoop
59 %define run_yarn %{run_hadoop_dirname}/%{name}-yarn
60 %define run_hdfs %{run_hadoop_dirname}/%{name}-hdfs
61 %define run_https %{run_hadoop_dirname}/%{name}-https
62 %define run_kms %{run_hadoop_dirname}/%{name}-https
63 %define run_mapreduce %{run_hadoop_dirname}/%{name}-mapreduce
64 %define state_hadoop_dirname /var/lib
65 %define state_hadoop %{state_hadoop_dirname}/hadoop
66 %define state_yarn %{state_hadoop_dirname}/%{name}-yarn
67 %define state_hdfs %{state_hadoop_dirname}/%{name}-hdfs
68 %define state_mapreduce %{state_hadoop_dirname}/%{name}-mapreduce
69 %define state_https %{state_hadoop_dirname}/%{name}-https
70 %define state_kms %{state_hadoop_dirname}/%{name}-kms
71 %define bin_hadoop %{_bindir}
72 %define man_hadoop %{_mandir}
73 %define doc_hadoop %{_docdir}/%{name}-%{hadoop_version}
74 %define doc_hadoop_mr1 %{_docdir}/hadoop-0.20-mapreduce
75 %define https_services https
```

ADH ↑

# diff Bigtop.sh ADH.sh




```
604
605 internal_versionless_symlinks \
606     ${HADOOP_DIR}/hadoop-*.jar \
607     ${HDFS_DIR}/hadoop-*.jar \
608     ${YARN_DIR}/hadoop-*.jar \
609     ${MAPREDUCE_DIR}/hadoop-*.jar \
610     ${MAPREDUCE_MR1_DIR}/hadoop-*.jar \
611     ${MAPREDUCE_MR1_DIR}/contrib/**/hadoop-*.jar \
612     ${PREFIX}/usr/lib/hadoop-kms/webapps/kms/WEB-INF/lib/hadoop-kms*.jar
613
614 # Creating an explicit hadoop-examples.jar symlink in MR1 for backwards compatibility
615 ln -s hadoop-examples-mr1.jar ${MAPREDUCE_MR1_DIR}/hadoop-examples.jar
616
617 external_versionless_symlinks 'hadoop' \
618     ${PREFIX}/usr/lib/hadoop-0.20-mapreduce \
619     ${PREFIX}/usr/lib/hadoop-0.20-mapreduce/lib \
620     ${PREFIX}/usr/lib/hadoop-mapreduce \
621     ${PREFIX}/usr/lib/hadoop-mapreduce/lib \
622     ${PREFIX}/usr/lib/hadoop-yarn/lib \
623     ${PREFIX}/usr/lib/hadoop/client \
624     ${PREFIX}/usr/lib/hadoop/client-0.20 \
625     ${PREFIX}/usr/lib/hadoop/lib
626
627 external_versionless_symlinks 'hadoop-https' \
628     ${PREFIX}/usr/lib/hadoop-https/webapps/webhdfs/WEB-INF/lib
629
630 external_versionless_symlinks 'hadoop-kms' \
631     ${PREFIX}/usr/lib/hadoop-kms/webapps/kms/WEB-INF/lib
632
633 # For adding extra native libs required by libhadoop.so (e.g. newer version of libcrypto.so) at run time
634 install -d -m 0755 ${PREFIX}/var/lib/hadoop/extra/native
635
```

ADH



```
381
382 # Create log, var and lib
383 install -d -m 0755 $PREFIX/var/{log,run,lib}/hadoop-hdfs
384 install -d -m 0755 $PREFIX/var/{log,run,lib}/hadoop-yarn
385 install -d -m 0755 $PREFIX/var/{log,run,lib}/hadoop-mapreduce
386
387 # Remove all source and create version-less symlinks to offer integration point with other projects
388 for DIR in ${HADOOP_DIR} ${HDFS_DIR} ${YARN_DIR} ${MAPREDUCE_DIR} ${HTTPFS_DIR} ; do
389     (cd $DIR &&
390         rm -fv *-sources.jar
391         rm -fv lib/hadoop-*.jar
392         for j in hadoop-*.jar; do
393             if [[ $j =~ hadoop-(.*)-${HADOOP_VERSION}.jar ]]; then
394                 name=${BASH_REMATCH[1]}
395                 ln -s $j hadoop-$name.jar
396             fi
397         done)
398 done
399
400 # Now create a client installation area full of symlinks
401 install -d -m 0755 ${CLIENT_DIR}
402 for file in `cat ${BUILD_DIR}/hadoop-client.list` ; do
403     for dir in ${HADOOP_DIR}/{lib,} ${HDFS_DIR}/{lib,} ${YARN_DIR}/{lib,} ${MAPREDUCE_DIR}/{lib,} ; do
404         [ -e $dir/$file ] && \
405             ln -fs ${dir#PREFIX}/$file ${CLIENT_DIR}/${file} && \
406             ln -fs ${dir#PREFIX}/$file ${CLIENT_DIR}/${file/-[:digit:]*}.jar && \
407             continue 2
408     done
409     exit 1
410 done
```

Apache



# Our rpm examples



```
hadoop-0.20-conf-pseudo-2.6.1+adh1.0.0+1-1.adh1.0.0.p0.7.el6.x86_64.rpm  
hadoop-0.20-mapreduce-2.6.1+adh1.0.0+1-1.adh1.0.0.p0.7.el6.x86_64.rpm  
hadoop-0.20-mapreduce-jobtracker-2.6.1+adh1.0.0+1-1.adh1.0.0.p0.7.el6.x86_64.rpm  
hadoop-0.20-mapreduce-jobtrackerha-2.6.1+adh1.0.0+1-1.adh1.0.0.p0.7.el6.x86_64.rpm  
hadoop-0.20-mapreduce-tasktracker-2.6.1+adh1.0.0+1-1.adh1.0.0.p0.7.el6.x86_64.rpm  
hadoop-2.6.1+adh1.0.0+1-1.adh1.0.0.p0.7.el6.x86_64.rpm  
hadoop-client-2.6.1+adh1.0.0+1-1.adh1.0.0.p0.7.el6.x86_64.rpm  
hadoop-conf-pseudo-2.6.1+adh1.0.0+1-1.adh1.0.0.p0.7.el6.x86_64.rpm  
hadoop-debuginfo-2.6.1+adh1.0.0+1-1.adh1.0.0.p0.7.el6.x86_64.rpm  
hadoop-doc-2.6.1+adh1.0.0+1-1.adh1.0.0.p0.7.el6.x86_64.rpm  
hadoop-hdfs-2.6.1+adh1.0.0+1-1.adh1.0.0.p0.7.el6.x86_64.rpm  
hadoop-hdfs-datanode-2.6.1+adh1.0.0+1-1.adh1.0.0.p0.7.el6.x86_64.rpm  
hadoop-hdfs-fuse-2.6.1+adh1.0.0+1-1.adh1.0.0.p0.7.el6.x86_64.rpm  
hadoop-hdfs-journalnode-2.6.1+adh1.0.0+1-1.adh1.0.0.p0.7.el6.x86_64.rpm  
hadoop-hdfs-namenode-2.6.1+adh1.0.0+1-1.adh1.0.0.p0.7.el6.x86_64.rpm  
hadoop-hdfs-nfs3-2.6.1+adh1.0.0+1-1.adh1.0.0.p0.7.el6.x86_64.rpm  
hadoop-hdfs-secondarynamenode-2.6.1+adh1.0.0+1-1.adh1.0.0.p0.7.el6.x86_64.rpm  
hadoop-hdfs-zkfc-2.6.1+adh1.0.0+1-1.adh1.0.0.p0.7.el6.x86_64.rpm  
hadoop-httplib-2.6.1+adh1.0.0+1-1.adh1.0.0.p0.7.el6.x86_64.rpm  
hadoop-kms-2.6.1+adh1.0.0+1-1.adh1.0.0.p0.7.el6.x86_64.rpm  
hadoop-kms-server-2.6.1+adh1.0.0+1-1.adh1.0.0.p0.7.el6.x86_64.rpm  
hadoop-libhdfs-2.6.1+adh1.0.0+1-1.adh1.0.0.p0.7.el6.x86_64.rpm  
hadoop-libhdfs-devel-2.6.1+adh1.0.0+1-1.adh1.0.0.p0.7.el6.x86_64.rpm  
hadoop-mapreduce-2.6.1+adh1.0.0+1-1.adh1.0.0.p0.7.el6.x86_64.rpm  
hadoop-mapreduce-historyserver-2.6.1+adh1.0.0+1-1.adh1.0.0.p0.7.el6.x86_64.rpm  
hadoop-yarn-2.6.1+adh1.0.0+1-1.adh1.0.0.p0.7.el6.x86_64.rpm  
hadoop-yarn-nodemanager-2.6.1+adh1.0.0+1-1.adh1.0.0.p0.7.el6.x86_64.rpm  
hadoop-yarn-proxyserver-2.6.1+adh1.0.0+1-1.adh1.0.0.p0.7.el6.x86_64.rpm  
hadoop-yarn-resourcemanager-2.6.1+adh1.0.0+1-1.adh1.0.0.p0.7.el6.x86_64.rpm
```

```
hive-1.1.0+adh1.0.0+1-1.adh1.0.0.p0.0.el6.noarch.rpm  
hive-hbase-1.1.0+adh1.0.0+1-1.adh1.0.0.p0.0.el6.noarch.rpm  
hive-hcatalog-1.1.0+adh1.0.0+1-1.adh1.0.0.p0.0.el6.noarch.rpm  
hive-jdbc-1.1.0+adh1.0.0+1-1.adh1.0.0.p0.0.el6.noarch.rpm  
hive-metastore-1.1.0+adh1.0.0+1-1.adh1.0.0.p0.0.el6.noarch.rpm  
hive-server-1.1.0+adh1.0.0+1-1.adh1.0.0.p0.0.el6.noarch.rpm  
hive-server2-1.1.0+adh1.0.0+1-1.adh1.0.0.p0.0.el6.noarch.rpm  
hive-webhcat-1.1.0+adh1.0.0+1-1.adh1.0.0.p0.0.el6.noarch.rpm  
hive-webhcat-server-1.1.0+adh1.0.0+1-1.adh1.0.0.p0.0.el6.noarch.rpm
```

```
llama-1.0.0+adh1.0.0+1-1.adh1.0.0.p0.0.el6.noarch.rpm  
llama-doc-1.0.0+adh1.0.0+1-1.adh1.0.0.p0.0.el6.noarch.rpm  
llama-master-1.0.0+adh1.0.0+1-1.adh1.0.0.p0.0.el6.noarch.rpm
```

```
hbase-1.0.0+adh1.0.0+1-1.adh1.0.0.p0.0.el6.x86_64.rpm  
hbase-doc-1.0.0+adh1.0.0+1-1.adh1.0.0.p0.0.el6.x86_64.rpm  
hbase-master-1.0.0+adh1.0.0+1-1.adh1.0.0.p0.0.el6.x86_64.rpm  
hbase-regionserver-1.0.0+adh1.0.0+1-1.adh1.0.0.p0.0.el6.x86_64.rpm  
hbase-rest-1.0.0+adh1.0.0+1-1.adh1.0.0.p0.0.el6.x86_64.rpm  
hbase-thrift-1.0.0+adh1.0.0+1-1.adh1.0.0.p0.0.el6.x86_64.rpm
```

```
parquet-1.5.0+adh1.0.0+1-1.adh1.0.0.p0.0.el6.noarch.rpm
```

```
parquet-format-2.1.0+adh1.0.0+1-1.adh1.0.0.p0.0.el6.noarch.rpm
```

```
hbase-1.0.0+adh1.0.0+1-1.adh1.0.0.p0.0.el6.x86_64.rpm  
hbase-doc-1.0.0+adh1.0.0+1-1.adh1.0.0.p0.0.el6.x86_64.rpm  
hbase-master-1.0.0+adh1.0.0+1-1.adh1.0.0.p0.0.el6.x86_64.rpm  
hbase-regionserver-1.0.0+adh1.0.0+1-1.adh1.0.0.p0.0.el6.x86_64.rpm  
hbase-rest-1.0.0+adh1.0.0+1-1.adh1.0.0.p0.0.el6.x86_64.rpm  
hbase-thrift-1.0.0+adh1.0.0+1-1.adh1.0.0.p0.0.el6.x86_64.rpm
```

## Why parquet/avro?



- 集群原使用cdh，新版本基于Hadoop patched re-distribution
- cdh 发行版 hadoop安装依赖parquet, avro等，否则yum update会报错
- 这部分在apache bigtop里是没有的，需要自己写一部分，而且包名要对应依赖的包名



## Key points



- cdh package基于bigtop 0.6.0, 除spec外不更新其他
- apache bigtop基于apache hadoop, 不包含cdh特有项目
- 融合cdh package和现版bigtop是一个比较有挑战的工程
- 理解apache bigtop架构, 基于apache 脚本重构仿cdh编译安装脚本
- Key of all keys: Over the GreatFireWall





# Thank You

---

向磊 : [xianglei@admaster.com.cn](mailto:xianglei@admaster.com.cn)  
Wechat: 880391

 [www.admaster.com.cn](http://www.admaster.com.cn)

 [Marketing@admaster.com.cn](mailto:Marketing@admaster.com.cn)

 400-8062-600

任何对本数据及报告的使用不得违反任何法律法规或侵犯任何第三方合法权益，任何场合下的转述或引述均需标明本数据及报告来源于AdMaster，非经AdMaster书面同意，不得将本数据及报告转卖/提供给任何第三方（除最终客户）。  
The use of data and report by AdMaster shall not violate any laws and regulations and infringe legitimate rights and interests of any third party. Any recital and quoting of such data and report shall be expressed the source is derived from AdMaster. Unless written consent by AdMaster, such data and report under this Agreement cannot be resell/provided to any third party (except the end client).