



Building Cutting Edge Big Data Platform with Apache Bigtop

Big Data Innovation Summit

PRESENTED BY Evans Ye | December 5, 2016

Who am I

- Software Engineer @ Yahoo! APAC Data Team
- Building personalized data products for...



- Apache Bigtop PMC member



Outline

- Quick Intro to Apache Bigtop
- Bigtop Provisioner
- Bigtop Sandbox
- Big Data Landscape
- Open Source Adoption Strategy
- Release Timeline

Quick Intro to Apache Bigtop

Linux Distributions



CentOS



ubuntu
linux for human beings

fedora 




debian




Hadoop Distributions






There're some other great
Hadoop ecosystem components..



How do I add patches?



From source code to packages



Bigtop feature set

Packaging



Testing



Deployment



Virtualization



for you to easily build your own Big Data Stack

Supported components



Configuration Matrix	centos-6	centos-7	fedora-20	ubuntu-14.04	ubuntu-16.04	debian-8	opensuse-42.1
apex	●	●	●	●	●	●	●
alluxio	●	●	●	●	●	●	●
bigtop-groovy	●	●	●	●	●	●	●
bigtop-jsvc	●	●	●	●	●	●	●
bigtop-tomcat	●	●	●	●	●	●	●
bigtop-utils	●	●	●	●	●	●	●
crunch	●	●	●	●	●	●	●
datafu	●	●	●	●	●	●	●
flume	●	●	●	●	●	●	●
flink	●	●	●	●	●	●	●
giraph	●	●	●	●	●	●	●
gpdb	●	●	●	●	●	●	●
hadoop	●	●	●	●	●	●	●
hama	●	●	●	●	●	●	●
hbase	●	●	●	●	●	●	●
hive	●	●	●	●	●	●	●
hue	●	●	●	●	●	●	●
iglu	●	●	●	●	●	●	●
karbon	●	●	●	●	●	●	●
kite	●	●	●	●	●	●	●
mahout	●	●	●	●	●	●	●
oozie	●	●	●	●	●	●	●
phoenix	●	●	●	●	●	●	●
pig	●	●	●	●	●	●	●
qfs	●	●	●	●	●	●	●
solr	●	●	●	●	●	●	●
spark1	●	●	●	●	●	●	●
spark	●	●	●	●	●	●	●
sqoop	●	●	●	●	●	●	●
sqoop2	●	●	●	●	●	●	●
tajo	●	●	●	●	●	●	●
tez	●	●	●	●	●	●	●
ycsb	●	●	●	●	●	●	●
zookeeper	●	●	●	●	●	●	●
zeppelin	●	●	●	●	●	●	●

Address dependency Issues

Bigtop early mission accomplished



All major Hadoop distros leverage Bigtop
to build its foundation

Leveraged by app providers...



What now?

Get out from the Apache dome

New focus and target end users

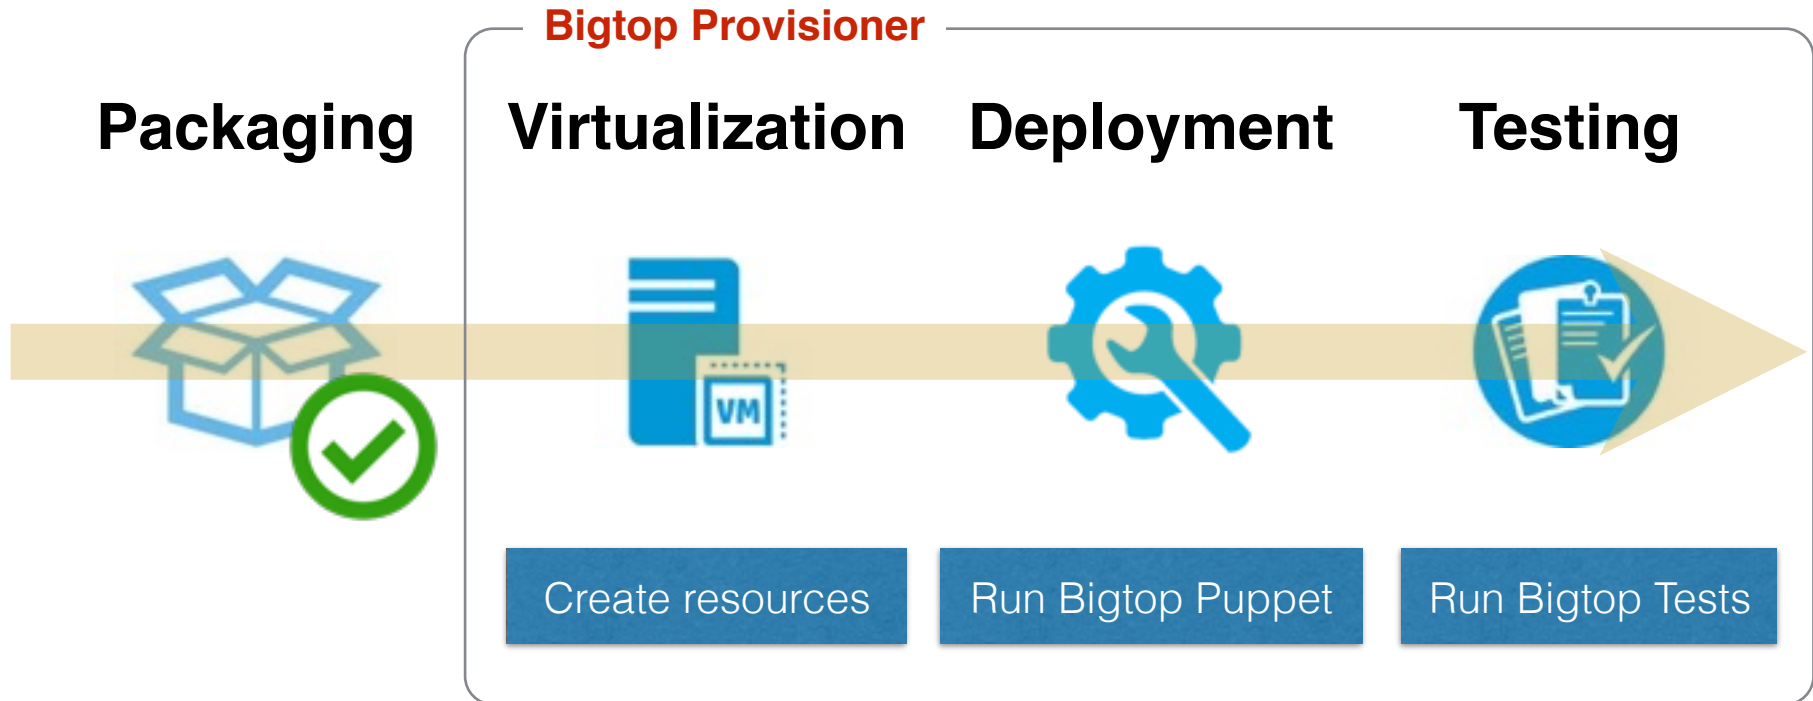
- Data engineers vs Distro. builders
- Reference implementation & stacks
- Solution diversity:
 - Streaming: Flink, Apex
 - In-memory cache: Alluxio, Ignite
- User/developer tools:
 - Bigtop Provisioner
 - Bigtop Sandbox



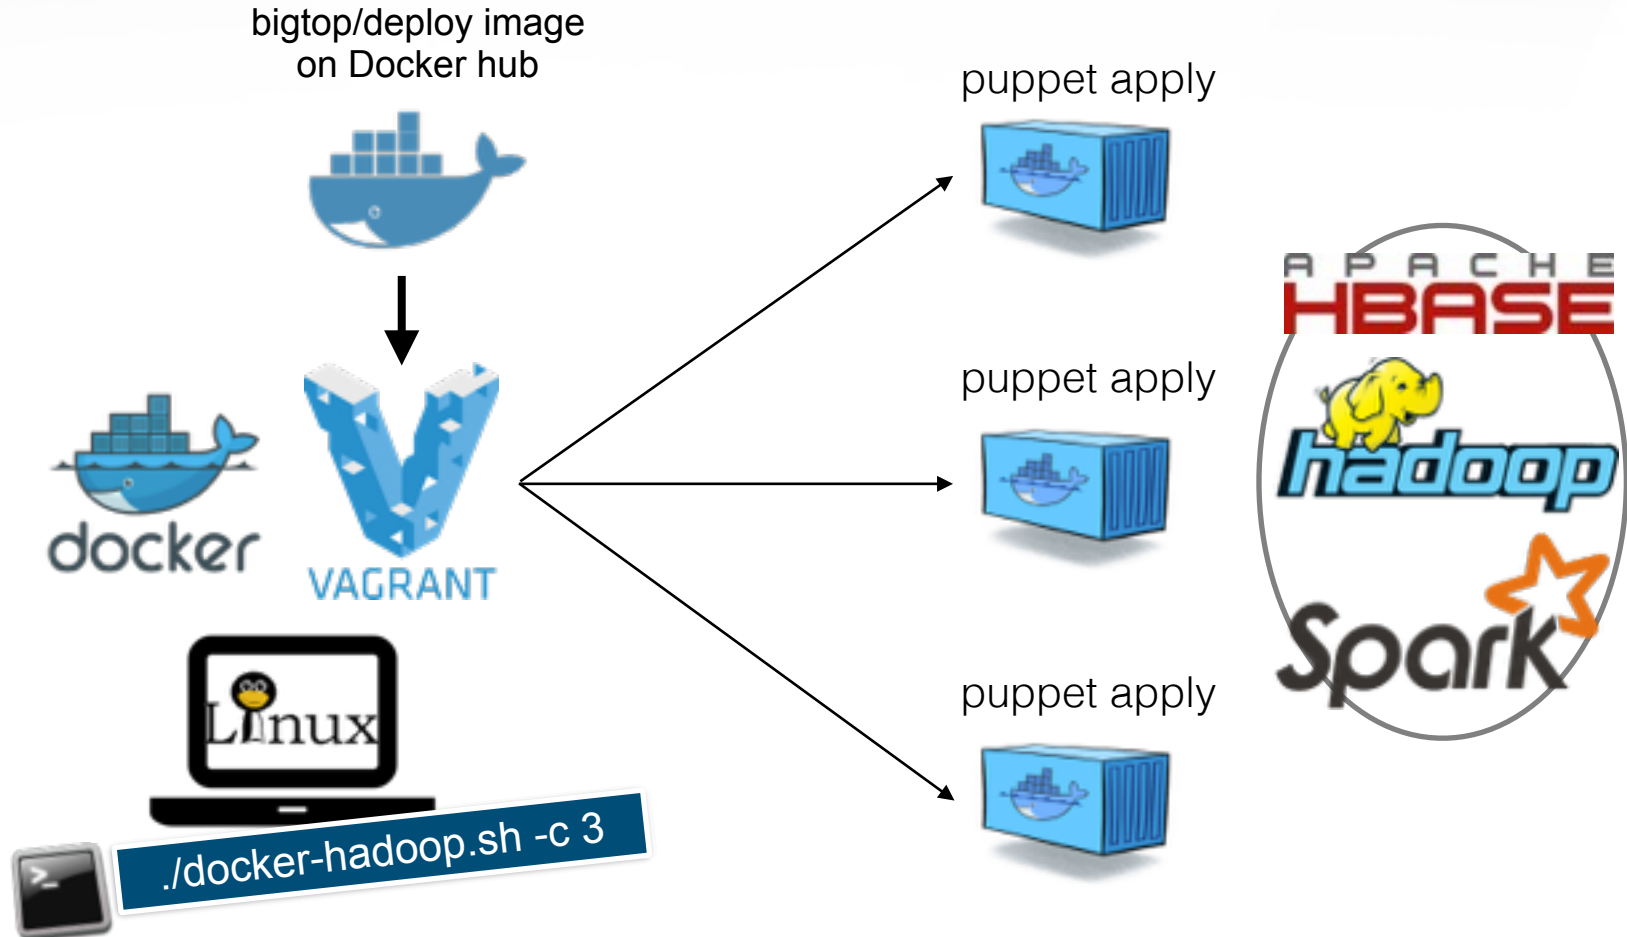
Bigtop Provisioner

Bigtop Provisioner

- A tool to demonstrate full life cycle of Bigtop



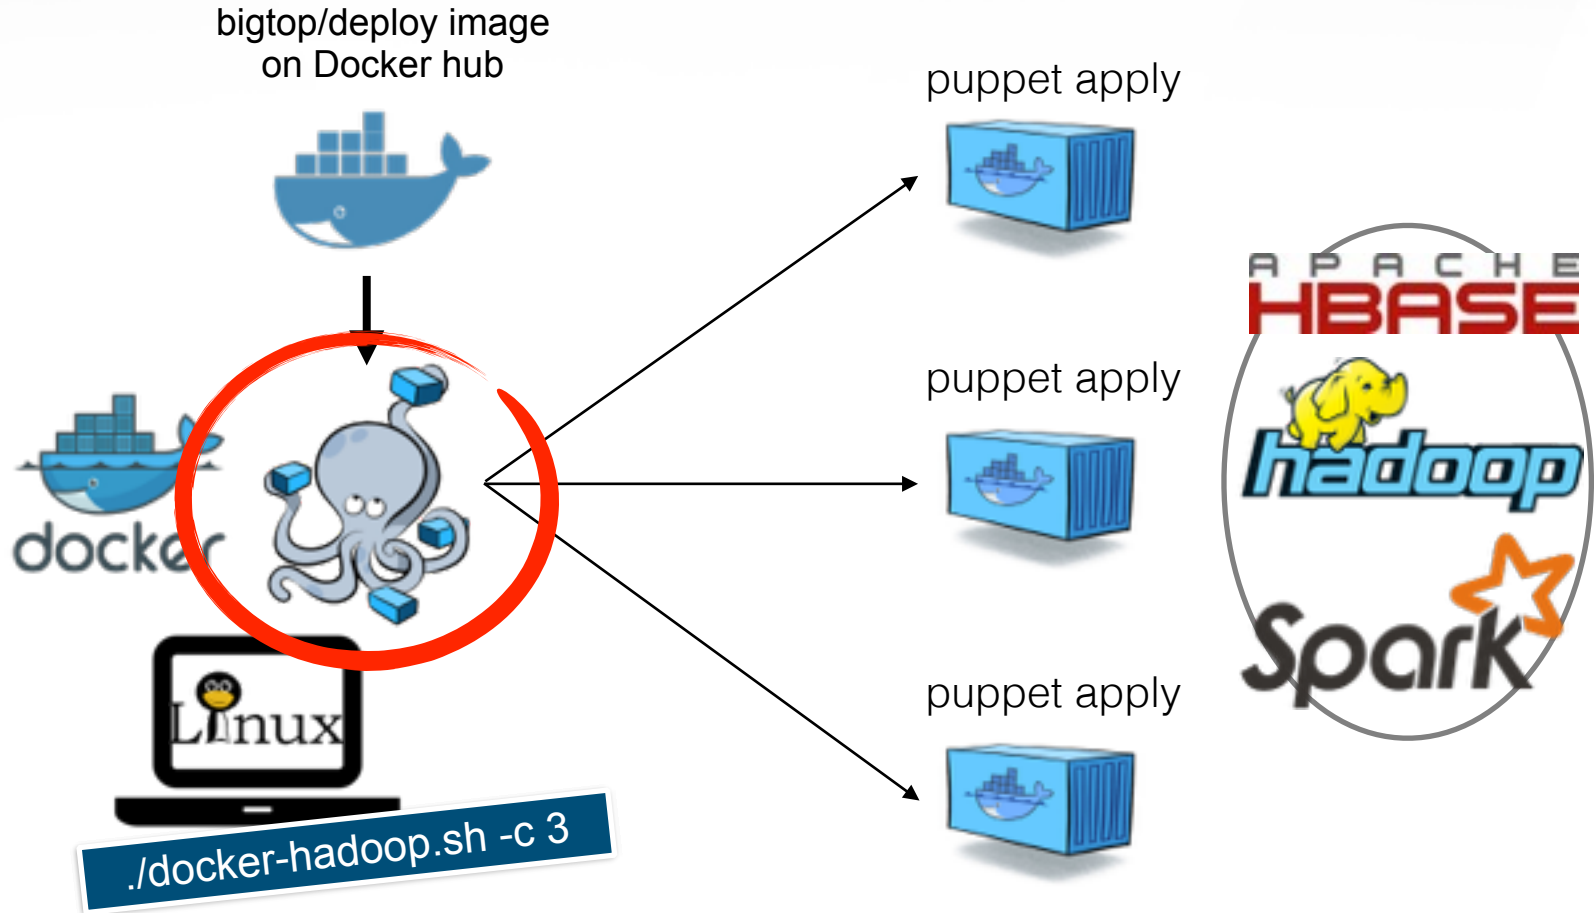
One click Hadoop provisioning (Bigtop 1.0.0)



What's the problem with Vagrant's Docker Provider?

- Need to add vagrant public key into docker images
- Too many issues with auto-created boot2docker VM
- A bug for docker provider keep opening for almost 2y
 - Waiting for machine to boot' hangs infinitely
- Can not share same code for different providers anyway
- Not all the docker options supported in Vagrantfile
- ^#?& slow

Impl. replaced by docker-compose (1.2.0-SNAPSHOT)



Advantages

- No need to create customized image beforehand
- Better compatibility with Docker's native solutions
- Clear, simple yaml file for orchestration settings
- Supports new features such as overlay network and named volume
- Fast —> better user experience

Bigtop Sandbox

Introducing Bigtop Docker Sandbox

- Docker images that has Bigtop stacks installed and configured
- Pseudo cluster up & running w/ zero installation/configuration
- Command-line tool to build your own stack

Docker Image layer Interface



Customized big data stack

Deploy & management tool

Base image (OS)

Docker Image layer

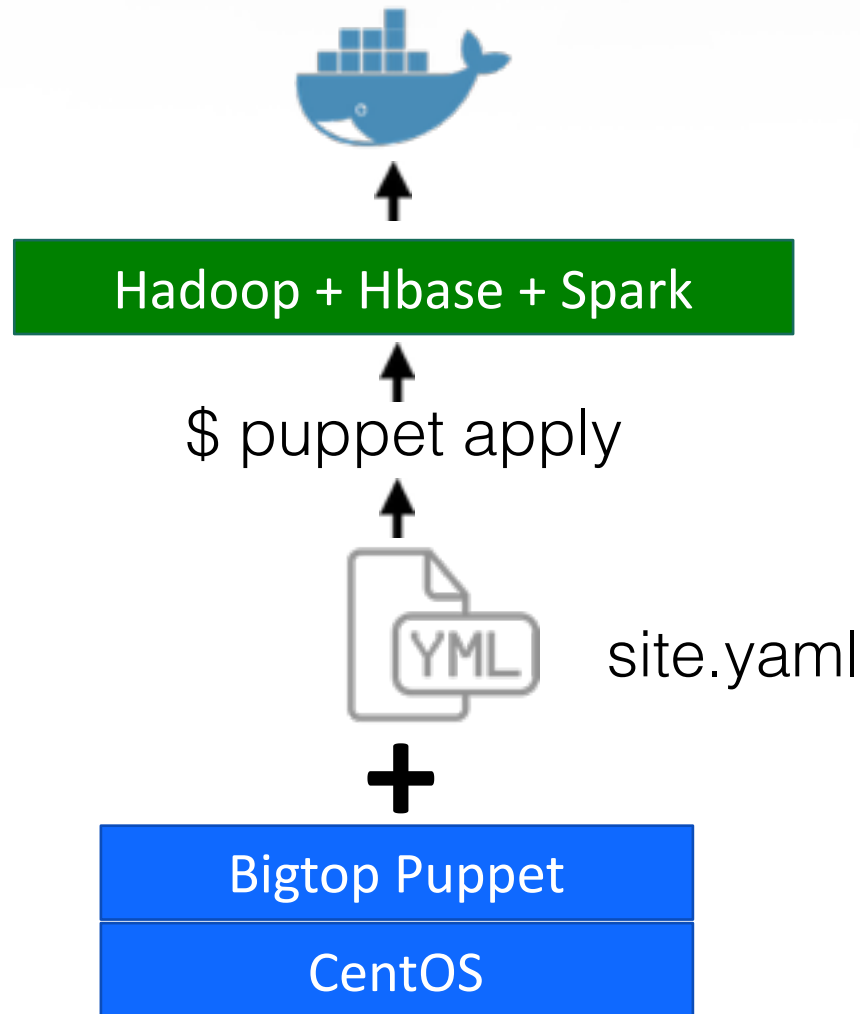
Concrete implementation

Hadoop + HBase + Spark

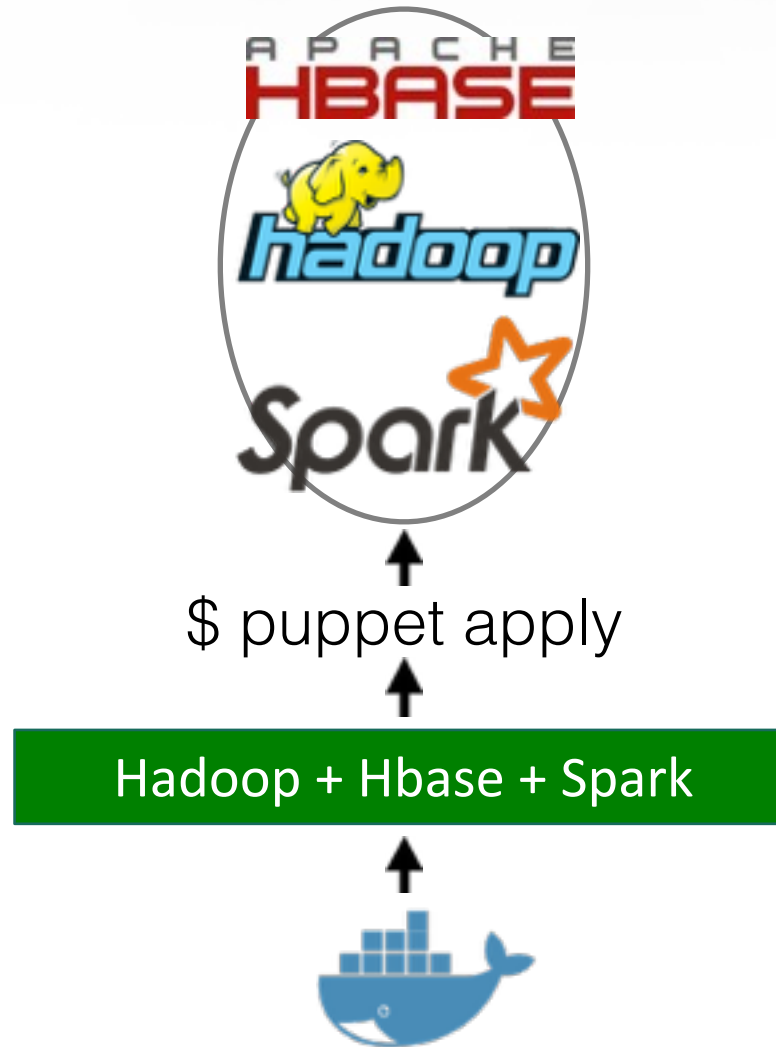
Bigtop Puppet

CentOS

Building images



Running images



Support integration tests in CI/CD

	Bigtop Provisioner	Bigtop Sandbox
Data engineers	Create multi-node cluster for testing	Build/use sandboxes for dev/test
Ops	Create multi-node cluster for testing	-
Contributors	Test packages, puppet recipes, test cases	-
Distro. Builders	Test packages, puppet recipes, test cases	Provide Sandboxes



Big Data Landscape



Hot topics in Apache conferences

- Spark dominates the big data world and the research area
 - 26/27 talks with Spark on the title
- Streaming is still a hot topic
 - 8/8 talks with steaming on the title

Innovations

- New features and major version upgrade in key projects
 - Spark 2.0
 - Hadoop 3.0
 - HBase 2.0
 - Cassandra 3.X
 - Kafka Streams & Connect
- Still, there're so many new projects
 - Helix, Calcite, Unomi, Samoa, Kudu, Streams, OODT, Tinkerpop, Kerby, Yetus, HTrace, REEF, Bahir, ...

Apache HBase

Region Replicas for better Availability

- Multiple Region Servers host each region
 - Primary + N read replicas (usually N=2)
 - Primary is authority on writes
 - Replicas tail replicate edits, offer TIMELINE view
- Client's choice
 - Read primary only for "classic" strong consistency
 - Fan-out reads for faster, potentially TIMELINE results
- Ref:
 - [Apache HBase: State of the Database, Nick Dimiduk](#)

Off-heap Caching

- On-heap LRU cache limited by Java heap size
- Off-heap bucket cache - solution for GC pause
- Serve data directly from off-heap cache since cells are backed by byte arrays
- Ref:
 - [Recent Development in HBase, Zhihong Yu](#)

Apache Cassandra

SStable Attached Secondary Index (SASI)

- Attach B+ tree on SStable
- Range scan on 2nd index made possible
- Ref:
 - [Cassandra 3.4 and Beyond, Jon Haddad](#)

Apache Kafka

Kafka Connect

- Focus on copying
- Standardize
- Parallelism
- Scale
- Ref:
 - [Kafka Connect: Real-time Data Integration at Scale with Apache Kafka, Ewen Cheslack-Postava](#)

Kafka Streams

- Available in Kafka since 0.10, May 20, 2016
- Powerful yet easy-to-use stream processing library
 - Event-at-a-time, Stateful
 - Windowing with out-of-order handling
 - Highly scalable, distributed, fault tolerant
- Ref:
 - [Introduction to Kafka Streams, Guozhang Wang](#)

Summary

- HBase
 - region replicas in 1.2, off-heap read in 2.0
- Cassandra
 - SASI in 3.4
- Kafka
 - Kafka Connectors available on Confluent
 - Kafka Streams in Kafka 0.10

Open Source Adoption Strategy

3 key aspects

- Functionality
- Community
- Maturity

Functionality

- Survey! Only you knows your requirement
 - Requires an architect /senior data engineer who has deep understanding as well as broader view
- Streaming for instance:
 - Exactly-once
 - Out-of-order handling
 - CEP
 - <1s SLA

Community

- How many contributors? (see on github)
- Employers of committers?
- When's the last release?
- What's the last code updated time?
- Subscribe to the mailing list and see how many discussions daily/weekly

Maturity

- Who already on-boarded?
- Founded by major big data players?
 - Cloudera, Hortonworks, MapR, Pivotal, AMPLab
- Donated by big tech companies?
 - Yahoo!, Facebook, Twitter
- Previously a commercial software?
 - For example: Apex, Ignite, Geode
- Backed by a company?
 - Databricks, dataArtisans, Confluent, DataTorrent, Gridgain

Release Timeline

New components Bigtop 1.1.0

- Released Feb, 2016:
 - Zeppelin 0.5.6
 - Tachyon 0.6.0
 - Hama 0.7.0



TACHYON



New Arch in Bigtop 1.1.0



- Ported 22 out of 24 Apache Bigtop stacks to POWER in two week
- Apache Bigtop has Dockerized the entire build environment

What's coming in 1.2 release?

Expected to be out in Feb, 2017

What's coming in 1.2 release?

- Featured upgrade:
 - Hadoop 2.7.3
 - Spark 2.0.2
 - Kafka 0.10.0
 - HBase 1.2.4

What's coming in 1.2 release?

- New components:
 - Apache Flink 1.1.3
 - Apache Apex 3.4.0
 - Greenplum Database 4.3.99.0
 - Quantcast File System 1.1.4

What's coming in 1.2 release?

- New features:
 - Juju bigtop charms
 - Bigtop Docker Sandbox
- Improvement:
 - Bigtop Docker Provisioner made faster

Reference

- Home page: <http://bigtop.apache.org/>
- Document: <https://cwiki.apache.org/confluence/display/BIGTOP/Index>
- Source code: <http://www.apache.org/dist/bigtop/bigtop-1.1.0/>
- Packages: <https://www.apache.org/dist/bigtop/bigtop-1.1.0/repos/>
- JIRA: <https://issues.apache.org/jira/browse/BIGTOP>
- Sandbox Preview: <https://youtu.be/yvmZu7Jbtag>





Thank you !

Questions?