

*Creating a Multinode Hadoop
cluster in 4 mins using docker
containers*

Rachit Arora

rachitar@in.ibm.com

IBM, India Software Labs

As a data scientist

- I want to run my analytics jobs
 - Social media analytics
 - Text analytics (Structure and Unstructured)
- I want to run queries on demand
- I want to run R scripts
- I want to submit Spark jobs

- I need a Hadoop cluster on demand to fulfil my jobs
- I do not want to spend time to setup the Hadoop cluster

Can someone give me a Cluster?

I have a Hadoop Cluster

- I do not want to manage it
- I do not want to Patch them for upgrades
- I do not need this cluster running all the time.
- I want to scale and descale the cluster based on job load

Can someone manage this Cluster?

Technical Interpretation of customer needs

- Multi node Hadoop cluster in minutes
- Elasticity – Add or remove data nodes on demand
- Economical
- Fully Managed
 - Choose what to install and keep it running all the time
- Repeatable, scalable & Highly available provisioning infra
- Minimize disruptions during patching
- Support for Service composition
- Auto Heal Services

Typical Hadoop Cluster - Setup

- Get the suitable hardware
- Prepare host machine
- Setup various networks
 - Private
 - Public
 - Management
- Fetch the binaries for the install
- Prepare the blueprint/config file for the install
- Start the install
- Many a times install fails, debug and retry again.

Bluemix

IBM's Open Cloud Architecture implementation based on the Cloud Foundry project

Bluemix is an open-standard, cloud-based platform for building, managing, and running applications of all types (web, mobile, big data, new smart devices, and so on).



Instant Environments

The developer can choose any language runtime or bring your own. Zero to production in one command.

APIs and Services

A catalog of IBM, third party, and open source API services allow the developer to stitch an application together in minutes.

On-Prem Integration

Build hybrid environments. Connect to on-premise assets plus other public and private clouds.

DevOps

Development, monitoring, deployment, and logging tools allow the developer to run the entire application.

Layered Security

IBM secures the platform and infrastructure and provides you with the tools to secure your apps.

Flexible Pricing

Sign up in minutes. Pay as you go and subscription models offer choice and flexibility.

Bluemix Services









Runtimes

Run an app in the language of your choice

 Liberty for Java™ IBM	 SDK for Node.js™ IBM	 XPages IBM	 ASP.NET Core Community <i>BETA</i>	 Go Community
 PHP Community	 Python Community	 Ruby Community	 Swift Community <i>BETA</i>	 Tomcat Community

Mobile

Quickly get started with your next app

 Mobile Application Content Manager IBM	 Mobile Client Access IBM	 Mobile Foundation IBM	 Mobile Quality Assurance IBM	 Push Notifications IBM
 Kinetise Third Party	 Testdroid Cloud Third Party	 Twilio Third Party		

Watson

Build cognitive apps that help enhance, scale, and accelerate human expertise

 AlchemyAPI IBM	 Concept Insights IBM	 Conversation IBM	 Dialog IBM	 Document Conversion IBM
 Language Translation IBM	 Natural Language Classifier IBM	 Personality Insights IBM	 Retrieve and Rank IBM	 Speech to Text IBM
 Text to Speech IBM	 Tone Analyzer IBM	 Tradeoff Analytics IBM	 Visual Recognition IBM	 Cognitive Commerce™ Third Party

Data and Analytics

Essential data services; limitless possibilities

HELP ME PICK

 Apache Spark IBM	 BigInsights for Apache Hadoop IBM	 Cloudbant NoSQL DB IBM	 dashDB IBM	 DataWorks IBM
 Elasticsearch by Compose IBM	 Geospatial Analytics IBM	 IBM DataStage on Cloud IBM	 IBM DB2 on Cloud IBM	 IBM Graph IBM
 IBM Master Data Management on Cloud	 Insights for Twitter IBM	 MongoDB by Compose IBM	 PostgreSQL by Compose	 Predictive Analytics IBM

BigInsights on cloud (Basic)

IBM Bluemix Ready? Try the new Bluemix | New! Try OpenWhisk

DASHBOARD SOLUTIONS CATALOG PRICING DOCS COMMUNITY 349

Back to All Categories

BigInsights for Apache Hadoop

PUBLISH DATE: 07/28/2016
AUTHOR: IBM
TYPE: Service
LOCATION: US South
[VIEW DOCS](#)

Develop analytics applications by using open source Apache Hadoop and BigInsights™ APIs without having to manage the platform. The service is managed and scaled for you by the Big Data experts at IBM®.

- Power of open source innovation**
This service provides 100% open source Apache Hadoop through the IBM Open Platform. It is updated regularly to provide the best versions of the Apache Hadoop ecosystem components, including, among others, Ambari, YARN, Spark, Knox, HBase, Hive, and an encrypted HDFS.
- Focus on delivering insights**
Key components of the platform, including the infrastructure, are proactively monitored by a 24x7 cloud operations team. Critical security fixes, high severity patches, and defect resolutions are promptly applied on the cluster. Take advantage of high value Hadoop analytics features such as Big SQL, BigDweets, Text Analytics, Big R, and Machine Learning to gain insight faster.

Monthly prices shown are for country or region: [United States](#)

Plan	Features	Price
<input checked="" type="checkbox"/> Basic (Beta)	Scalable Apache Hadoop clusters provisioned within minutes.	Free

Add Service

Space: blue

Service name: BigInsights for Apache Hadoop-xx

Selected Plan: Basic (Beta)

Enterprise

BigInsights for Apache Hadoop-e2

IBM BigInsights for Apache Hadoop

[OPEN](#)

The BigInsights™ for Apache Hadoop beta provides an exclusive early access to a nimble model for creating and using Hadoop clusters for analytics. You can spin up an open source Hadoop cluster within a few minutes. The clusters are based on IBM® Open Platform for Apache Hadoop.

During the beta period, you can have at most one cluster per service instance. You can start with only one data node and scale up to as many as five data nodes.

The service also allows you to integrate an instance of the Object Storage service on Bluemix with your Hadoop clusters. This integration provides a persistent storage for your data when the cluster is deleted, thereby enabling an on-demand usage pattern.

BigInsights Cluster List

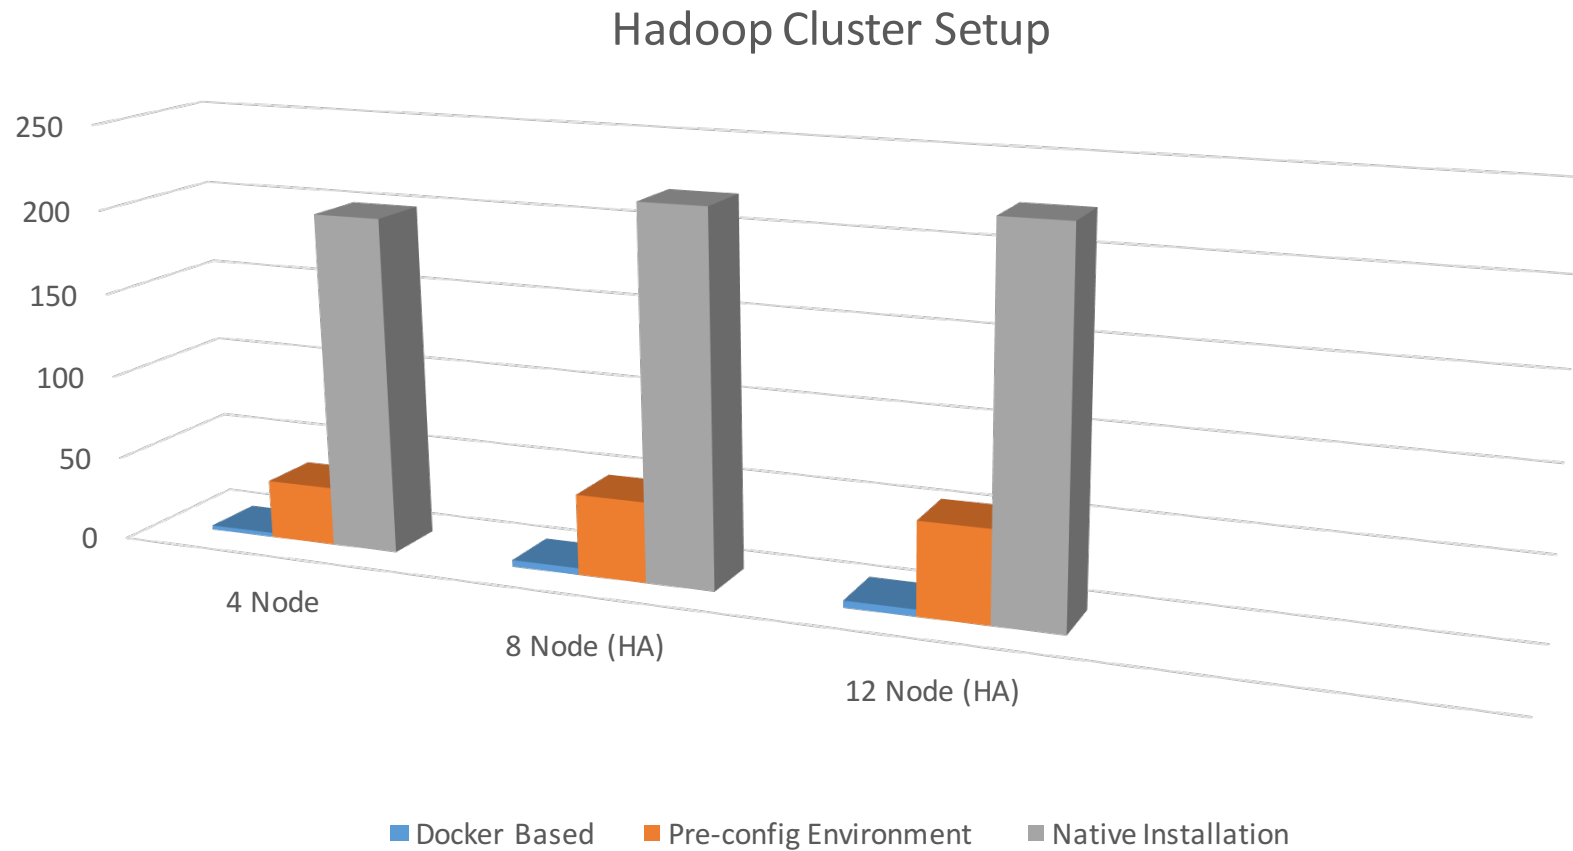
[Clusters](#) [Data Store](#)

[+ New Cluster](#) [Refresh](#)

Clusters (2)

Cluster Name	Management Nodes	Data nodes	High Availability	Status	Actions
rach-c2	1	1	No	● Active	
rach-c1	0	0	No	● Deleted	

Comparison of time to Setup Hadoop cluster



Earlier Experiments

Option	OS Provisioning	Config	Cluster Management / Updates
1	Bare metal	Chef	Chef
2	xCAT – Stateful(Create your own VMs)	PostScripts	xCAT - updateNode
3	xCAT – Sysclone(Image from current system)	Not Needed	xCAT - updateNode
4	Bare metal	PostScripts	xCAT - updateNode
5	Cloud Provider Specific Images	Not Needed	Manual/Scripts
6	Standard-ISO Image	Anaconda -post scripts	Manual/Scripts

Guiding Principles

- Virtualization helps repeatability, lesser failures & speed
 - Maintenance
 - Performance(Equivalent to Bare metal)
 - Use open source from an active community
 - Cloud-agnostic
-

Think Containers!!

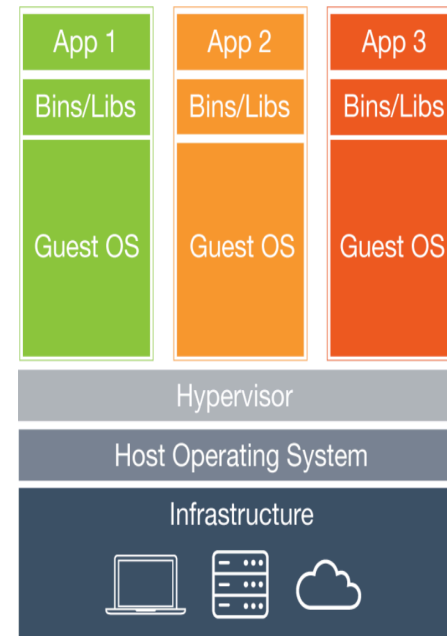
Docker : Thinking as VMs ? (mistake)

- **Key Concepts**

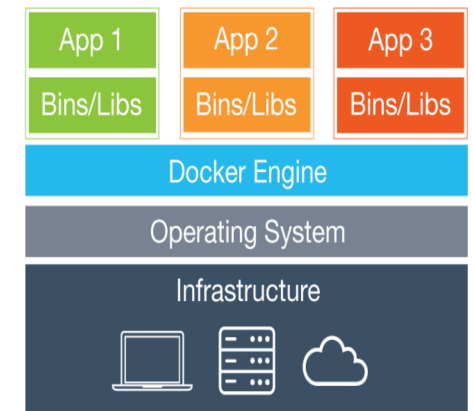
- Containers share host kernel
- Images & Image registry
- Build-able Images (make like)
- Images are layered & hence can be extended

- **Relevant concepts**

- Host FS directory can be mounted as 'volumes'
- IP specific port-forwarding



Virtual Machines

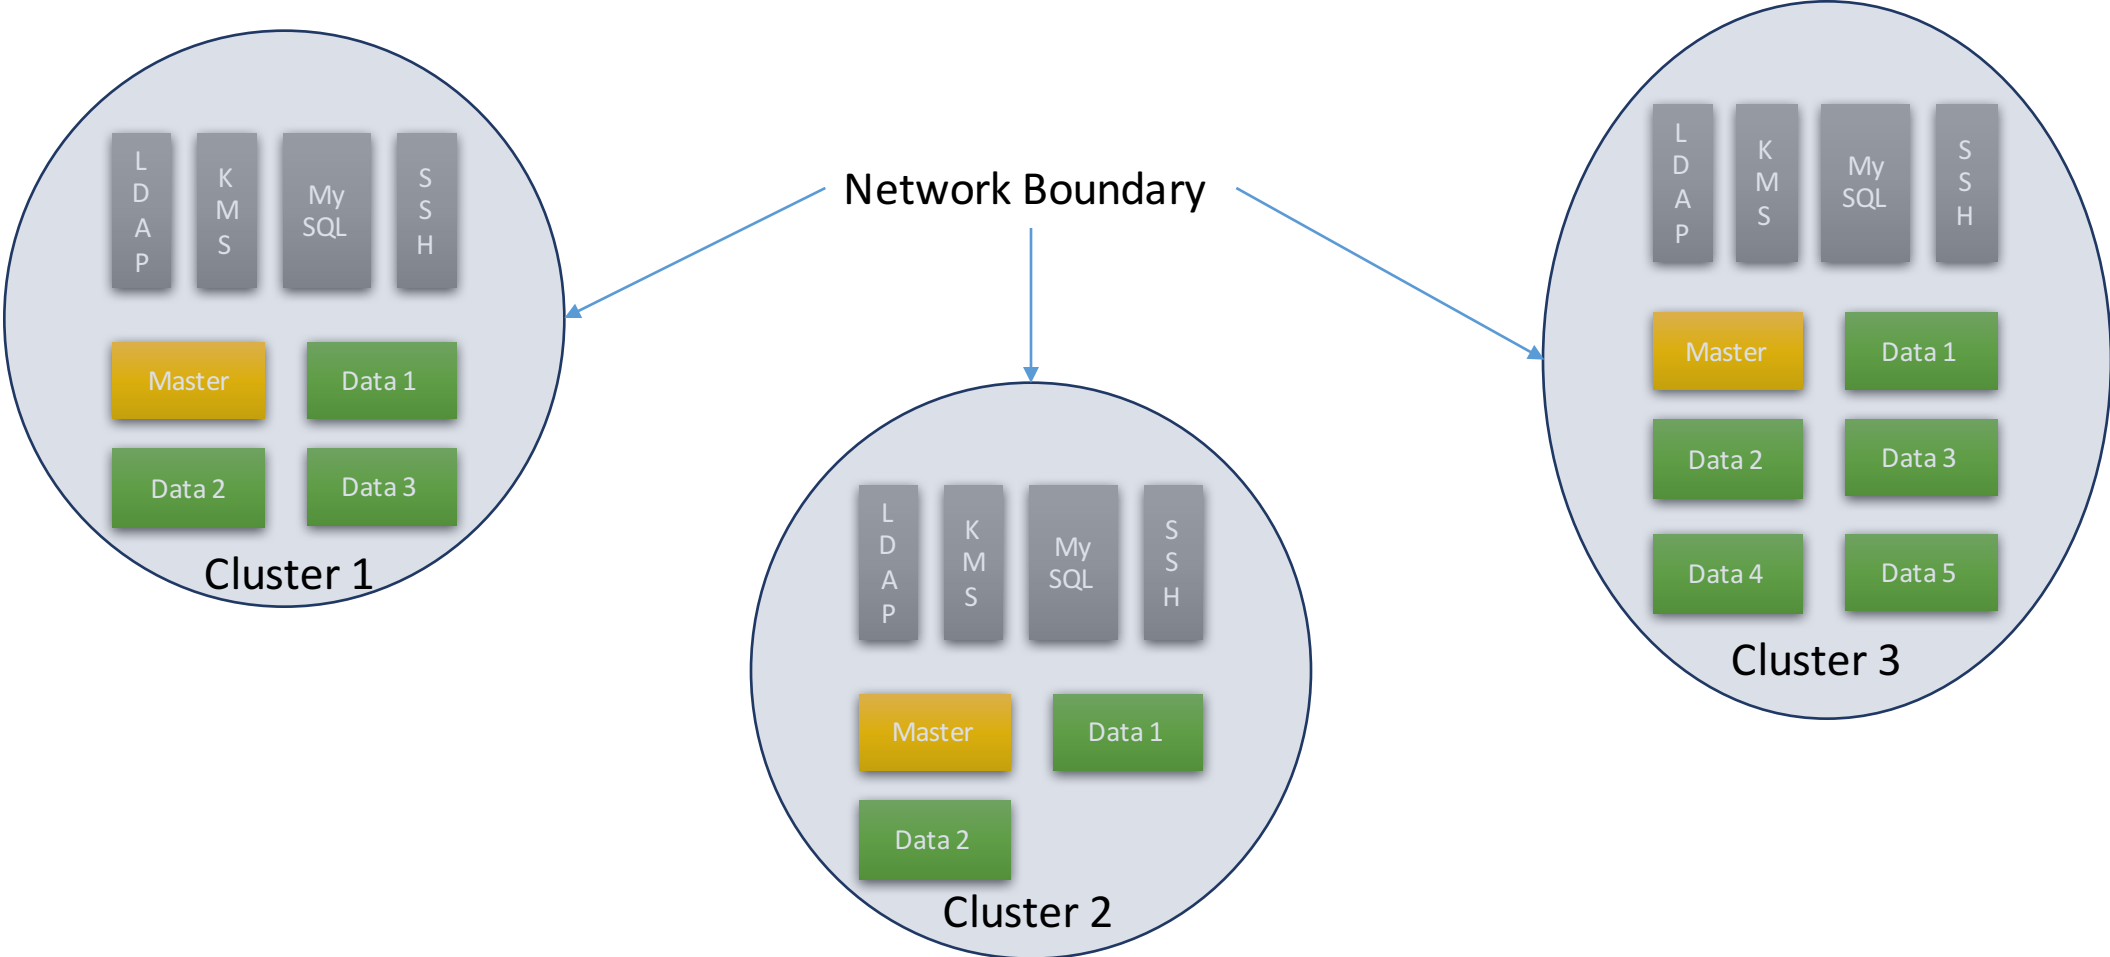


Containers

Docker in Hadoop Cluster on Cloud

- Each cluster node is a virtual node powered by Docker. Each node of the cluster is a Docker container
- Docker containers run on a bunch of **bare metal hosts** (Docker-hosts)
- Each Hadoop cluster will have multiple nodes/Docker containers spanning multiple hosts
- Docker
 - Container management - Custom
 - Multi host networking – Weave (Pluggable architecture to support other solutions)
 - Registry – Private

Typical Clusters



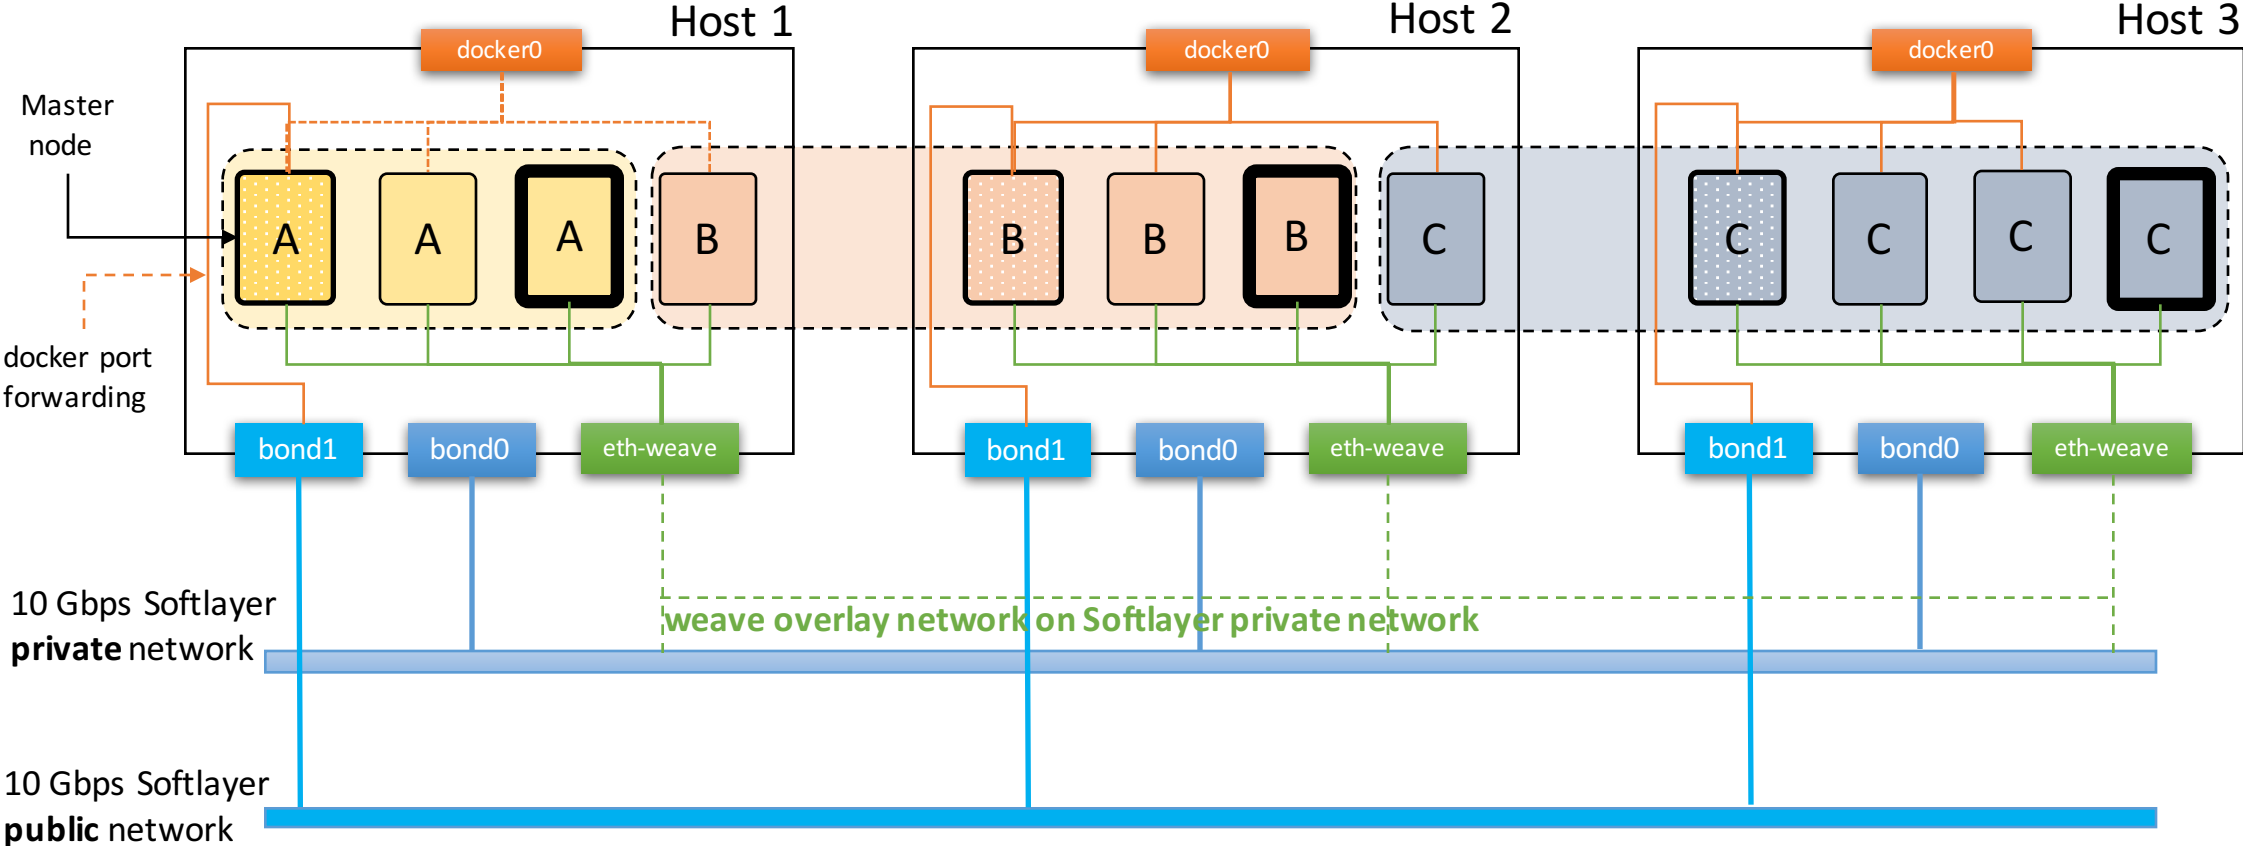
Docker Images

- Master node
- Data Node
- Edge Node
- Auxiliary service images
 - Ldap
 - Mysql
 - Ambari server
 - KMS


Multi host Docker networking with weave

- **Weave** based overlay network among IOP nodes
- One /26 private subnet per cluster (172.x.x.x)
- Master node to have a public IP – ports-forwarding
 - Portable public IPs
 - Network speed (shared with other masters)
- Edge node will be accessible using a public IP
 - User can SSH and run Hive, Hbase, Hadoop & Spark shells
- Private network
 - High Speed
 - Secure

Network Architecture



- * docker ICC=false (no inter container communication over docker0 network)
- * All inter-container communication is through weave network
- * One weave's private subnet per cluster (No communication across subnets)



Cluster Provisioning

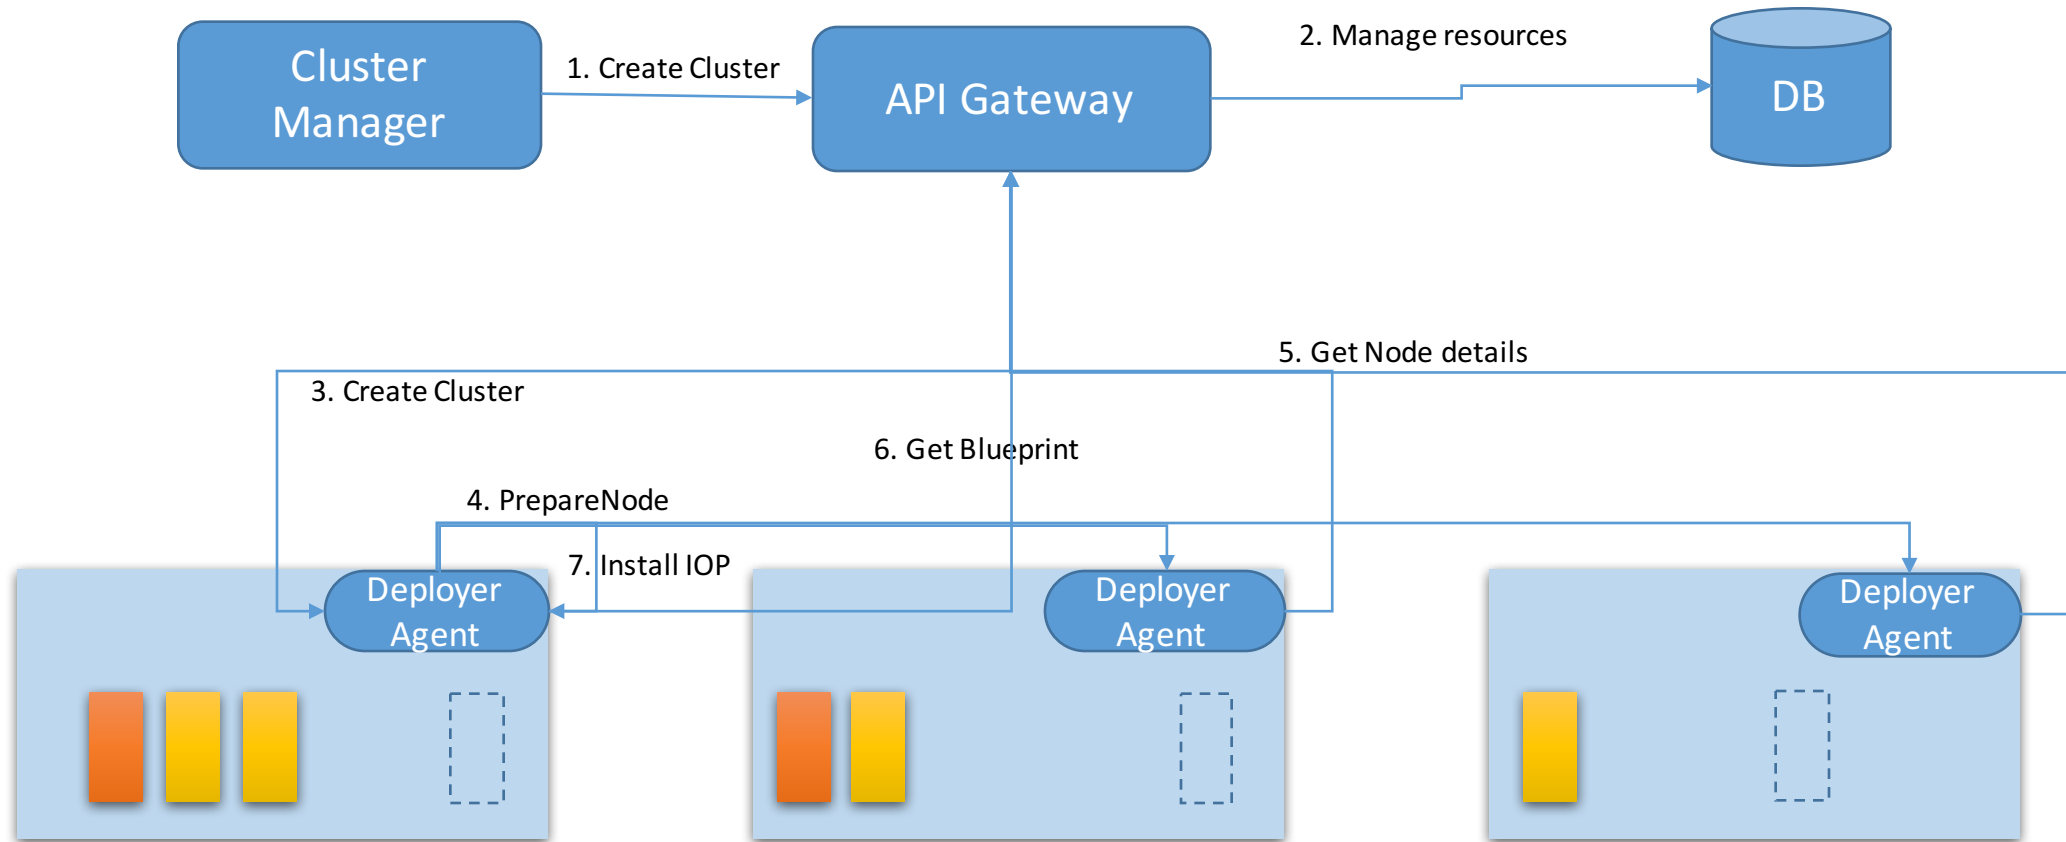
Provisioning Infrastructure

- Provisioning infrastructure consists of
 - Cluster Manager that provides REST API to create cluster
 - API Gateway application
 - Deployment agent
 - Deployer scripts that actually do all the work
 - And of course the database that holds all metadata related to clusters, hosts, nodes etc.

Phases involved

1. Acquire Hardware
2. Deploy Provisioning Infra
3. Add Hardware to Resource Pool
4. Prepare Host Machines
5. Orchestrate Cluster lifecycle
 1. Create Cluster, Add Nodes, Remove Nodes, Delete Cluster

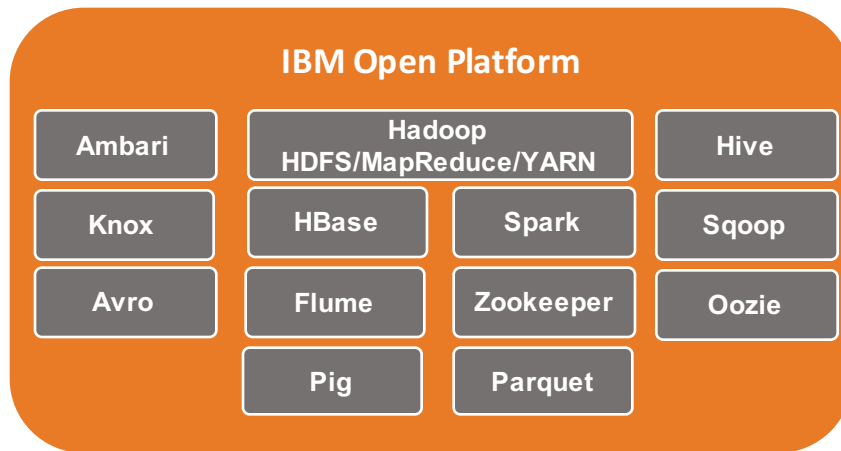
How is a cluster created?



Summary

- Cluster Creation in around 90 secs , 2-3 mins to start services
- Reliable, repeatable cluster provisioning
- Fully managed
- Highly Secure
- Close to zero failure rates
- Easy Patching
- Optimal use of resources
- Cost effective services

IBM BigInsights on Cloud – Basic Plan



Secure, flexible,
managed Apache
Hadoop clusters



Cheap – Pay for only what you need

Scalable – Add or remove nodes as per your workload

Modular – Compose and deploy applications using data and analytics services on Bluemix

Reliable – Durable SWIFT Object Store service always persists data

References

- Bluemix : <http://www.ibm.com/cloud-computing/bluemix>
- BigInsights On Cloud
 - <http://www-03.ibm.com/software/products/en/ibm-biginsights-for-apache-hadoop>
 - <https://developer.ibm.com/clouddataservices/docs/biginsights/get-started-in-bluemix/what-is-basic-plan/>
 - <https://www.youtube.com/watch?v=0DN1cDEs6ME>
- Tutorials
 - <https://www.youtube.com/watch?v=S3n9L2X91xM>
 - https://www.youtube.com/watch?v=t1Nuy_zrL7U
- IOP : <http://www.ibm.com/analytics/us/en/technology/hadoop/>
- Docker
- Weave