

2015 Spark 技术峰会

SuperVessel: Enabling Spark as a
Service with OpenStack and Docker

Guan Cheng (G.C.) Chen

IBM Research - China

weibo.com/parallelabs



程序员最可信赖的求职帮手

SuperVessel Cloud

- Public cloud built on the POWER7/POWER8 servers with OpenStack
- It provides **free** access for students, researchers, developers across the world, and helps grow OpenPOWER ecosystem (used in 30+ universities now)
- It provides advanced technology services such as **Spark as a Service**, Docker Services, Cognitive Computing Service, IoT Service, Accelerator as a Service (FPGA and GPU)



程序员最可信赖的求职帮手



清华大学
Tsinghua University



北京大学
PEKING UNIVERSITY



上海交通大学
SHANGHAI JIAO TONG UNIVERSITY



北京邮电大学
Beijing University of Posts and Telecommunications



Spark as a Service

Try it at: www.ptopenlab.com

3 steps to launch a Spark cluster, easy!

SuperVessel Big Data Service(Beta version)

Login

User Name: chengc@cn.ibm.com

Password:
Login

Recommend to use Firefox and Chrome browser

SuperVessel Big Data Service(Beta version)

Create

Create

Select Big Data Service Type: Spark

Select Cluster Node Size: 3

Select Extra Disk Size: 20G

Need Blueprint:30/day

Create Now

SuperVessel Big Data Service(Beta version)

Status

cluster1

Big Data Service Type:	Map Reduce
Node number:	3
Disk size per node:	20G
Mounted disk path:	/mnt (only on master node)
master_server_floating_ip:	172.16.9.193
master_server_internal_ip:	192.168.100.5
slave_node01_internal_ip:	192.168.100.4
slave_node1_internal_ip:	192.168.100.2

Map Reduce

HDFS (need to connect to VPN)

Step 1: Login

Step 2: Create

Step 3: Ready!

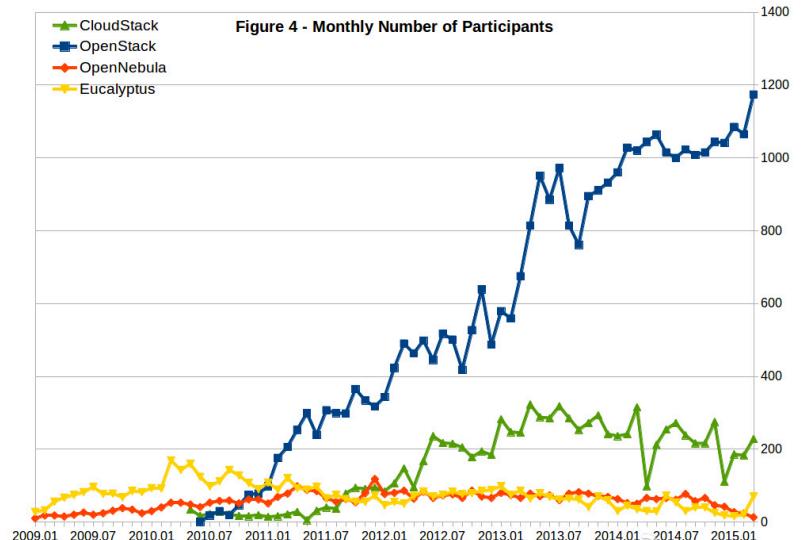


Why OpenStack?



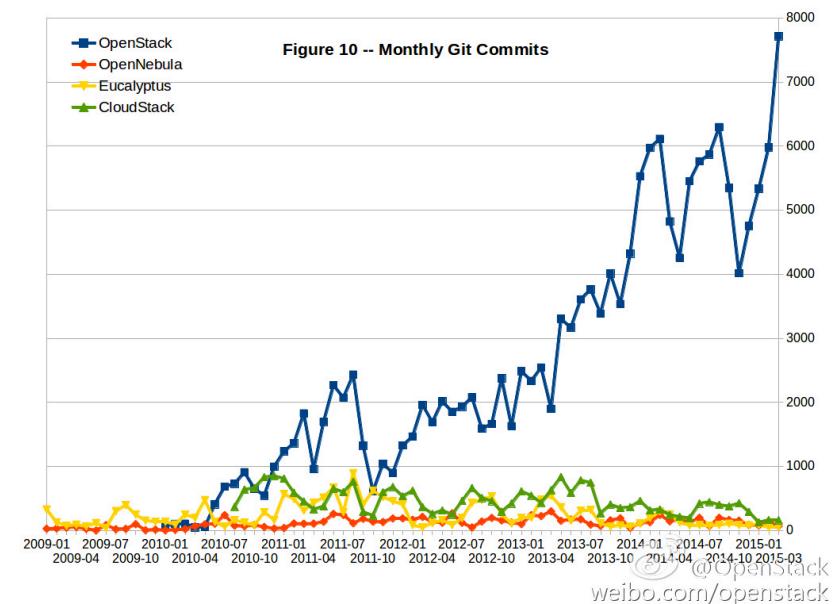
程序员最可信赖的求职帮手

- Most popular IaaS software
- Supports Docker*
- Heat can orchestra docker containers easily – good for provision a Spark cluster



 @OpenStack
weibo.com/openstack

Picture source: <http://www.gvjohn.net/?p=3801>

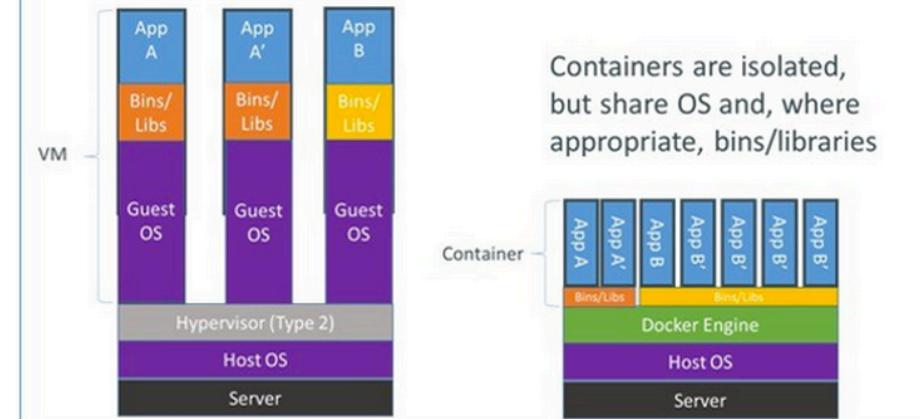


OpenStack 101

Why Docker?

- Less resource consumption than KVM
 - We can provision more Spark clusters!
- Boot faster than KVM
 - Users like fast provision!
- Incrementally build, revert and reuse your container
 - We love Git and AUFS!
- However, Docker is not officially supported in OpenStack yet
 - Nova Docker is an external component of OpenStack
- Port Docker to POWER Architecture (ppc64 and ppc64le)
 - Ubuntu 15.04 includes Docker for POWER8 ppc64le

Containers vs. VMs



Picture source: <http://www.zdnet.com/article/what-is-docker-and-why-is-it-so-darn-popular/>



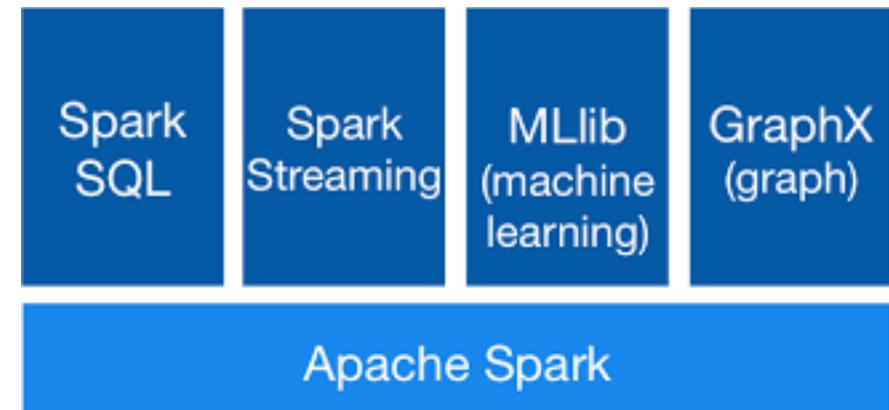
Why Spark?

- Fast
- Unified
- Ecosystem



Porting to POWER

- Bugfix submitted to the community
- Spark 1.3 works smoothly! ☺



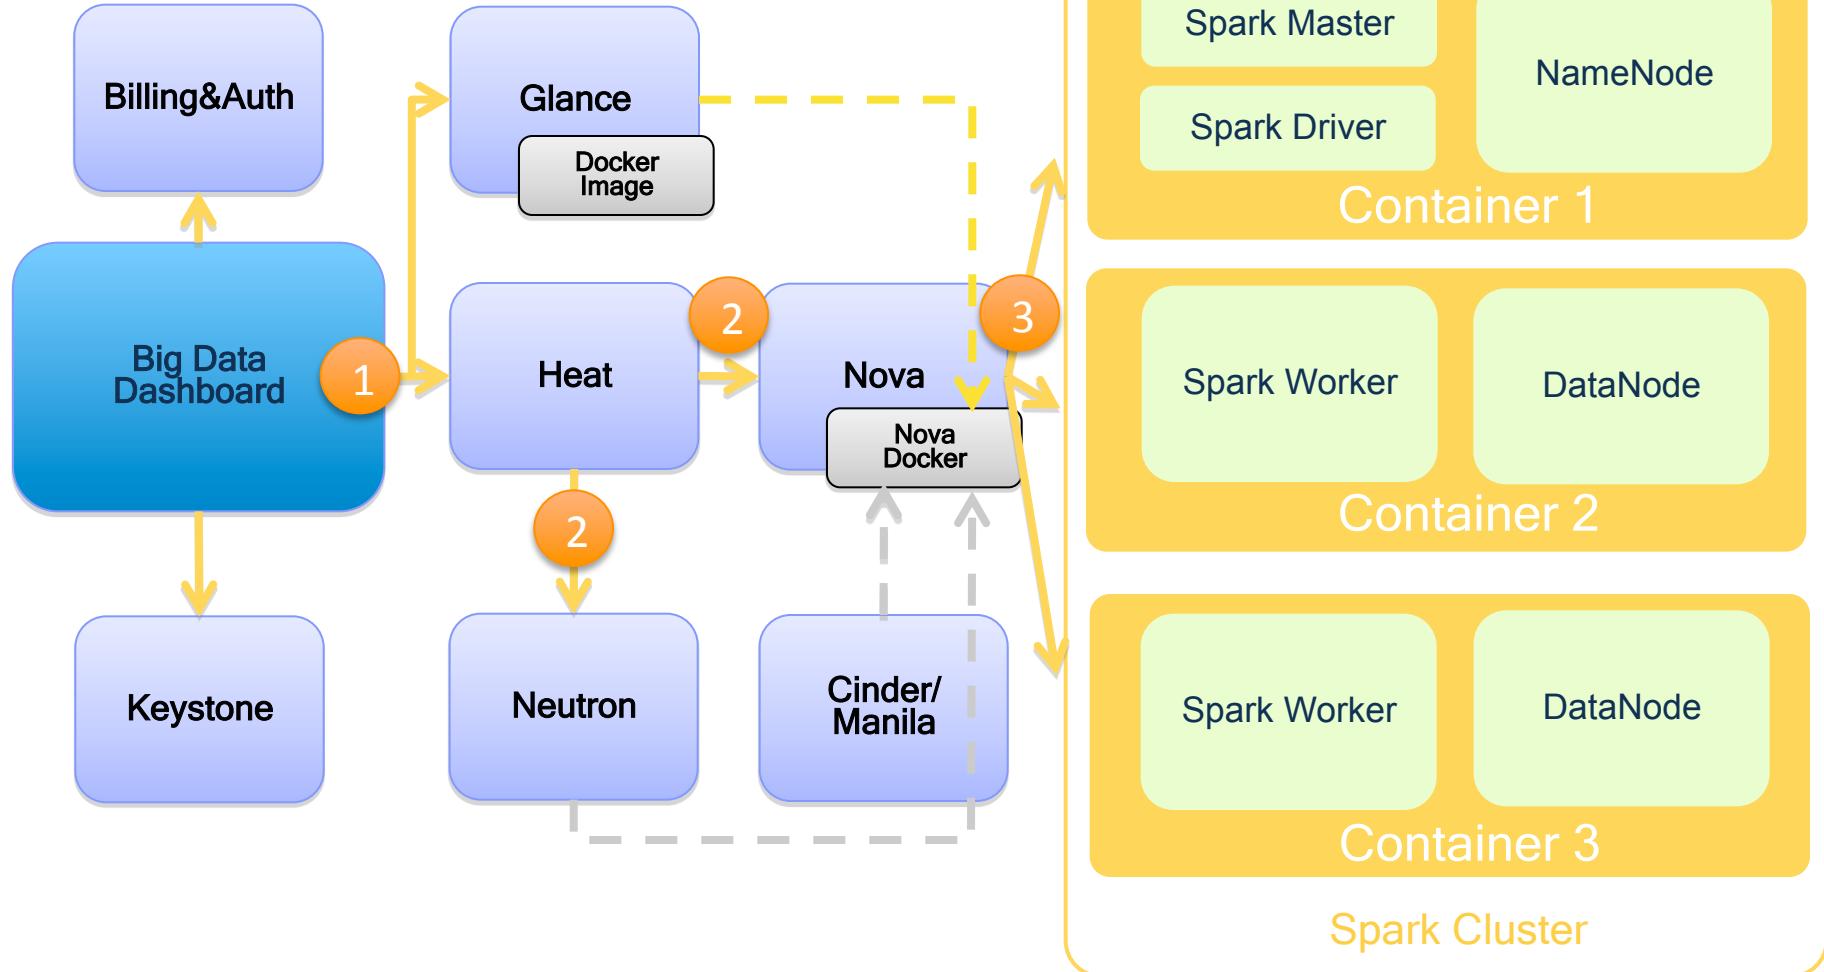
Why not Sahara?

Sahara is a component for Hadoop/Spark as a Service in OpenStack

- We started from OpenStack Icehouse ...
- Dockerization
 - Better service deployment and isolation for Big Data dashboard server
- Customization
 - Waiting for Sahara's improvements is somehow ***SLOW***
 - Docker, user analytics, Spark 1.4, Spark IDE, scheduling, data visualization etc.



Architecture Design



Dockerize everything!

- We use containers to
 - run **applications**
 - run **other containers**
 - run **OpenStack python daemons**
 - run **OpenStack services inside OpenStack** (with a special trunk link in neutron/OVS)
 - run **OpenStack inside OpenStack**
- Good for multi-site expansion

Heat template design

Heat is a component for orchestration in OpenStack

- Parameters
 - Cinder/Malina/Neutron uuid
 - Size of Cinder/Malina resources
- Resources
 - Master/slave node
 - Neutron
 - Cinder/Manila
 - Need to modify nova-docker to mount the cinder/ manila resources when booting the docker container

Spark Docker Image

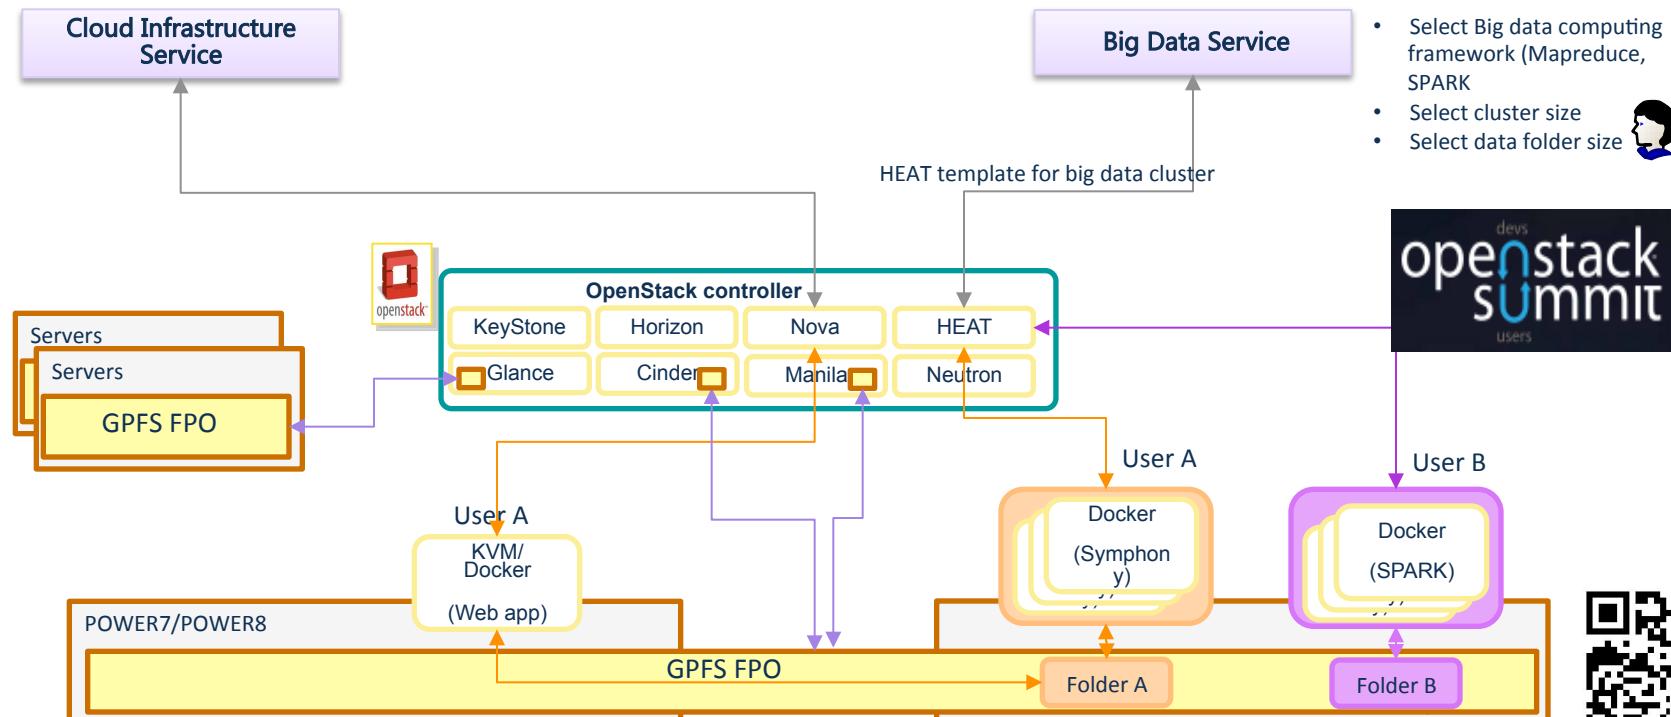
- Built from Ubuntu 14.04.1
- All Spark nodes use the same image, with different initialization scripts by using cloudinit
- Initialization scripts will
 - Sync /etc/hosts across all nodes
 - Set HDFS and Spark configurations accordingly
 - Format HDFS and launch HDFS
 - Launch Spark

Big Data Dashboard Development

- 2 Developers (front end + backend)
 - Online in 2 months
 - Separates Heat related stuff and Dashboard
 - Restful API (for billing and authentication etc)
 - Dockerize the big data dashboard server
 - Separates development and production environment

Where should I put the data?

Shared File System for Cloud and Spark as a Service

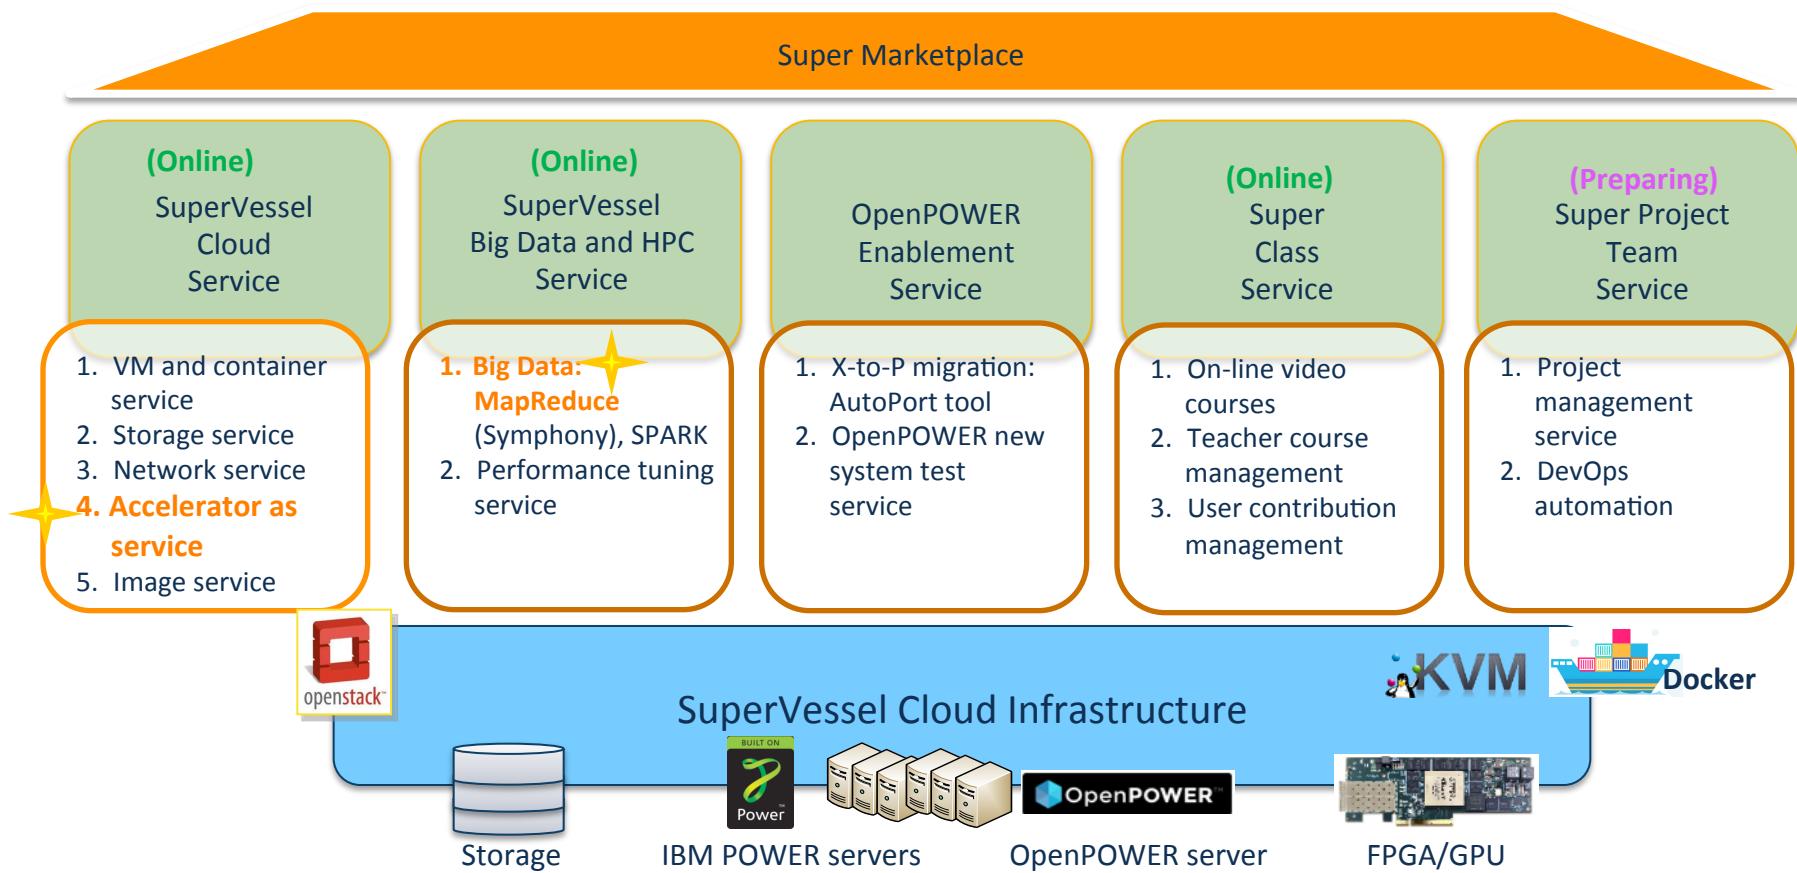


- HEAT will orchestrate docker instances, subnet and data folder based on user's request
- Manila provides the NFS service using GPFS as backend, and the folder will be mounted via nova-docker (with -v support)
- Folder created by Manila could be accessed by the KVM/docker instances created for big data and other purpose



程序员最可信赖的求职助手

SuperVessel Services Roadmap



Summary

- Spark + OpenStack + Docker works very well on OpenPOWER servers
- Dockerized services made DevOps easier
- Docker issues
 - Zombie process
 - Can't dynamically attach volume
 - Commit, attach operations is not supported on Nova-docker
- Monitoring everything (API, resources etc)
 - key for operating a cloud
- TODO: Spark 1.4, Spark IDE, SWIFT, Data Visualization



Join Us!



SuperVessel
WeChat group



程序员最可信赖的求职帮手

www.ptopenlab.com



QQ group: SuperVessel



@冠诚
chengc@cn.ibm.com