

MESOS(DC/OS)入门与实践

龙辉 2018.4

微博、微信: @huihoo

MESOS是什么？

一个架构模式，一种编程模型

Linux：操作系统内核

Mesos：分布式操作系统内核

分布式系统的内核

Mesos: 数据中心操作系统的内核

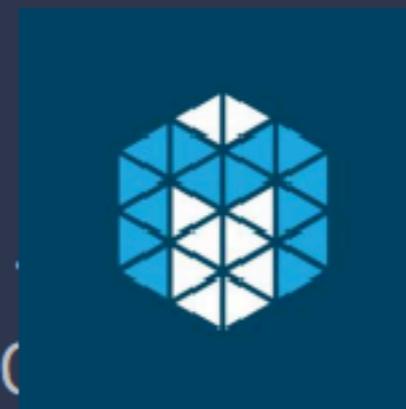
- 传统的操作系统内核

- 资源管理： 单机CPU, 内存, 磁盘, ...
- 编程抽象： POSIX API：进程, 线程,
- 安全与隔离： 虚拟内存, 用户, ...



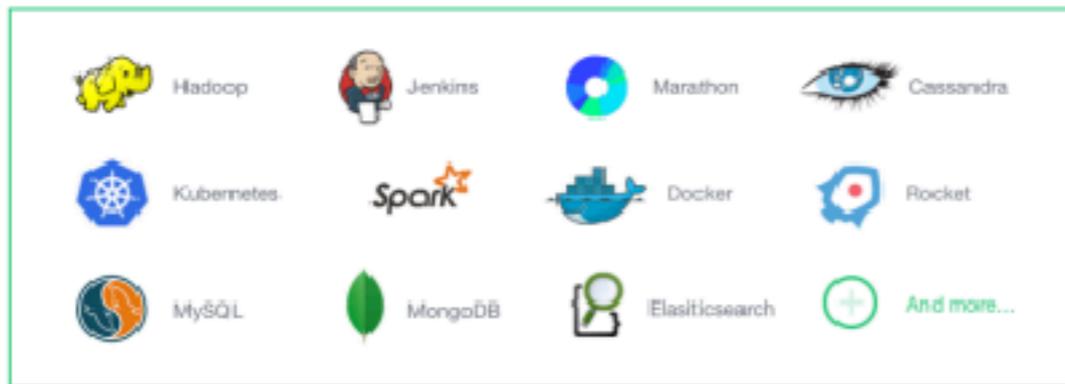
- 数据中心操作系统的内核

- 资源管理： 集群CPU, 内存, 磁盘, ...
- 编程抽象： Mesos API：Resource, ...
- 安全与隔离： 容器化



DC/OS是什么？

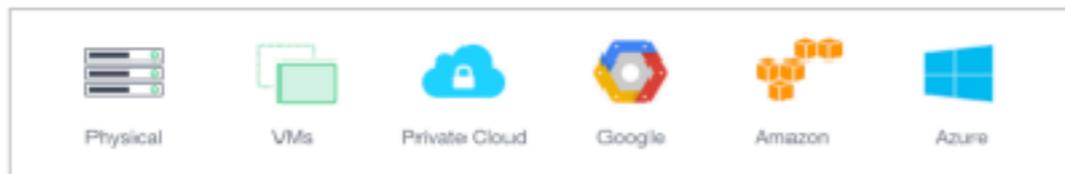
Mesosphere DC/OS - Mesos的一个发行版



容器和服务



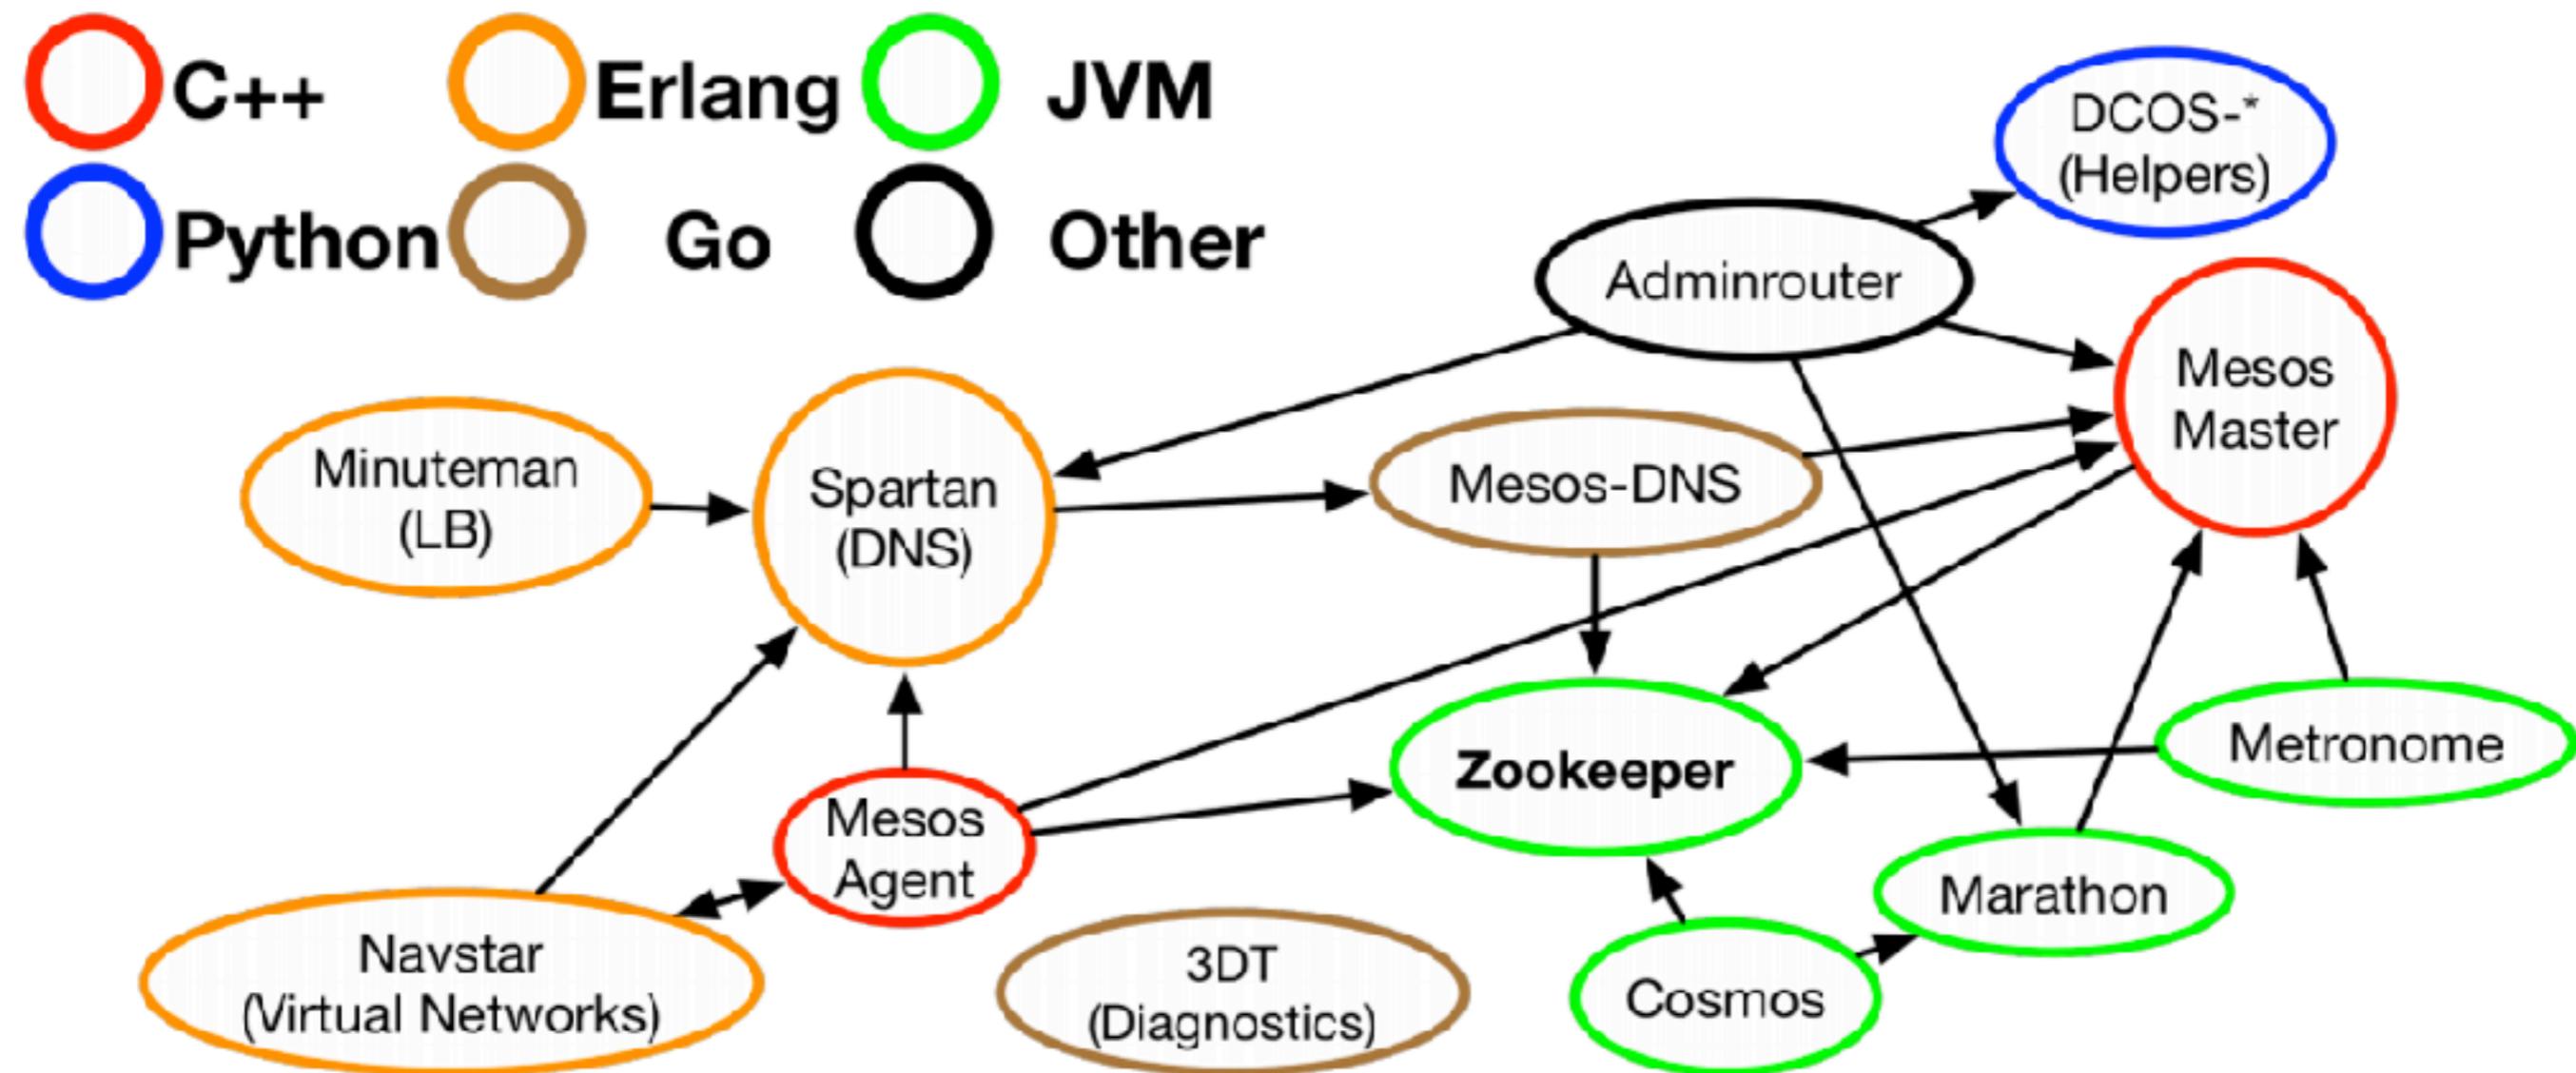
Mesosphere DC/OS



任意基础架构
公有云
私有云

- 只有内核是不够的！
- DC/OS: Mesos最好的发行版
 - 命令行和图形界面
 - 包裹管理
 - 服务发现
 - 负载均衡
 - 调试，日志和监控
 - 用户管理，安全性加强
 - Framework SDK
- 是开源的！

DC/OS架构



MESOS简介

分布式、集群管理系统和框架

Mesos是分布式操作系统内核，Mesos是以与Linux内核同样的原则而创建的，不同点仅仅是在于抽象的层面。Mesos内核运行在每一个机器上，同时通过API为各种应用提供跨数据中心和云的资源管理调度能力。这些应用包括Hadoop、Spark、Kafka、Elasticsearch。还可配合框架 Marathon 来管理大规模的Docker等容器化应用。

开发模式转变

从针对一台PC、一个OS、一个Database、一个Web Server编程；

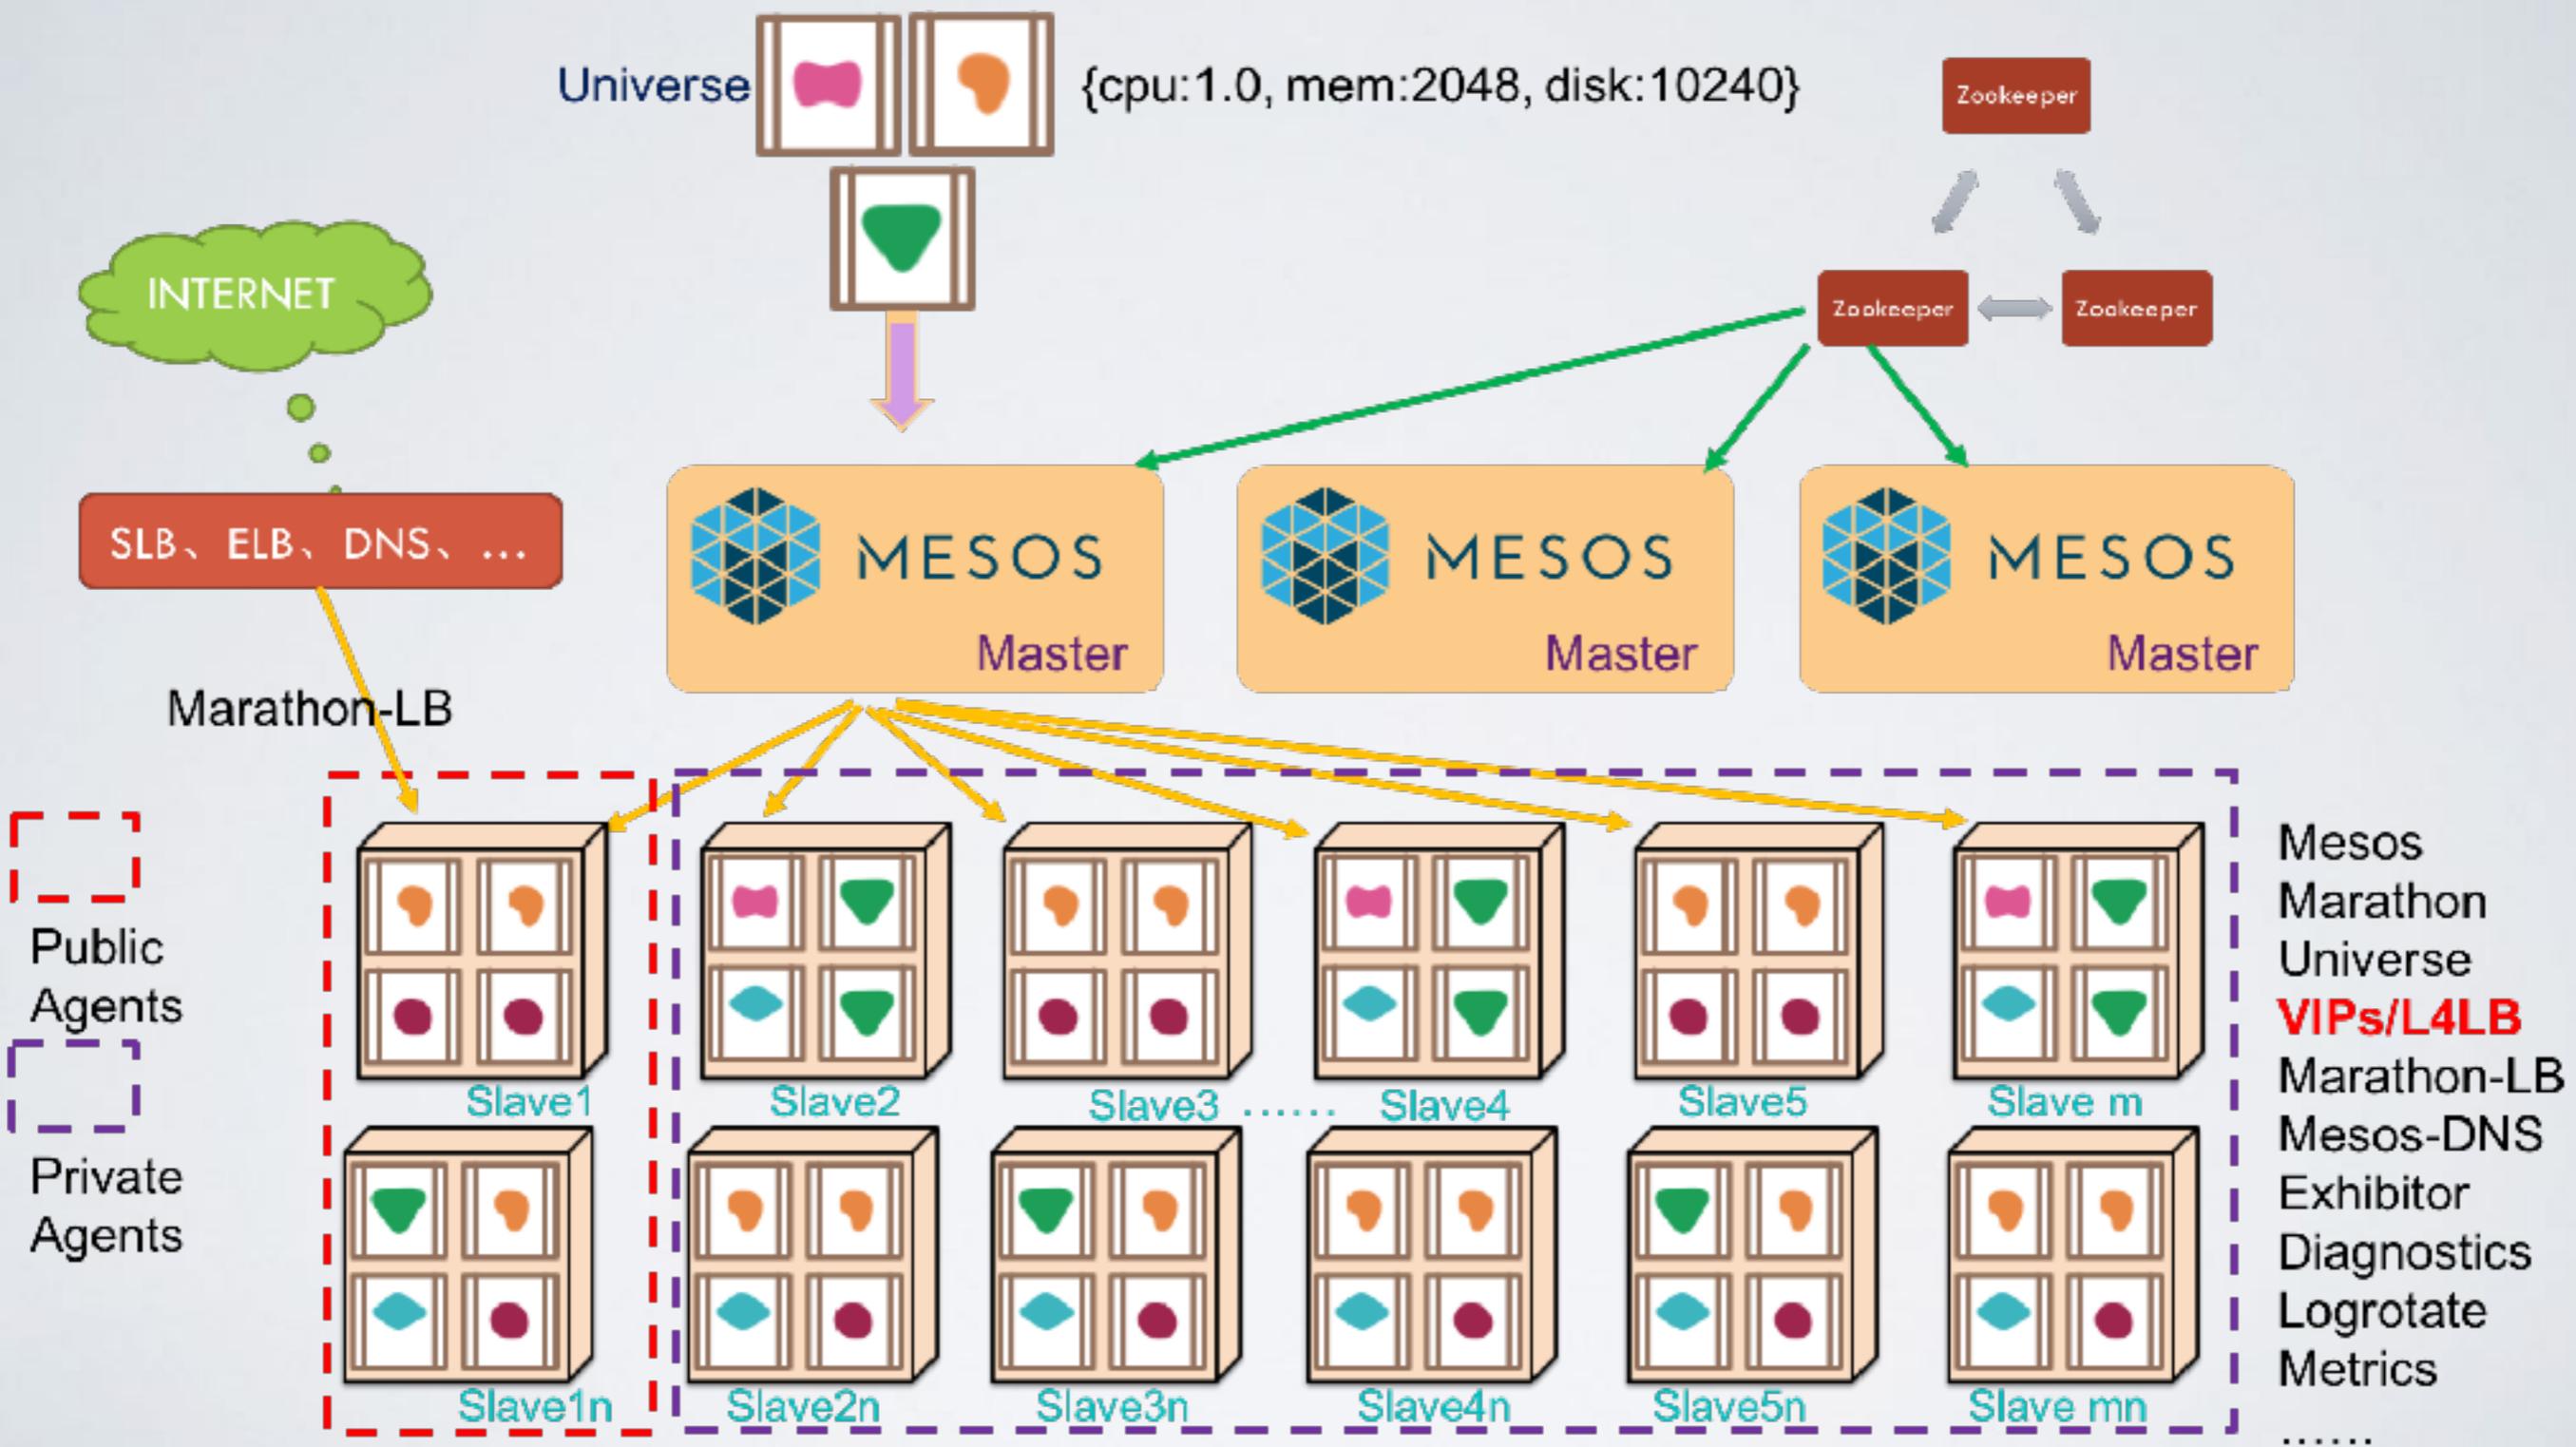
到针对一堆服务器、集群、云平台、数据中心编程；

Mesos将PC编程转变为对Data Center编程，对资源分配、任务调度编程。

MESOS特性

- * 可扩展到10000个节点；
- * 使用 ZooKeeper 实现 Master 和 Slave 的容错；
- * 支持 Docker 容器；
- * 使用 Linux 容器实现本地任务隔离；
- * 多资源调度能力（内存，CPU、磁盘、端口）；
- * 提供 Java, Python, C++等多种语言 APIs；
- * 通过 Web 界面查看集群状态；
- * 一个快速增长的生态系统。

MESOS集群内网/外网

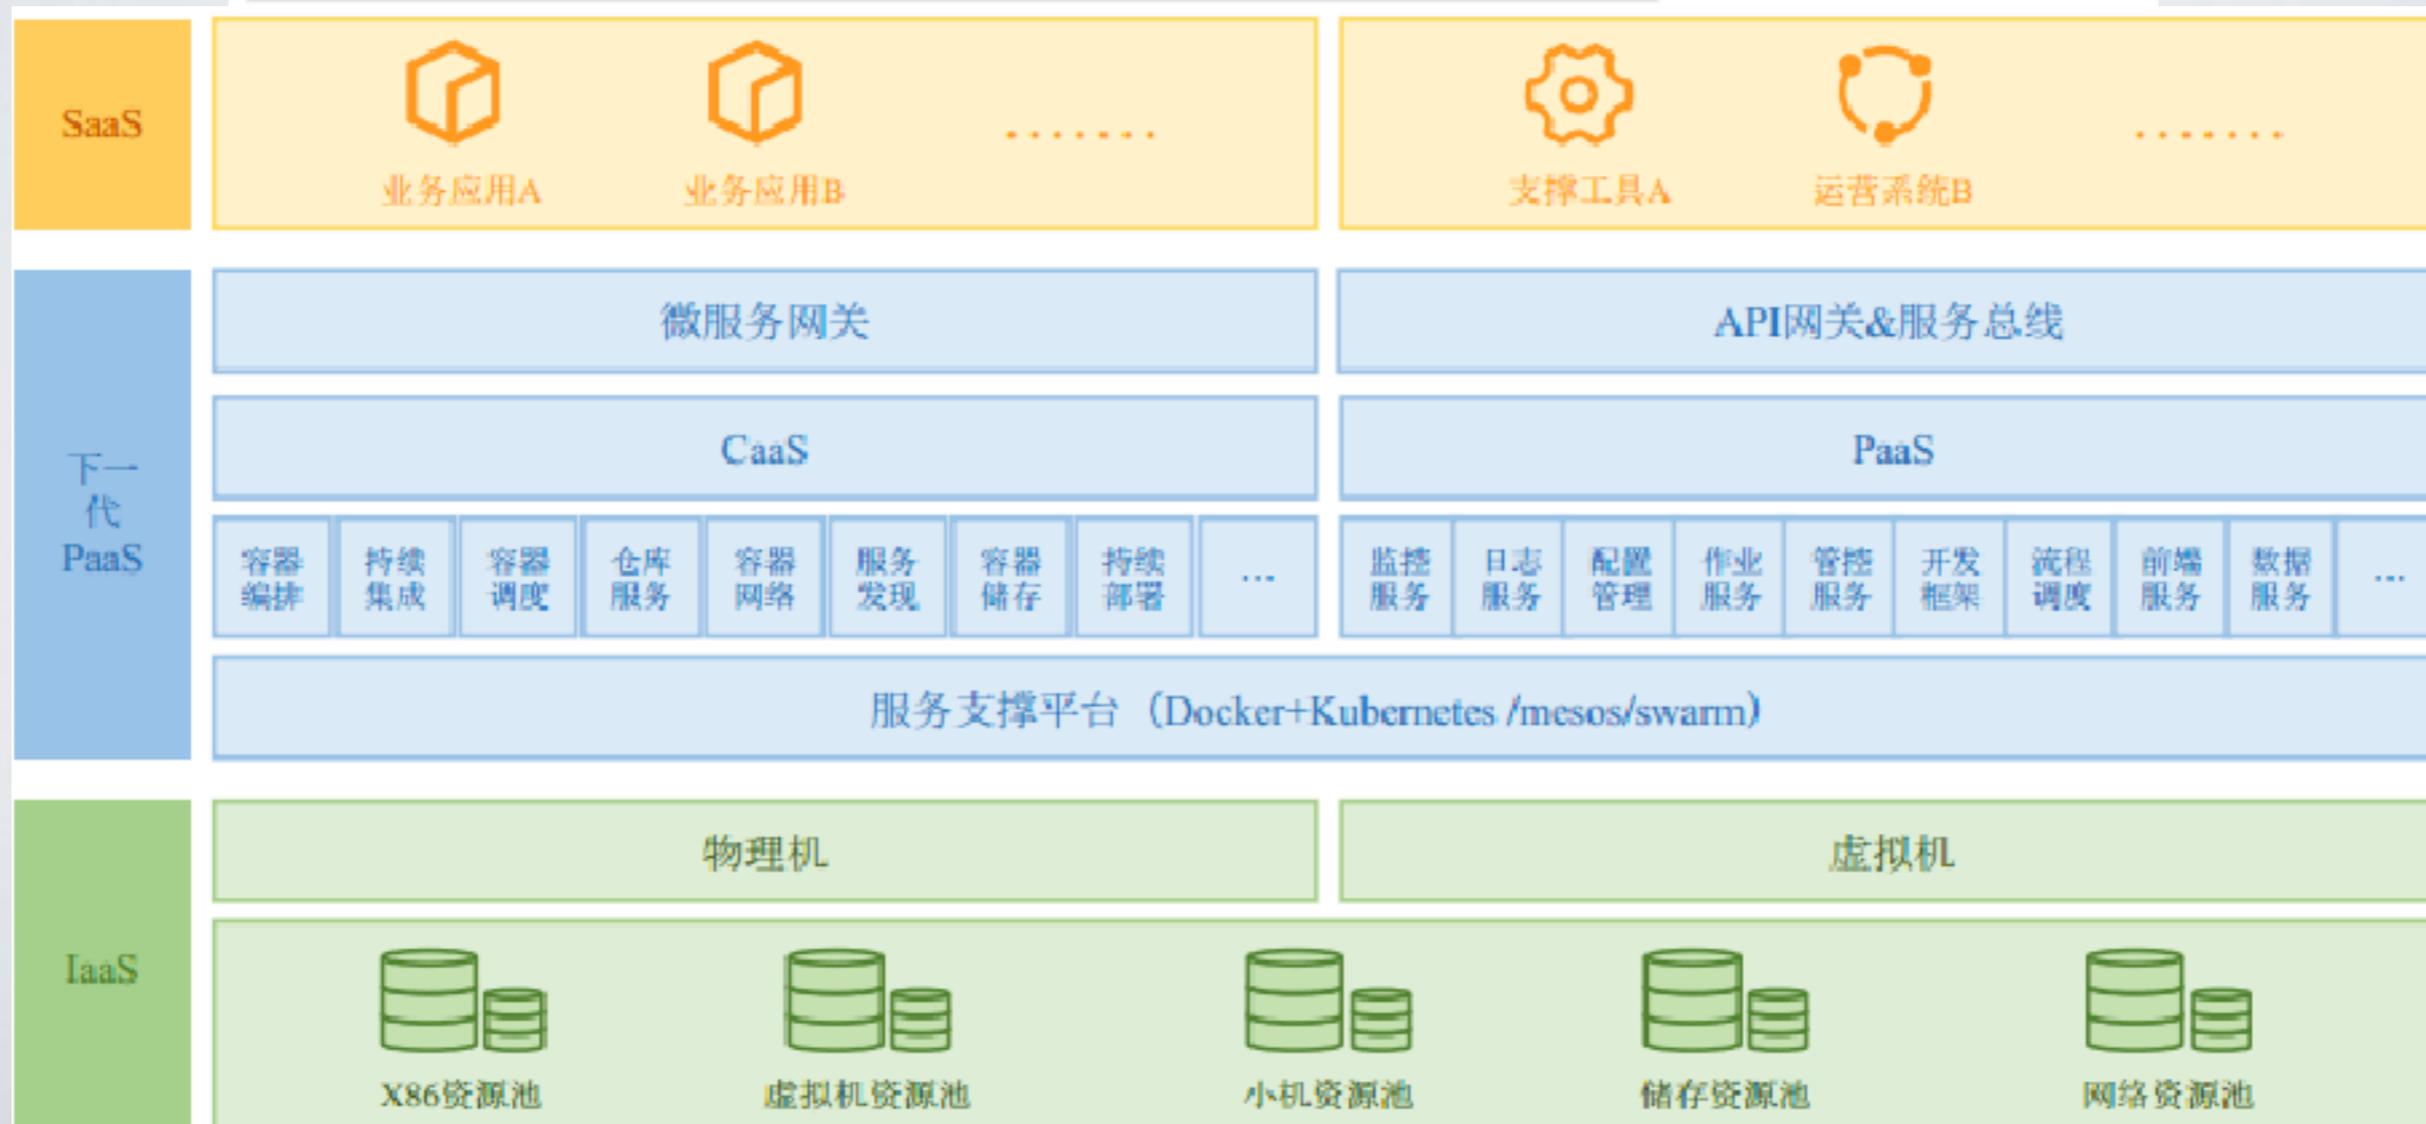


为什么是MESOS(DC/OS)

腾讯蓝鲸容器管理平台



- 容器编排
- 弹性扩缩容
- 服务发现
- 镜像管理
- 自动构建
- 一键部署
- 跨区调度



小米弹性调度平台

Core Stack

Marathon

Mesos

Docker

Flannel

Calico

PaaS Service

ELB

Mysql

Memcached

Redis

Other

Stateless APP

Code

OS

Build

Docker Image

Deploy

App

CI

CD

Infrastructure

CMDB

Naming

ELK

Monitor

LB

Relay Login

dynamics Auth

Key Center

distributed storage

Target

Depoy & Migration

Auto Scaling

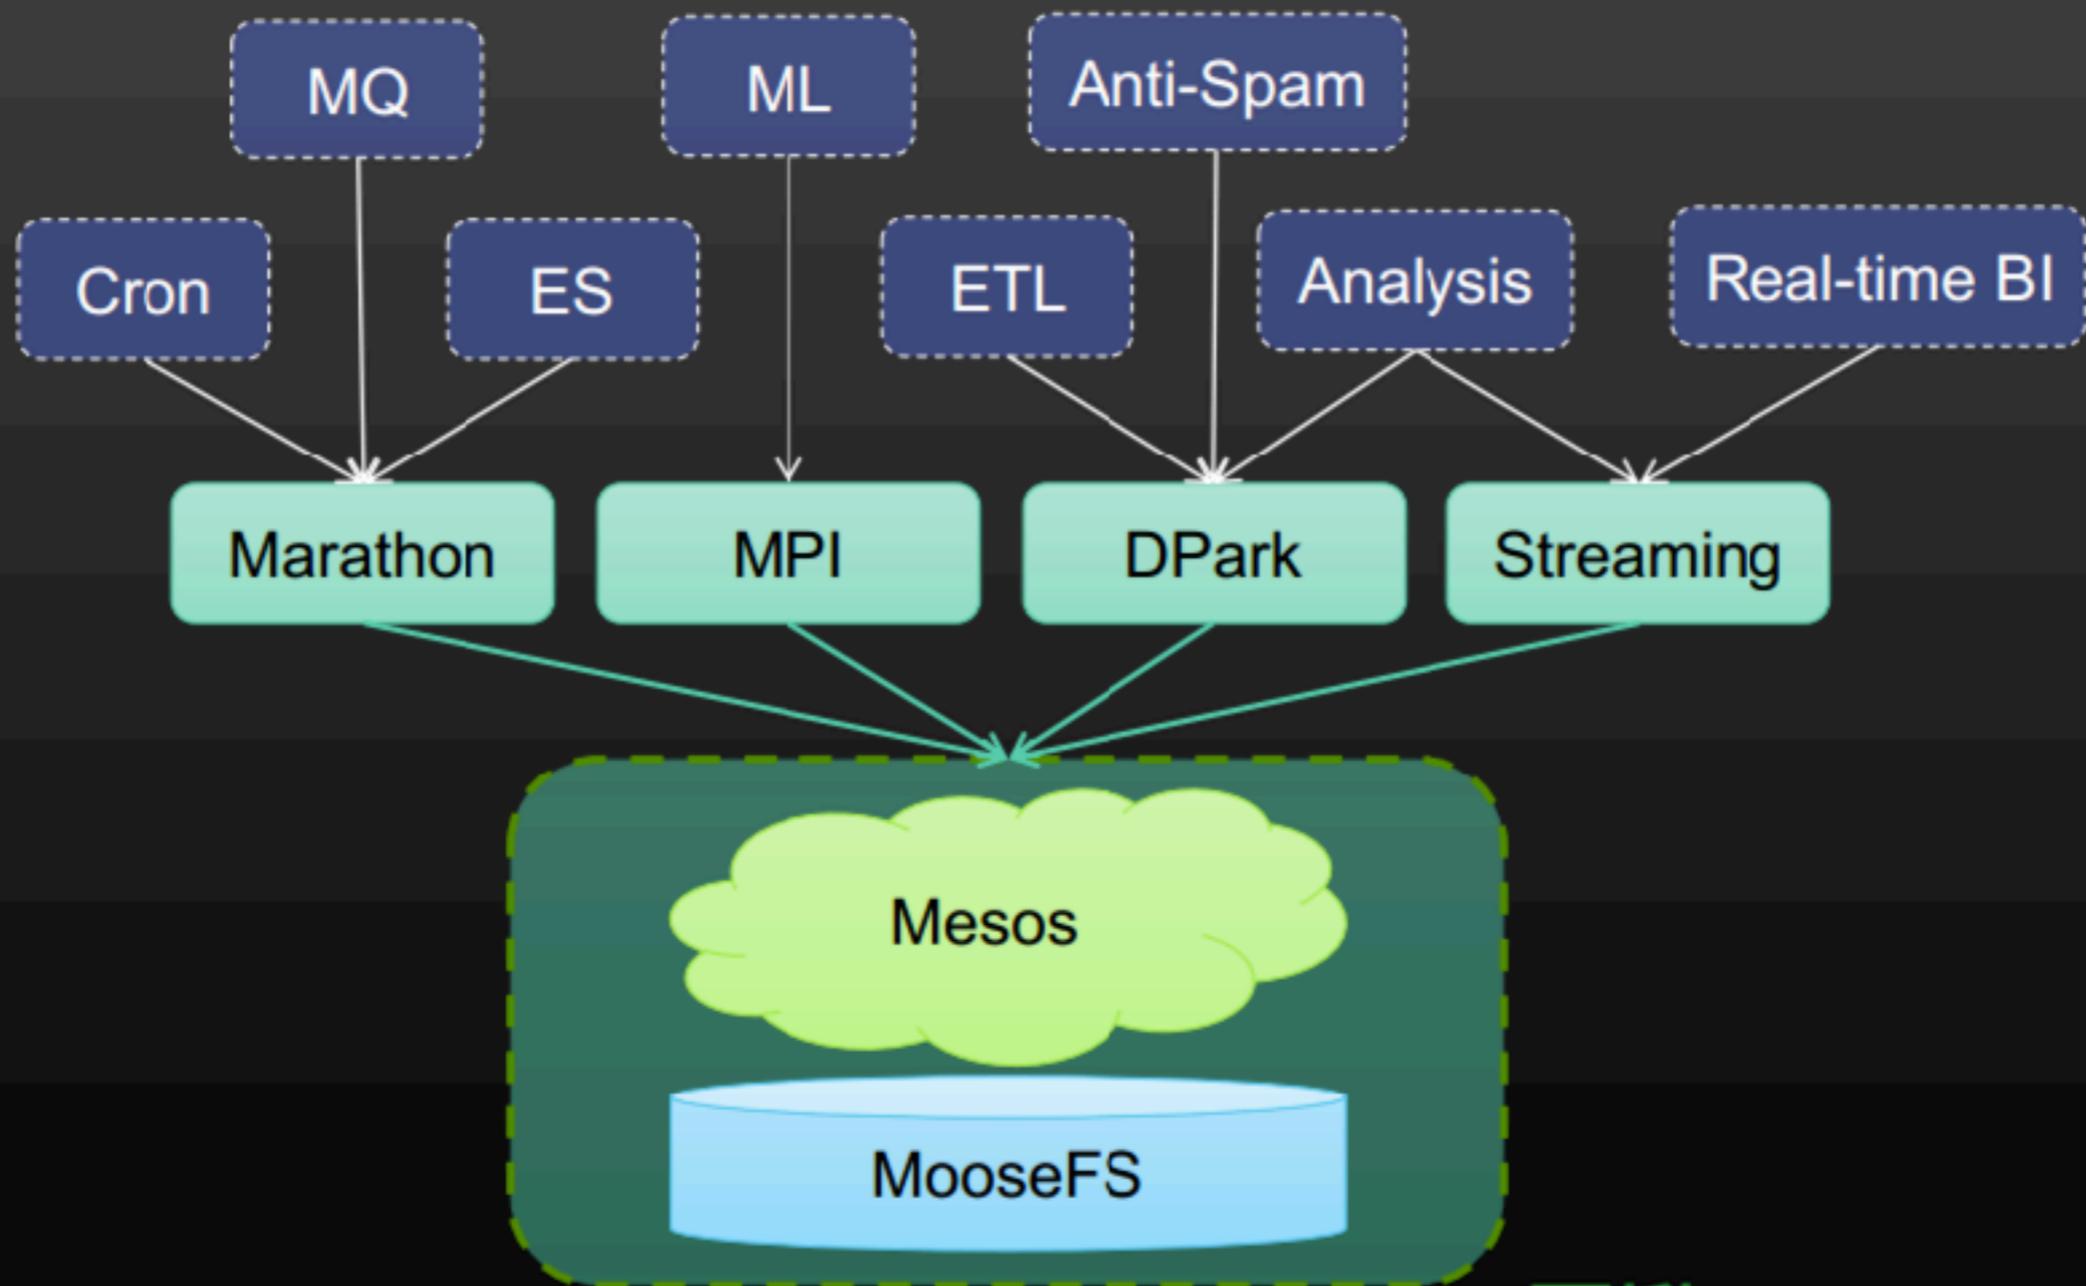
Auto PipeLine

Auto FailOver

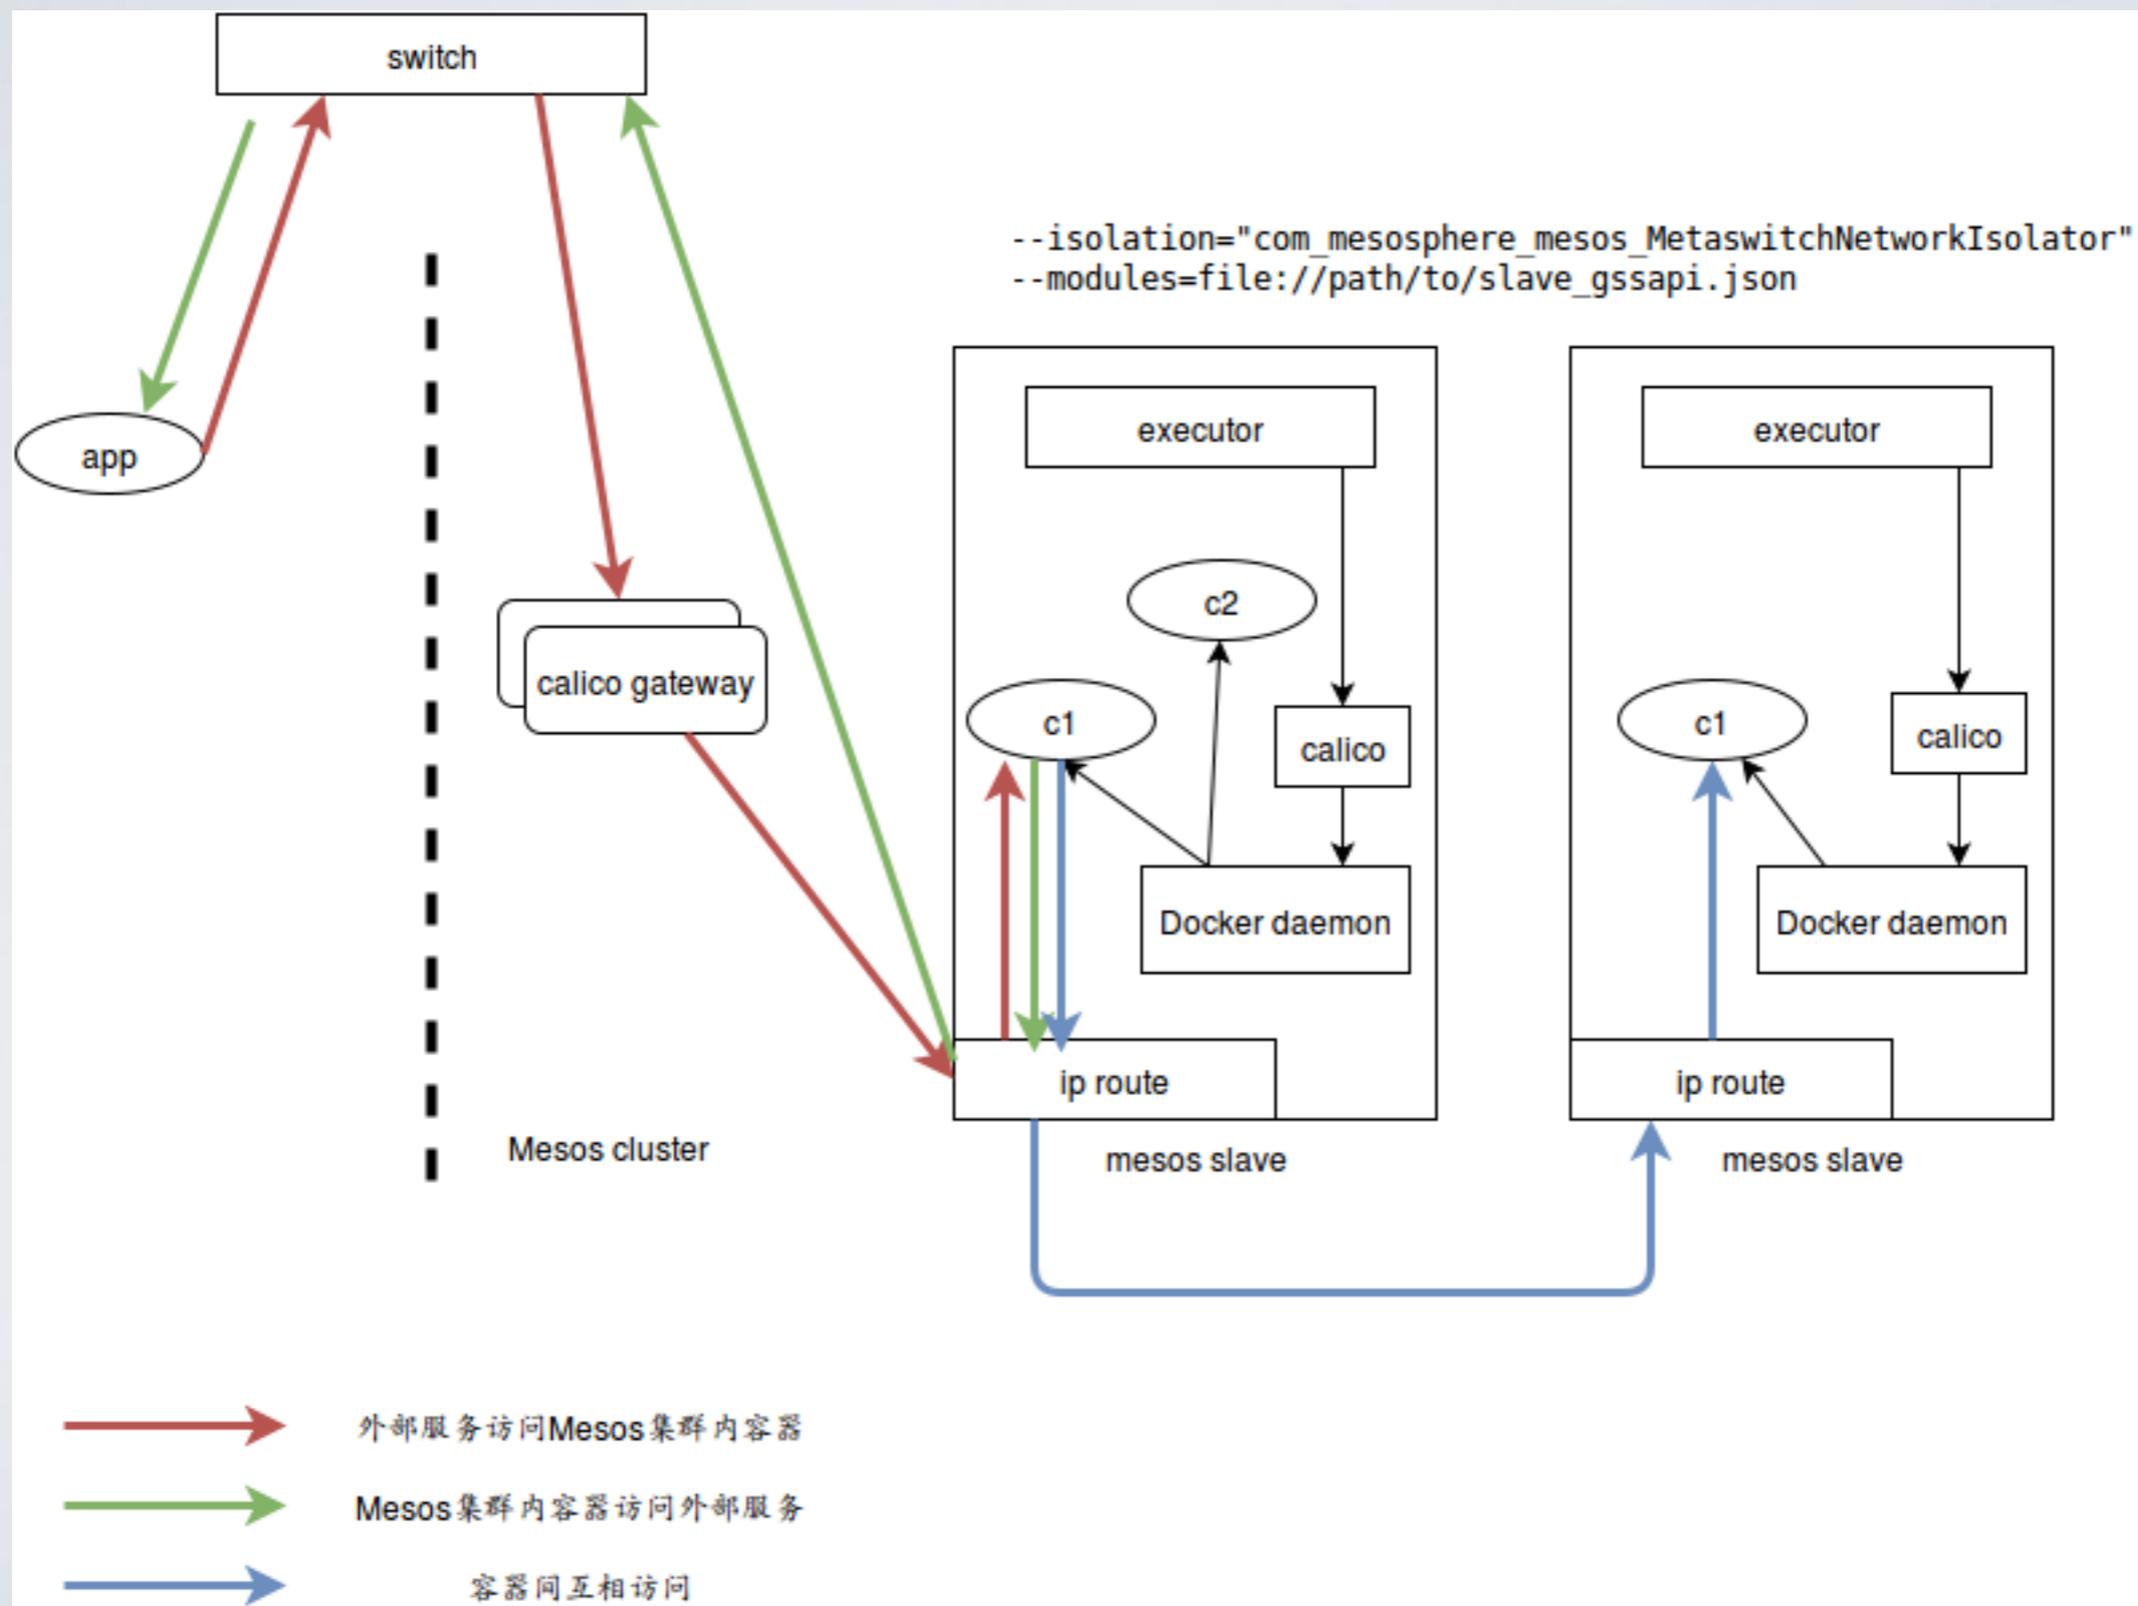
Dependen cy Chain

豆瓣基于MESOS的大数据基础架构

Big Picture



去哪儿



去哪儿架构演进

1) OpenStack + nova-docker + VLAN

2) Mesos + Marathon + Docker(--net=host) + 随机端口

3) Mesos + Marathon + Docker + Calico

联通CU-DCOS平台

CU-DCOS | 技术选型Mesosphere DC/OS



Open DC/OS

Management

Admin Router

ZooKeeper

Mesos-DNS

Minuteman

Distributed
DNS Proxy

Auth

GUI

Cosmos

docker

docker

docker

Myriad

Kafka

Spark

Marathon

Frameworks

MESOS

Open DC/OS 是mesosphere2016年上旬发布的企业数据中心操作系统，主要是围绕开源的Apache Mesos附加了一些面向企业级应用的管理功能，核心内容还是两支，一边是面向微服务治理的marathon，一边是通过framework来管理大数据类的一些应用，其附加功能：

Admin Router：一个反向代理用于引导用户请求到相应的服务，服务引流的功能。

Zookeeper：或者是Exhibitor，为DCOS其上的分布式系统提供协调服务

Mesos-DNS：服务注册、服务发现

Minuteman：内部网络第四层负载均衡

Distributed DNS Proxy：分布式DNS代理

Auth：用户权限管理

GUI：图形化的为监控、配置界面

Cosmos：框架安装包私库

MESOS中国生态系统

China Ecosystem



DC/OS系统服务

/ETC/SYSTEMD/SYSTEM/DCOS.TARGET.WANTS/

DC/OS由一组系统服务组成，包括大家所熟知的Mesos、Marathon、Zookeeper及Rexray等。

DC/OS容器服务

提供容器集群服务

支持K8S容器集群管理

Marathon支持两种运行Docker容器镜像的方式：

- 1、Docker容器化，使用Docker引擎运行容器镜像；
- 2、Mesos容器化，使用通用容器运行时(UCR)运行容器镜像。

MESOS-DNS

MESOS-DNS为DC/OS集群中的应用提供了服务发现功能

通过Calico项目给每个容器分配一个IP

EXHIBITOR服务 (ZOOKEEPER监控)

EXHIBITOR是是NETFLIX开源的一个用于ZOOKEEPER监控和配置管理的系统。

特性

实例监控、日志清理、备份/还原、集群配置、集群滚动升级、自动实例管理...

ADMIN ROUTER

Admin Router服务是DC/OS内部核心的负载均衡服务。Admin Router是一个定制的
Nginx，它在端口80上代理所有内部的服务。

研发

PIPELINES

PaaS、PM、CI、CD、DevOps、BPM

JAVA JVM ON DC/OS

JAVA JVM ON DC/OS

Microservice Framework

Microservices

Cluster Orchestrator

Anwendungen

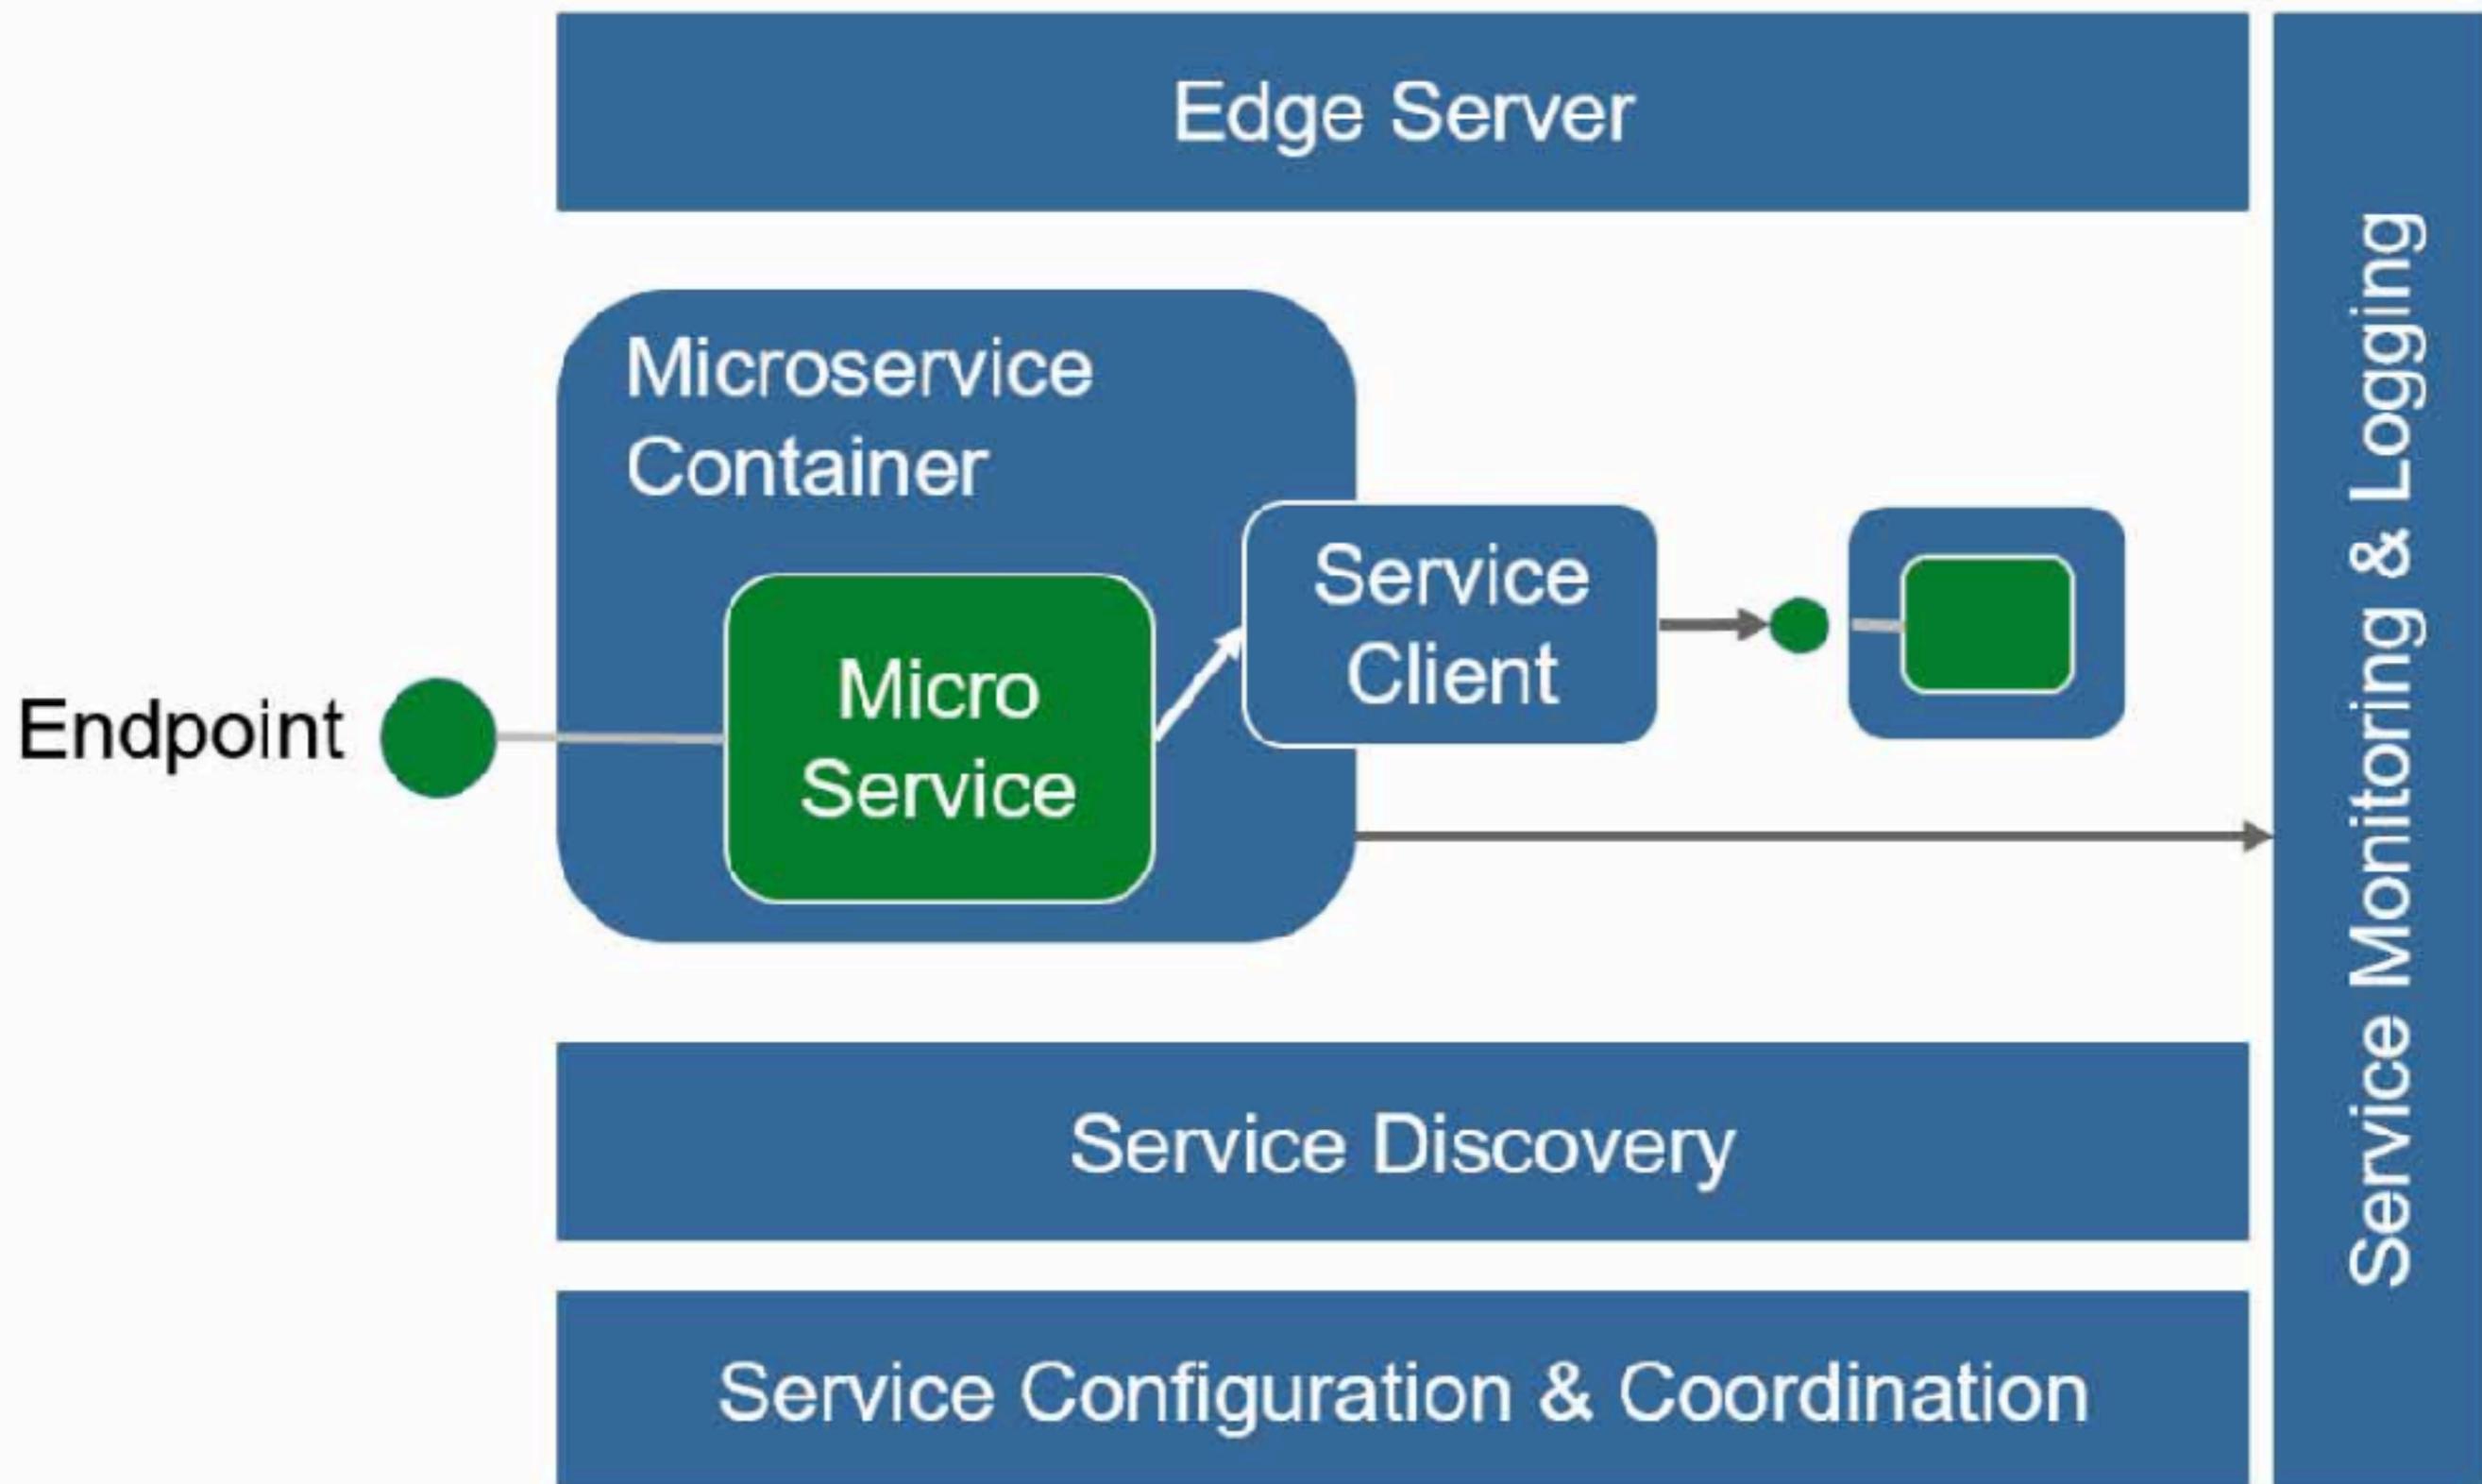
Cluster Scheduler

Container

Ressourcen

Distributed
Systems Services

JAVA JVM ON DC/OS



JAVA JVM ON DC/OS



Spring Cloud

NETFLIX | OSS

Microservice Framework



kubernetes
by Google

Cluster Orchestrator



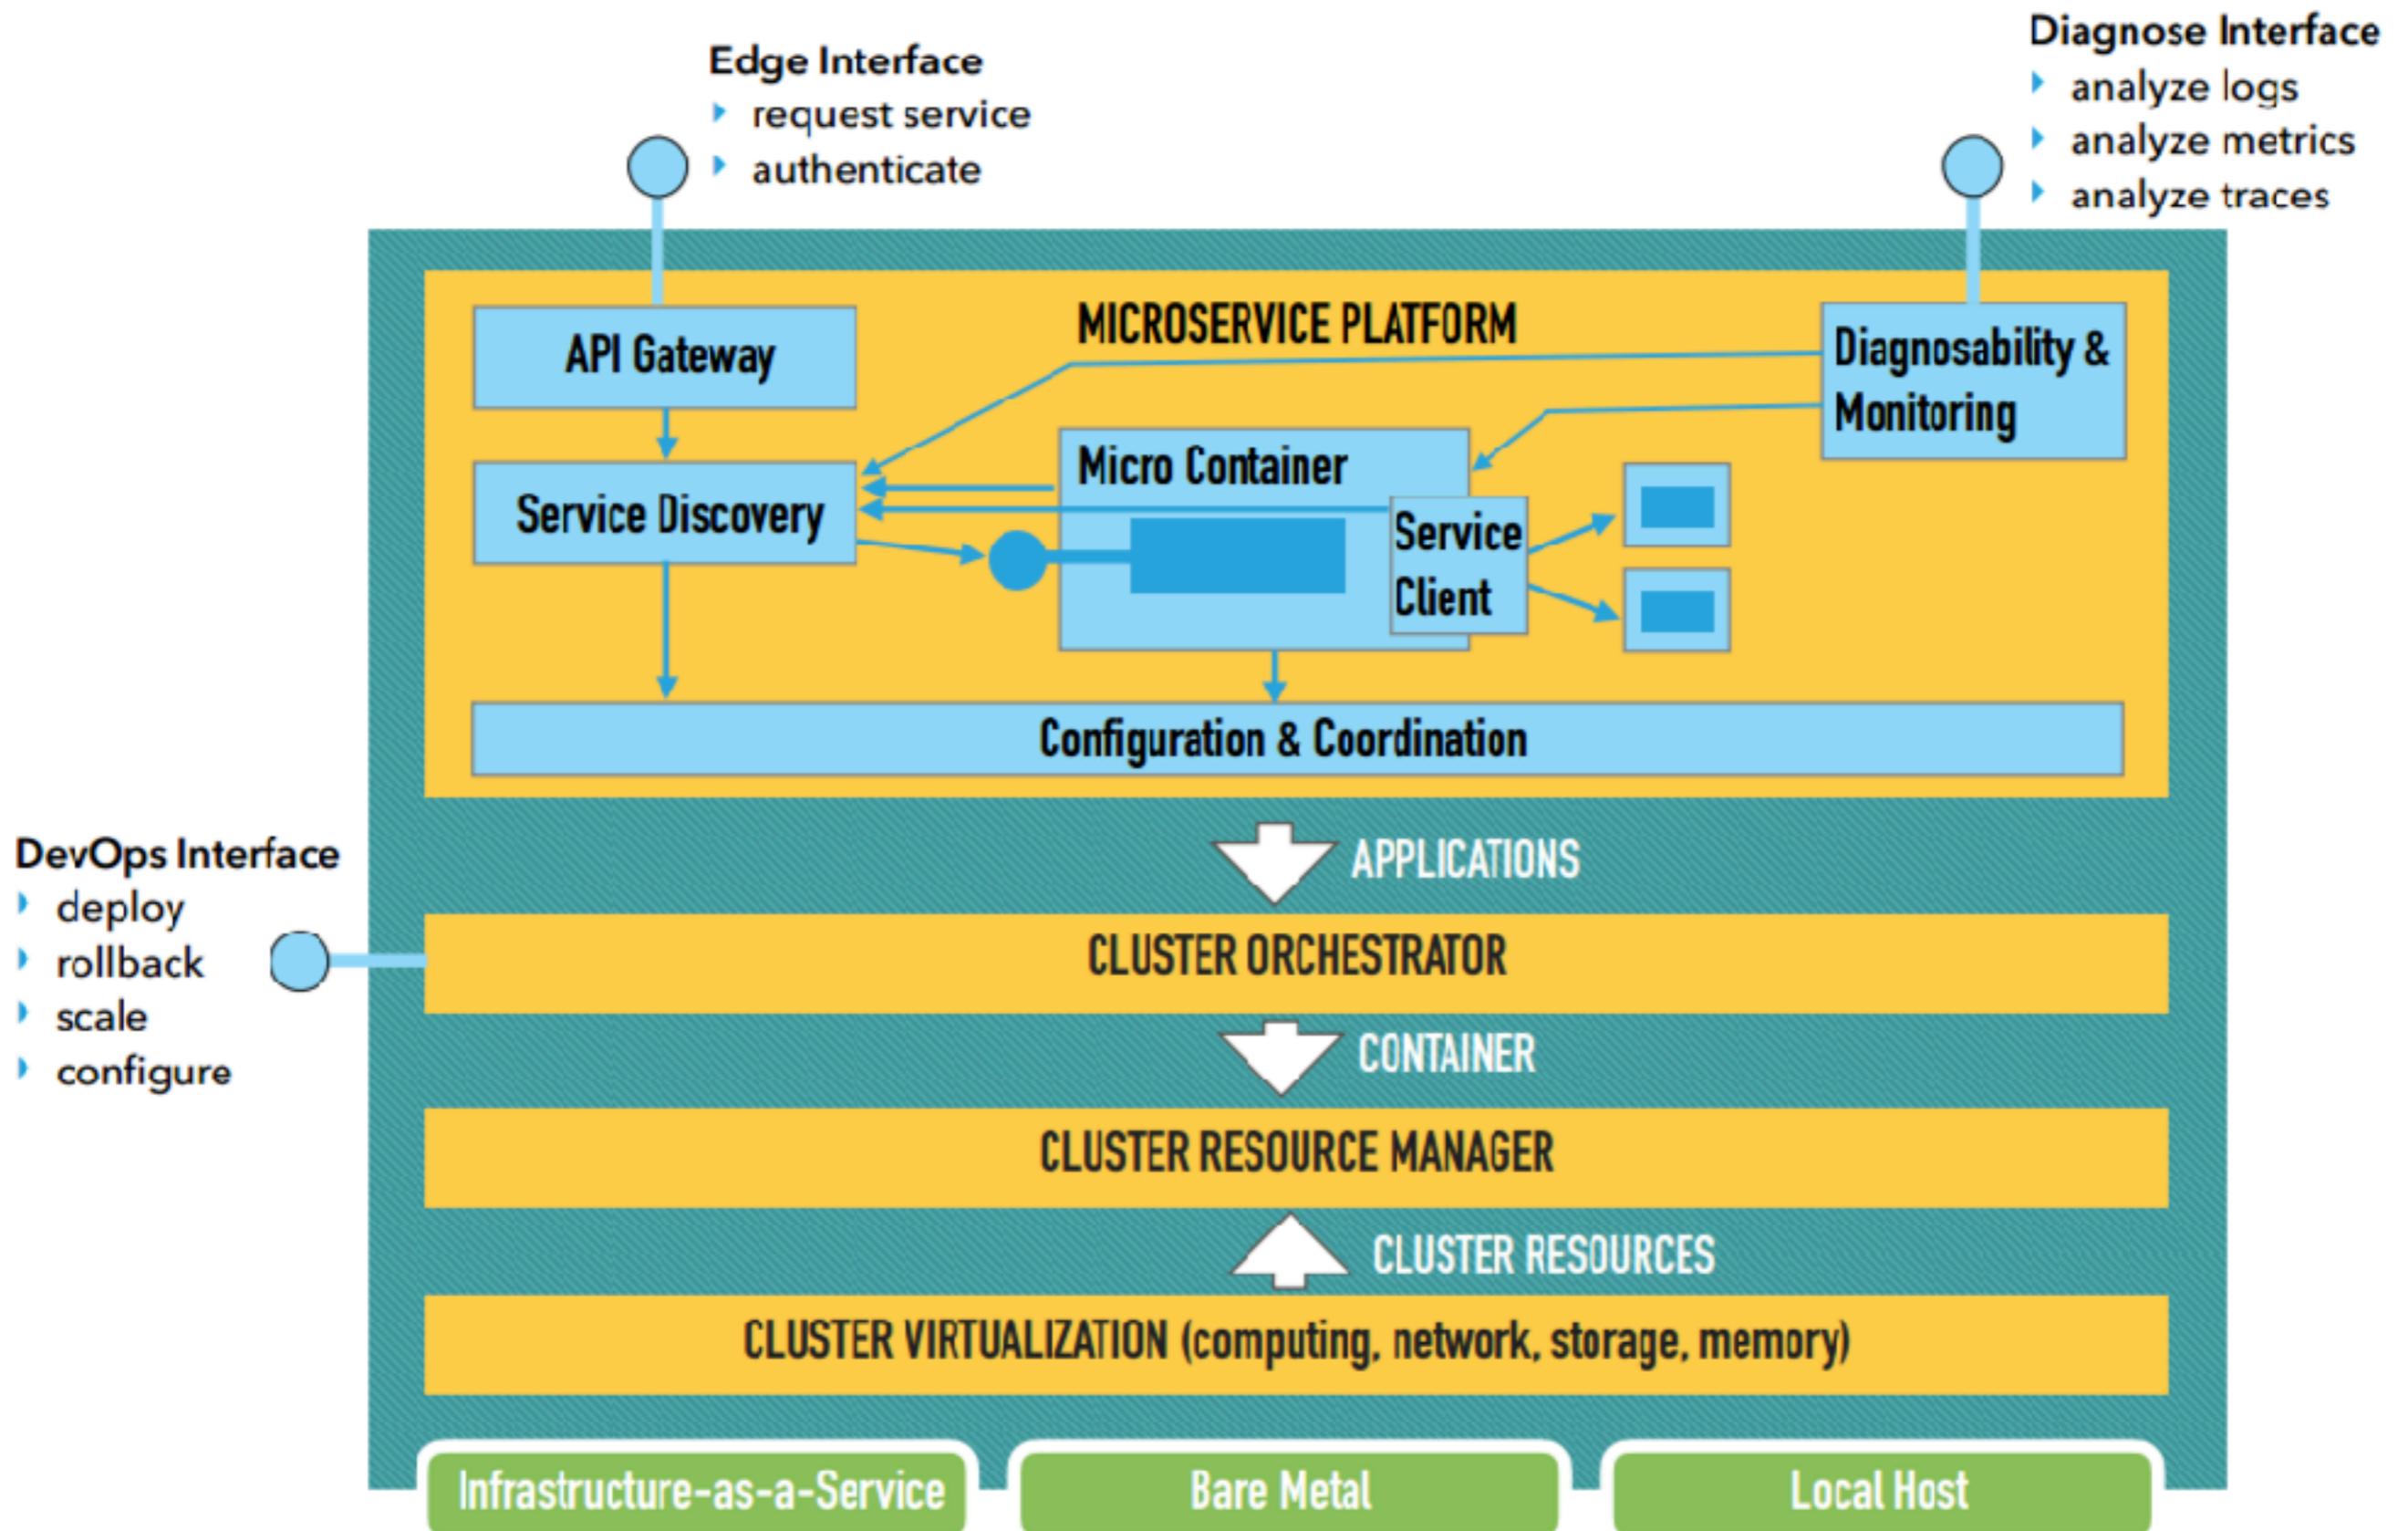
MESOS

Cluster Scheduler

JAVA JVM ON DC/OS

	JEE	Cloud Native
Applikation	WAR	Microservices
App-Framework	JEE, Spring	Spring Cloud
App-Server	WildFly, Tomcat, ...	Kubernetes
Betriebssystem	Linux	Mesos

JAVA JVM ON DC/OS



SPRING BOOT/CLOUD

▼ Core			
<input type="checkbox"/> Security	<input type="checkbox"/> AOP	<input type="checkbox"/> Atomikos (JTA)	<input type="checkbox"/> Bitronix (JTA)
<input type="checkbox"/> Narayana (JTA)	<input type="checkbox"/> Cache	<input type="checkbox"/> DevTools	<input type="checkbox"/> Configuration Processor
<input type="checkbox"/> Validation	<input type="checkbox"/> Session	<input type="checkbox"/> Retry	<input type="checkbox"/> Lombok
▼ Web			
<input type="checkbox"/> Web	<input type="checkbox"/> Websocket	<input type="checkbox"/> Web Services	<input type="checkbox"/> Jersey (JAX-RS)
<input type="checkbox"/> Ratpack	<input type="checkbox"/> Vaadin	<input type="checkbox"/> Rest Repositories	<input type="checkbox"/> HATEOAS
<input type="checkbox"/> Rest Repositories HAL Browser	<input type="checkbox"/> Mobile	<input type="checkbox"/> REST Docs	
▼ Template Engines			
<input type="checkbox"/> Freemarker	<input type="checkbox"/> Velocity	<input type="checkbox"/> Groovy Templates	<input type="checkbox"/> Thymeleaf
<input type="checkbox"/> Mustache			
▼ SQL			
<input type="checkbox"/> JPA	<input type="checkbox"/> JOOQ	<input type="checkbox"/> MyBatis	<input type="checkbox"/> JDBC
<input type="checkbox"/> H2	<input type="checkbox"/> HSQLDB	<input type="checkbox"/> Apache Derby	<input type="checkbox"/> MySQL
<input type="checkbox"/> PostgreSQL	<input type="checkbox"/> SQL Server	<input type="checkbox"/> Flyway	<input type="checkbox"/> Liquibase
▼ NoSQL			
<input type="checkbox"/> MongoDB	<input type="checkbox"/> Cassandra	<input type="checkbox"/> Couchbase	<input type="checkbox"/> Neo4j
<input type="checkbox"/> Redis	<input type="checkbox"/> Gemfire	<input type="checkbox"/> Solr	<input type="checkbox"/> Elasticsearch
▼ Social			
<input type="checkbox"/> Facebook	<input type="checkbox"/> LinkedIn	<input type="checkbox"/> Twitter	
▼ I/O			
<input type="checkbox"/> Batch	<input type="checkbox"/> Integration	<input type="checkbox"/> Activiti	<input type="checkbox"/> Apache Camel
<input type="checkbox"/> JMS (ActiveMQ)	<input type="checkbox"/> JMS (Artemis)	<input type="checkbox"/> JMS (HornetQ)	<input type="checkbox"/> AMQP
<input type="checkbox"/> Kafka	<input type="checkbox"/> Mail	<input type="checkbox"/> LDAP	
▼ Ops			
<input type="checkbox"/> Actuator	<input type="checkbox"/> Actuator Docs	<input type="checkbox"/> Remote Shell	

SPRING BOOT/CLOUD

▼ Cloud Core

Cloud Connectors

Cloud Bootstrap

Cloud Security

Cloud OAuth2

Cloud Task

▼ Cloud Config

Config Client

Config Server

Zookeeper Configuration

Consul Configuration

▼ Cloud Discovery

Eureka Discovery

Eureka Server

Zookeeper Discovery

Cloud Foundry Discovery

Consul Discovery

▼ Cloud Routing

Zuul

Ribbon

Feign

▼ Cloud Circuit Breaker

Hystrix

Hystrix Dashboard

Turbine

Turbine Stream

▼ Cloud Tracing

Sleuth

Zipkin Client

Sleuth Stream

Zipkin Stream

Zipkin UI

Zipkin Server

▼ Cloud Messaging

Cloud Bus AMQP

Cloud Bus Kafka

Stream Rabbit

Stream Kafka

▼ Cloud AWS

AWS Core

AWS JDBC

AWS Messaging

▼ Cloud Data Flow

Local Data Flow Server

Data Flow Shell

▼ Cloud Cluster

Cluster Redis

Cluster Zookeeper

Cluster Hazelcast

Cluster Etcd

▼ Cloud Contract

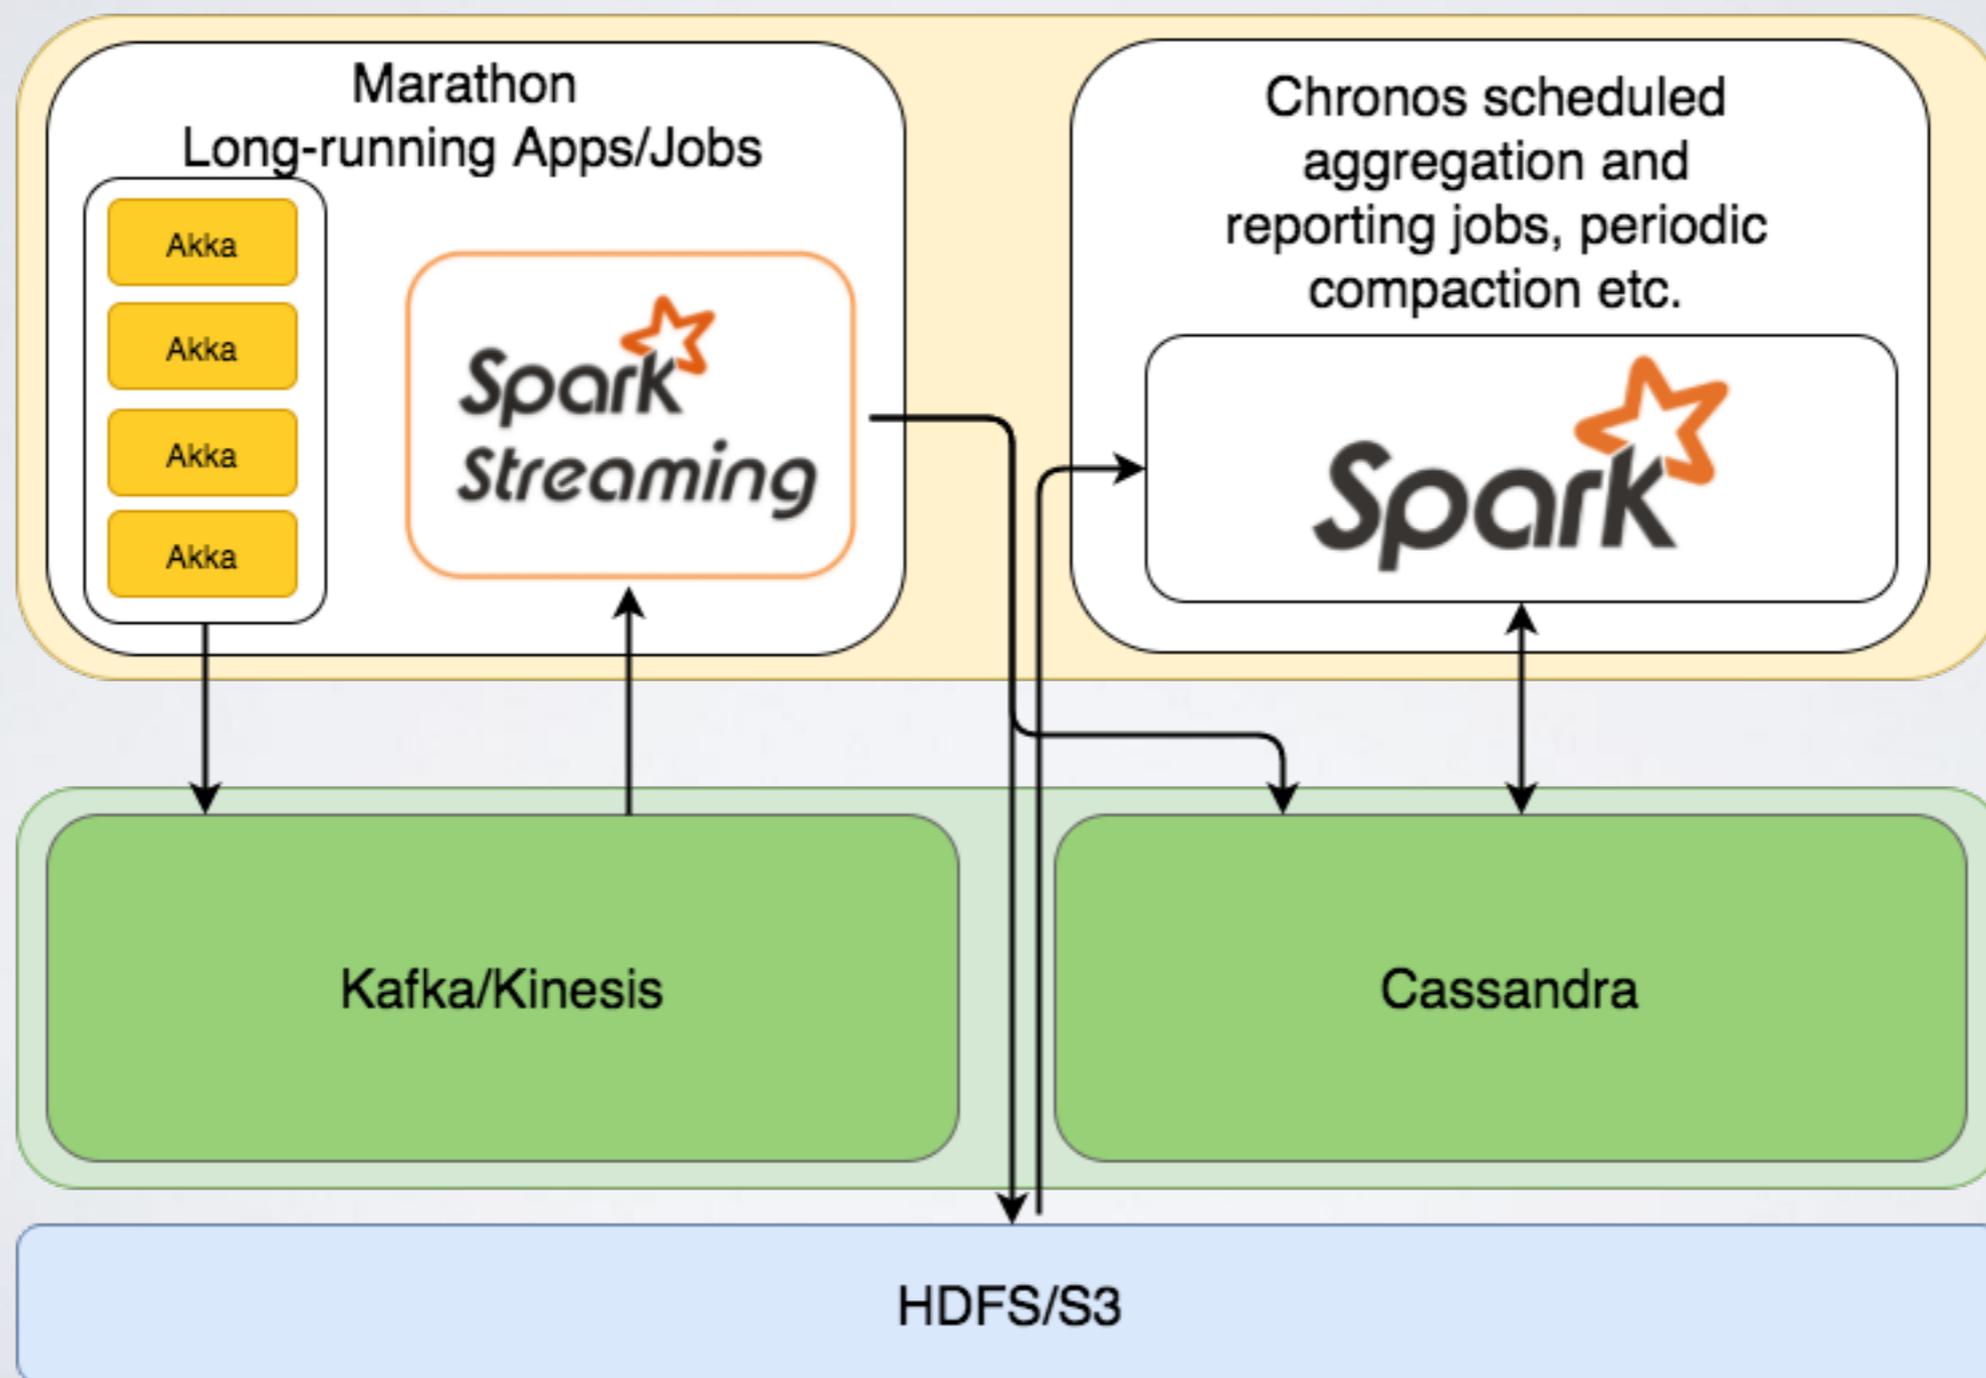
Cloud Contract Verifier

Cloud Contract Stub Runner

Cloud Contract WireMock

▶ Pivotal Cloud Foundry

SMARK堆栈



CI/CD

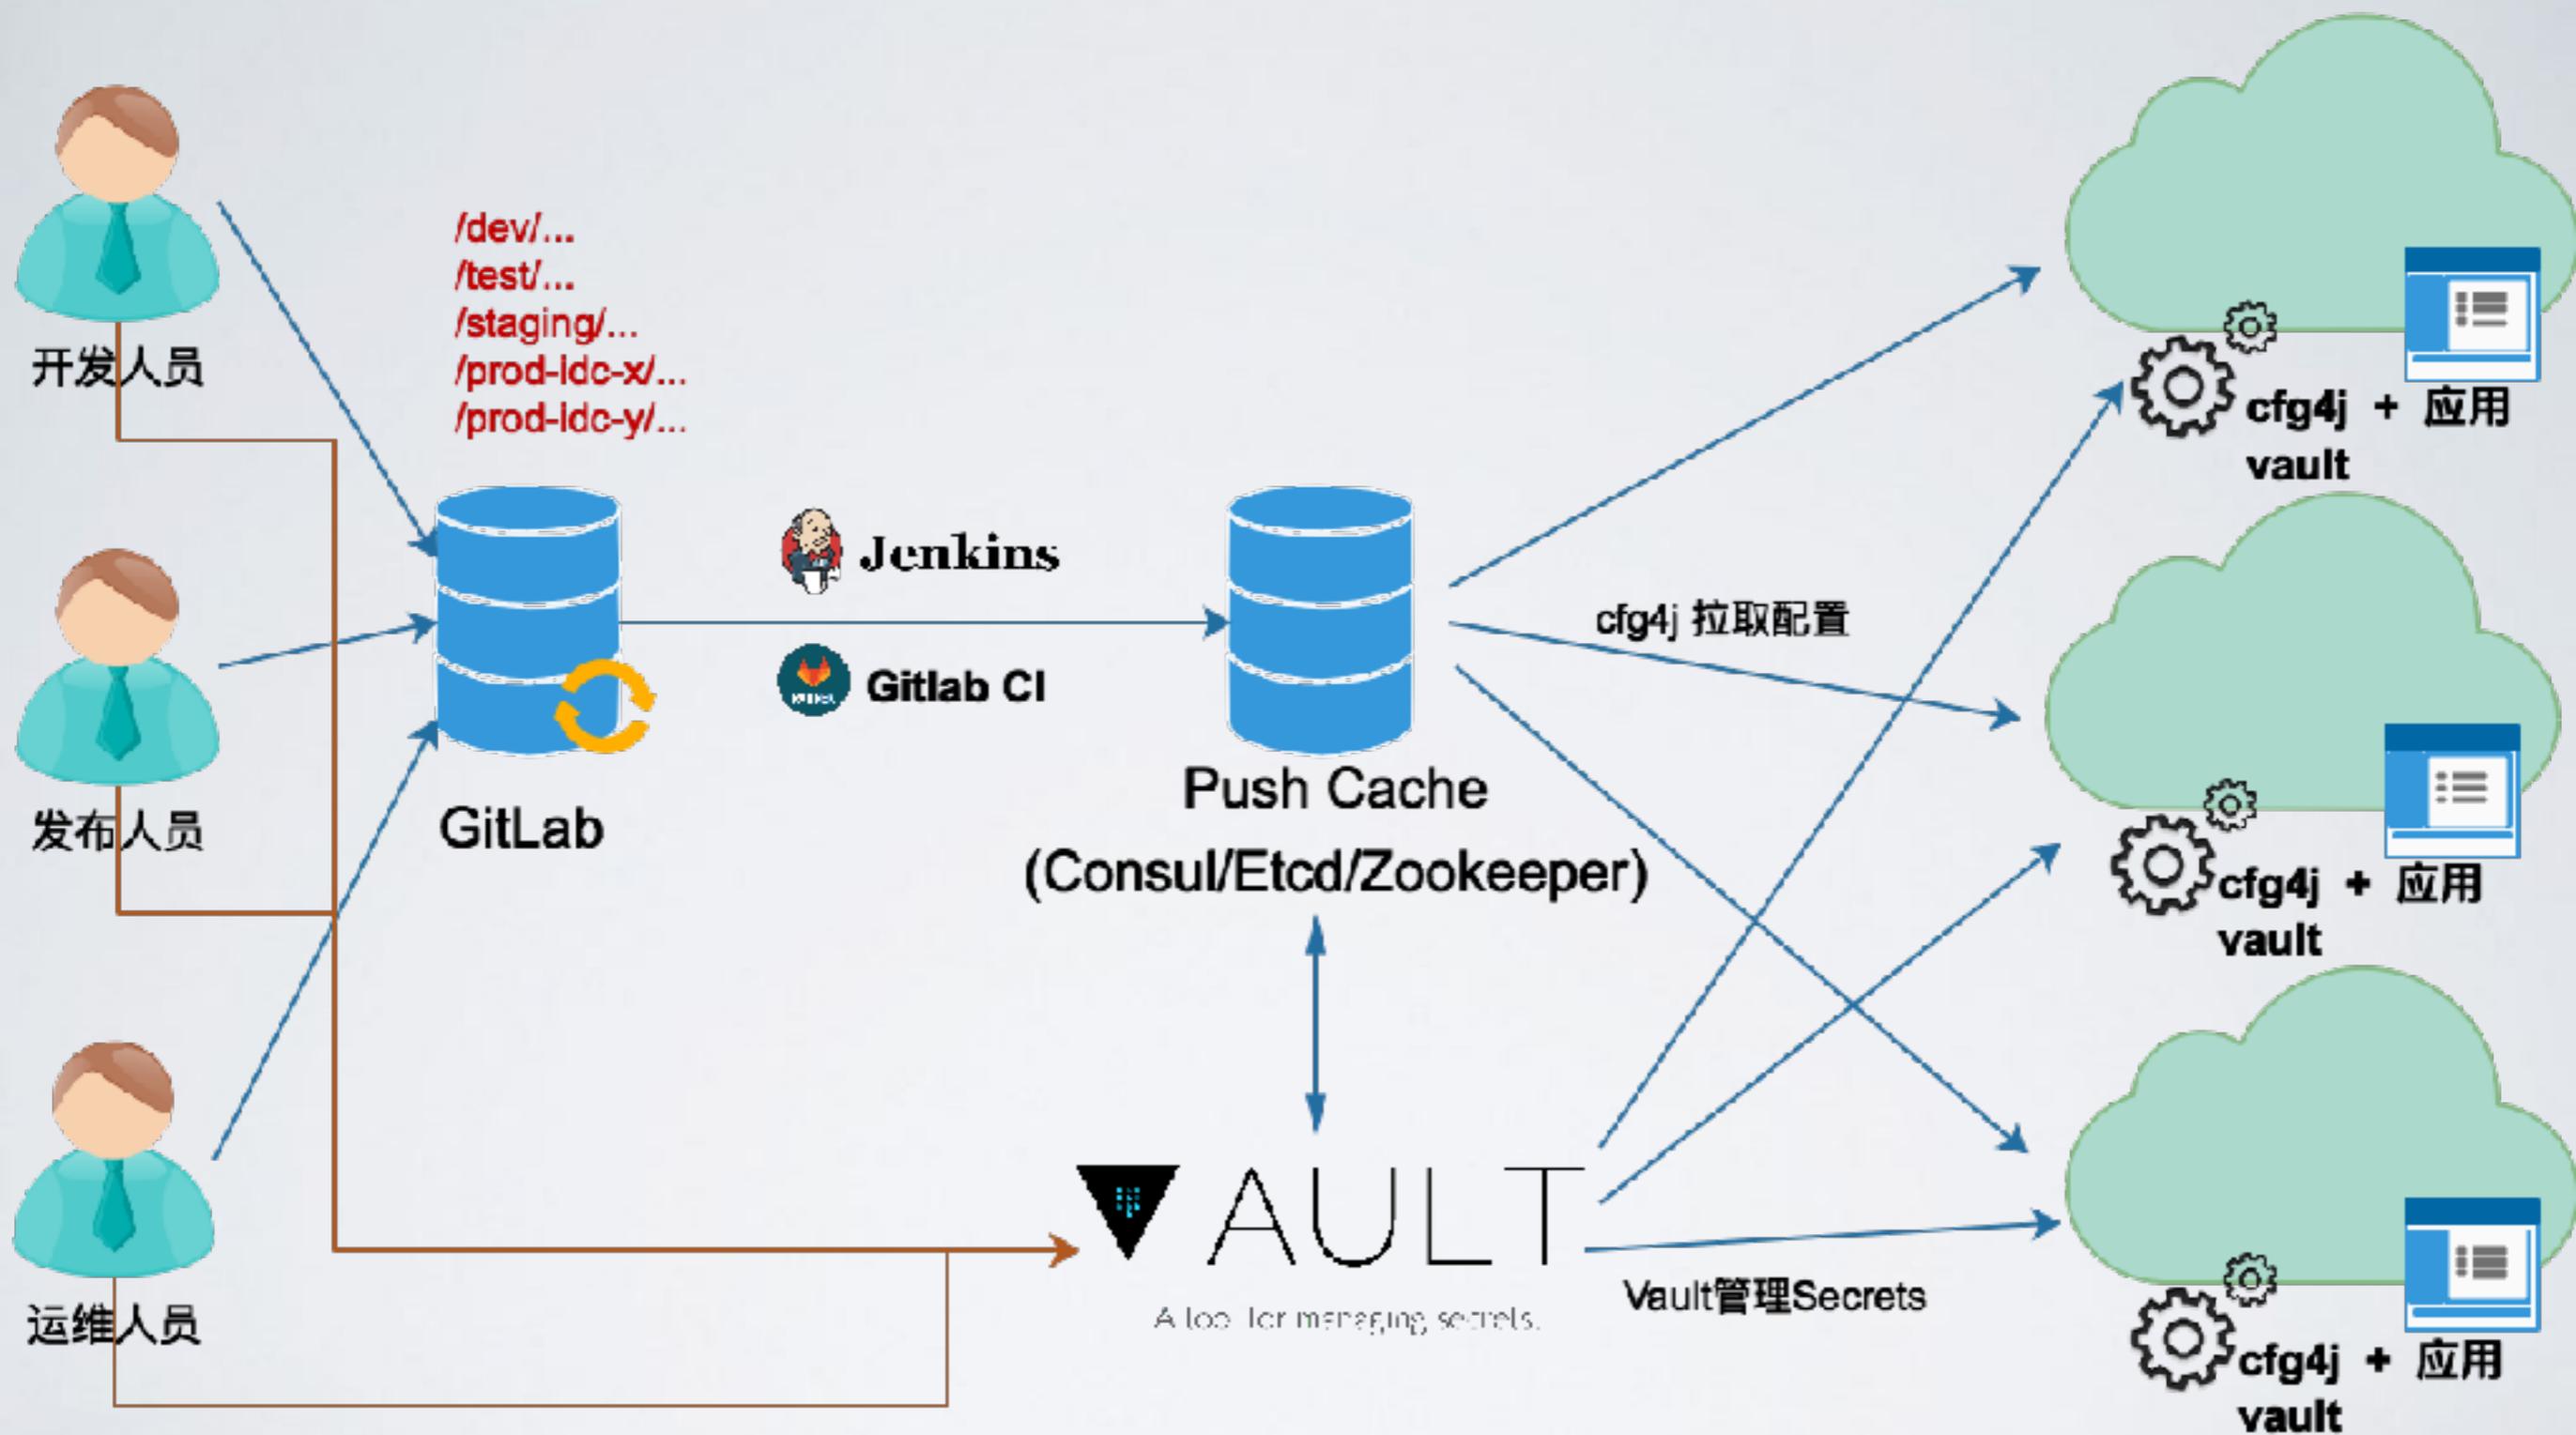
JENKINS ON DC/OS WITH VAMP

Average stage times:

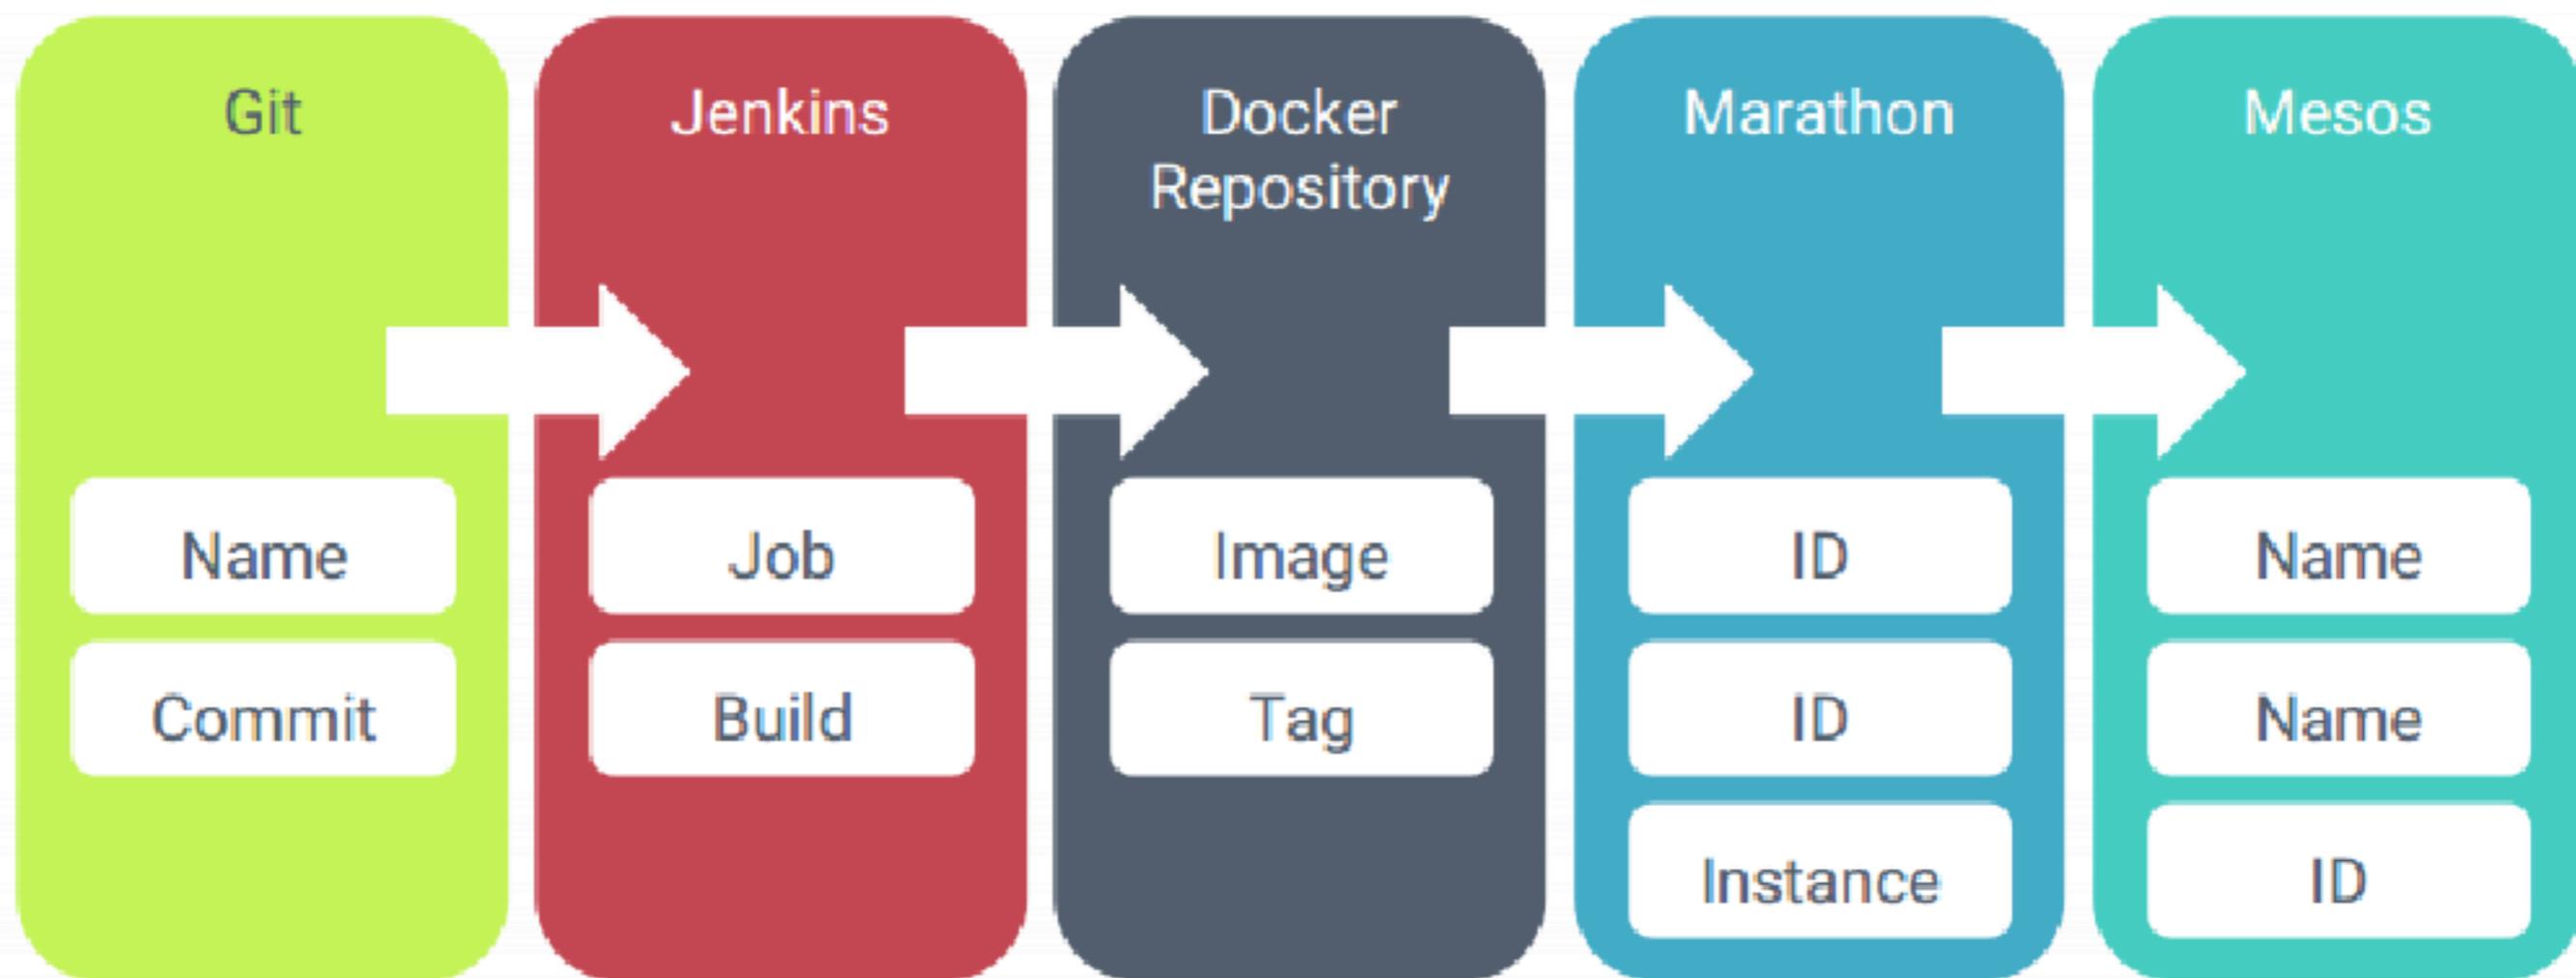
Checkout	Install	Test	Build Docker image	Push Docker image	Deploy to Vamp	Undeploy	Cleanup
3s	7s	1s	11s	14s	2min 3s	1s	309ms
1s	2s	1s	9s	13s	2min 36s	1s	309ms

39
3ep 18
17:09
No
Changes

CI/CD配置管理



部署流



CI/CD实施步骤

产品服务、研发运营、IT基础架构和商业流程的持续改善和自动化

1、提升PM/CI/CD/DevOps流程和效率 2、PaaS/API 3、私有公有混合云

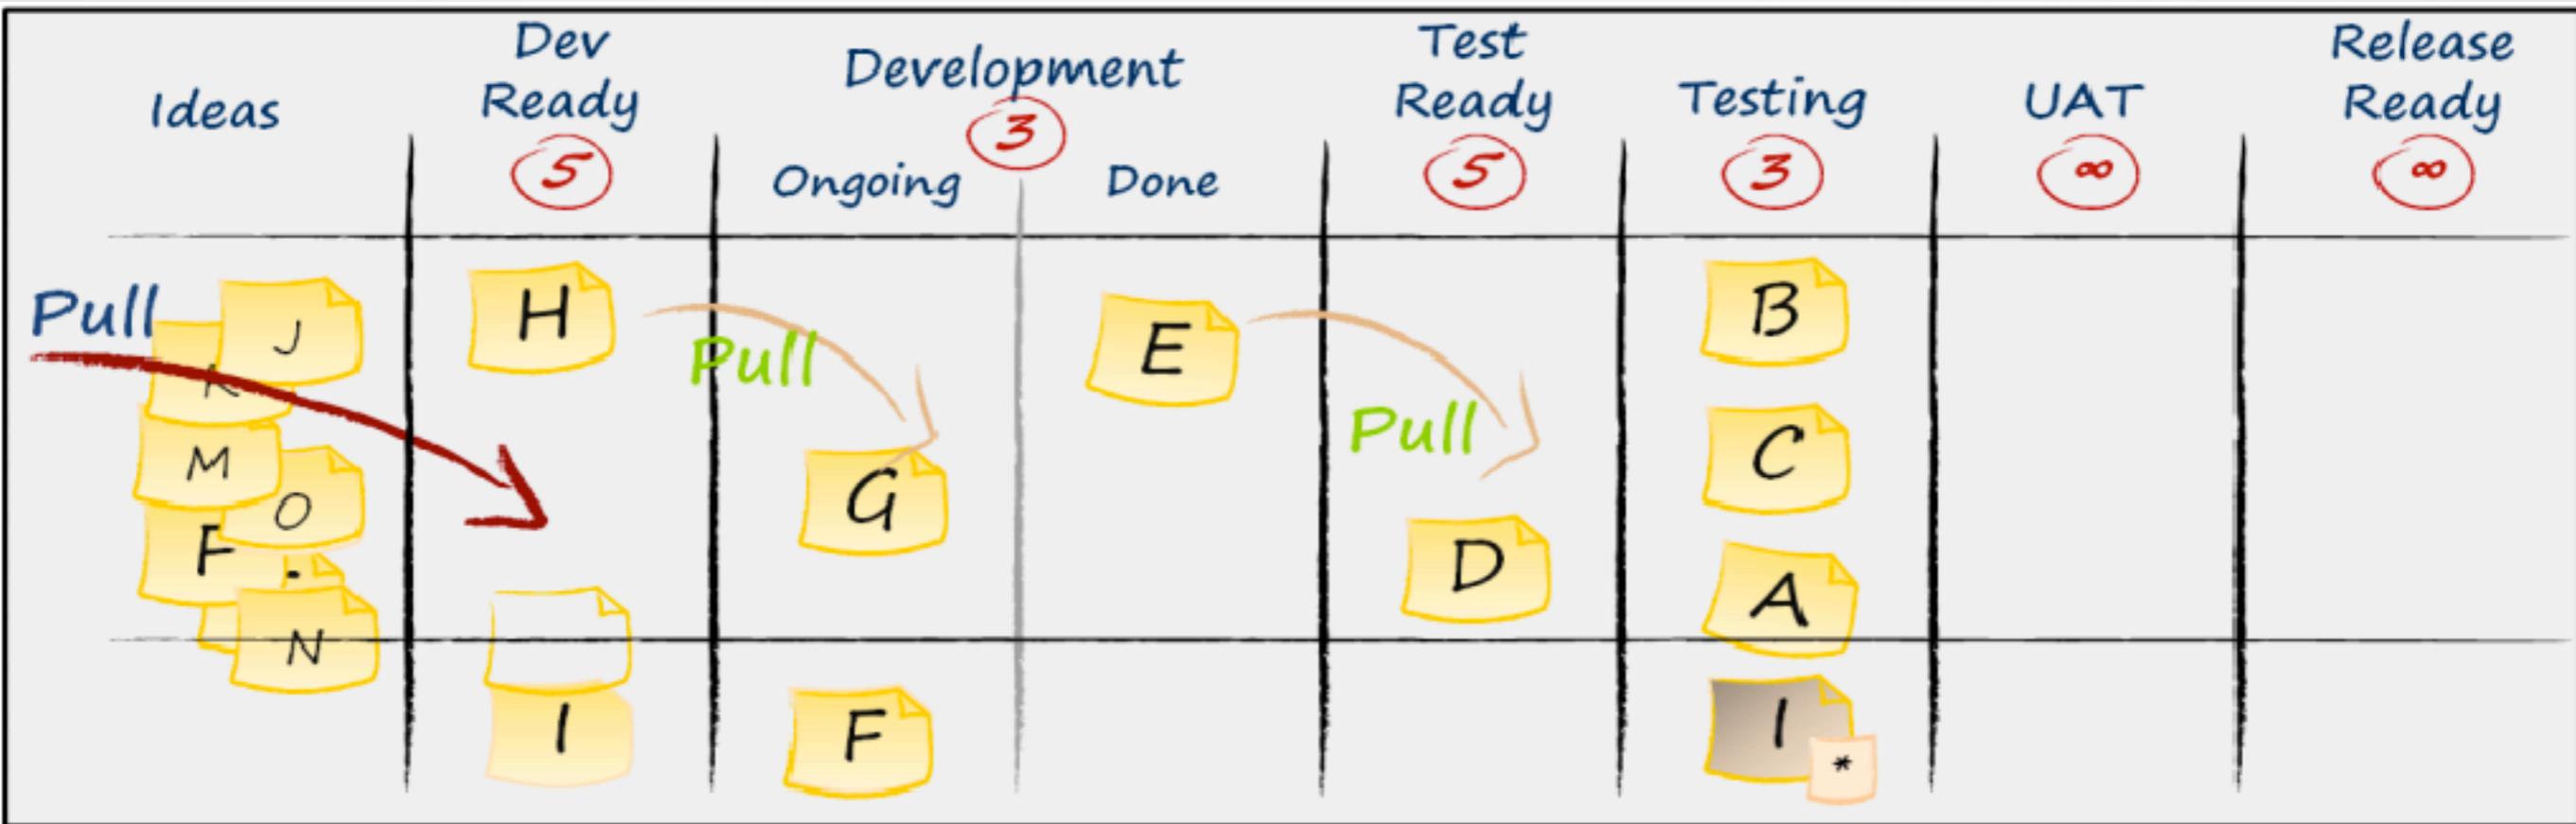
自动化

IT/商业的规范、适应和进化

1、Phabricator和GitLab项目协作 2、CI/CD/DevOps/BPM自动化 3、监控基础设施

看板方法

采用精益的思维范式，将软件开发视为一个价值流，并且基于PULL模式来驱动



高效协作

精益方法论 + PHABRICATOR

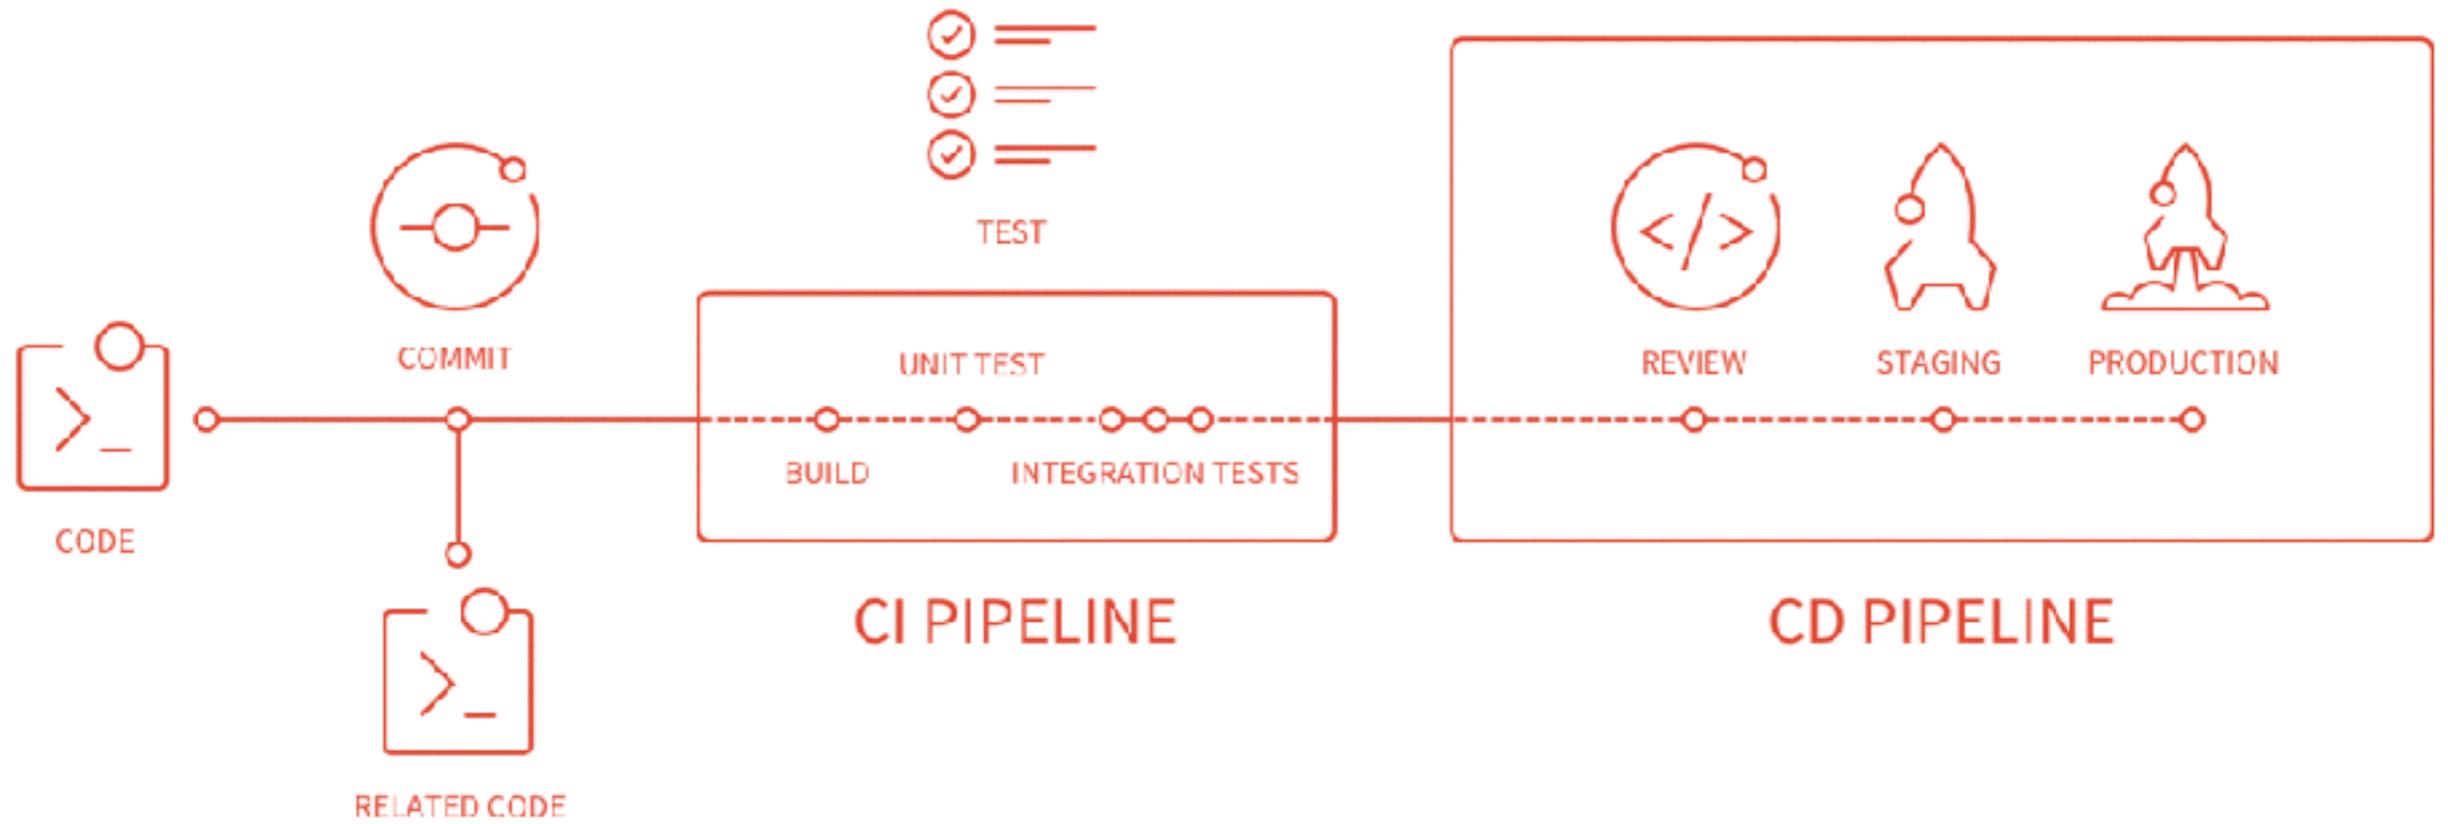
← → ↻ 安全 https://phabricator.wikimedia.org/tag/analytics-kanban/ ☆ [Browser Extensions] [Search] [Log In]

Phabricator [Natural] [Open Tasks] [Settings]

Projects > Analytics-Kanban > Workboard

- Parent Tasks** (16 | 0)
 - T185233** Eventlogging of the Future
 - T175344** Move away from jmxtrans in favor of prometheus jmx_exporter
User-Elukey [Patch-For-Review]
 - T90759** Create Daily & Monthly pageview dump with country data and Visualize on UI
[Analytics-Wikistats]
 - T131280** Making geowiki data public
 - T159170** Find an alternative query interface for eventlogging on analytics cluster that can replace MariaDB
[Analytics-EventLogging]
 - T143927** Per Family Unique Devices Counts
[Analytics-EventLogging]
 - T164019** Webrequest tagging and
- Next Up (Default)** (44 | 79)
 - T185776** Reduction of stat1005's disk space usage
 - T185602** Monitor and alert if no new data from JsonRefine jobs
[Analytics-EventLogging]
 - T185335** Wikistat Beta: expand topic explorer by default
[Analytics-Wikistats]
 - T185130** Hive EventLogging tables not updating since January 26
[Analytics-EventLogging]
 - T185124** Wikistats Beta: How about historic data
[Analytics-Wikistats]
 - T185122** Wikistats Beta Localization on the UI
[Analytics-Wikistats]
 - T185121** Remaining reports
- Paused** (12 | 14 / 10)
 - T185419** Mediaccounts missing top1000 files after 2018-01-01
[Analytics-Cluster] [Datasets-Archiving] [Datasets-Webstatscollector] [Patch-For-Review]
 - T174386** Remove sensitive fields from whitelist for QuickSurvey schemas (end of Q2)
[Patch For Review]
 - T174110** Private data access for non-person user that calculates metrics
[Discovery] [Discovery-Analysis]
 - T185009** Handle long-term retention of ChangesListFilters and ChangesListFilterGrouping schemas
 - T182718** SEO-friendly HTML titles for Wikistats 2.0
[Analytics-Wikistats]
- In Progress**
 - T181064** Purge re after 90 days
[Analytics-F] [Patch-For-Review]
 - T184759** Sqoop co geowiki
[Analytics-F]
 - T186510** Eventlog forwarder/zmq-leg over time
[User-Elukey]
 - T185952** EventLog
[Analytics-E] [Beta-Cluster-Inf] [User-Elukey]
 - T184768** Bug beha for quantileBounds
[Wikipedi] [Discover] [Mobile]

GITLAB CI/CD



[Project](#) [Activity](#) [Repository](#) [Pipelines](#) [Graphs](#) [Issues](#) 2 [Merge Requests](#) 1 [Wiki](#)

[Pipelines](#) [Builds](#) [Environments](#) [Cycle Analytics](#)

All 34 [Running](#) 0 [Branches](#) [Tags](#)

[Run pipeline](#)

[CI lint](#)

Status	Pipeline	Commit	Stages	Duration
passed	#34 by latest	develop 895db605 Revert: feat(all): 后台微信支付		00:15 7 minutes ago
passed	#31 by	develop 3615c759 feat(Payment): 支持商家自主选择可用...		16 minutes ago

JENKINS CI/CD

- Back to Dashboard
- Status
- Changes
- Build Now
- Delete Pipeline
- Configure
- Move
- Full Stage View

Pipeline tasks-cd-pipeline

 [Recent Changes](#)

 [add description](#)



Build History [trend](#)

find

- #4 Jun 14, 2016 6:28 AM
- #3 Jun 14, 2016 6:25 AM
- #1 Jun 14, 2016 5:37 AM

[RSS for all](#) [RSS for failures](#)

Stage View

Average stage times: (Average full run time: ~3min 40s)

	Build	Test and Analysis	Push to Nexus	Deploy DEV	Deploy STAGE
#4 Jun 14 12:28 1 commits	51s 410037b09d9	1min 40s 410037b09d9	21s 410037b09d9	44s 410037b09d9	8s 410037b09d9
#3 Jun 14 12:25 1 commits	31s 3d99248035	1min 0s 3d99248035 failed			
#1 Jun 14 11:37 No Changes	1min 15s 145eb047318	1min 7s 145eb047318	35s 145eb047318	28s 145eb047318	7s 145eb047318

 [Latest Test Result](#) (no failures)

JENKINS CI/CD



GITLAB ON DC/OS

 Catalog > gitlab



gitlab

COMMUNITY 1.0-9.1.0

REVIEW & RUN

By deploying you agree to the [terms and conditions](#)

Description

Collaboration and source control management: code, test, and deploy together! Installation Guide: <https://github.com/dcos/examples/tree/master/1.8/gitlab>

Pre-Install Notes

Having Marathon-lb installed and setting a virtual host in the config is required for being able to access GitLab. Setting the pinned hostname is required if using the single-node install (the default).

Information

SCM: <https://gitlab.com/gitlab-org/gitlab-ce>

Maintainer: support@gitlab.com

Licenses

MIT License: <https://gitlab.com/gitlab-org/gitlab-ce/raw/master/LICENSE>

EE License: <https://gitlab.com/gitlab-org/gitlab-ee/raw/master/LICENSE>

JENKINS ON DC/OS

 Catalog > jenkins



jenkins

CERTIFIED 3.4.0-2.89.2

REVIEW & RUN

By deploying you agree to the [terms and conditions](#)

Description

Jenkins is an award-winning, cross-platform, continuous integration and continuous delivery application that increases your productivity. Use Jenkins to build and test your software projects continuously making it easier for developers to integrate changes to the project, and making it easier for users to obtain a fresh build. It also allows you to continuously deliver your software by providing powerful ways to define your build pipelines and integrating with a large number of testing and deployment technologies.

Pre-Install Notes

WARNING: If you didn't provide a value for `storage.host-volume` (either using the CLI or via the Advanced Install dialog), YOUR DATA WILL NOT BE SAVED IN ANY WAY.

Information

SCM: <https://github.com/mesosphere/dcos-jenkins-service.git>

Maintainer: support@mesosphere.io

Licenses

Apache License Version 2.0: <https://github.com/mesosphere/dcos-jenkins-service/blob/master/LICENSE>

WAMP ON DC/OS

 Catalog > vamp



vamp

COMMUNITY 0.9.4-0.0.1

REVIEW & RUN

By deploying you agree to the [terms and conditions](#)

Description

Canary test/release and autoscaling platform for DC/OS

Pre-Install Notes

This DC/OS Service is currently in preview. In order to run Vamp you need to specify connection parameters to Elasticsearch. See installation instruction at <http://vamp.io/documentation/installation/dcos/>

Information

SCM: <https://github.com/magneticio/vamp.git>

Maintainer: vamp@magnetic.io

Licenses

Apache License Version 2.0: <https://raw.githubusercontent.com/magneticio/vamp/master/LICENSE>

Media

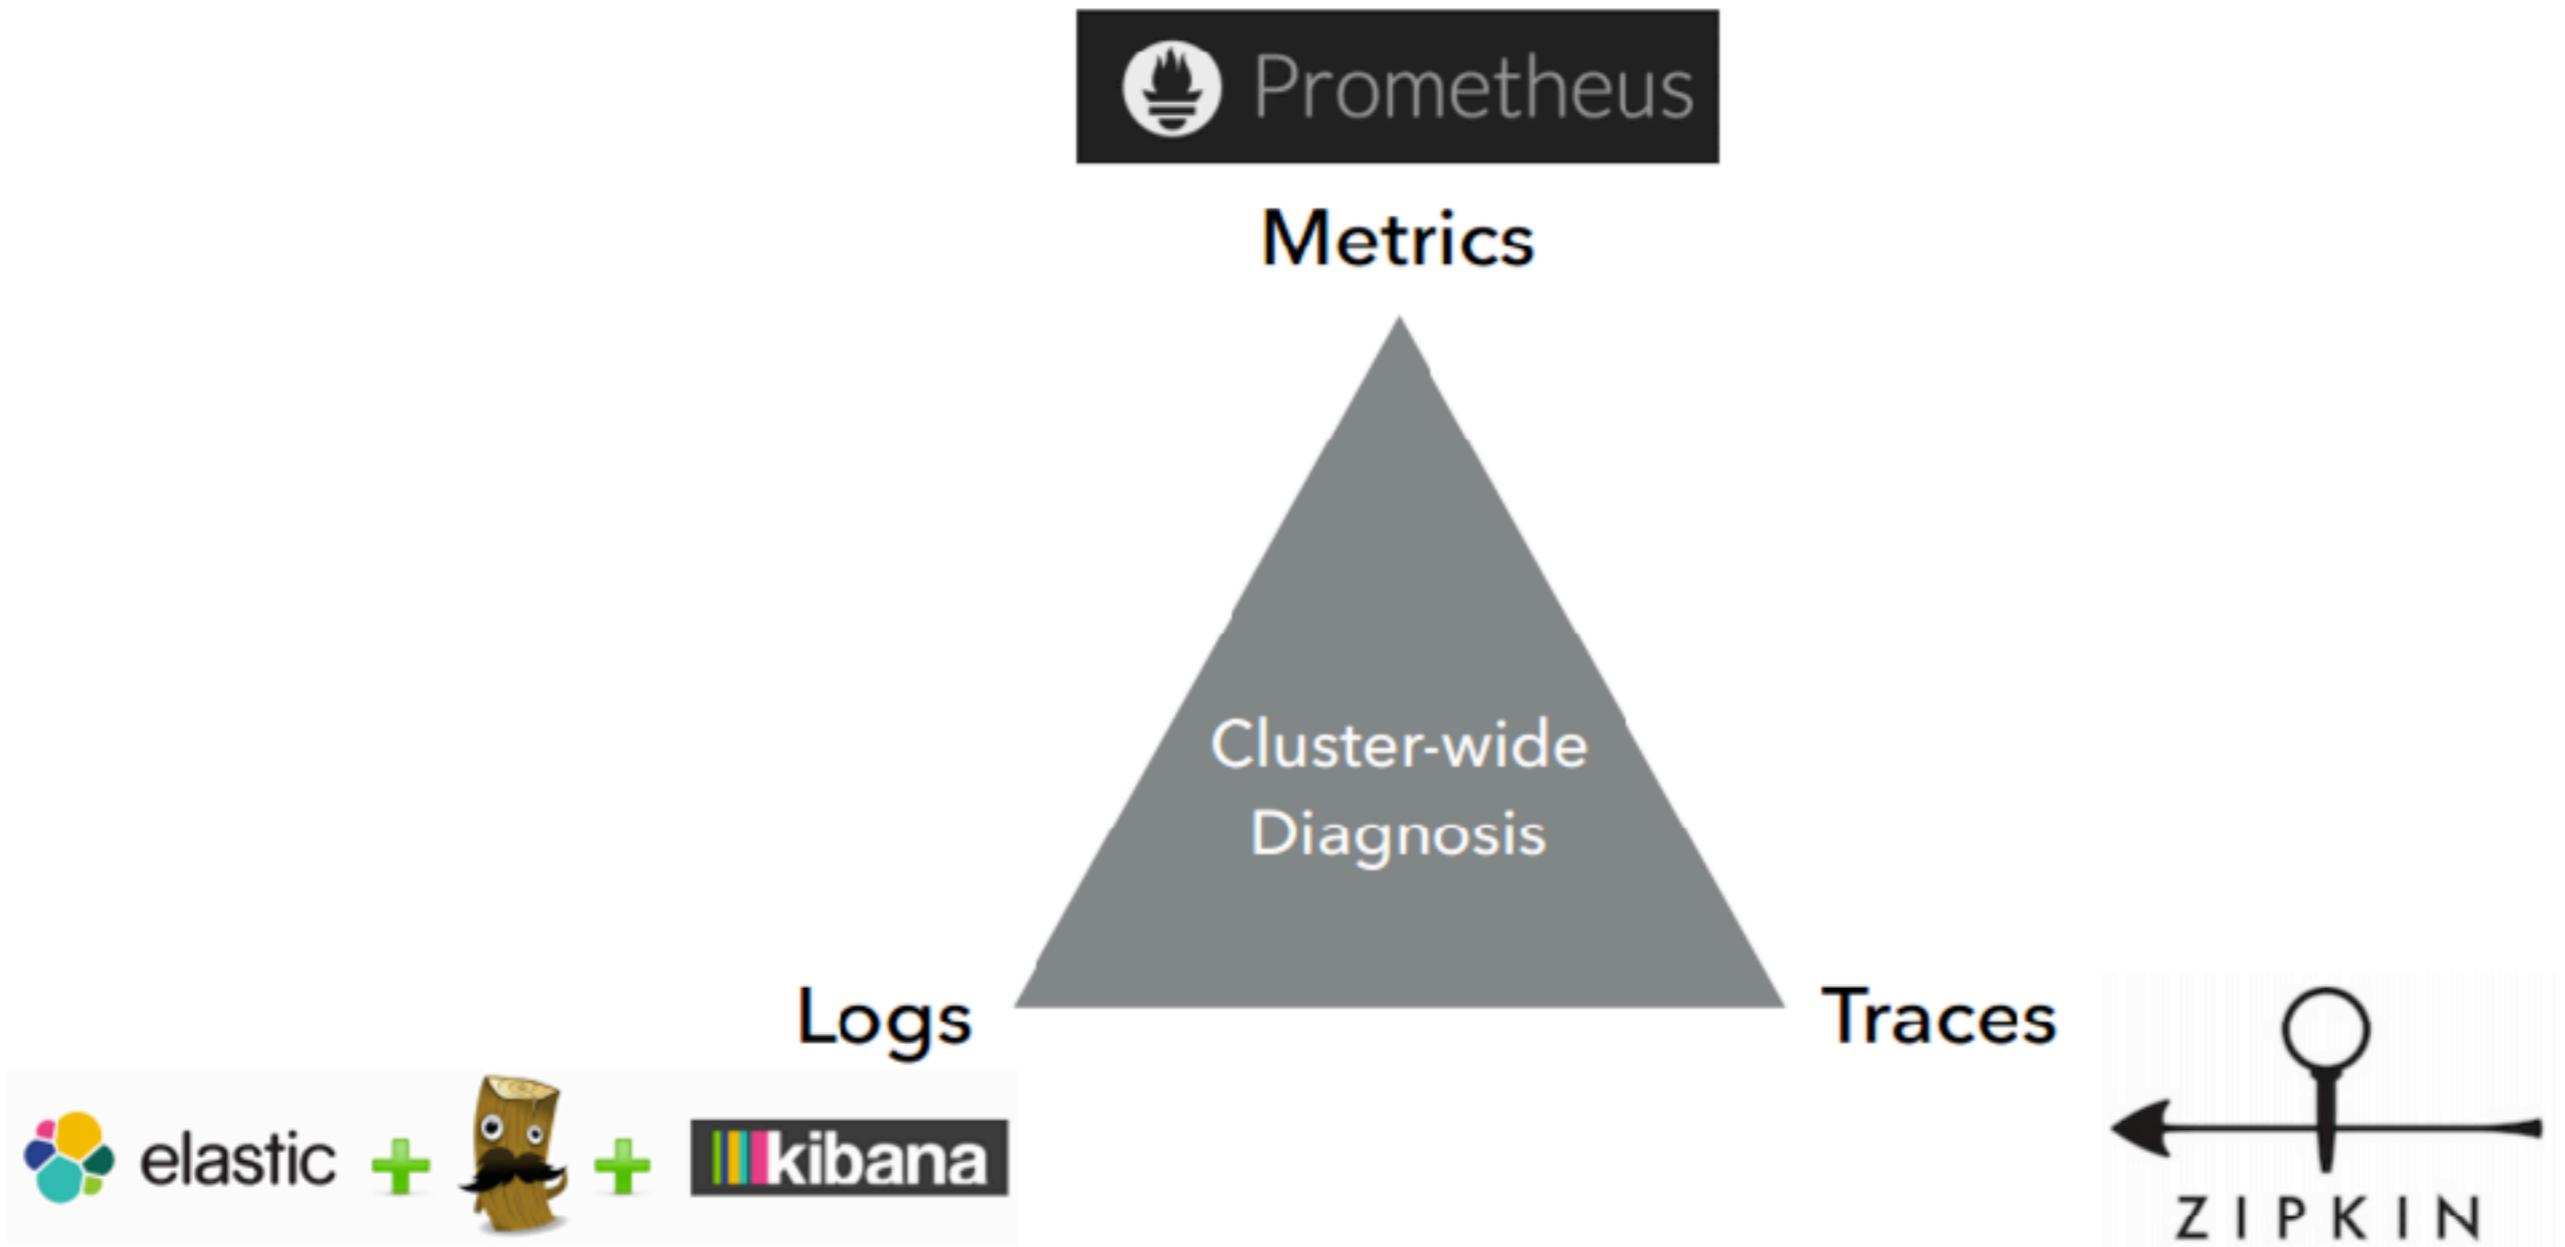


通过WAMP安装JENKINS

JENKINS ON DC/OS WITH VAMP

1、安装Vamp on DC/OS 2、Jenkins on DC/OS with Vamp 3、集成GitLab and GitLab CI

MONITORING & DIAGNOSABILITY



监控系统

监控所有，基于时间序列数据进行有效报警

系统将收集时间序列信息作为监控系统的首要工作，通过对时间序列信息进行操作(同时发展成PromQL信息查询语言)，将数据转化为图表和报警取代以前的探针脚本。

Prometheus: A Next-Generation Monitoring System

EXPORTER主/代理指标

PROMETHEUS MESOS EXPORTER

https://github.com/mesosphere/mesos_exporter

知识星球



龙辉

灰狐的朋友们



请使用微信扫描二维码加入星球

谢谢

欢迎订阅灰狐会员服务：灰狐的朋友们