Automatic datacenter and service deployments based on capacity planning artifacts

Xiaoxiang(XX) JIAN
Apsara Infrastructure - Alibaba

Alibaba Group

About Me

- Xiaoxiang(XX) JIAN
  - aka @jianxx
  - Apsara Infrastructure team in Alibaba
  - Leader of Cluster Infrastructure in Alibaba
- Middleware (2007~2011)
- Virtualization & Software-defined Data Center (2011~2014)
- Technical Infrastructure (2014~Now)

Alibaba Group

How to deploy a datacenter for cloud service?

Alibaba Group

Difficulties to deploy a datacenter

- Capacity planning
- Server placement
- Network cabling
- Network bootstrap
- OS bootstrap
- Service bootstrap
- ……

Our Goal

- Capacity planning by business requirement
- Automatic deployment
- Auto-healing for hardware failure

Our Solution

- Immutable bare-metal infrastructure
- Describe the datacenter with artifacts

Alibaba Group

IMMUTABLE BARE-METAL INFRASTRUCTURE

Traditional Deployment Method

- Build-in place approaches
  - Start from the foundation
  - Follow the delivery pipeline

Alibaba Group

Traditional Deployment Method

- Pros
  - Easy to learn for the beginners
- Cons
  - Impossible to rollback
  - Difficult to update
  - Difficult to lock down

Better solution: let’s lock it down

- Make the configuration before deployment
- Make the runtime state-driven
- A good example: Kubernetes

Immutable Infrastructure: Image based Deployment

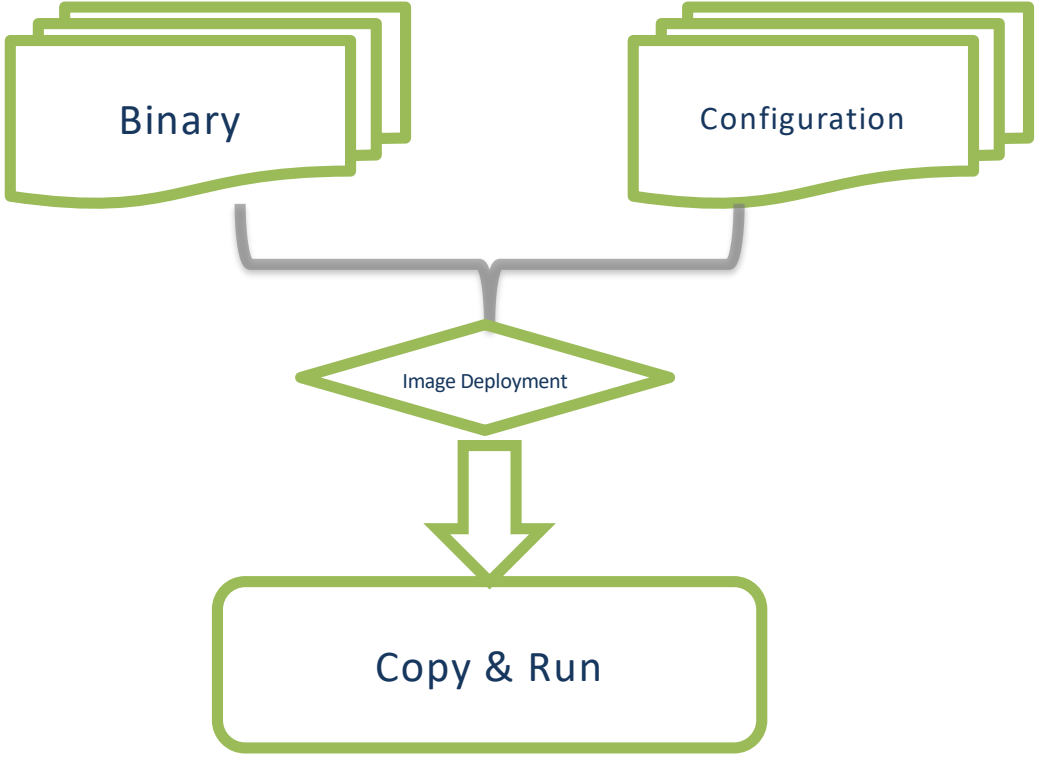

Alibaba Group

Image building process

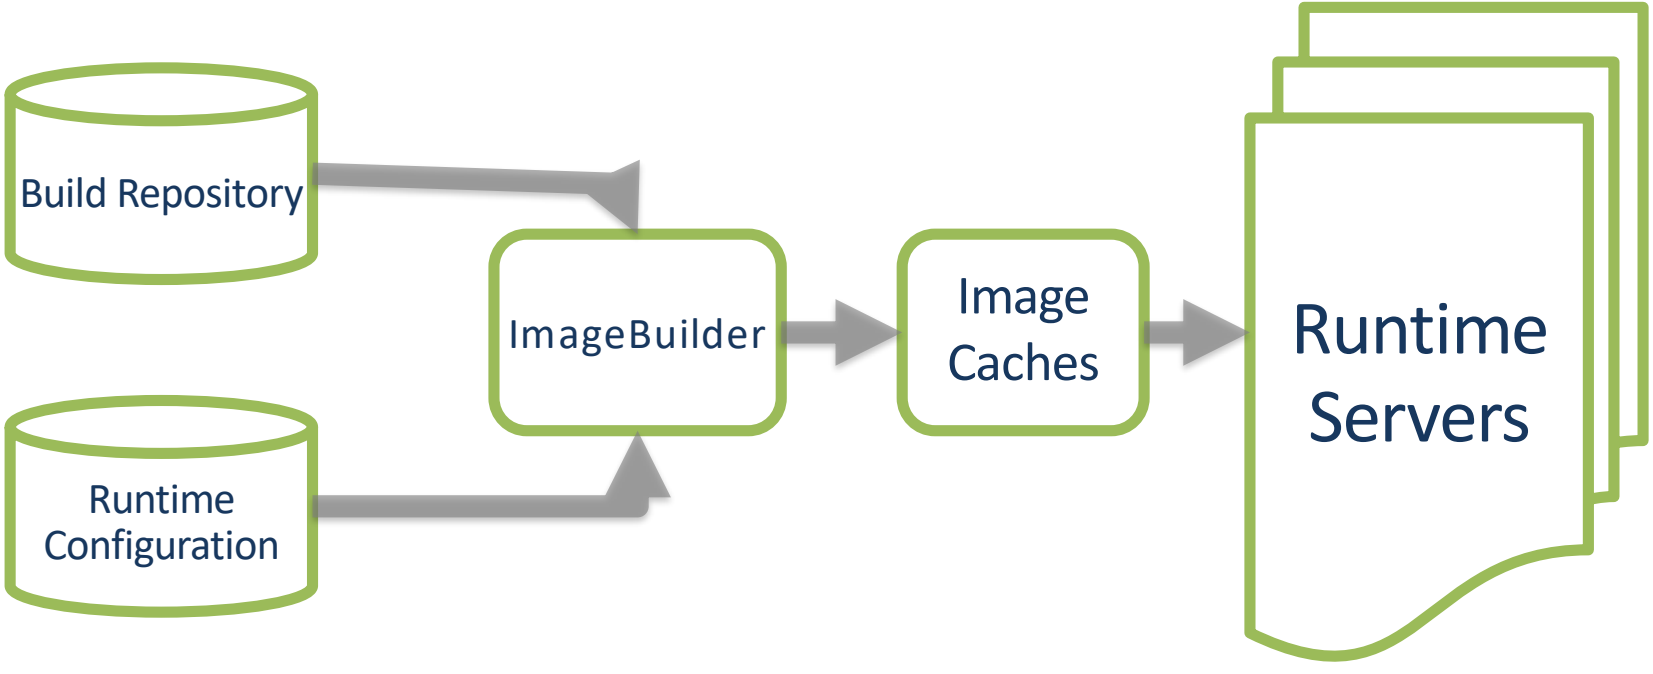

Normal immutable infrastructure

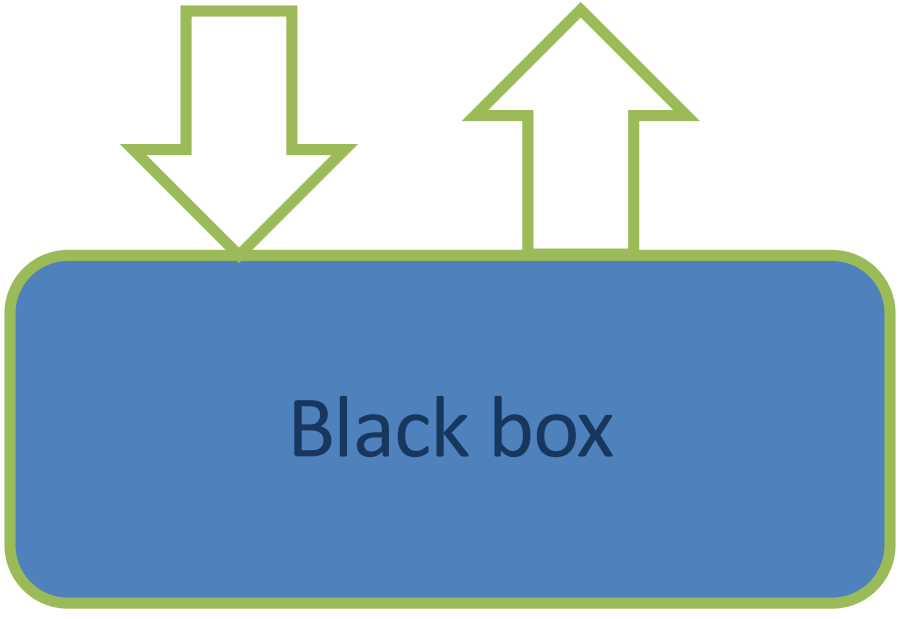

- Infrastructure is treated as a black box
- Changes are applied based on state description

Immutable infrastructure in a bare-metal environment

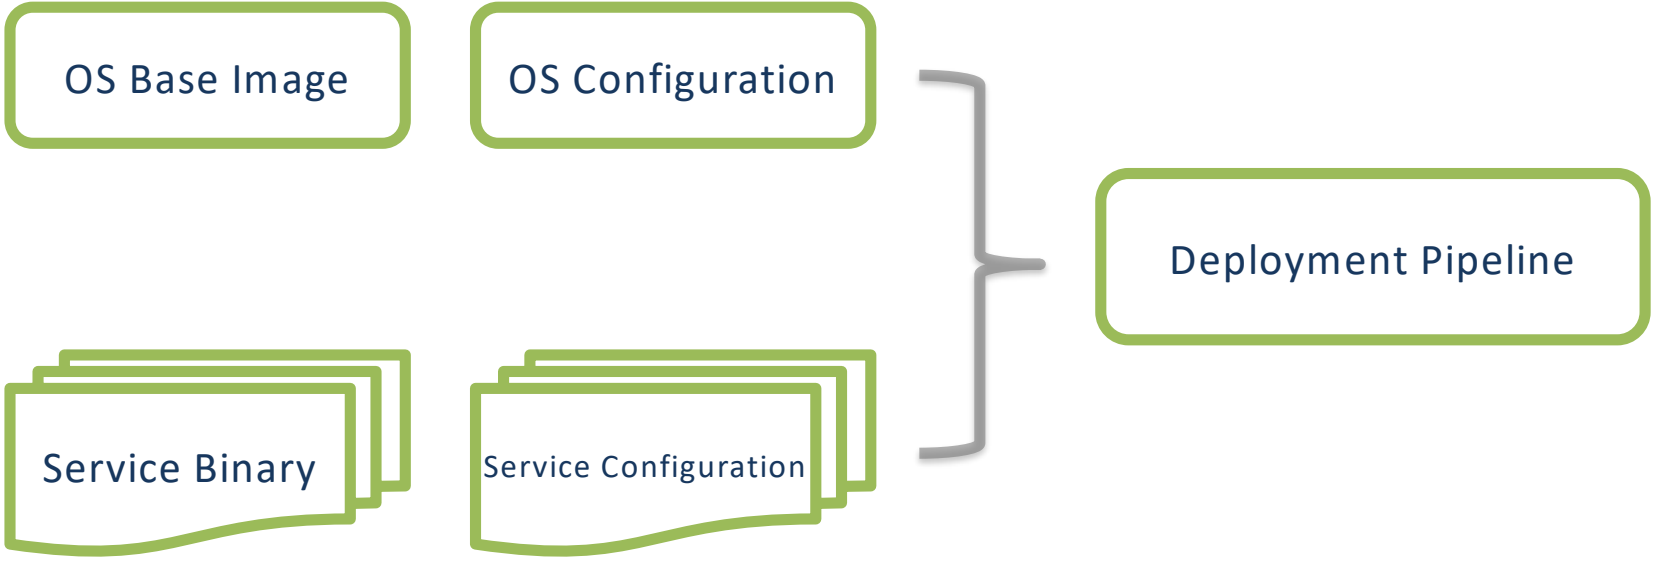

DESCRIBE THE DATACENTER WITH ARTIFACTS

Our approach to deploy a datacenter with cloud service

Business planning phase

- Categorize the services
- Capacity planning based on service measures
- Network planning

Delivery planning phase

- IDC planning
- Network configuration configuration
- OS/Service configuration generation

Alibaba Group

Concepts in the artifacts

- Product
  - The final deliverable for business.
- Service
  - Software concept deployed on a cluster.
- Application
  - The real thing runs as process on a server.

Product

- Products are composed by services.
- Product describes the deployment topology.
- Product describes the exposed features.
- Example: Alibaba Elastic Compute Service

Product – Services

- Id or name
- Version.
- Configuration templates.

Product – Deployment topology

- Cluster descriptions
  - Services
  - Capacity measures
- Service deployment topology

Product - Features

- Deployment requirement.
- Capacity measures.

Service

- Service is composed by serverroles.
- Service might have dependencies.
- ServerRoles might be in groups.
- Example: BlockStorage

Service - ServerRoles

- Types
- Server requirements.
- OS requirements.
- Applications.

Service - Dependencies

- Dependency on a node.
- Dependency on a cluster.

Service - ServerRoleGroups

- How to place the serverroles on a cluster?

Application

- Normal application.
- Docker application.
- ……

Data center described in artifacts

- IDC/Racks
- Network
- Server clusters
  - OS
  - Services

What do we have get now?

- An immutable bare-metal infrastructure
- Artifacts which describe the datacenter.
  - Generated in planning phase
- Datacenter bootstrap
  - From 2 months to 24 hours.
- Available Product
  - Apsara Stack from Alibaba Cloud

Summary

- Immutable bare-metal infrastructure
  - State-driven
  - Describe everything with state-based configuration files
  - Image based deployment
- Artifacts which describe the datacenter.
  - Model concepts for software services
  - Model concepts for the datacenter

Questions?

THANK YOU!