



NOVEMBER 2 & 3, 2015 | BERLIN

CLOUD **FOUNDRY**
SUMMIT



NOVEMBER 2 & 3, 2015 | BERLIN

CLOUD **FOUNDRY**
SUMMIT

Building of an external CPI for CloudStack


Guillaume Berche & Pierre Oblin






NOVEMBER 2 & 3, 2015 | BERLIN
CLOUD **FOUNDRY**
SUMMIT




Introduction

- Who we are
- Goal of this talk
- Why did we build a CPI for cloudStack ?



 @gberche
 gberche-orange
guillaume.berche
 [at] orange.com



 poblin-orange
 pierre.oblin
 [at] orange.com



NOVEMBER 2 & 3, 2015 | BERLIN
CLOUD **FOUNDRY**
SUMMIT

Outline

- Bosh & CPI intro
- CloudStack CPI design tradeoffs
- Status, future work, learnings



NOVEMBER 2 & 3, 2015 | BERLIN

CLOUD **FOUNDRY**
SUMMIT

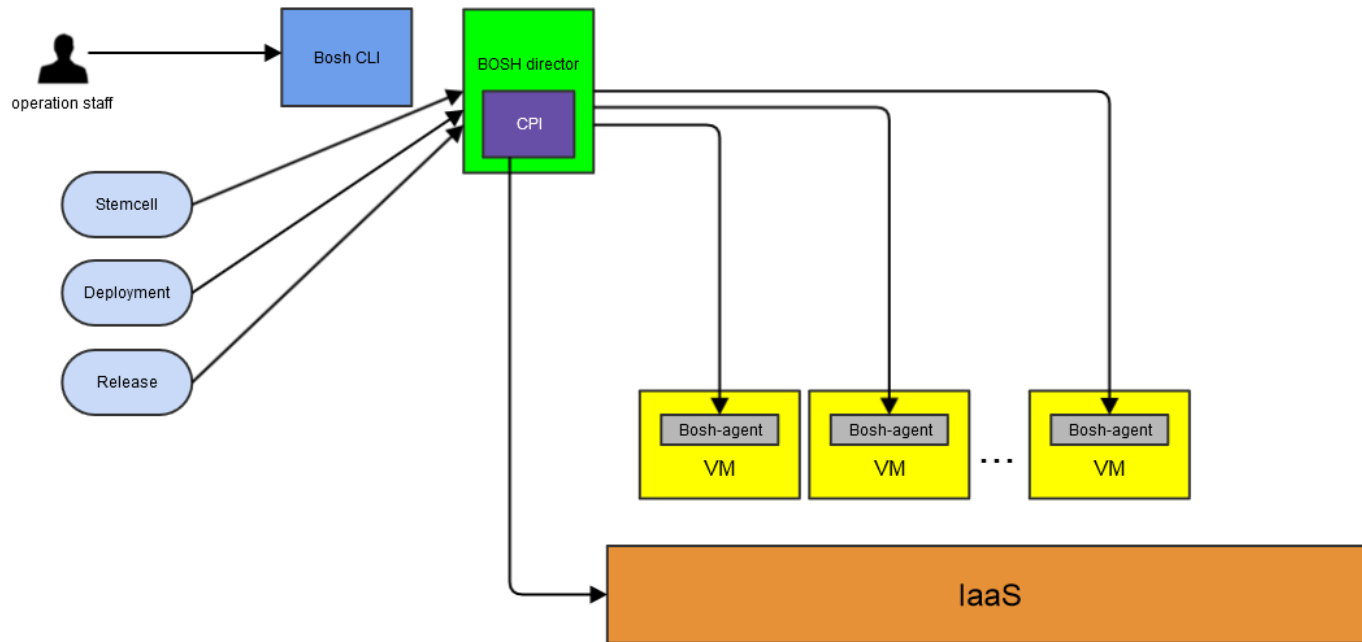
Bosh and CPI overview





NOVEMBER 2 & 3, 2015 | BERLIN
CLOUD **FOUNDRY**
SUMMIT

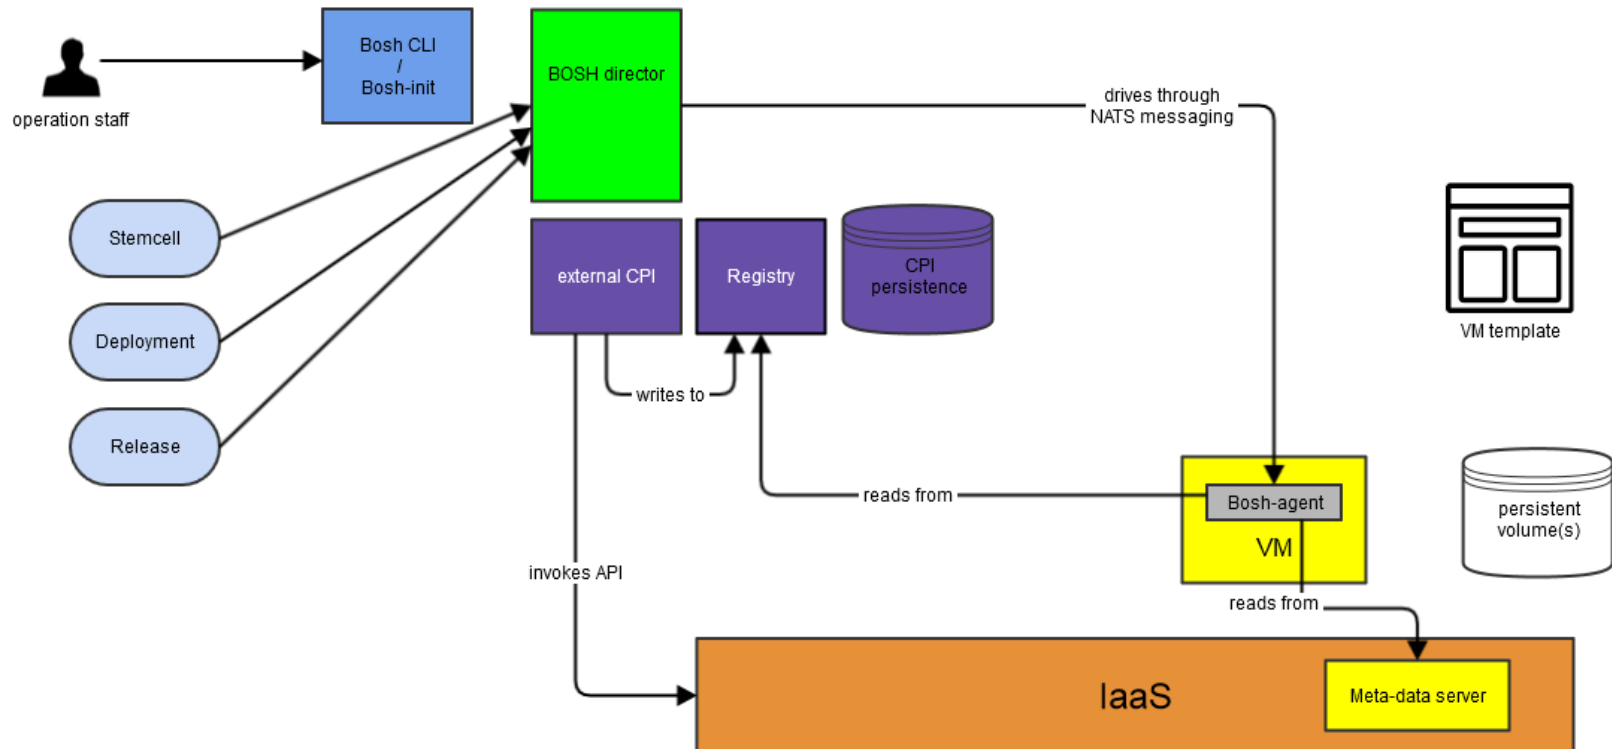
Bosh intro





NOVEMBER 2 & 3, 2015 | BERLIN
CLOUD **FOUNDRY**
SUMMIT

Bosh CPI overview





NOVEMBER 2 & 3, 2015 | BERLIN
CLOUD **FOUNDRY**
SUMMIT

CPI API V1 overview

- Stemcell management
 - ``create stemcell``
 - ``delete stemcell``
- VM management
 - ``create vm``
 - ``delete vm``
 - ``has vm``
 - ``reboot vm``
 - ``set vm metadata``
 - ``configure networks``
- Disk management
 - ``delete disk``
 - ``has disk``
 - ``attach disk``
 - ``detach disk``
 - ``get disks``
- Disk snapshots
 - ``snapshot disk``
 - ``delete snapshot``
 - ``current vm id``

Stemcell management

`create_stemcell`

Creates a reusable VM image in the IaaS from the `stemcell` image. It's used later for creating VMs. For example AWS CPI creates an AMI and returns AMI ID.

See [Stemcell Building](#) for more details.

Arguments

- `image_path` [String]: Path to the stemcell image extracted from the stemcell tarball on a local filesystem.
- `cloud_properties` [Hash]: Cloud properties hash extracted from the stemcell tarball.

Example

```
[
  "/tmp/extracted-stemcell-348754vdsn87fr/image",
  {
    "name": "bosh-openstack-esxi-ubuntu-trusty-go_agent",
    "version": "2972",
    "infrastructure": "openstack",
    "hypervisor": "esxi",
    "disk": 3072,
    "disk_format": "ovf",
    "container_format": "bare",
    "os_type": "linux",
    "os_distro": "ubuntu",
    "architecture": "x86_64",
    "auto_disk_config": true
  }
]
```

Returned

- `stemcell_cid` [String]: Cloud ID of the created stemcell (e.g. stemcells in AWS CPI are made into AMIs so `cid` would be `ami-83fd1f`)

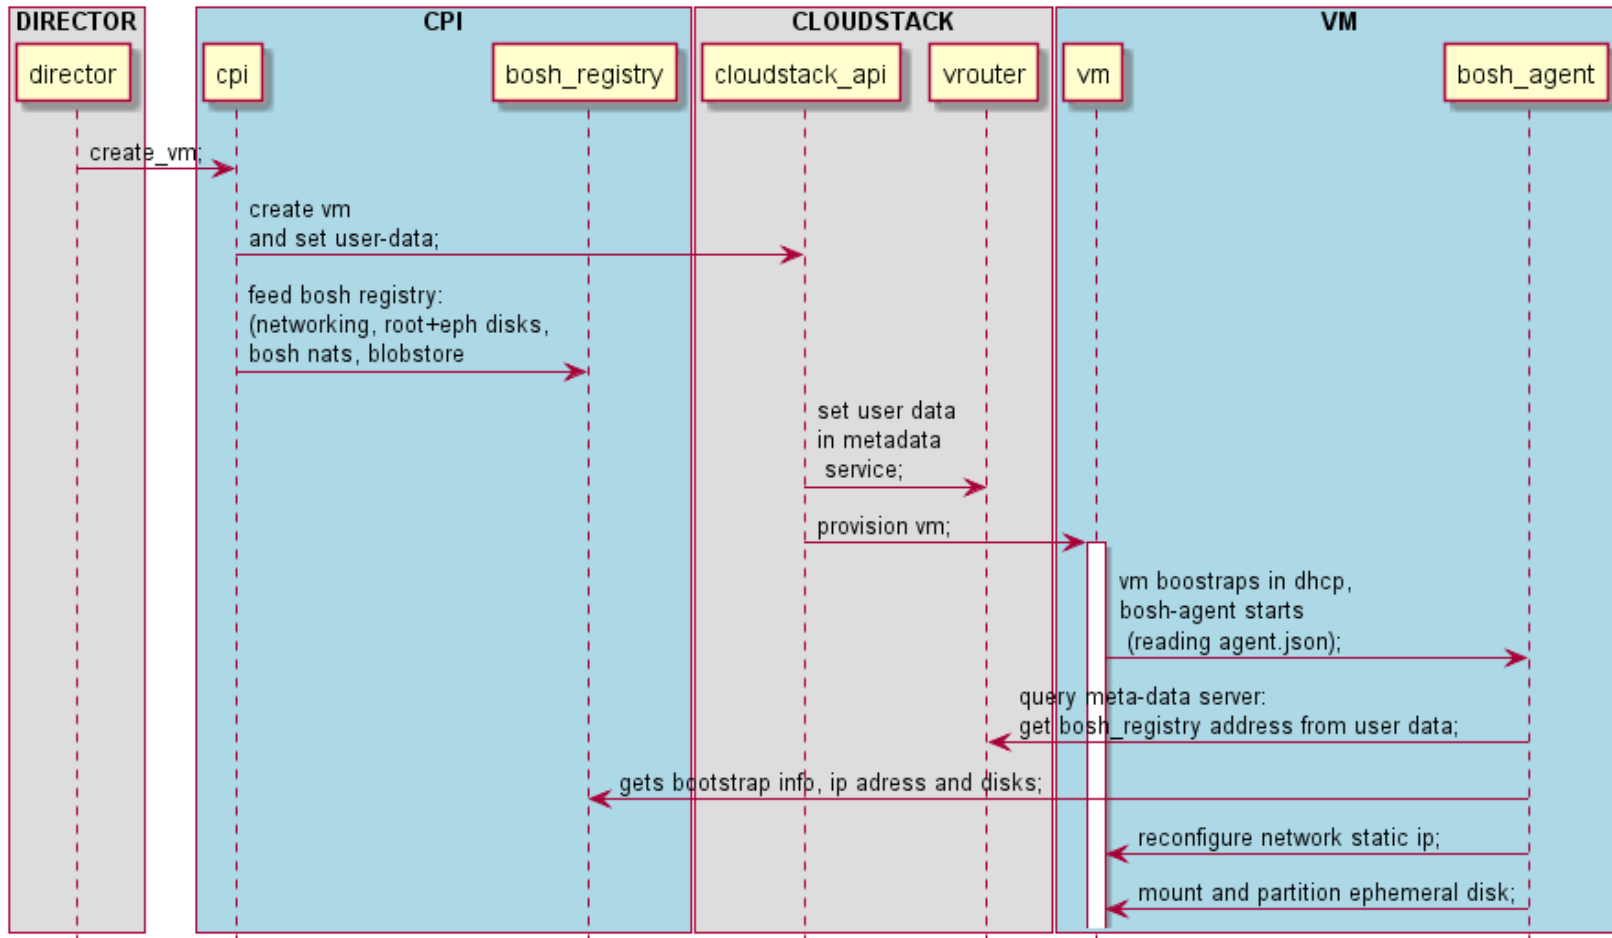


Bosh agent configuration overview

- bootstrapping from stemcell provided config
- iaas-specific initial networking & registry lookup
 - available methods: meta-data server, config drive, cdrom
- initial registry config loading
- dynamic configs (network reconfig, disk resize...) through NATs messages



VM instantiation illustrated





NOVEMBER 2 & 3, 2015 | BERLIN

CLOUD **FOUNDRY**
SUMMIT

CloudStack CPI design choices





NOVEMBER 2 & 3, 2015 | BERLIN
CLOUD **FOUNDRY**
SUMMIT

Context and constraints

- Provide flexibility for innovation and optimization while remaining opensource
- Chicken and egg: CPI and stemcell
- Bosh is in active development
- Lack of public infrastructure enabling community contributions





NOVEMBER 2 & 3, 2015 | BERLIN
CLOUD **FOUNDRY**
SUMMIT

Design tradeoffs

- Programming language: ruby, go, java ?
- CloudStack api client
- CPI process model
 - per invocation process or daemon mode ?
 - stdin/stdout or REST facade ?
- Stemcell to cloudstack vm templates
 - insider or outsider ?
 - from volume or webdav exposed image ?
 - a light stemcell to kick things off
- CS custom metadata URL





NOVEMBER 2 & 3, 2015 | BERLIN

CLOUD **FOUNDRY**
SUMMIT

Status,
future work,
learnings





NOVEMBER 2 & 3, 2015 | BERLIN
CLOUD **FOUNDRY**
SUMMIT

Status


- Supported
 - stemcell upload
 - advanced zone networking: static/dynamic ips
- Limitations
 - single persistent disk support
 - single nic support
 - no yet floating ip/vip





Future short-term work

- Stemcell: Xen (+KVM), Ubuntu+Centos
 - another infrastructure in bosh/stemcell_builder ?
 - patch one of the official stemcell ?
- CI (concourse): BATS + CPI lifecycle tests + ...
- Local storage for persistent disks
 - snapshot
 - live migration
 - vm placement to existing local storage
- Happy to collaborate and support the community.
 - Any public CloudStack instance for public ci ?
- Apply for CF incubator process ?






NOVEMBER 2 & 3, 2015 | BERLIN
CLOUD **FOUNDRY**
SUMMIT

Challenges met & potential bosh improvements

- documentation:
 - bdd/living doc
 - unit tests
 - reference implementations
- CPI daemon support in bosh-init
- bosh-agent cloud-stack style volume support
- laas-specific registry impl
 - => use NATS for disk mount support ?
- bosh-init optim for CPI development
 - (external blobstore)





NOVEMBER 2 & 3, 2015 | BERLIN

CLOUD **FOUNDRY**
SUMMIT

Thanks, Questions ?



@gberche



gberche-orange
guillaume.berche



[at] orange.com



poblin-orange



pierre.oblin



[at] orange.com

