Hadoop Backup and Scaling in Hybrid Environment

Paweł Leszczyński, Robert Mroczkowski, Mariusz Strzelecki
Allegro Group

allegro.tech

Agenda

- Allegro Data Hub.
- Defining a problem:
  - Single Hadoop cluster in a single DC.
- Applied solutions in details
  - Live Backup in S3.
  - Scaling Hadoop in hybrid environment.

Applications and Device Logs

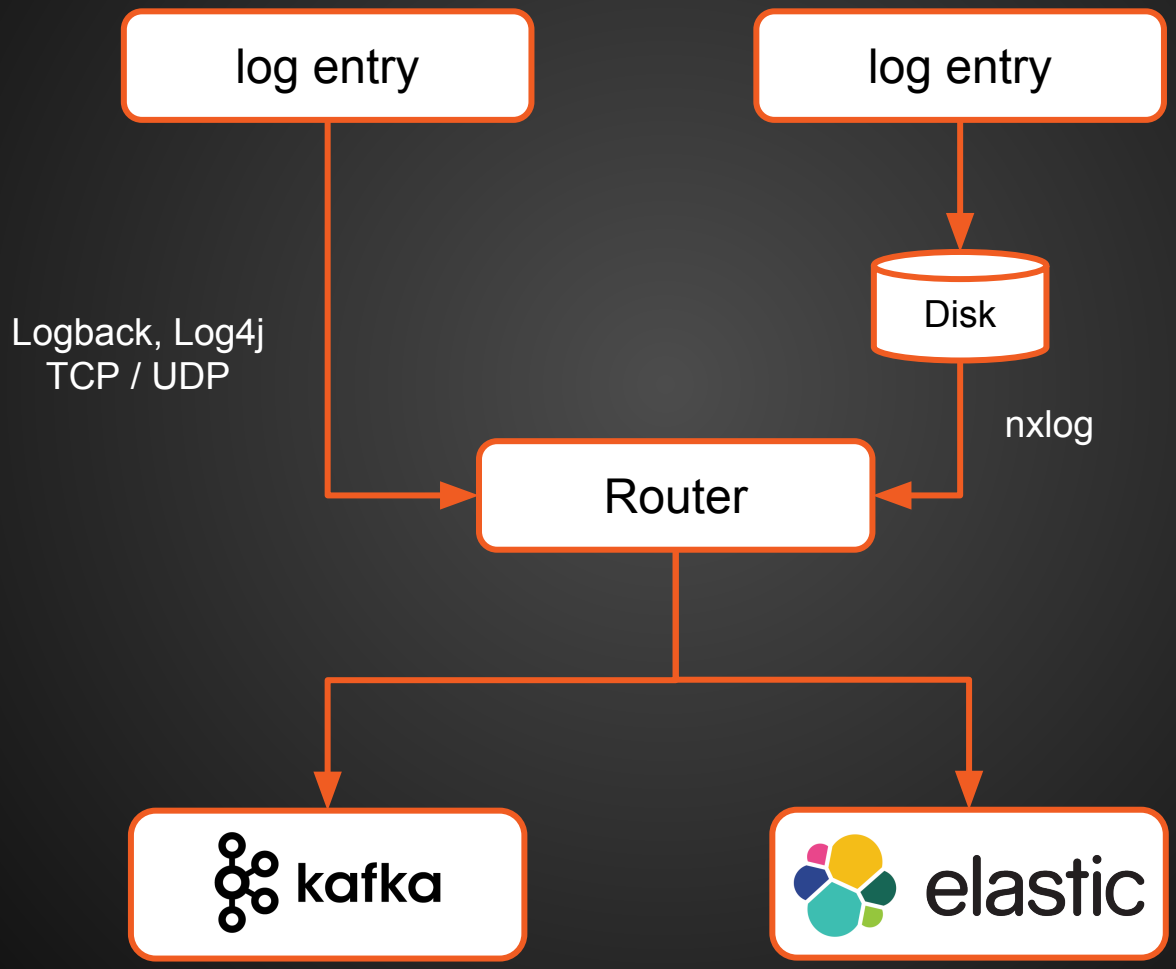

Microservices' messaging

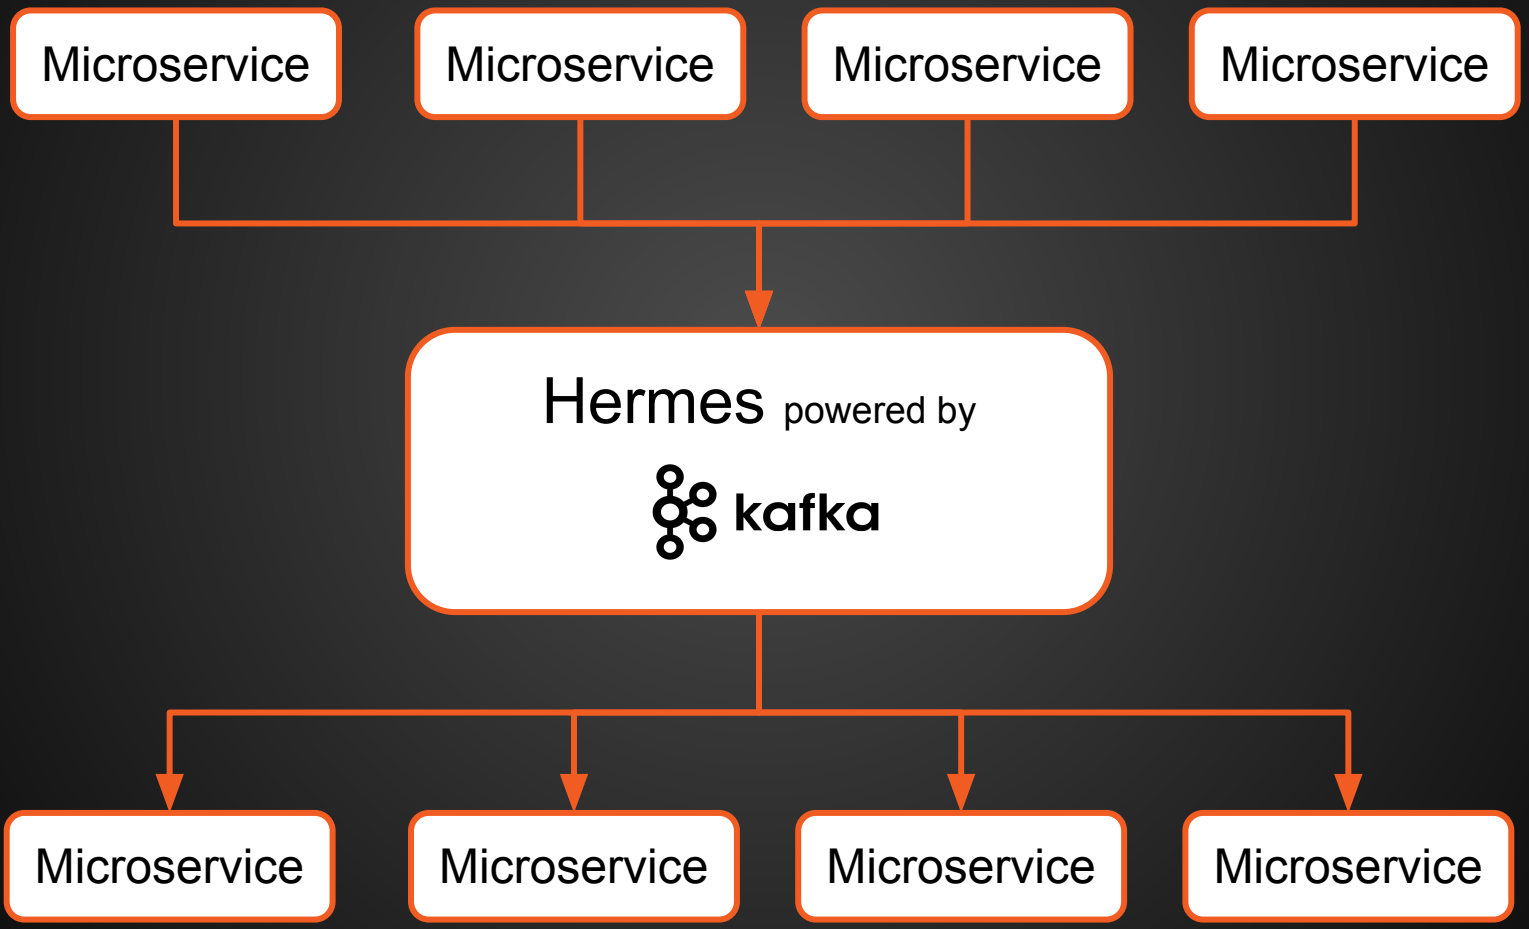

Clickstream data

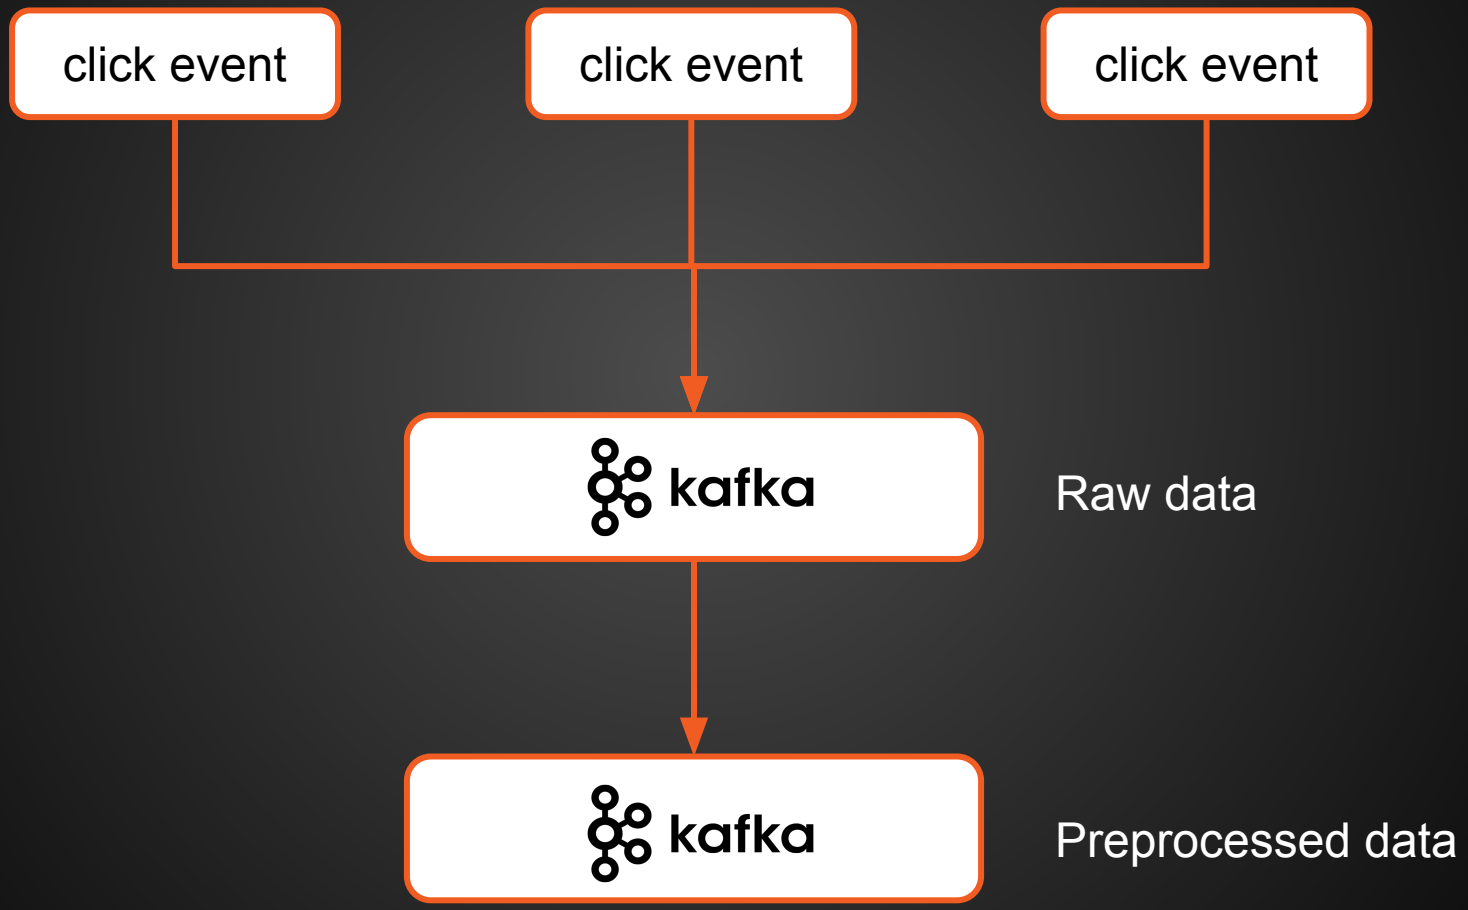

Data hub's backbone

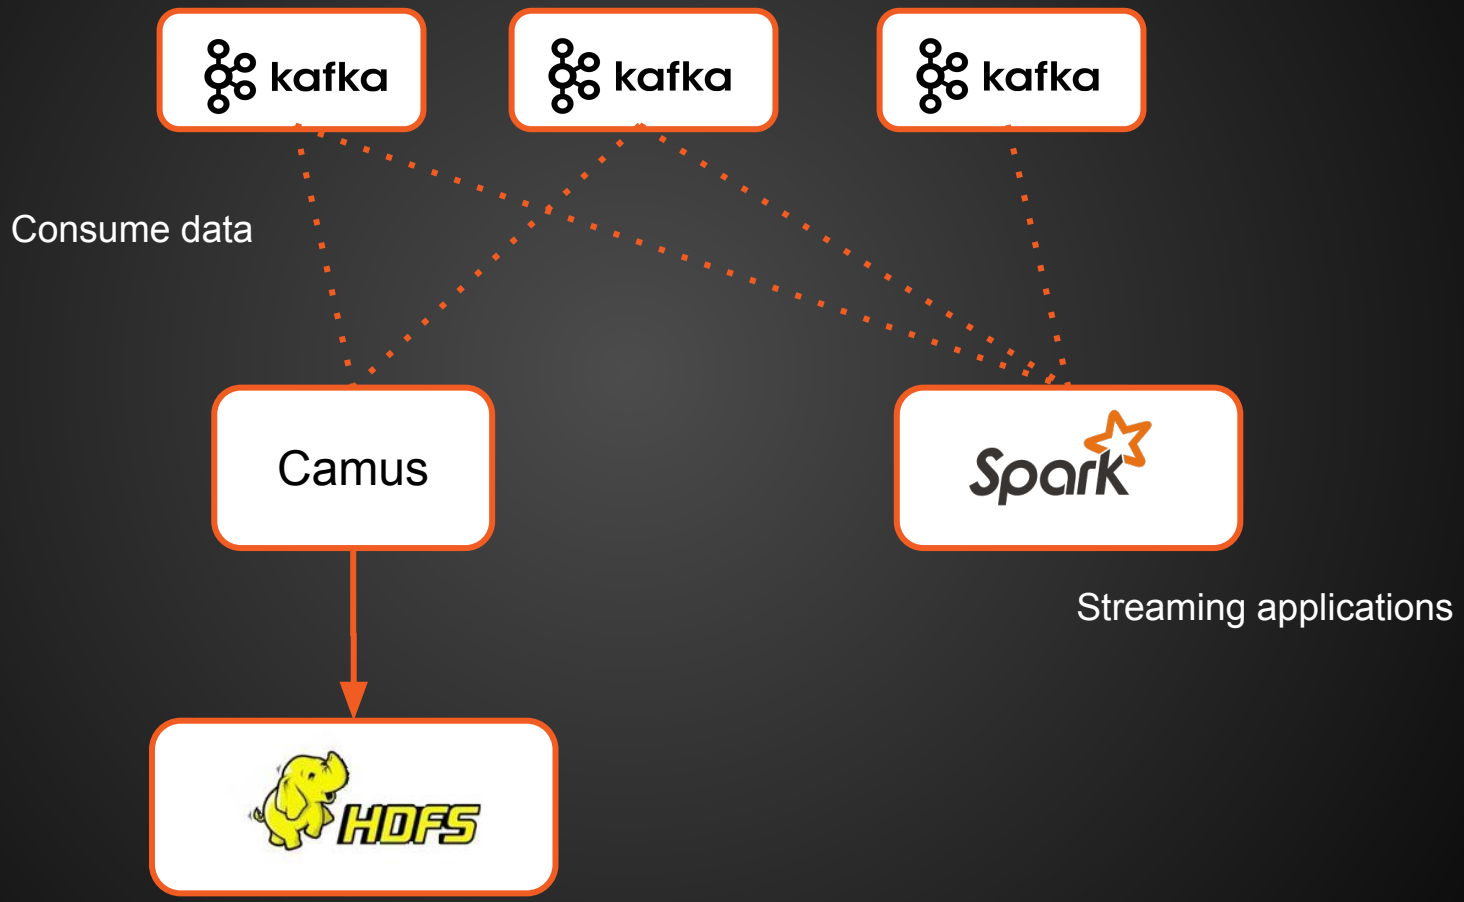

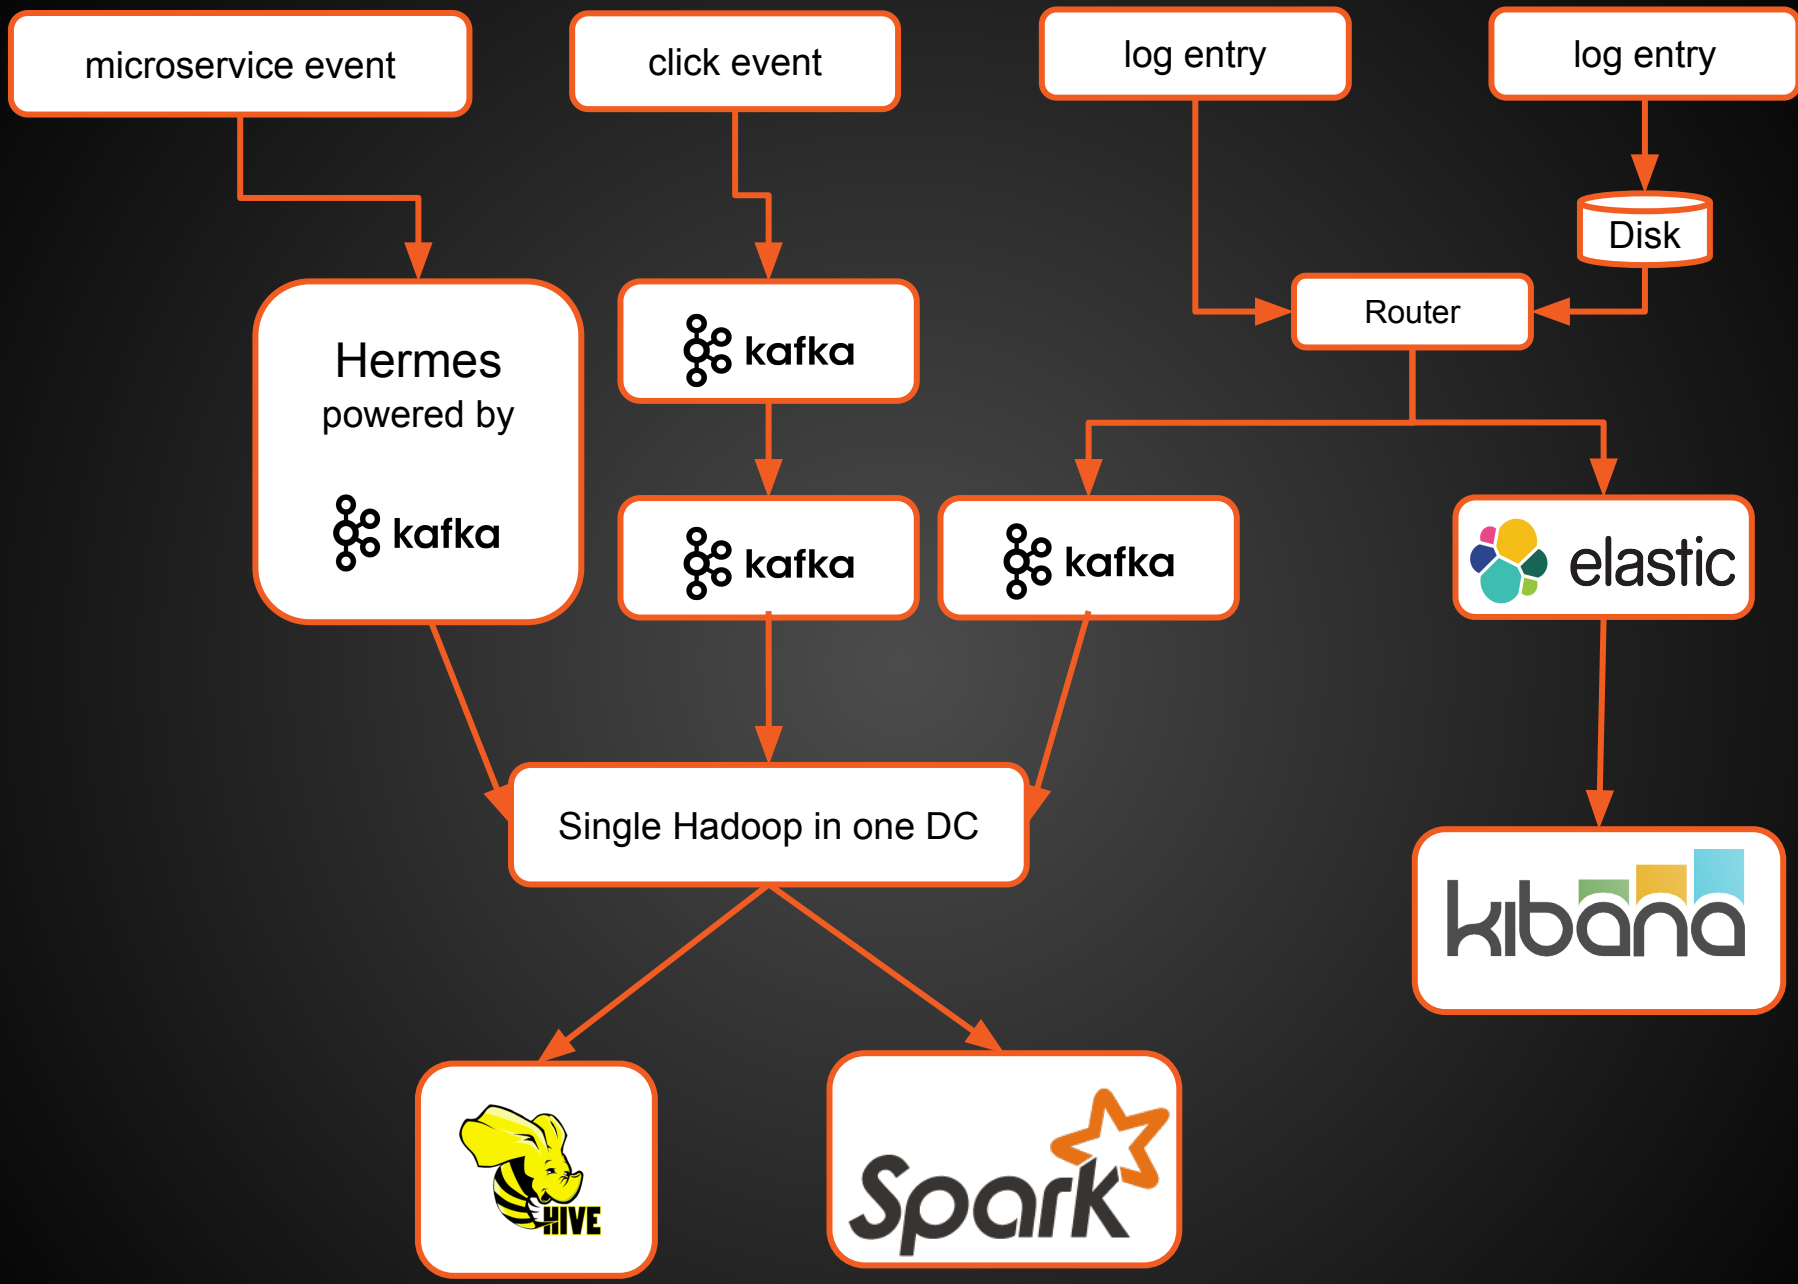
microservice event
click event
log entry
log entry
Disk
Router
Hermes
powered by
kafka
kafka
kafka
kafka
kafka
elastic
Single Hadoop in one DC
kibana
HIVE
Spark

Allegro Hadoop cluster in numbers

4	2	47
terabytes ram	petabytes disk space	nodes

HDFS Backup?

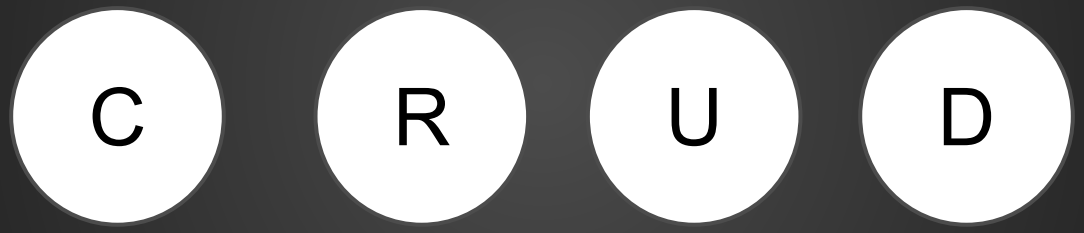

Data Hub for the rescue!

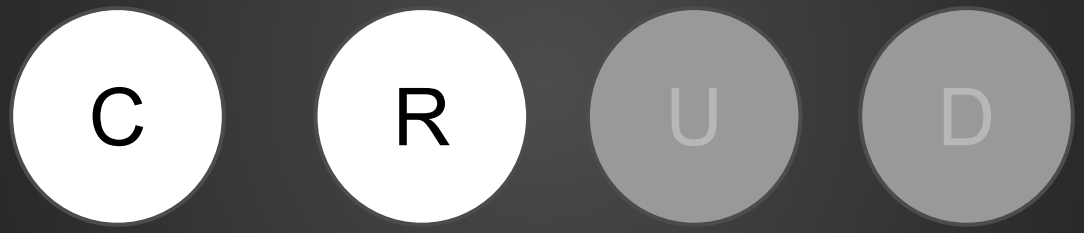

ETL system - Camus

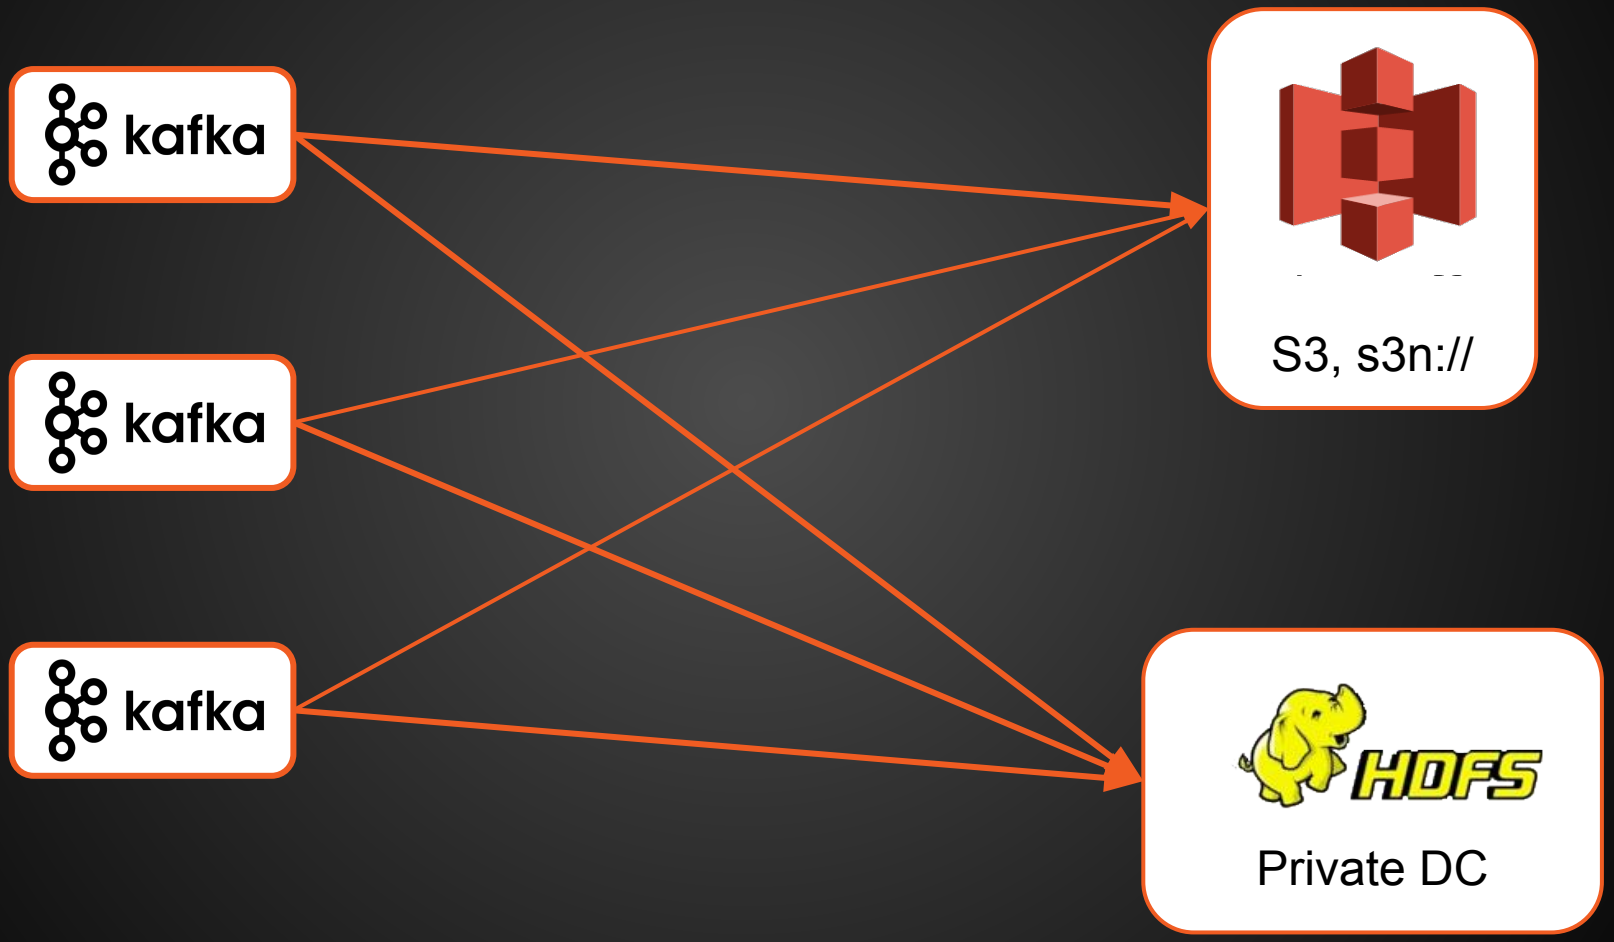

Data Integrity Verification

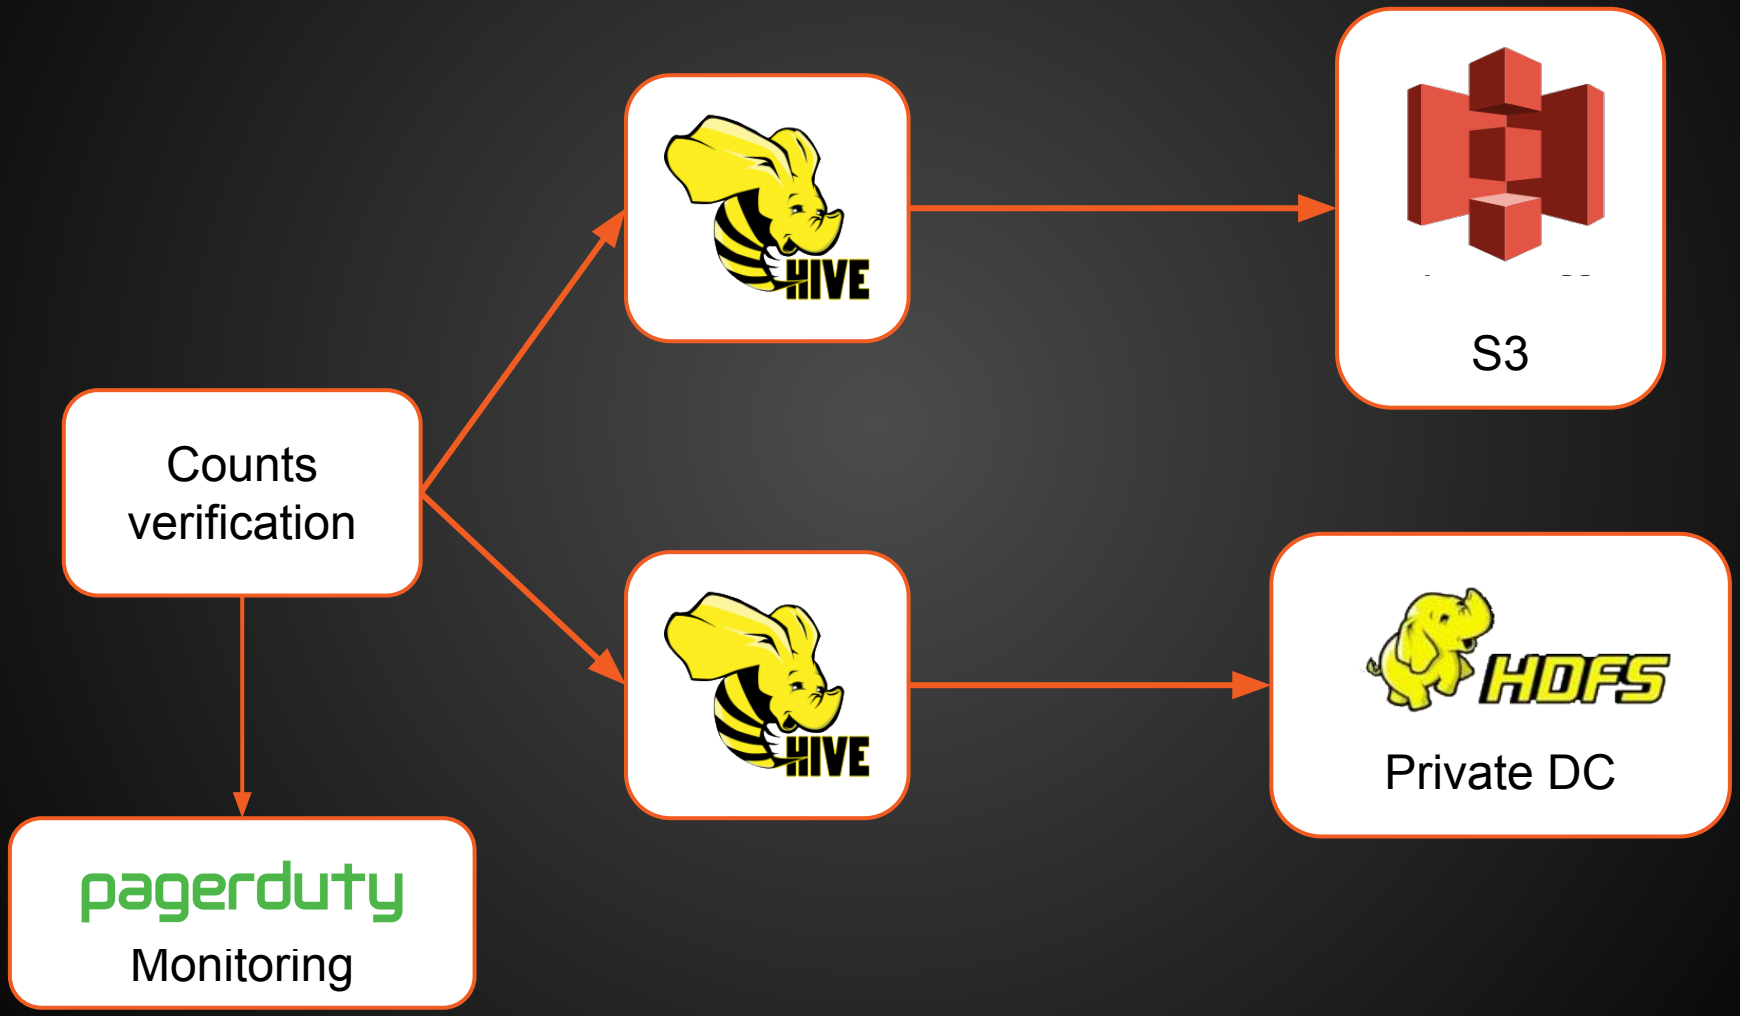

Data Flow Monitoring

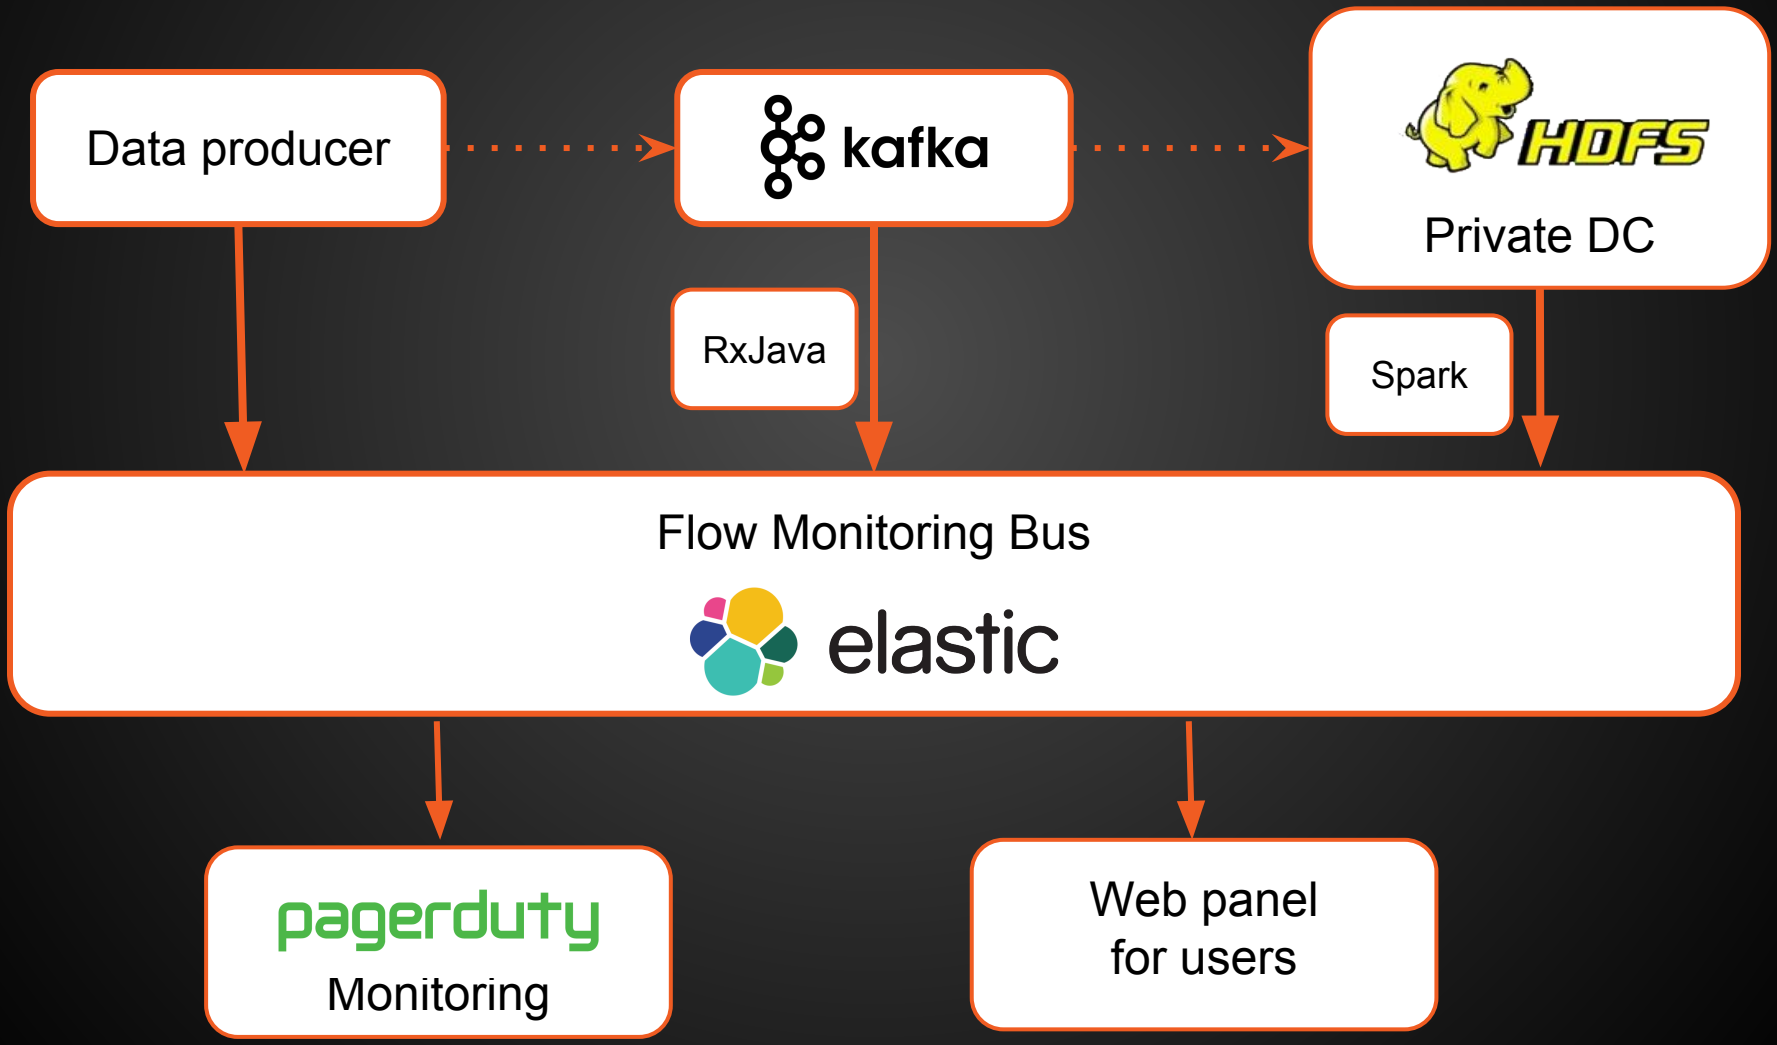

Backup and HA

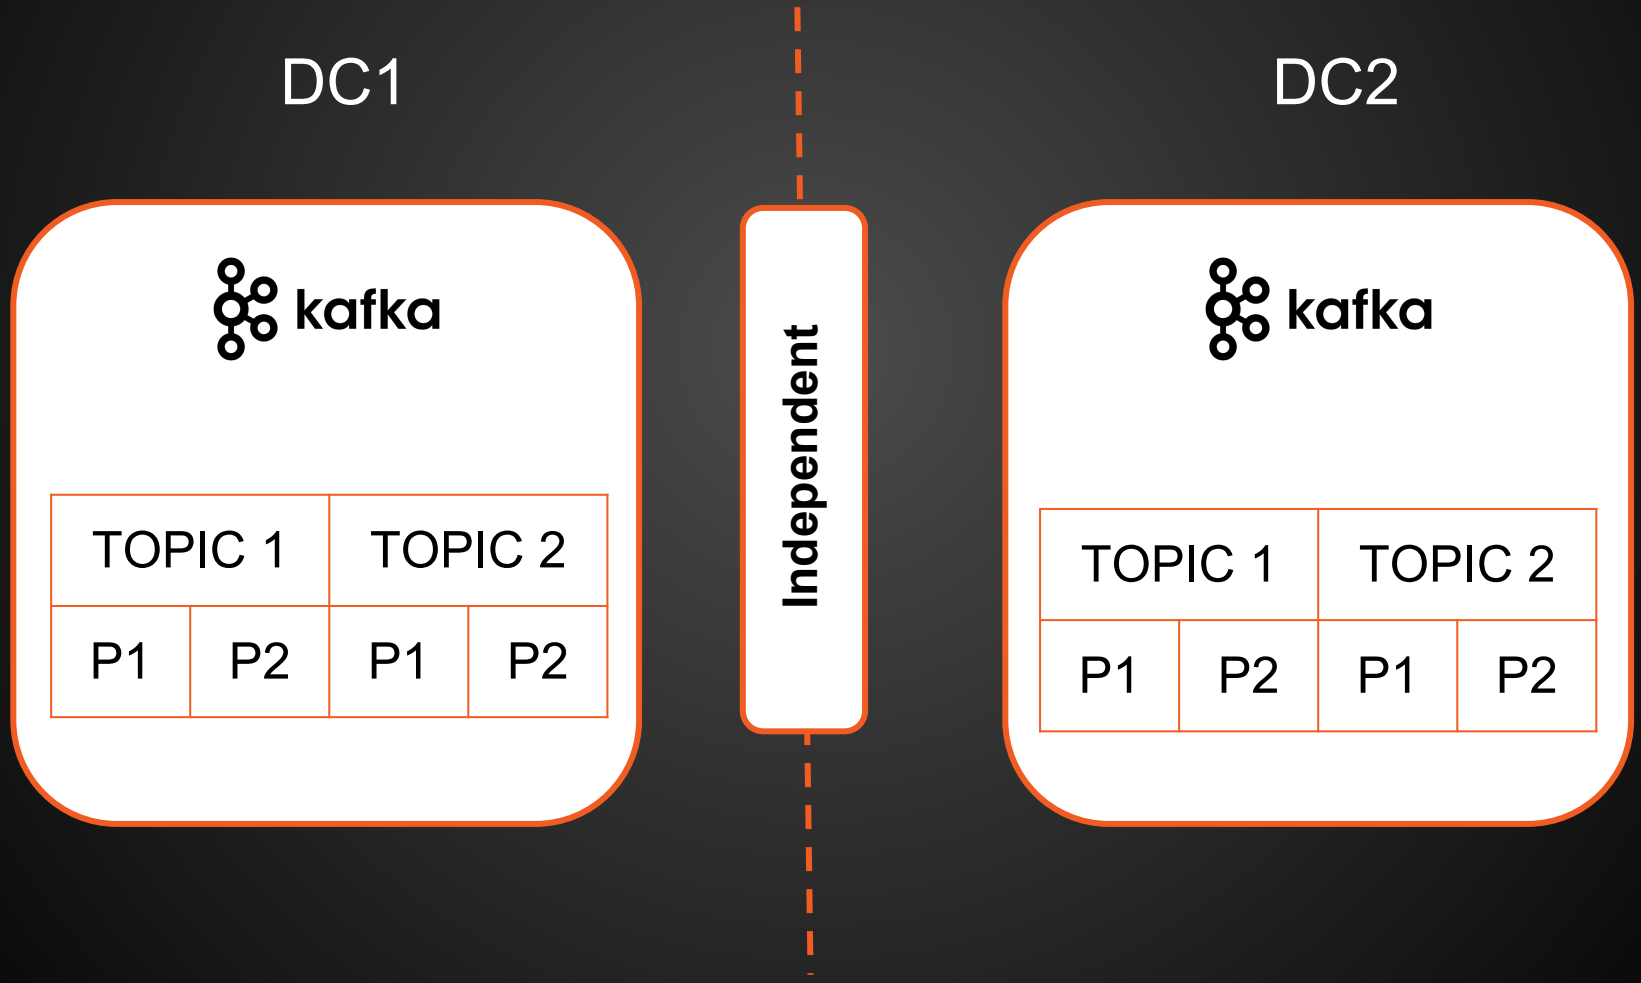

Backup and HA

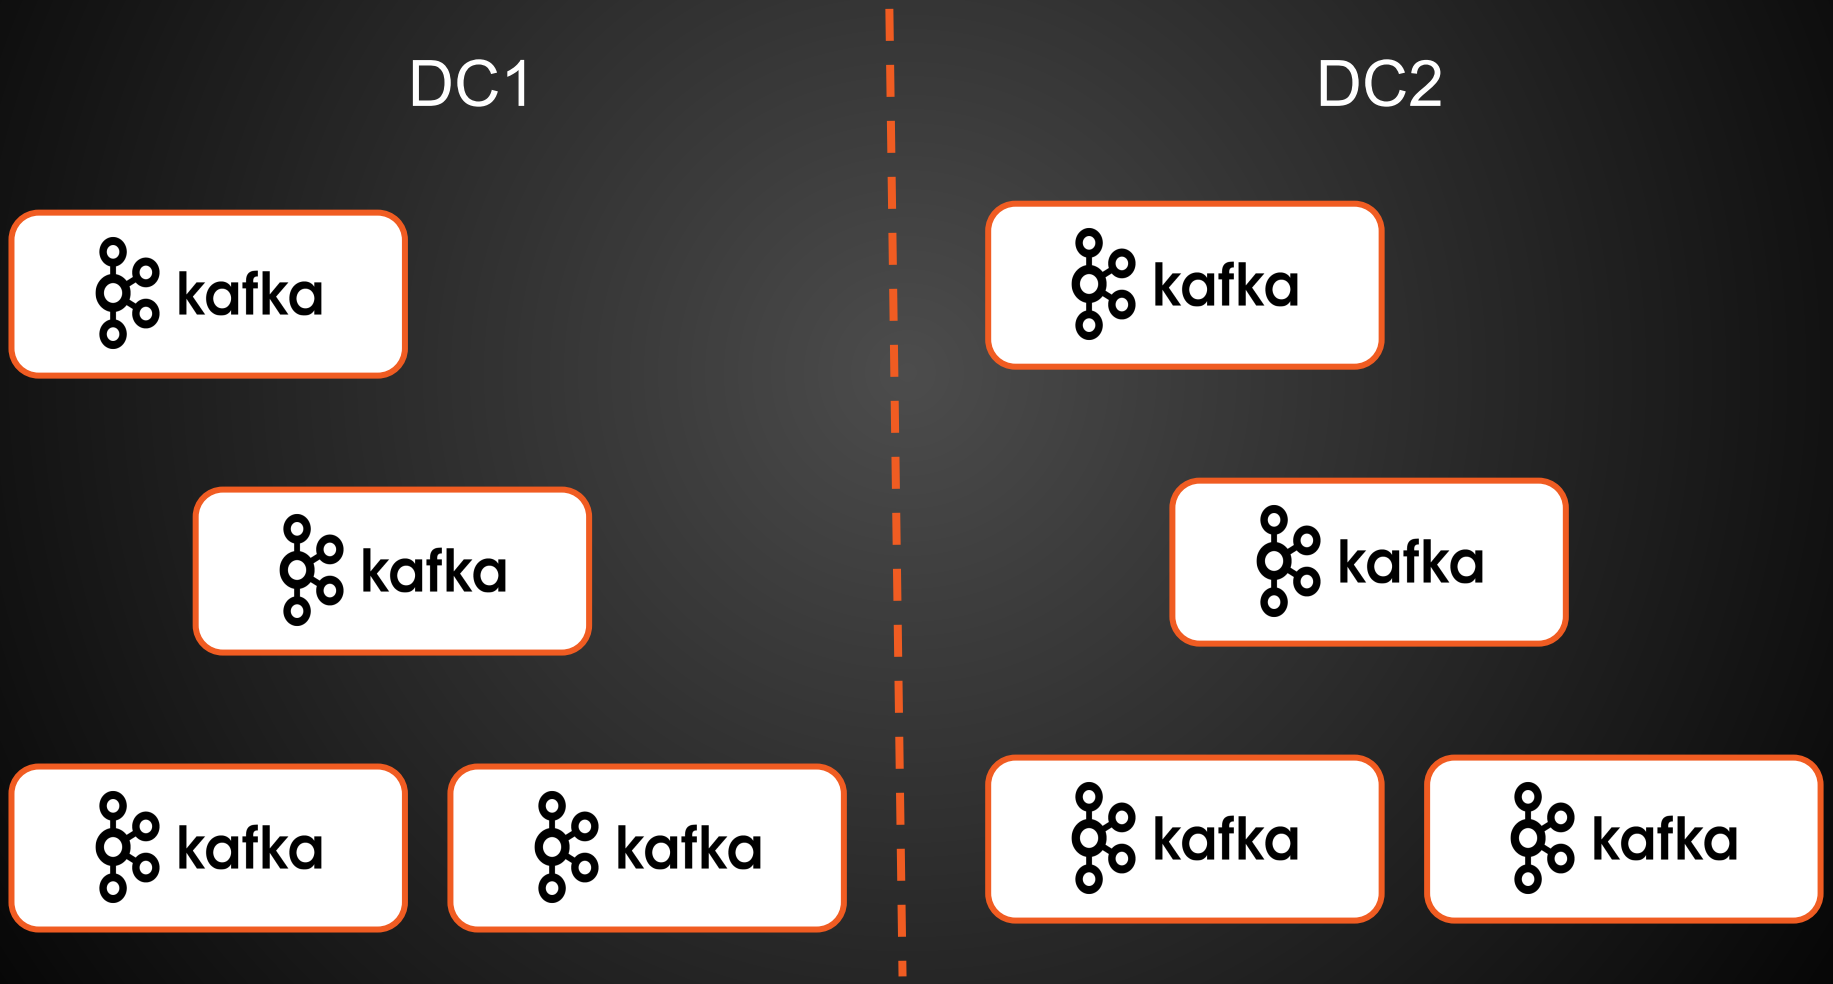

Automatization - Assumptions

- Fully functional production cluster
- One-click deploy
- Provide infrastructure in AWS and OpenStack
- Dynamic resources scaling
- Dynamic configuration
- Scalability and dynamicity on Bare Metal cluster
- Vendor independent
- Simplicity
- Fast delivery

Automatization - technology stack

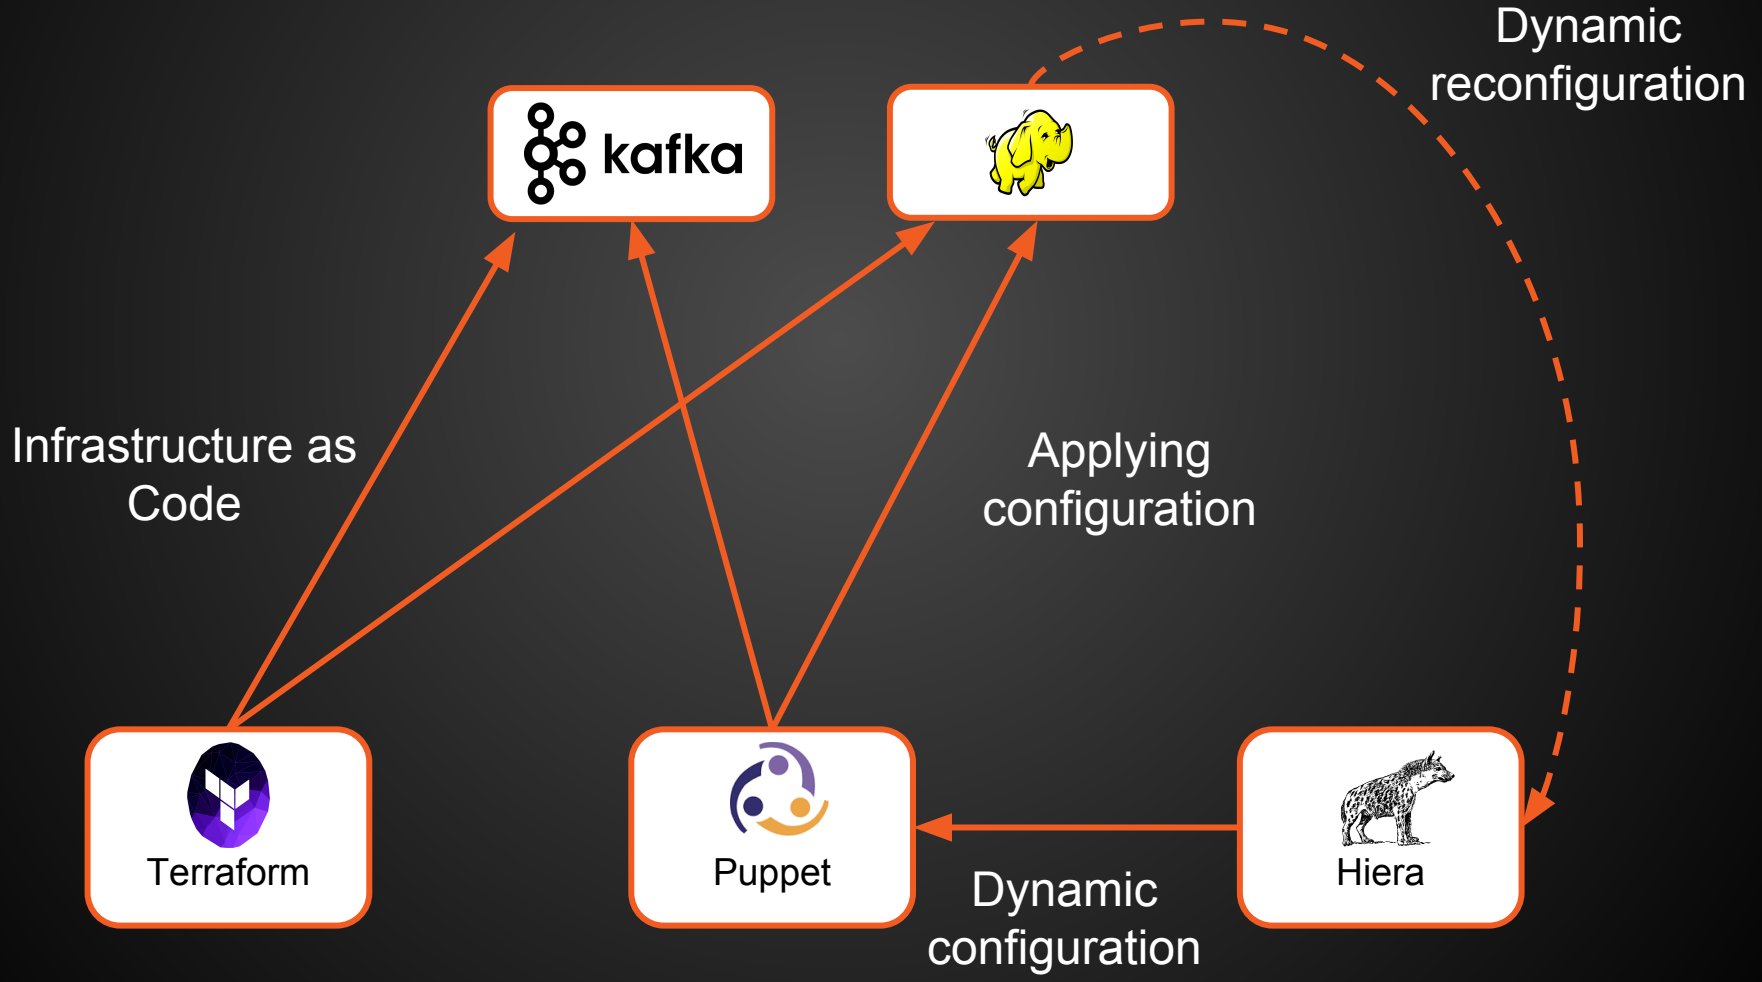

Terraform

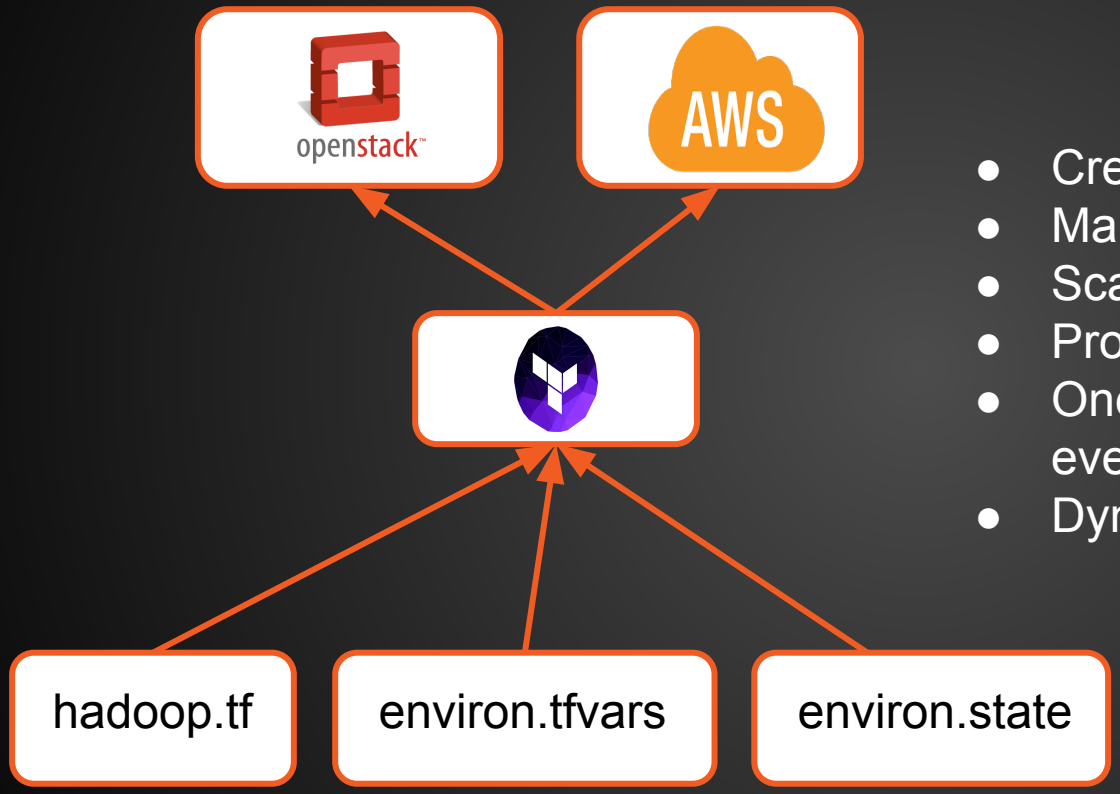

- Create infrastructure
- Manage network
- Scale hosts and storage resources
- Provides zone / region awareness
- One code provides multiple clusters everywhere
- Dynamic configuration

Terraform

Name	Availability Zone	Instance State
aloefreedom-datanode	eu-west-1a	running
aloefreedom-datanode	eu-west-1c	running
aloefreedom-datanode	eu-west-1b	running

terraform.clustername.tfvars:

```
metanode_count = "4"
storage_force_detach = "true"
metanode_instance_type = "m3.large"
metanode_storage_size = "50"
metanode_storage_type = "gp2"
metanode_storage_termination = "true"
metanode_user_data = "hadoop_prod::claster"
datanode_count = "3"
datanode_instance_type = "m3.2xlarge"
datanode_storage_size = "100"
datanode_storage_type = "standard"
datanode_storage_termination = "true"
datanode_user_data = "hadoop_prod::claster"
```

Name	Availability Zone	State	Attachment Information
aloefreedom-datanode_volume0	eu-west-1a	in-use	i-b5826919 (aloefreedom-datanode) :/dev/xvdbb
aloefreedom-datanode_volume1	eu-west-1b	in-use	i-8c66062d (aloefreedom-datanode) :/dev/xvdbc
aloefreedom-datanode_volume2	eu-west-1c	in-use	i-b8afc615 (aloefreedom-datanode) :/dev/xvdbd
aloefreedom-datanode_volume3	eu-west-1a	in-use	i-b5826919 (aloefreedom-datanode) :/dev/xvdce

Puppet

- Manage OS
- Setup roles
- Configure services
- Dynamicity via templates
- Configuration backends
- Scale

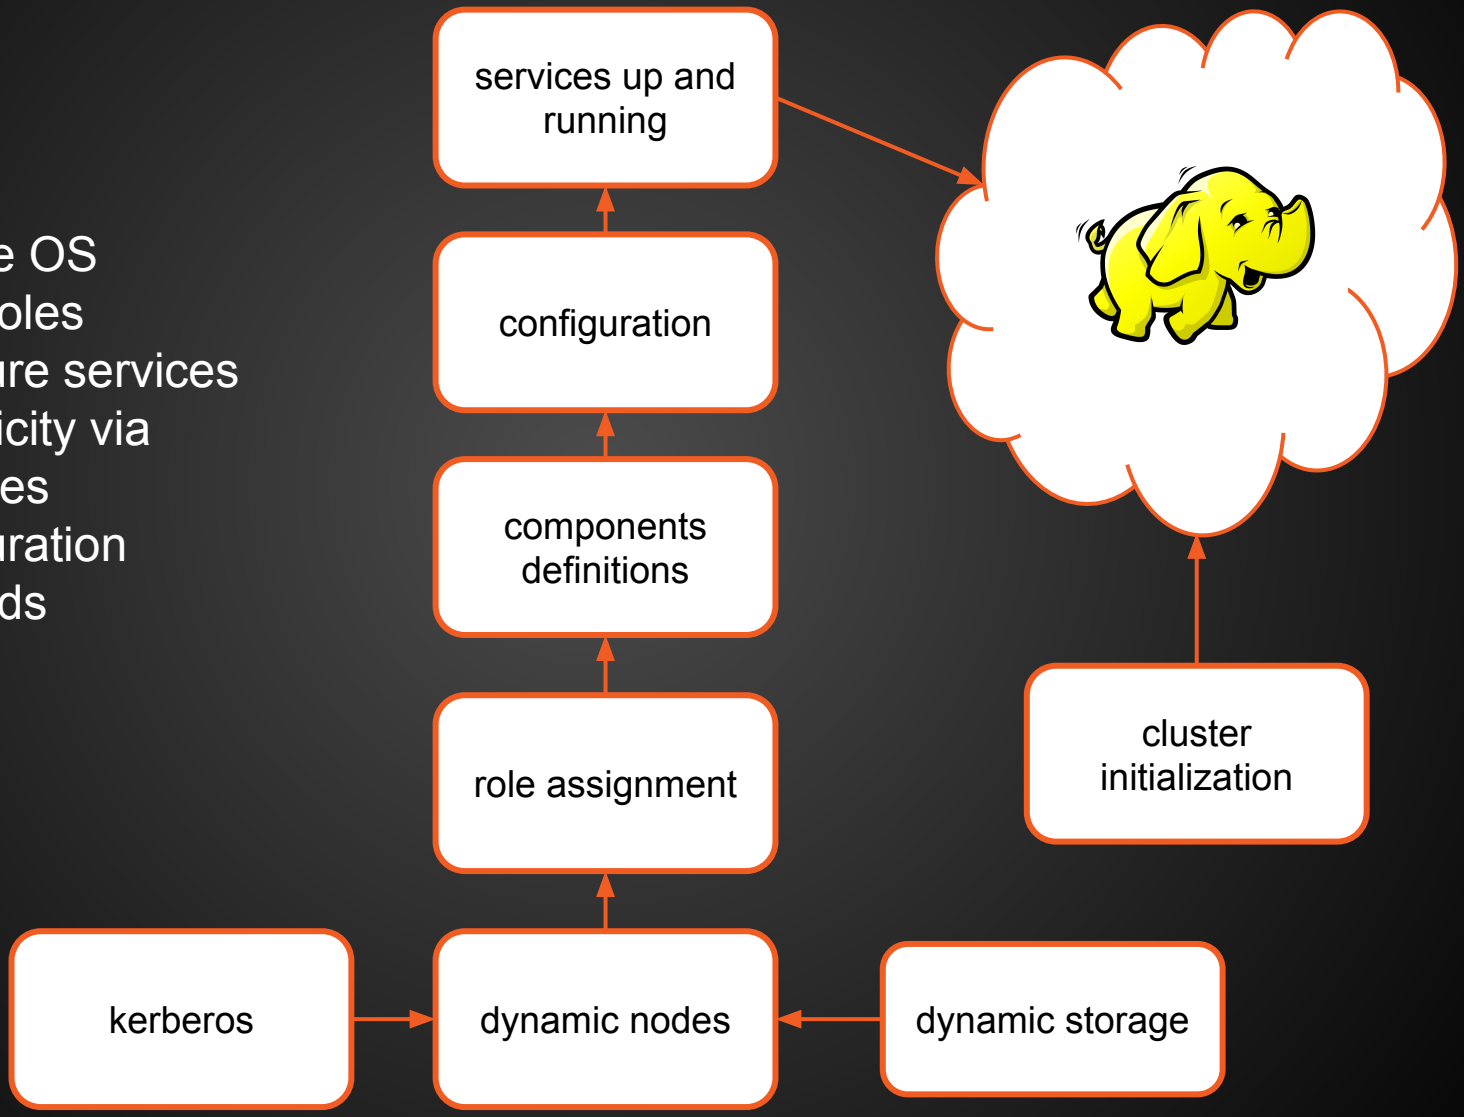

Configuration Service: Hiera

Instance specific

```
hadoop_prod/name/cfg:
  format: false
  bootstrap_standby: false
  init_filesystem: false
  cluster_version: "cdh5.4.4"
```

Defaults

```
hadoop_prod/default/cfg/hdfs:
  dfs.ha.fencing.methods:
     "sshfence(hdfs)"
  dfs.block.size: 128m
  dfs.replication: 3
```

Roles in cluster

```
hadoop_prod/clusters/members:
  server1.eu-west-1.aws.ourdc.net:
    cluster_name: "some_cluster"
    services:
      - namenode
      - zkfc
      - journalnode
  server2.eu-west-1.aws.ourdc.net:
    cluster_name: "some_cluster"
    services:
      - journalnode
      - resourcemanager
      - client
```

Summary

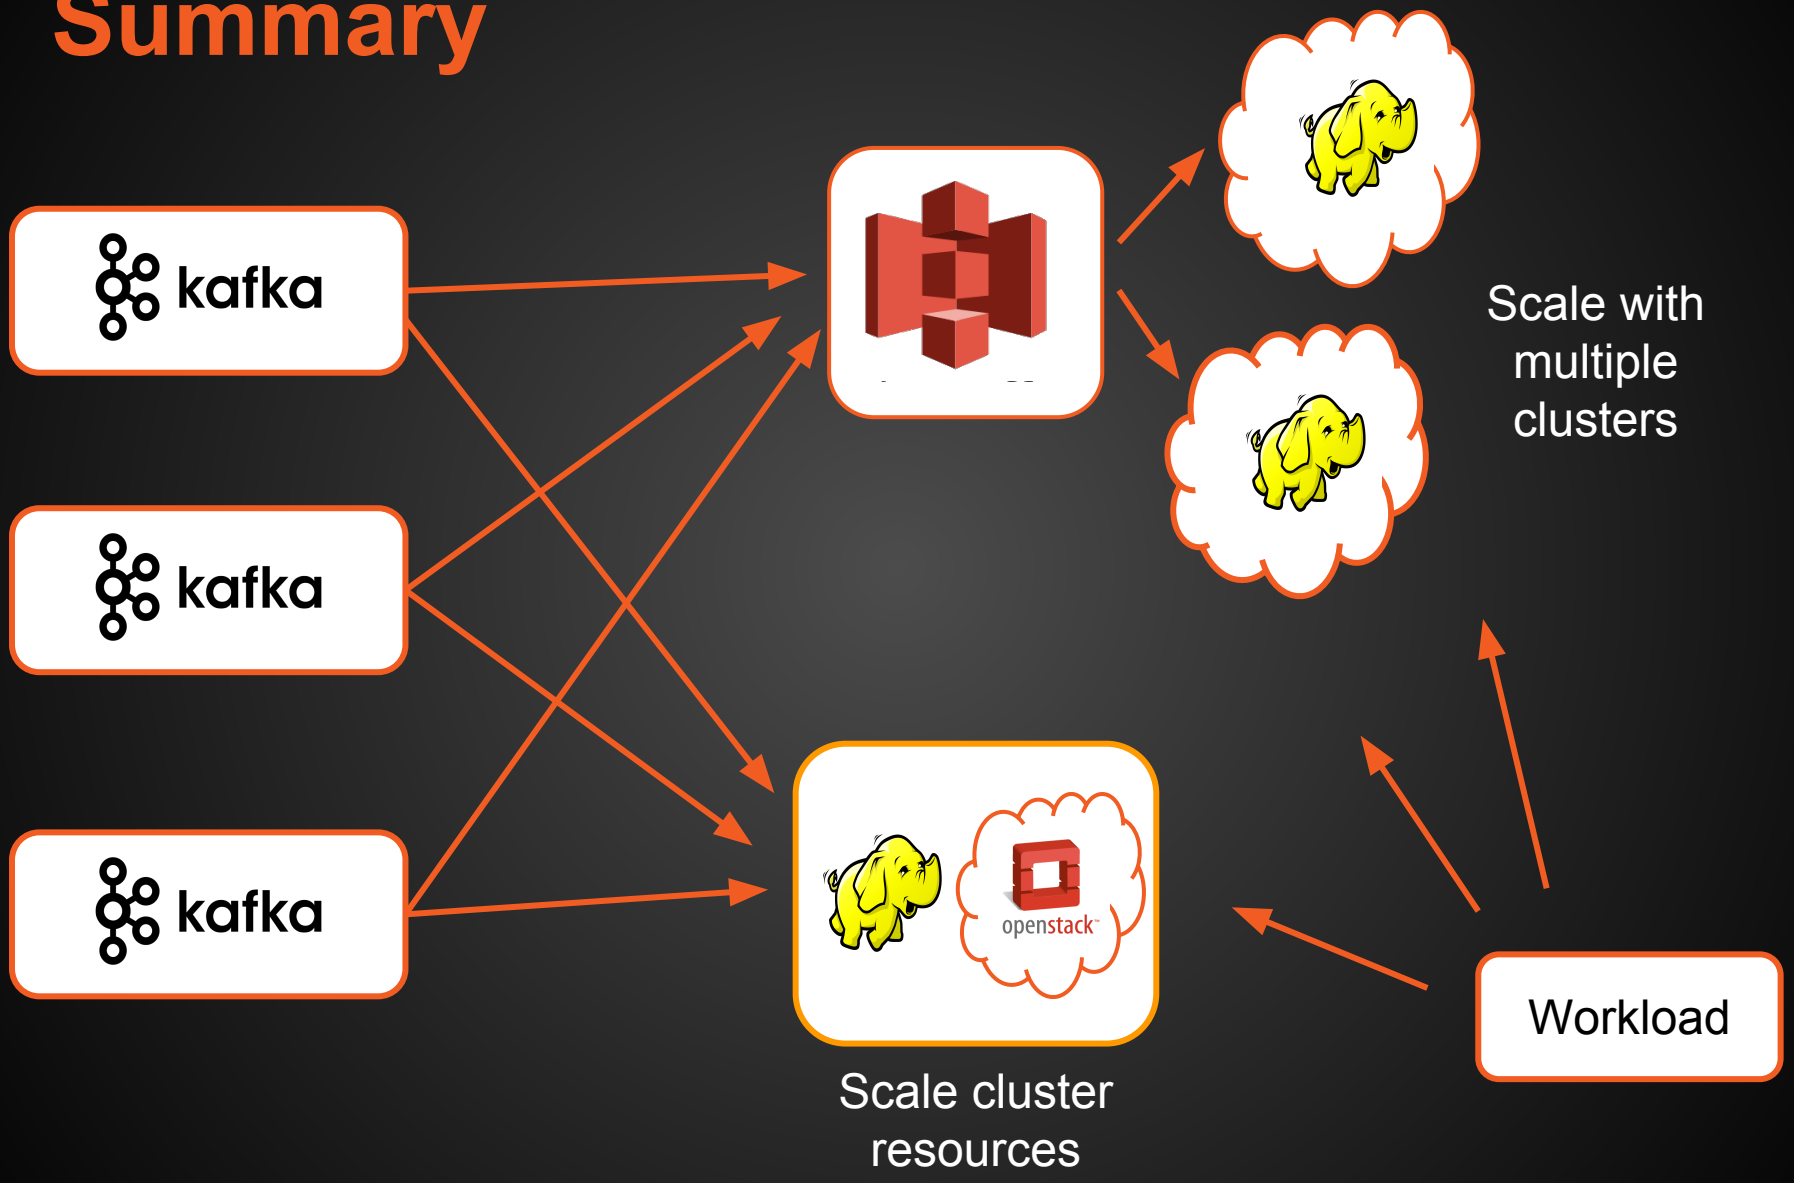
kafka
kafka
kafka
openstack
Scale with multiple clusters
Scale cluster resources
Workload

Thank you!

Questions?