

Pivotal®

Greenplum for Containers

Greenplum on Pivotal Container Service



gpcloud@pivotal.io

Postgres Conf, April 2018

Speaker: Goutam Tadi

Agenda

- **Greenplum** Architecture
- **Greenplum** Data Platform
- Kubernetes on PCF
- **Greenplum** on Kubernetes
- Demo
- Q+A






Pivotal
Greenplum®

WHAT IS GREENPLUM?

**MASSIVELY PARALLEL
DATA WAREHOUSE**



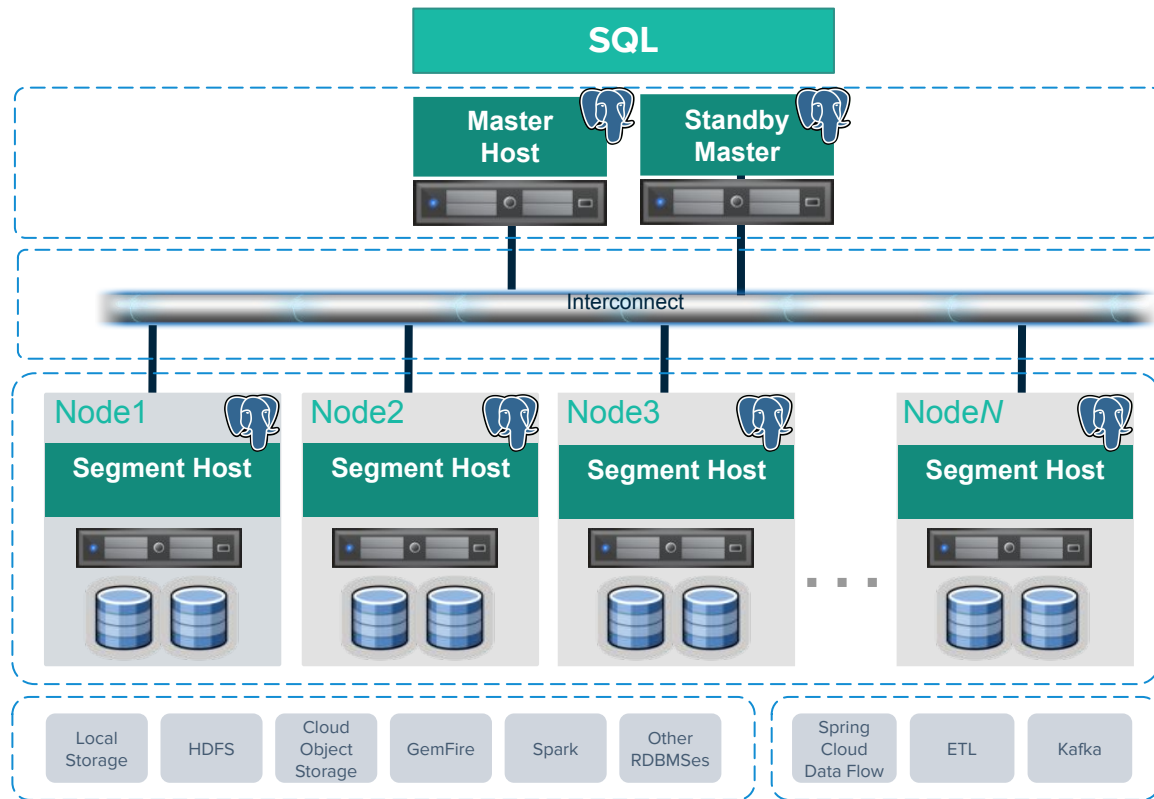
Greenplum = Massively Parallel Postgres for Analytics

Master Servers
Query planning and dispatch

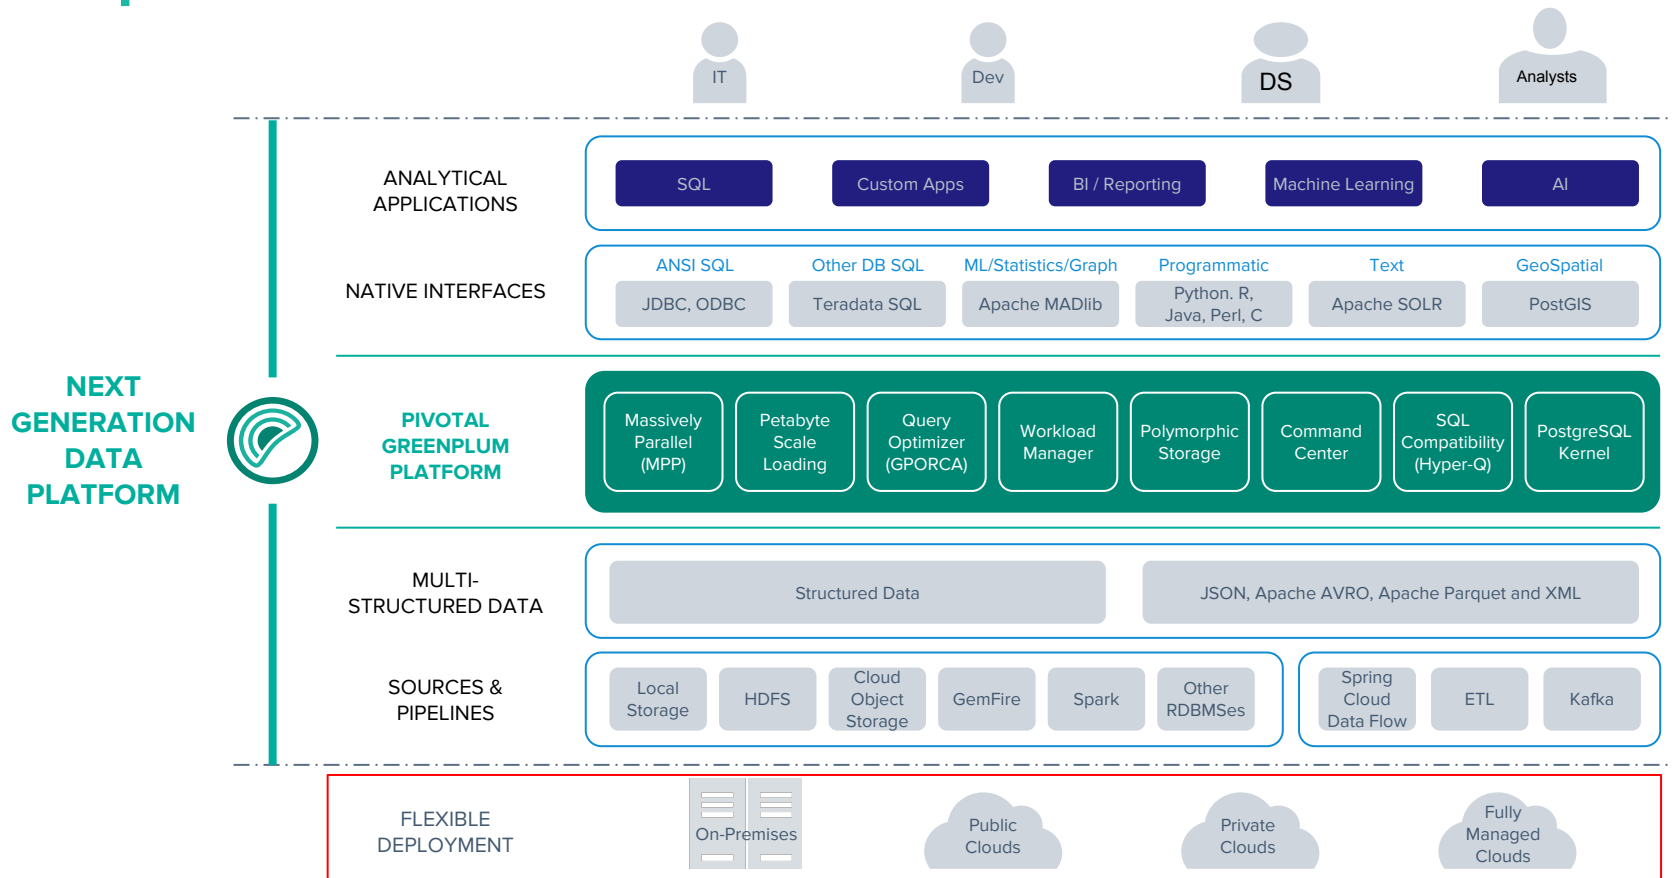
Interconnect

Segment Servers
Query processing and data storage

External Sources & Pipelines
Parallel loading and streaming



Greenplum Data Platform



Faster Deployments... How?

Infrastructure-Agnostic

Bare-Metal



Private Cloud



Public Cloud



Microsoft Azure



Google Cloud Platform

Have you

experienced these before with any database?

Have you...

- Ran out of disk space ?
- Been able to provision more than 100 postgres instances in few minutes?
- Faced issues in recovering failures ?
- Faced issues in expanding the database?



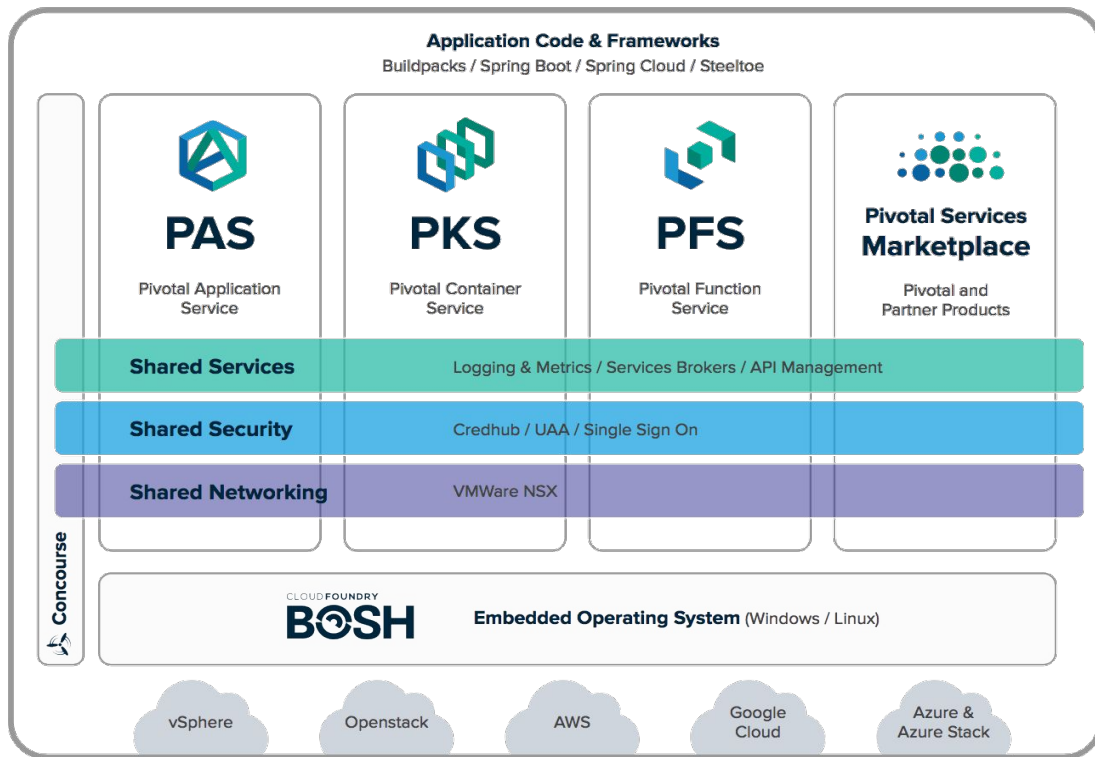
Pivotal
Container Service™

WHAT IS PKS?

**RELIABLY DEPLOY AND RUN
CONTAINERIZED WORKLOADS.**

Kubernetes on Pivotal Cloud Foundry

Continuously deliver any app to every major private and public cloud with a single platform.



Faster Deployments... How?

Infrastructure-Agnostic

Bare-Metal



Private Cloud



Public Cloud



Microsoft Azure



Google Cloud Platform

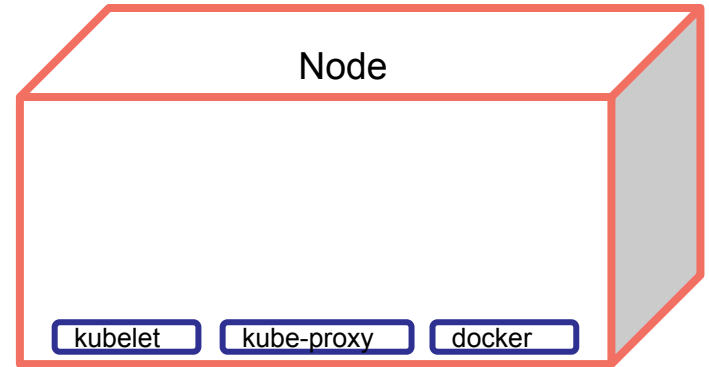
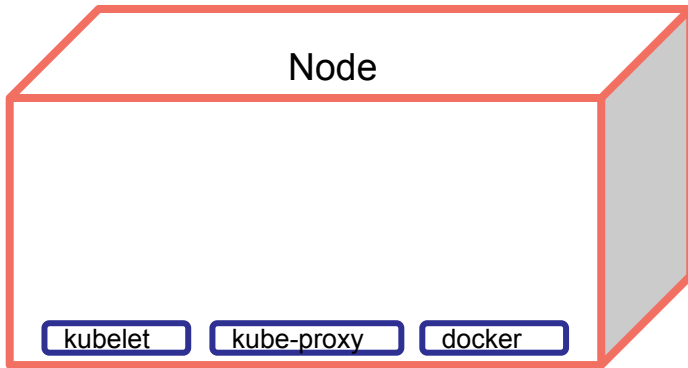
Greenplum Data Platform + PKS





Kubernetes 101






Kubernetes 101



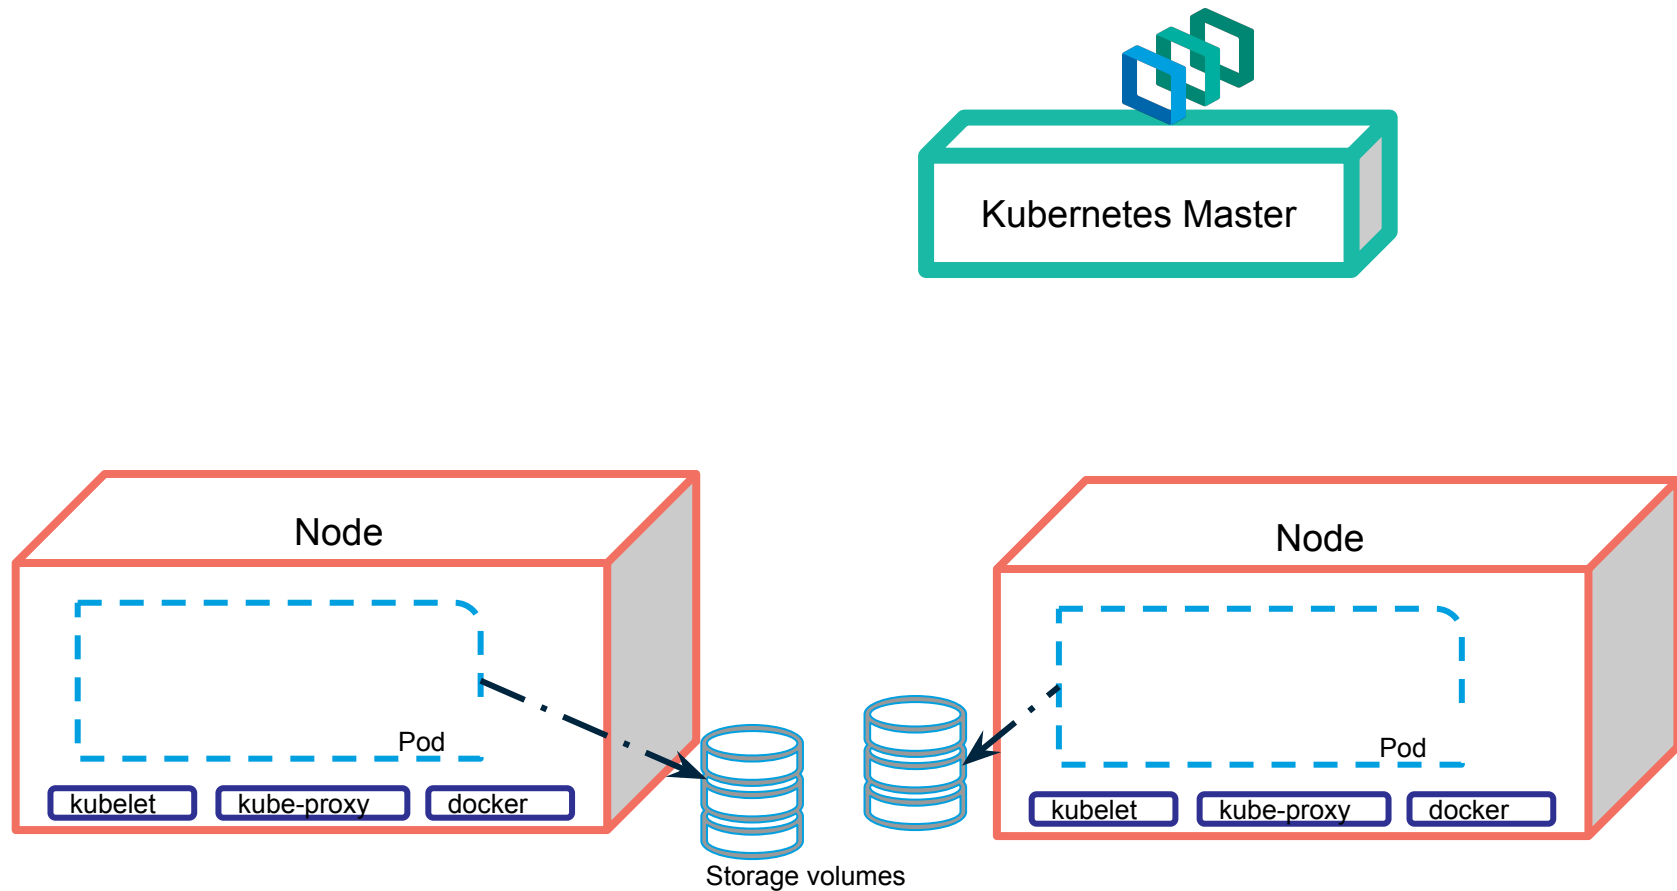
Kubernetes 101



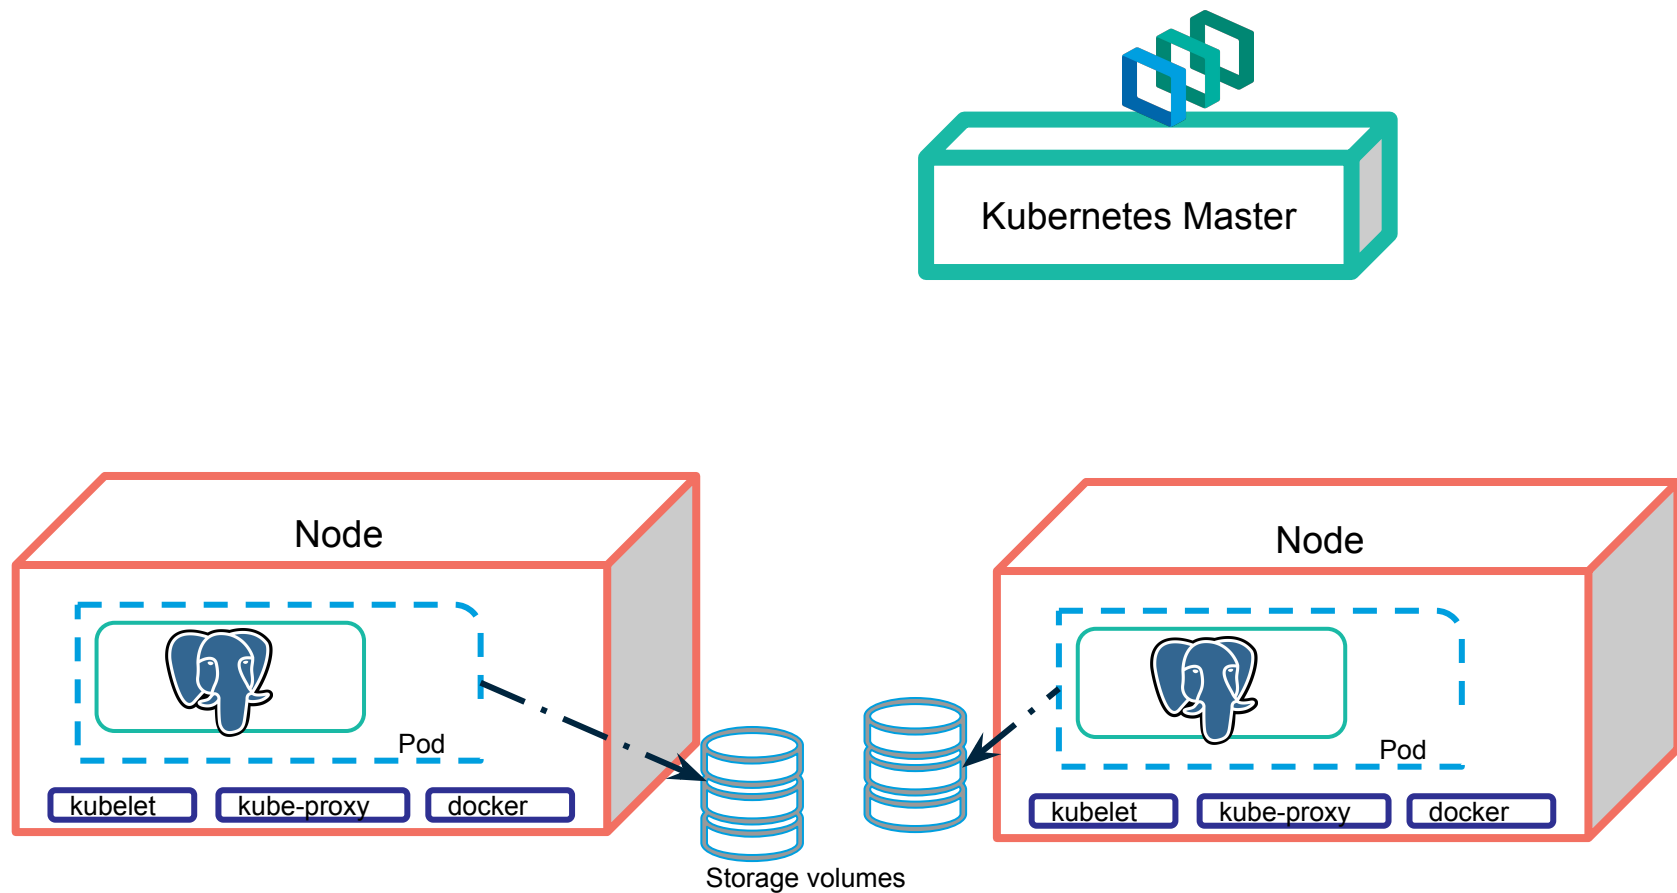
Kubernetes 101



Kubernetes 101




Kubernetes 101

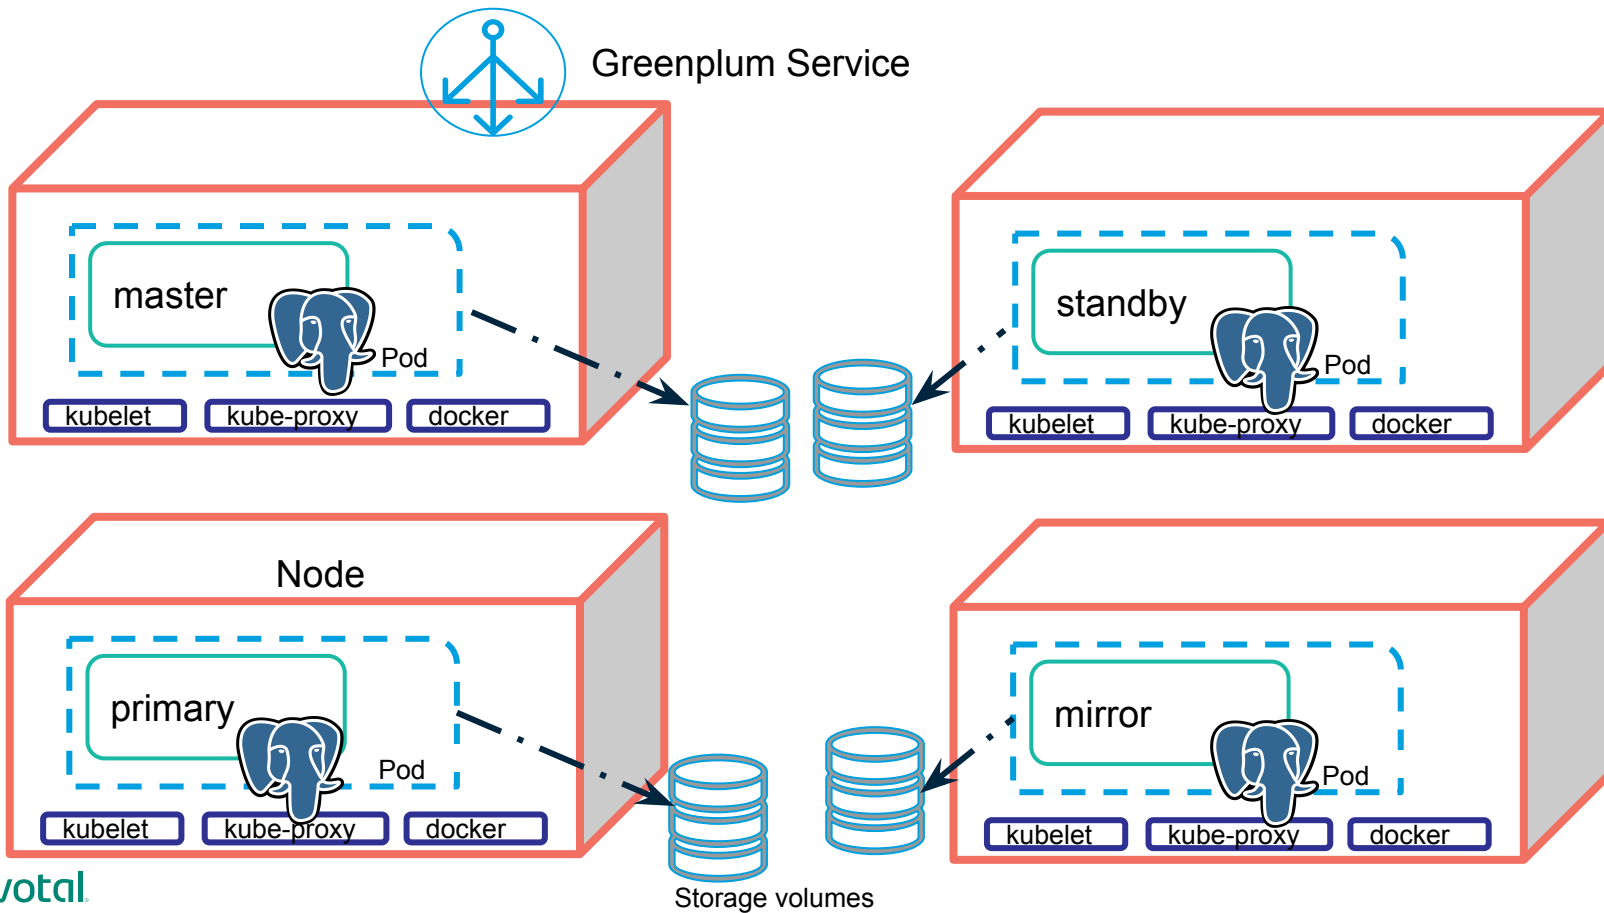


Kubernetes 101

Load Balancer Service



Greenplum on Kubernetes




Greenplum on PKS

Benefits

1. On Demand Cluster Provisioning



Alana




PKS

1. On Demand Cluster Provisioning



Alana

Give me a Greenplum
Cluster

A blue outline arrow pointing from the user icon towards the cluster icon.


PKS

1. On Demand Cluster Provisioning



Alana

Give me a Greenplum
Cluster



PKS

Cluster Alana



1. On Demand Cluster Provisioning




Alana

Give me a Greenplum
Cluster



gpdb-alana:5432



PKS



Cluster Alana




1. On Demand Cluster Provisioning




Alana

psql gpdb-alana:5432



Cluster Alana




PKS

1. On Demand Cluster Provisioning




Alana

psql gpdb-alana:5432



Cluster Alana



PKS




Dev Team

1. On Demand Cluster Provisioning



Alana


psql gpdb-alana:5432



Cluster Alana



Dev Team




PKS

1. On Demand Cluster Provisioning



Alana


psql gpdb-alana:5432



Cluster Alana



Dev Team



PKS

Cluster Dev




1. On Demand Cluster Provisioning



Alana



psql gpdb-alana:5432



Cluster Alana



Dev Team



PKS



Cluster Dev




1. On Demand Cluster Provisioning




Alana

psql gpdb-alana:5432



Cluster Alana



PKS



Dev Team

psql gpdb-dev:5432



Cluster Dev



2. Service Discovery

We can always discover a container by DNS.

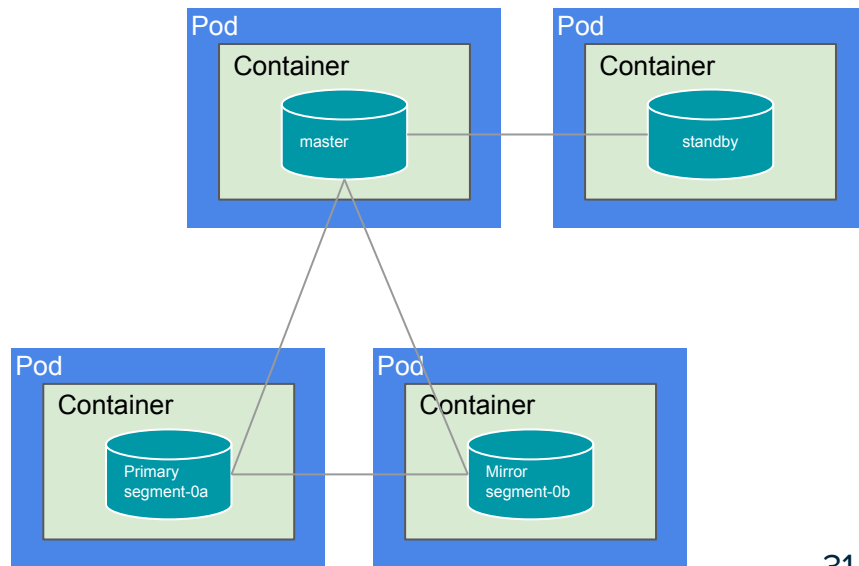
For example, DNS address for different roles:

master.greenplum.svc.cluster.local


standby.greenplum.svc.cluster.local

segment-0a.greenplum.svc.cluster.local

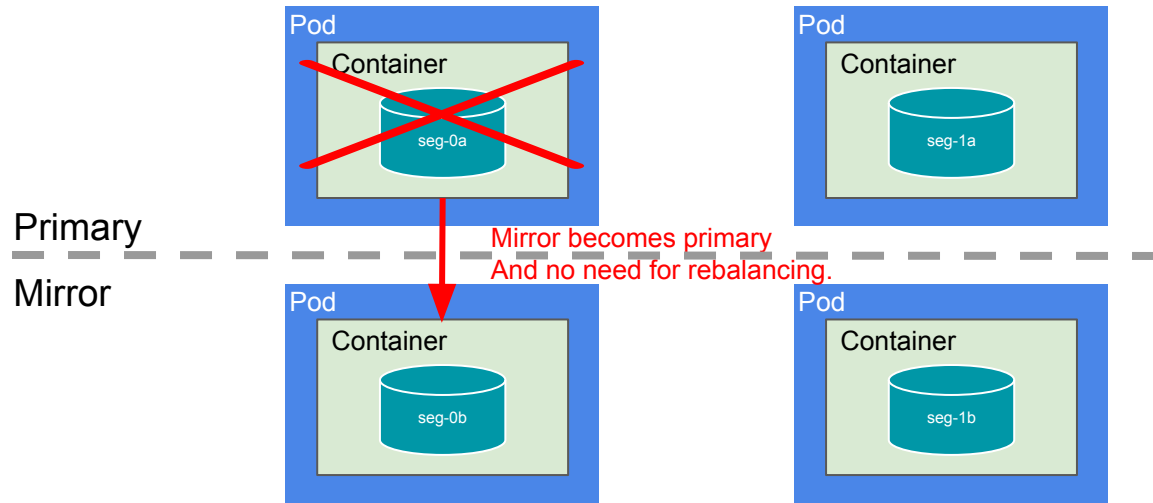
segment-0b.greenplum.svc.cluster.local



3. HA without Rebalancing



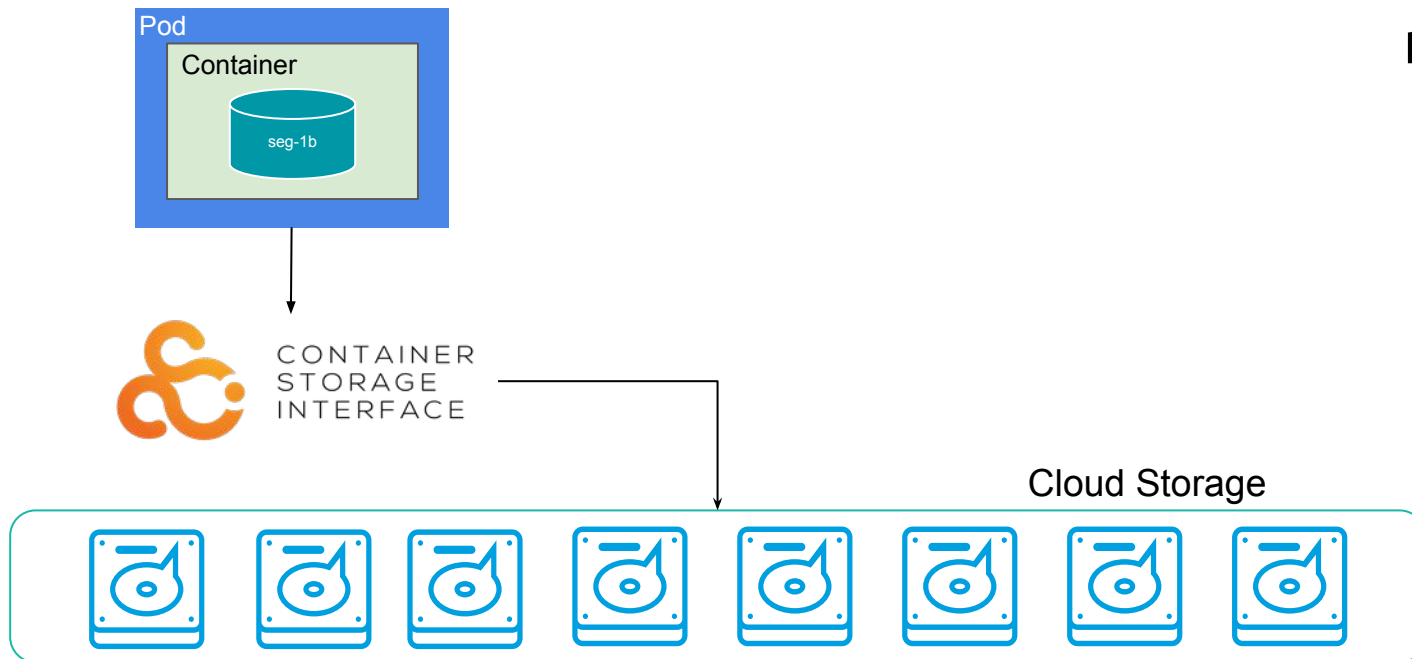
3. HA without Rebalancing



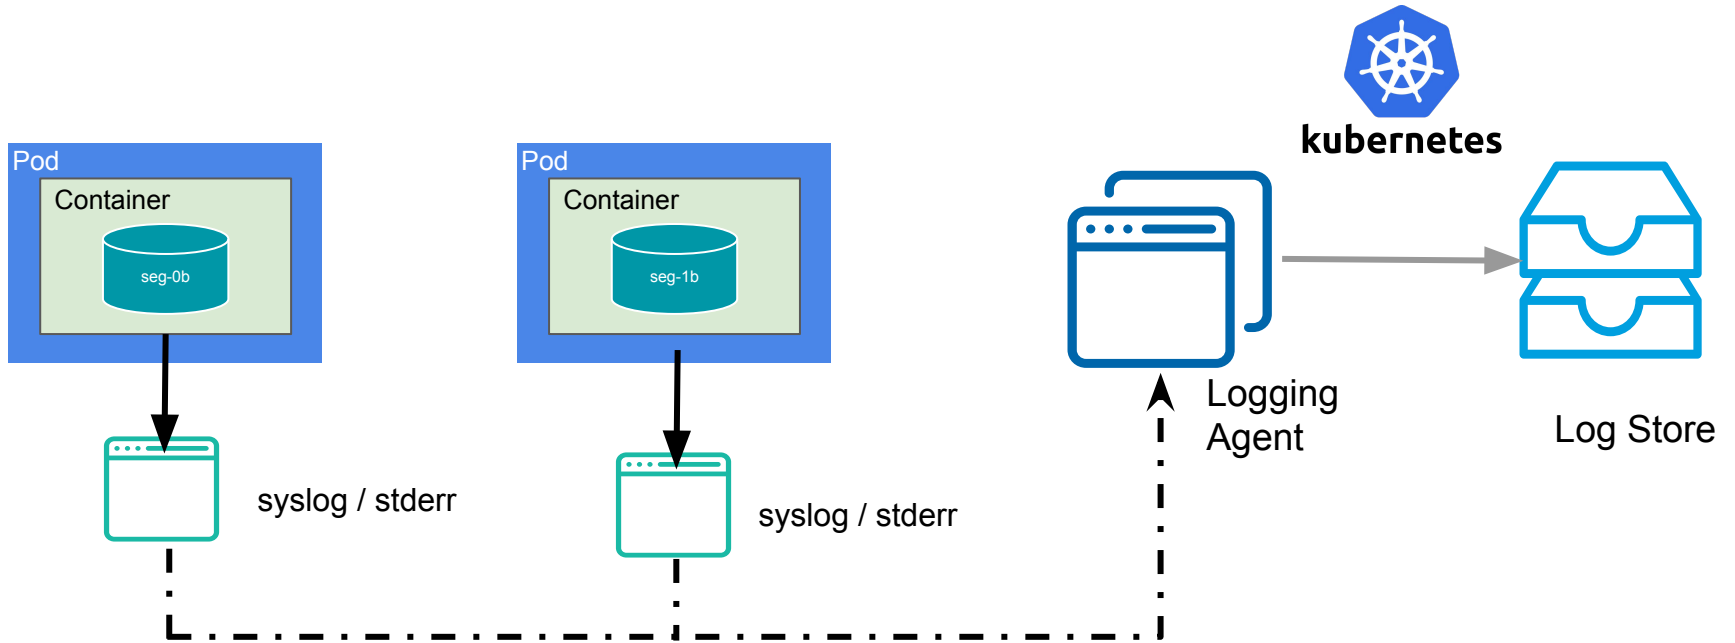
4. Kubernetes Plugins Support : Container Storage Interface



kubernetes



4. Kubernetes Plugins Support : Logging



GREENPLUM ON PKS DEMO

**HEY PKS! GIVE ME A
GREENPLUM CLUSTER OF “N”
SEGMENTS**

Demo

Deploy Greenplum on PKS

```
20180409:00:15:27:002418 gpstop:master:gpadmin-[INFO]:-Signalling all postmaster processes to reload
20180409:00:15:28:002457 gpinitstandby:master:gpadmin-[INFO]:-Validating environment and parameters for standby initialization.
20180409:00:15:29:002457 gpinitstandby:master:gpadmin-[INFO]:-Checking for filespace directory /greenplum/data-1 on standby
20180409:00:15:29:002457 gpinitstandby:master:gpadmin-[INFO]:-----
20180409:00:15:29:002457 gpinitstandby:master:gpadmin-[INFO]:-Greenplum standby master initialization parameters
20180409:00:15:29:002457 gpinitstandby:master:gpadmin-[INFO]:-----
20180409:00:15:29:002457 gpinitstandby:master:gpadmin-[INFO]:-Greenplum master hostname                = master
20180409:00:15:29:002457 gpinitstandby:master:gpadmin-[INFO]:-Greenplum master data directory           = /greenplum/data-1
20180409:00:15:29:002457 gpinitstandby:master:gpadmin-[INFO]:-Greenplum master port                    = 5432
20180409:00:15:29:002457 gpinitstandby:master:gpadmin-[INFO]:-Greenplum standby master hostname       = standby
20180409:00:15:29:002457 gpinitstandby:master:gpadmin-[INFO]:-Greenplum standby master port           = 5432
20180409:00:15:29:002457 gpinitstandby:master:gpadmin-[INFO]:-Greenplum standby master data directory = /greenplum/data-1
20180409:00:15:29:002457 gpinitstandby:master:gpadmin-[INFO]:-Greenplum update system catalog        = On
20180409:00:15:29:002457 gpinitstandby:master:gpadmin-[INFO]:-----
20180409:00:15:29:002457 gpinitstandby:master:gpadmin-[INFO]:- Filespace locations
20180409:00:15:29:002457 gpinitstandby:master:gpadmin-[INFO]:-----
20180409:00:15:29:002457 gpinitstandby:master:gpadmin-[INFO]:-pg_system -> /greenplum/data-1
20180409:00:15:29:002457 gpinitstandby:master:gpadmin-[INFO]:-Syncing Greenplum Database extensions to standby
20180409:00:15:29:002457 gpinitstandby:master:gpadmin-[INFO]:-The packages on standby are consistent.
20180409:00:15:29:002457 gpinitstandby:master:gpadmin-[INFO]:-Adding standby master to catalog...
20180409:00:15:29:002457 gpinitstandby:master:gpadmin-[INFO]:-Database catalog updated successfully.
20180409:00:15:29:002457 gpinitstandby:master:gpadmin-[INFO]:-Updating pg_hba.conf file...
20180409:00:15:30:002457 gpinitstandby:master:gpadmin-[INFO]:-pg_hba.conf files updated successfully.
20180409:00:15:32:002457 gpinitstandby:master:gpadmin-[INFO]:-Updating filespace flat files...
20180409:00:15:32:002457 gpinitstandby:master:gpadmin-[INFO]:-Filespace flat file updated successfully.
20180409:00:15:32:002457 gpinitstandby:master:gpadmin-[INFO]:-Starting standby master
20180409:00:15:32:002457 gpinitstandby:master:gpadmin-[INFO]:-Checking if standby master is running on host: standby in direct
ory: /greenplum/data-1
20180409:00:15:33:002457 gpinitstandby:master:gpadmin-[INFO]:-Cleaning up pg_hba.conf backup files...
20180409:00:15:34:002457 gpinitstandby:master:gpadmin-[INFO]:-Backup files of pg_hba.conf cleaned up successfully.
20180409:00:15:34:002457 gpinitstandby:master:gpadmin-[INFO]:-Successfully created standby master on standby
```

Demo

Expand Greenplum on PKS

```
20180321:21:08:55:083489 gpexpand:master:gpadmin-[INFO]:--Checking database template1 for tables with unique indexes
...
20180321:21:08:55:083489 gpexpand:master:gpadmin-[INFO]:--Checking database postgres for tables with unique indexes.
..
20180321:21:08:55:083489 gpexpand:master:gpadmin-[INFO]:--Checking database gpadmin for tables with unique indexes..
20180321:21:08:56:083489 gpexpand:master:gpadmin-[INFO]:--Heap checksum setting consistent across cluster
20180321:21:08:56:083489 gpexpand:master:gpadmin-[INFO]:--Syncing Greenplum Database extensions
20180321:21:08:57:083489 gpexpand:master:gpadmin-[INFO]:--The packages on segment-1a are consistent.
20180321:21:08:59:083489 gpexpand:master:gpadmin-[INFO]:--The packages on segment-1b are consistent.
20180321:21:09:01:083489 gpexpand:master:gpadmin-[INFO]:--Creating segment template
20180321:21:09:01:083489 gpexpand:master:gpadmin-[INFO]:--VACUUM FULL on the catalog tables
20180321:21:09:06:083489 gpexpand:master:gpadmin-[INFO]:--Starting copy of segment dbid 1 to location /greenplum/gpe
xpend_03212018_3489
20180321:21:09:08:083489 gpexpand:master:gpadmin-[INFO]:--Copying postgresql.conf from existing segment into templat
e
20180321:21:09:09:083489 gpexpand:master:gpadmin-[INFO]:--Copying pg_hba.conf from existing segment into template
20180321:21:09:11:083489 gpexpand:master:gpadmin-[INFO]:--Adding new segments into template pg_hba.conf
20180321:21:09:11:083489 gpexpand:master:gpadmin-[INFO]:--Creating schema tar file
20180321:21:09:13:083489 gpexpand:master:gpadmin-[INFO]:--Distributing template tar file to new hosts
20180321:21:09:26:083489 gpexpand:master:gpadmin-[INFO]:--Configuring new segments (primary)
20180321:21:09:27:083489 gpexpand:master:gpadmin-[INFO]:--Configuring new segments (mirror)
20180321:21:09:29:083489 gpexpand:master:gpadmin-[INFO]:--Backing up pg_hba.conf file on original segments
20180321:21:09:30:083489 gpexpand:master:gpadmin-[INFO]:--Copying new pg_hba.conf file to original segments
20180321:21:09:31:083489 gpexpand:master:gpadmin-[INFO]:--Configuring original segments
20180321:21:09:31:083489 gpexpand:master:gpadmin-[INFO]:--Cleaning up temporary template files
20180321:21:09:33:083489 gpexpand:master:gpadmin-[INFO]:--Starting Greenplum Database in restricted mode
20180321:21:09:51:083489 gpexpand:master:gpadmin-[INFO]:--Stopping database
20180321:21:10:06:083489 gpexpand:master:gpadmin-[INFO]:--Checking if Transaction filespace was moved
20180321:21:10:06:083489 gpexpand:master:gpadmin-[INFO]:--Checking if Temporary filespace was moved
20180321:21:10:06:083489 gpexpand:master:gpadmin-[INFO]:--Configuring new segment filespaces
20180321:21:10:06:083489 gpexpand:master:gpadmin-[INFO]:--Cleaning up databases in new segments.
20180321:21:10:06:083489 gpexpand:master:gpadmin-[INFO]:--Starting master in utility mode
20180321:21:10:10:083489 gpexpand:master:gpadmin-[INFO]:--Stopping master in utility mode
20180321:21:10:17:083489 gpexpand:master:gpadmin-[INFO]:--Starting Greenplum Database in restricted mode
20180321:21:10:36:083489 gpexpand:master:gpadmin-[INFO]:--Creating expansion schema
```


Demo

Greenplum Segment Failover

```

20180409:00:27:01:003141 gprecoverseg:master:gpadmin-[INFO]:-Starting gprecoverseg with args: -a
20180409:00:27:01:003141 gprecoverseg:master:gpadmin-[INFO]:-local Greenplum Version: 'postgres (Greenplum Database) 5.7.0 build
d f7c6eb5-oss'
20180409:00:27:01:003141 gprecoverseg:master:gpadmin-[INFO]:-master Greenplum Version: 'PostgreSQL 8.3.23 (Greenplum Database 5
.7.0 build f7c6eb5-oss) on x86_64-pc-linux-gnu, compiled by GCC gcc (Ubuntu 5.4.0-6ubuntu1-16.04.9) 5.4.0 20160609, 64-bit comp
iled on Mar 30 2016 19:36:48'
20180409:00:27:01:003141 gprecoverseg:master:gpadmin-[INFO]:-Checking if segments are ready to connect
20180409:00:27:01:003141 gprecoverseg:master:gpadmin-[INFO]:-Obtaining Segment details from master...
20180409:00:27:01:003141 gprecoverseg:master:gpadmin-[INFO]:-Obtaining Segment details from master...
20180409:00:27:02:003141 gprecoverseg:master:gpadmin-[INFO]:-Heap checksum setting is consistent between master and the segment
s that are candidates for recoverseg
20180409:00:27:02:003141 gprecoverseg:master:gpadmin-[INFO]:-Greenplum instance recovery parameters
20180409:00:27:02:003141 gprecoverseg:master:gpadmin-[INFO]:-----
20180409:00:27:02:003141 gprecoverseg:master:gpadmin-[INFO]:-Recovery type                - Standard
20180409:00:27:02:003141 gprecoverseg:master:gpadmin-[INFO]:-----
20180409:00:27:02:003141 gprecoverseg:master:gpadmin-[INFO]:-Recovery 1 of 1
20180409:00:27:02:003141 gprecoverseg:master:gpadmin-[INFO]:-----
20180409:00:27:02:003141 gprecoverseg:master:gpadmin-[INFO]:- Synchronization mode                = Incremental
20180409:00:27:02:003141 gprecoverseg:master:gpadmin-[INFO]:- Failed instance host                 = segment-0a
20180409:00:27:02:003141 gprecoverseg:master:gpadmin-[INFO]:- Failed instance address              = segment-0a
20180409:00:27:02:003141 gprecoverseg:master:gpadmin-[INFO]:- Failed instance directory            = /greenplum/data
20180409:00:27:02:003141 gprecoverseg:master:gpadmin-[INFO]:- Failed instance port                 = 40000
20180409:00:27:02:003141 gprecoverseg:master:gpadmin-[INFO]:- Failed instance replication port     = 50000
20180409:00:27:02:003141 gprecoverseg:master:gpadmin-[INFO]:- Recovery Source instance host       = segment-0b
20180409:00:27:02:003141 gprecoverseg:master:gpadmin-[INFO]:- Recovery Source instance address    = segment-0b
20180409:00:27:02:003141 gprecoverseg:master:gpadmin-[INFO]:- Recovery Source instance directory  = /greenplum/mirror
/data
20180409:00:27:02:003141 gprecoverseg:master:gpadmin-[INFO]:- Recovery Source instance port       = 50000
20180409:00:27:02:003141 gprecoverseg:master:gpadmin-[INFO]:- Recovery Source instance replication port = 6001
20180409:00:27:02:003141 gprecoverseg:master:gpadmin-[INFO]:- Mirror Target                       = In-place
20180409:00:27:02:003141 gprecoverseg:master:gpadmin-[INFO]:-----
20180409:00:27:02:003141 gprecoverseg:master:gpadmin-[INFO]:-1 segment(s) to recover
20180409:00:27:02:003141 gprecoverseg:master:gpadmin-[INFO]:-Ensuring 1 failed segment(s) are stopped

```

gprecoverseg a

Future Work

**More Components,
StatefulSets,
Custom Resource Definitions**

QUESTIONS?

gpcloud@pivotal.io

The background of the slide is a teal-tinted image of the Golden Gate Bridge, showing its iconic towers and suspension cables. The bridge spans across the frame from the bottom right towards the top left.

Pivotal[®]



Transforming How The World Builds Software