

POSTGRES ON KUBERNETES

Álvaro Hernández

ONGRES

POSTGRES ON KUBERNETES



`whoami`

- Founder & CEO, OnGres
- 20+ years PostgreSQL user and DBA
- Mostly doing R&D to create new, innovative software on Postgres
- Frequent speaker at PostgreSQL, database conferences
- Principal Architect of ToroDB
- Founder and President of the NPO *Fundación PostgreSQL*
- AWS Data Hero



Álvaro Hernández

<aht@ongres.com>

@ahachete

ONGRES

POSTGRESQL ON KUBERNETES



WHAT IS KUBERNETES

(non-classical intro)



//WHAT IS KUBERNETES?

An HTTP Server

ONGRES

POSTGRES ON KUBERNETES





//WHAT IS KUBERNETES?

An **object** HTTP Server

ONGRES

POSTGRES ON KUBERNETES





//WHAT IS KUBERNETES?

An object HTTP Server

That executes actions on containers when
create/modify/delete operations are run
on the HTTP-exposed objects

ONGRES

POSTGRES ON KUBERNETES



//WHAT IS KUBERNETES?

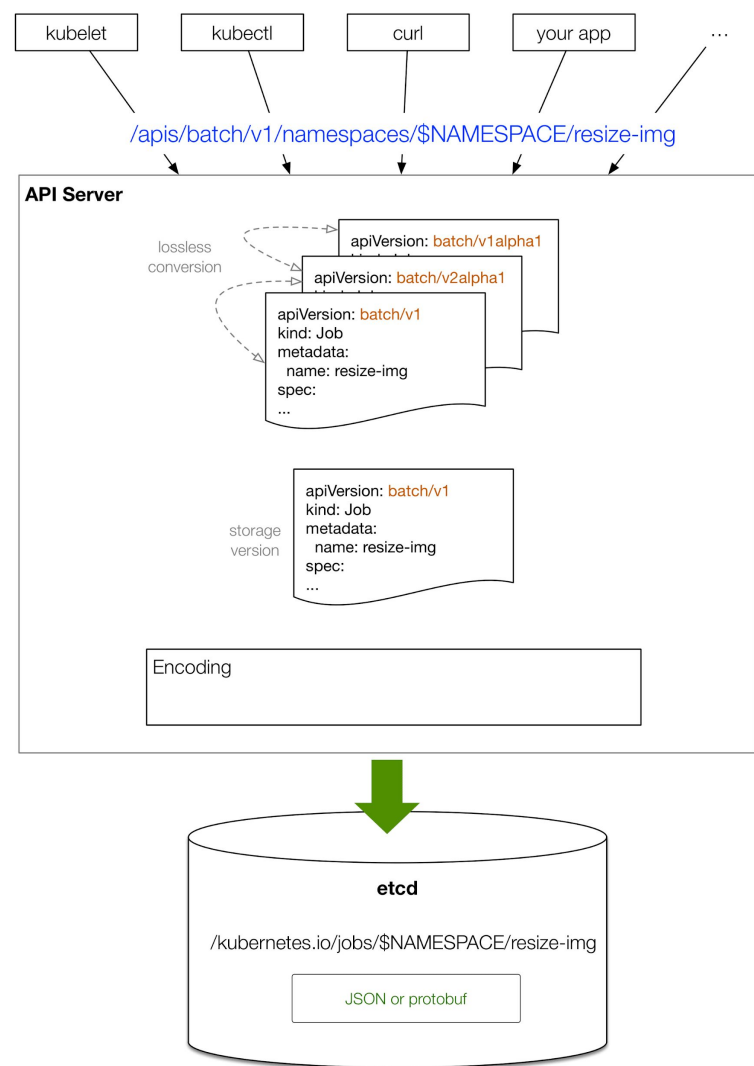
An object HTTP Server

Where the URLs of the objects and their fields form an API façade that represents the actions to be taken on containers

ONGRES

POSTGRES ON KUBERNETES

//WHAT IS KUBERNETES?



<https://blog.openshift.com/kubernetes-deep-dive-api-server-part-2>

//WHAT IS KUBERNETES?

WORKLOADS

Container v1 core
CronJob v1beta1 batch
DaemonSet v1 apps
Deployment v1 apps
Job v1 batch

Pod v1 core

Write Operations
Read Operations
Status Operations
Proxy Operations
Misc Operations

ReplicaSet v1 apps
ReplicationController v1 core
StatefulSet v1 apps

DISCOVERY & LOAD BALANCING

Endpoints v1 core
Ingress v1beta1 extensions
Service v1 core

CONFIG & STORAGE

ConfigMap v1 core
Secret v1 core
PersistentVolumeClaim v1 core
StorageClass v1 storage.k8s.io
Volume v1 core
VolumeAttachment v1beta1 storage.k8s.io

Pod v1 core

Group	Version	Kind
core	v1	Pod

! It is recommended that users create Pods only through a Controller, and not directly. See [Controllers: Deployment, Job, or StatefulSet](#).

Pod is a collection of containers that can run on a host. This resource is created by clients and scheduled onto hosts.

i Appears In:

- PodList core/v1

Field	Description
<code>apiVersion</code> <i>string</i>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#resources
<code>kind</code> <i>string</i>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#types-kinds
<code>metadata</code> <i>ObjectMeta</i>	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#metadata

kubectl

curl

Pod Config to print "Hello World".

```
apiVersion: v1
kind: Pod
metadata:
  name: pod-example
spec:
  containers:
  - name: ubuntu
    image: ubuntu:trusty
    command: ["echo"]
    args: ["Hello World"]
```

//WHY KUBERNETES?

- K8s is “the JVM” of the architecture of distributed systems: **an abstraction layer & API to deploy and automate infrastructure.**
- K8s provides APIs for nodes and IPs discovery, secret management, network proxying and load balancing, storage allocation, etc
- A PostgreSQL deployment can be fully automated!

UNGRES

POSTGRES ON KUBERNETES



DEMO - POSTGRES on K8S

("show the code")



TURNING POSTGRES CLOUD NATIVE

UNGRES

POSTGRES ON KUBERNETES



//CLOUD NATIVE

Cloud native applications are:

- designed to be packaged in containers
- scale and can be orchestrated for high availability

And follow cloud-native **best practices** including:

- Single-process hierarchy per container
- Sidecar containers to separate concerns
- Design for mostly ephemeral containers

ONGRES

POSTGRES ON KUBERNETES

//CONTAINERS ARE NOT SLIM VMS

- A container is an abstraction over a process hierarchy, with its own network, process namespaces and virtualized storage.
- But it is just a process hierarchy. Not many processes!
- No kernel, kernel modules, device drivers, no init system, bare minimum OS.
- **Should be just the binary of your process and its dynamic libraries and support files it needs.**

//IS POSTGRES FOR CONTAINERS?

- Overhead is minimal (1-2%): it is just a wrapper over the processes!
- Containers are as ephemeral as the process hierarchy they wrap.
- Advantage: they can be restarted somewhere if they fail.
- It's easier with stateless apps. But storage can be easily decoupled from containers: there are many storage persistence technologies.
- The entrypoint problem is typically solved by the container orchestration layer.

ONGRES

POSTGRES ON KUBERNETES

//MINIMAL CONTAINER IMAGE

- It's not about disk space or I/O.
It's about security and good design principles.
- PostgreSQL binaries are minimal: container image cannot be huge.
Remove:
 - Non-essential PostgreSQL binaries
 - Docs, psql
 - OS non system tools --all but */bin, /sbin, /lib**
 - Init system if any!

//LEVERAGE THE SIDECAR PATTERN

If a container should only have a single process hierarchy, how can we add support daemons like monitoring or HA agents?

- In K8s a **pod** is a set of 1+ containers that share the same namespaces, and run side-by-side on the same host.
- Sidecar pattern: deploy side functionality (like agents) to side containers (sidecars) on the same pod as PostgreSQL's container.
- Sidecars have the same IP and port space; process space (can send kill signals to processes), see the same persistent volume mount.

UNGRES

POSTGRES ON KUBERNETES

//HIGH AVAILABILITY (HA)

- HA is a native concept of *cloud native*.
- K8s provides mechanisms for leader election and HA.
But are not good for PostgreSQL!
- Leader election needs to be replication lag and topology aware.
- Also need to run operations after {fail,switch}over.
- Use PostgreSQL-specific HA mechanisms.
- Use K8s to automatically restart pods if they fail, and scale replicas.

//CENTRALIZED LOGGING

- A pattern that is not exclusive to containers, but reinforced in K8s.
- DBAs need not to “login” to every container to check logs.
- Centralized logs allow to:
 - Correlate events across multiple servers (leader / replicas).
 - Manage logs persistence once.
 - Run periodic reporting and alerting processes (like pgBadger).
 - Correlate with centralized monitoring (like Prometheus).

//K8S OPERATORS: AUTOMATE POSTGRES SQL OPS!

- Operators are just applications, developed for K8s
- **Understand PostgreSQL operations**
- Call K8s APIs to execute the operations
- Automate:
 - Minor version upgrades (rolling strategy)
 - Explicit vacuums
 - Repacks / reindex
 - Health checks


STACKGRES



UNGRES

POSTGRES ON KUBERNETES

//STACKGRES: CLOUD NATIVE POSTGRESQL



Running on Kubernetes. Embracing multi-cloud and on-premise.

Carefully selected and tuned a set of surrounding PostgreSQL components.

DB-as-a-Service without vendor lock-in. Root access.

Enterprise-grade PostgreSQL stack.

UNGRES

POSTGRESQL ON KUBERNETES

//A POSTGRES DISTRIBUTION (I)

Postgres installation is 25MB. Would you run this as-is on production?

OS, filesystem, Postgres tuning	Backups and DR
Connection pooling	High Availability
Centralized logging	Log parsing / alerting
Monitoring	Health checks
Repack	Logical replication software

UNGRES

POSTGRES ON KUBERNETES

//A POSTGRES DISTRIBUTION (II)

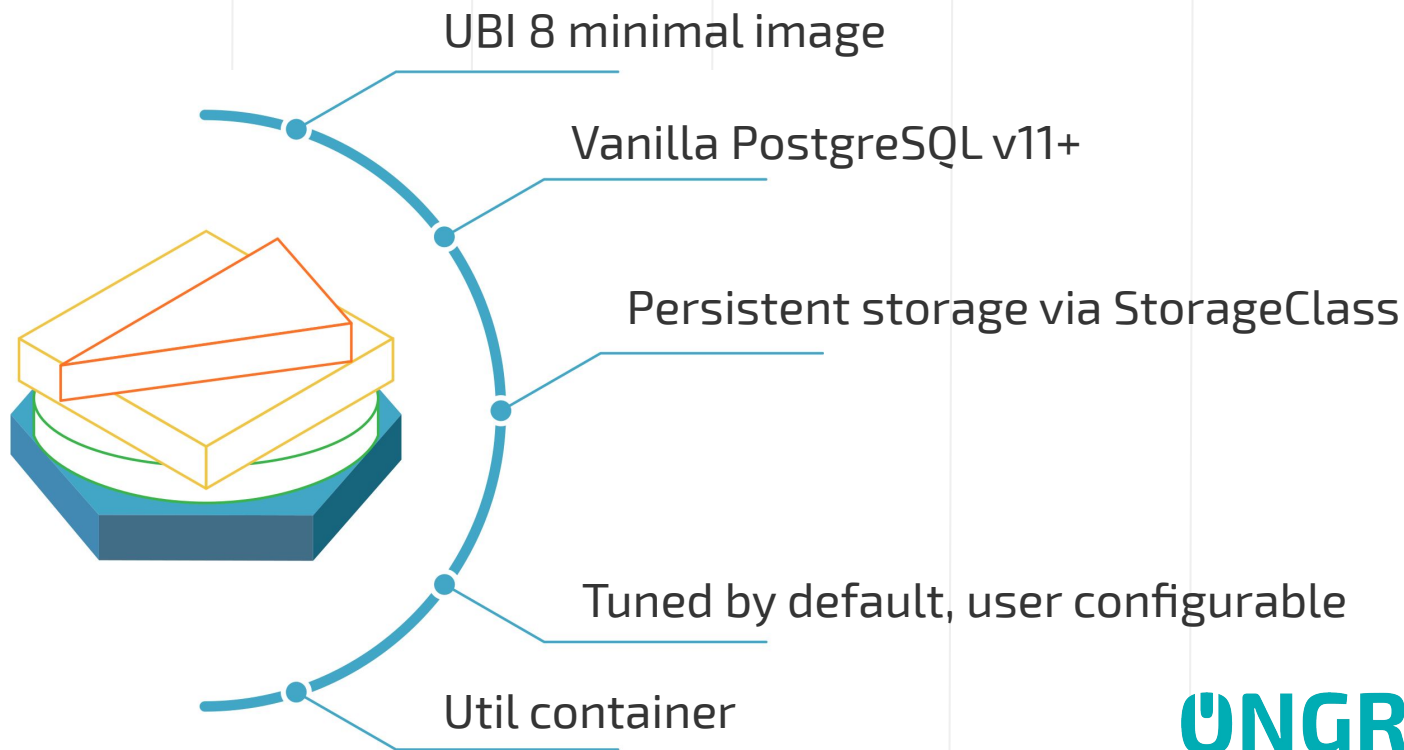
- StackGres = PostgreSQL + agg(all necessary PG ecosystem tools)
- Packaged as a single, easy deployable unit.
- Strongly opinionated stack (we picked the best tools for you).
- Based on Red Hat's UBI 8 (Universal Base Image).
- Basic PostgreSQL container is just 100MB.

The kernel is to Linux, what PostgreSQL is to StackGres

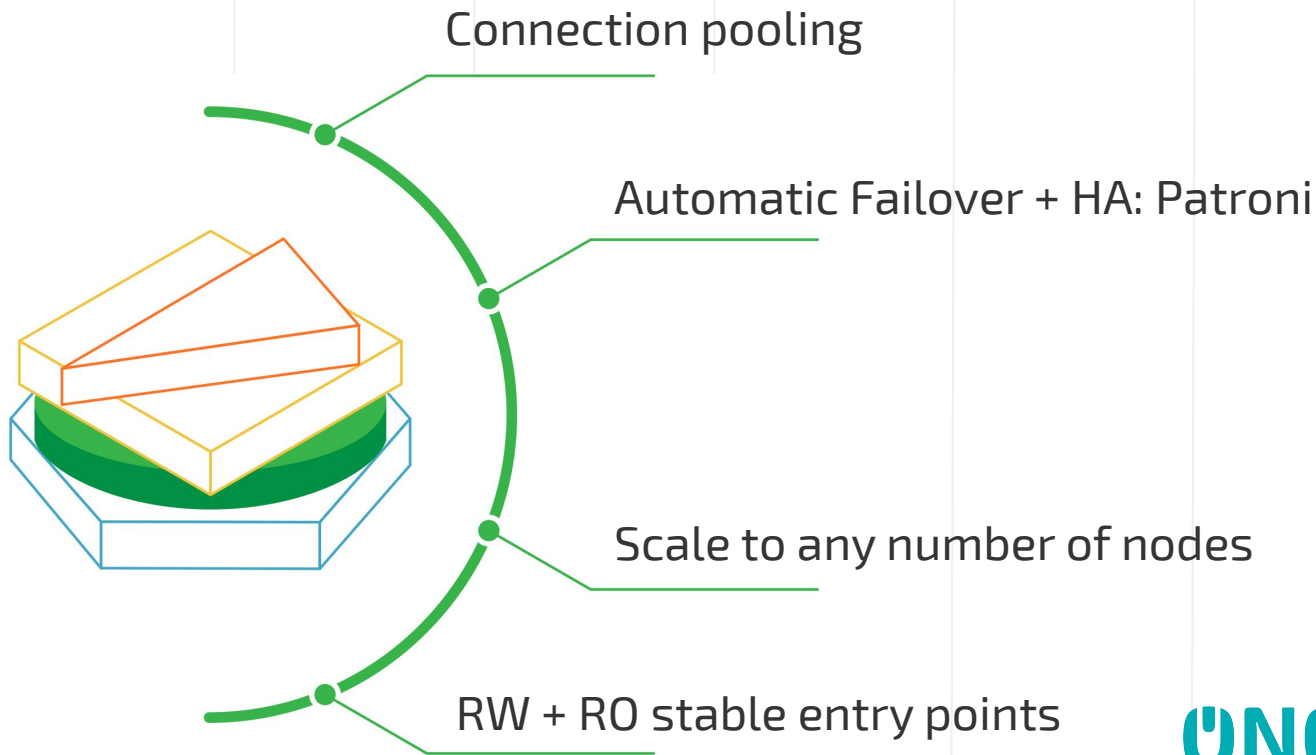
UNGRES

POSTGRES ON KUBERNETES

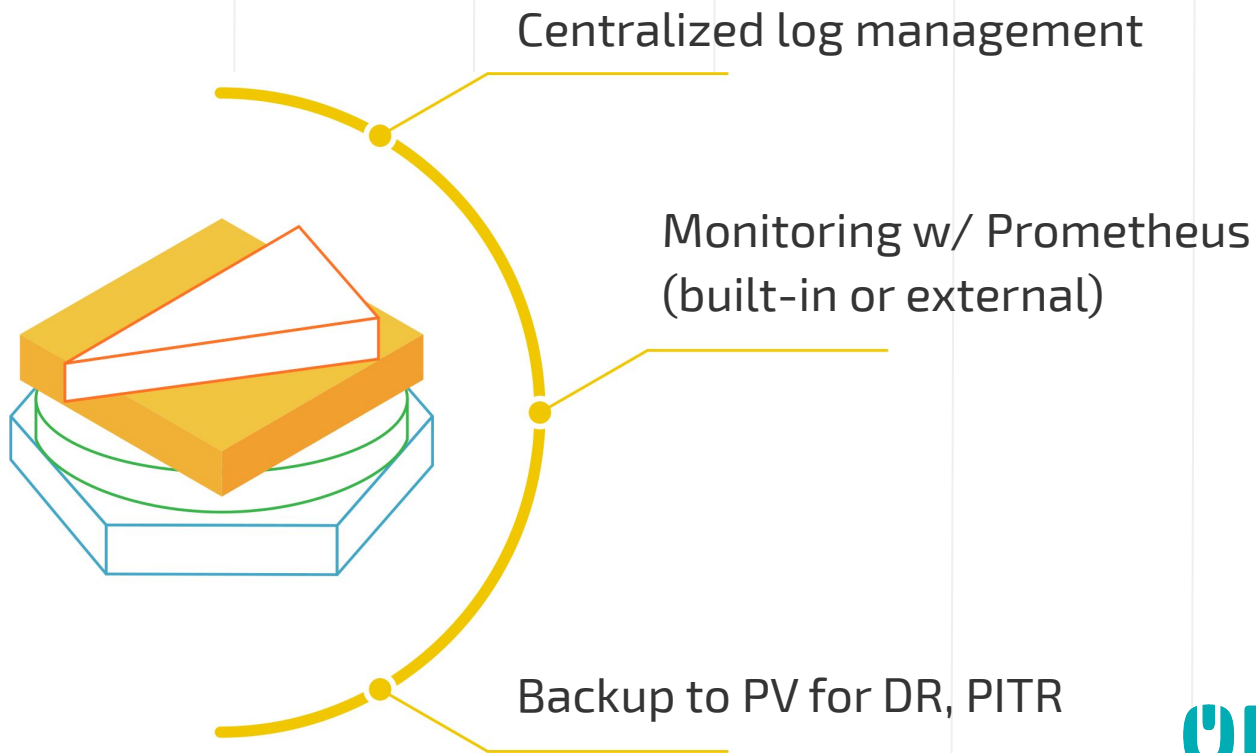
//THE STACKGRES STACK (I)



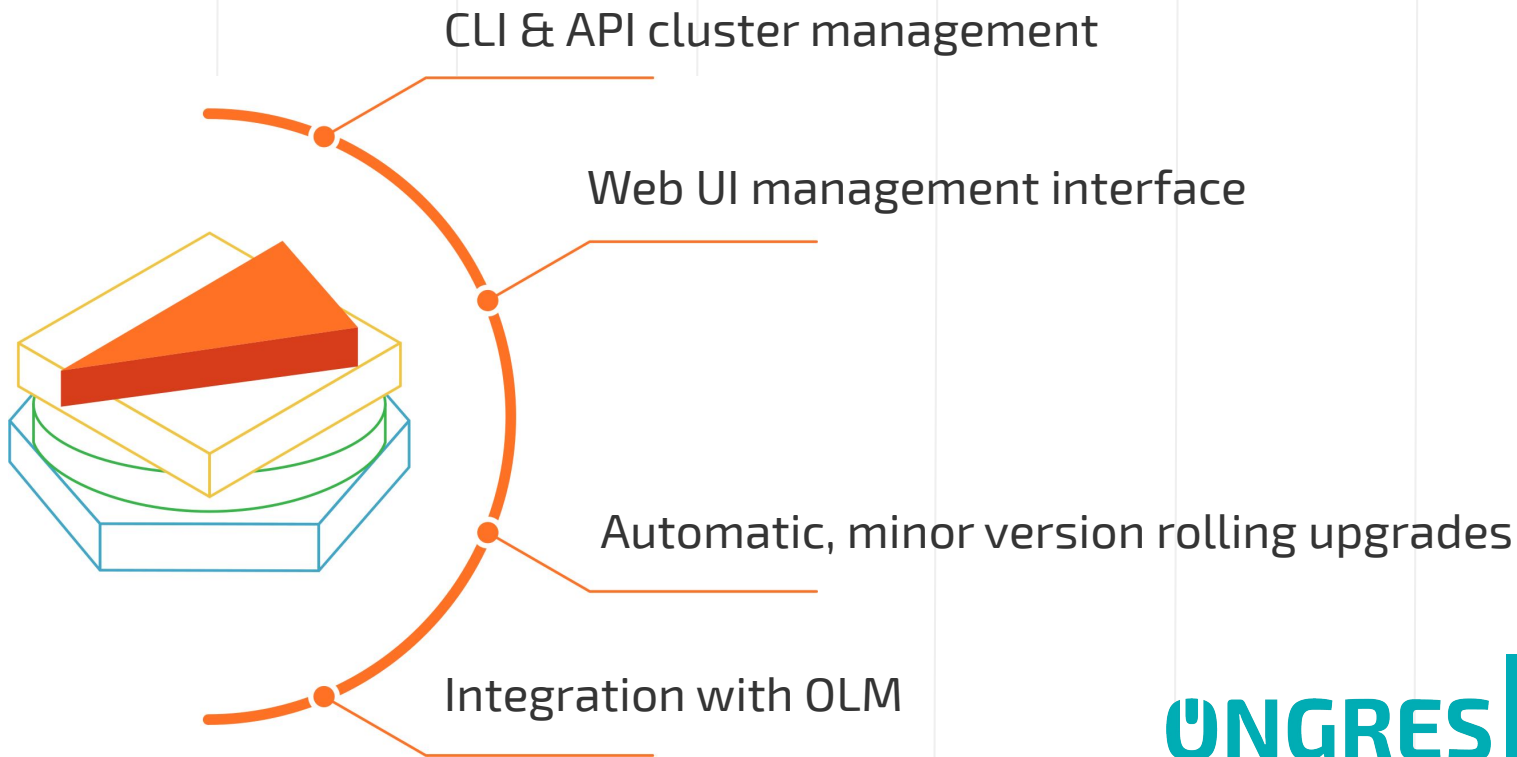
//THE STACKGRES STACK (II)




//THE STACKGRES STACK (III)



//THE STACKGRES STACK (IV)



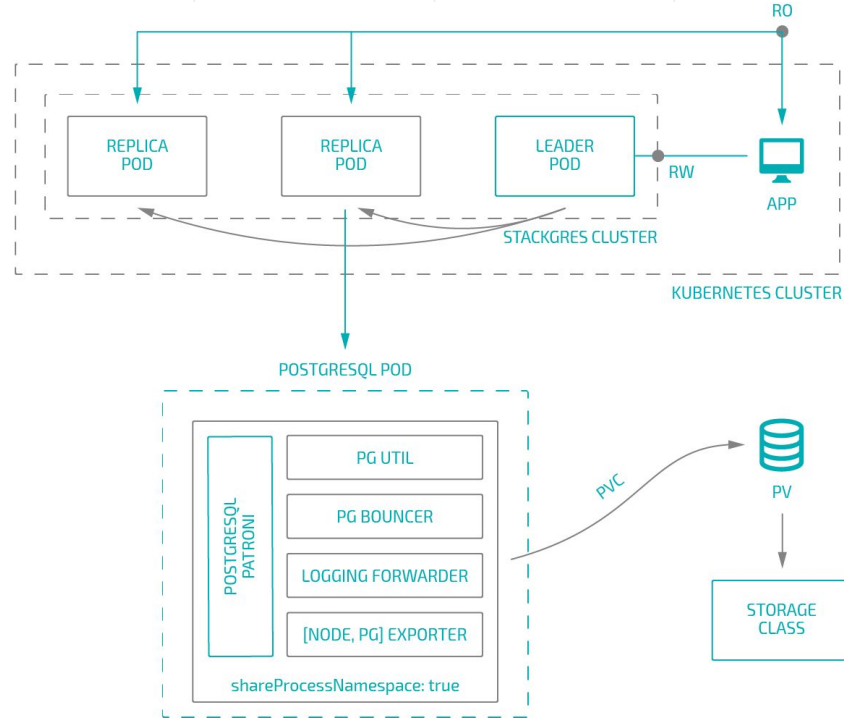
//STACKGRES ARCHITECTURE



ONGRES

STACKGRES: CLOUD NATIVE POSTGRESQL ON KUBERNETES

//STACKGRES ARCHITECTURE

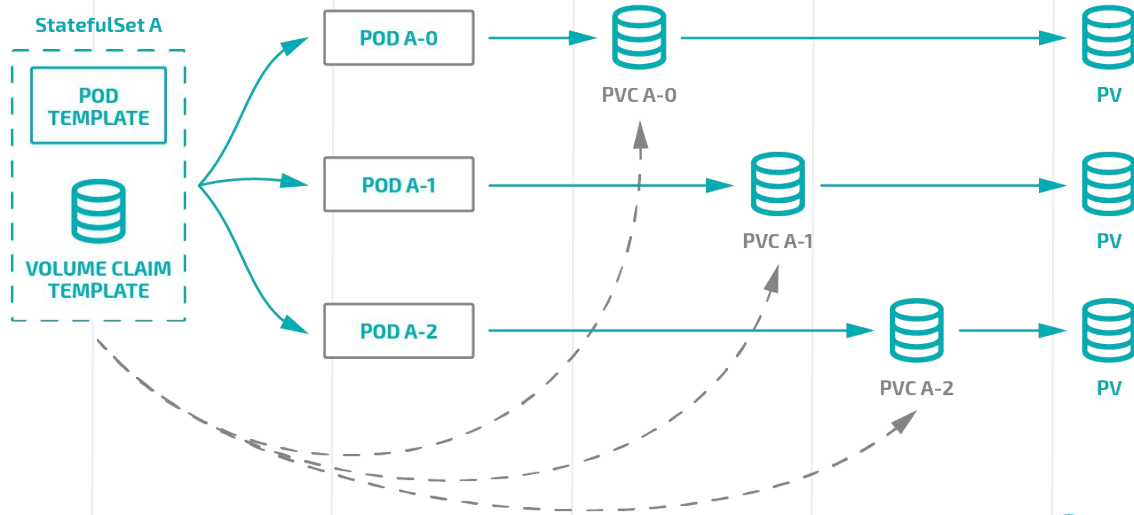


ONGRES

POSTGRES ON KUBERNETES

//STACKGRES ARCHITECTURE

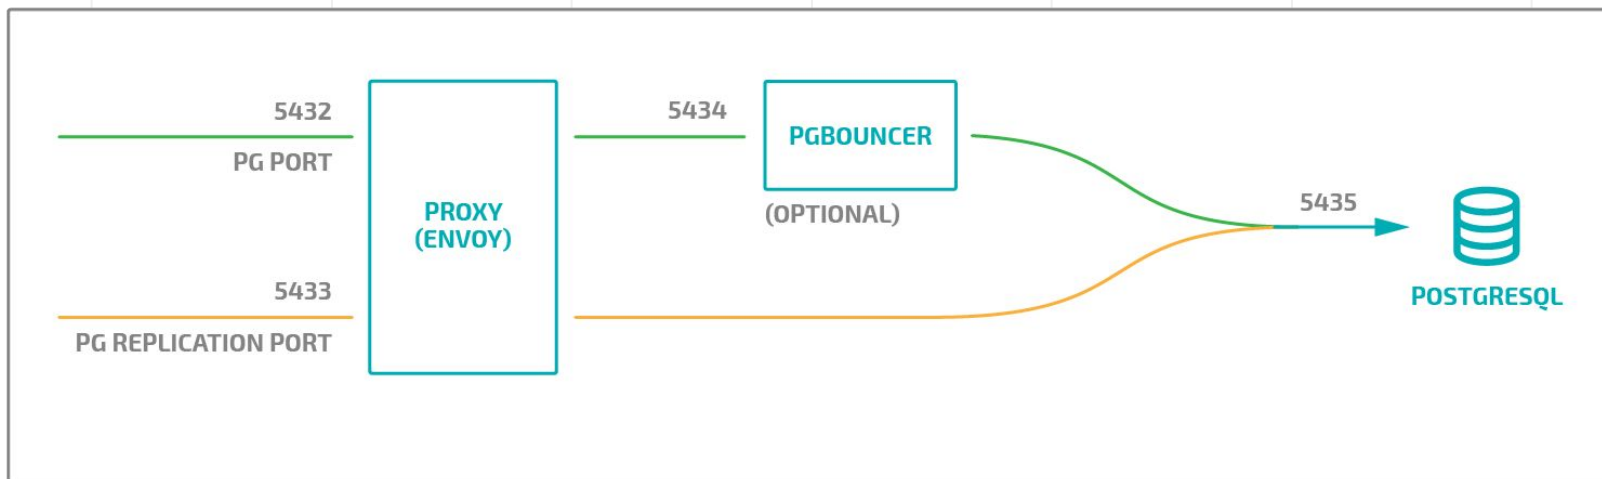
- Storage Class behavior:



//STACKGRES ARCHITECTURE

- Networking (via Envoy)

K8S POD



UNGRES

POSTGRES SQL ON KUBERNETES



DEMO

UNGRES



STACKGRES: CLOUD NATIVE POSTGRESQL ON KUBERNETES



//THANK YOU

QUESTIONS?

Álvaro Hernández
<aht@ongres.com>

@ahachete / www.ongres.com

ONGRES

CLOUD NATIVE POSTGRES SQL ON KUBERNETES

