# **DEPLOYING POSTGRESQL AS A STATEFULSET IN KUBERNETES**



There are different types of applications, from single large applications to <u>microservices-based</u> applications that cater to different needs. When it comes to the states of those applications, there are two states:

- **Stateless applications** can be run independently in isolation without any knowledge of past transactions.
- Stateful applications have full knowledge of the past information (state).

Most applications we use are stateful applications, and their state data may consist of user preferences, recent activity, database transactions, credentials, settings, etc.

Kubernetes provides StatefulSets when creating a stateful application in a Kubernetes cluster. Managing states within a <u>containerized environment</u> has become even more significant with the popularity of deploying database clusters in Kubernetes.

In this article, we will focus on how to deploy a <u>PostgreSQL database</u> on a Kubernetes cluster using StatefulSets.

(This article is part of our <u>Kubernetes Guide</u>. Use the right-hand menu to navigate.)

## What is Kubernetes StatefulSets?

StatefulSet is a Kubernetes workload API object that can be used to manage stateful applications.

In a normal deployment, the user does not care how the pods are scheduled as long as it doesn't have a negative impact on the deployed application. However, there is the need to properly identify

pods to preserve the state in stateful applications with persistent storage.

StatefulSet provides this functionality by <u>creating pods</u> with a persistent identifier that will pertain to its value across rescheduling. This way, a pod will get correctly mapped to the storage volumes even if it is recreated, and the application state will be preserved.

#### StatefulSets use cases

There are several use cases for StatefulSets.



### **Ordered deploying & scaling**

When an application relies on multiple containers, the ordered approach to scaling ensures that dependent containers are created in an orderly manner at deployments and scaling scenarios.

### Ordered automated rolling updates

Updating applications or microservices that are dependent should also be updated in an orderly manner. Additionally, an update should not affect the functionality.

Therefore, users can decide the order in which the applications or microservices should be updated by using a StatefulSet.

#### Mapping persistent storage

When considering <u>databases</u>, persistent storage is the most critical part as applications need to store data. With a StatefulSet, users can:

- 1. Define which pods correspond to each persistent storage
- 2. Create resilient application deployments

#### Using unique network identifiers to create persistent network connectivity

With unique identifiers, network users can manage and route traffic for specific pods without worrying about IP changes at rescheduling. This provides greater control over the network communications between pods by providing the ability to configure persistent routing, policies, and security configs for desired pods.

Even with these benefits, StatefulSets do not provide a solution for all requirements. For instance, StatefulSets are not interchangeable with <u>deployments</u> or <u>ReplicaSets</u>—these are instead geared to stateless configurations.

## **Drawbacks of StatefulSets**

StatefulSets also come with a set of limitations that users should be aware of before deploying the application.

- The storage for a StatefulSet must be provisioned either by a PersistentVolume Provisioner based on the storage class or pre-provisioned.
- Scaling or deleting pods will not affect the underlying persistent storage in order to ensure data safety. The provisioned volumes will remain within Kubernetes.
- The user needs to create a <u>headless service</u> manually to ensure network identity in StatefulSets.
- SatefulSets does not guarantee the termination of current pods when the StatefulSet is deleted. So, best practice is that the user implement an SOP to scale the StatefulSet to zero pods before deleting.
- Rolling updates with the default pod management policy may cause issues when deploying if a pod is broken (due to an application config error, bad binary, etc.). In such instances, users need to:
  - Manually revert to a previous deployment template.
  - $\,\circ\,$  Delete the broken Pods before attempting to rerun the new updates.

## Setting up a StatefulSet in a Kubernetes cluster

Now that we have a basic understanding of a StatefulSet, let's look at a sample StatefulSet deployment.

StatefulSets are ideal for database deployments. In this example, we will create a PostgreSQL deployment as a StatefulSet with a persistent storage volume.

#### postgresql.yaml

```
# PostgreSQL StatefulSet
apiVersion: apps/v1
kind: StatefulSet
metadata:
   name: postgresql-db
spec:
   serviceName: postgresql-db-service
   selector:
```

```
matchLabels:
    app: postgresql-db
replicas: 2
template:
  metadata:
    labels:
      app: postgresql-db
  spec:
    containers:
      - name: postgresql-db
        image: postgres:latest
        volumeMounts:
          - name: postgresql-db-disk
            mountPath: /data
        env:
          - name: POSTGRES PASSWORD
            value: testpassword
          - name: PGDATA
            value: /data/pgdata
# Volume Claim
volumeClaimTemplates:
  - metadata:
      name: postgresql-db-disk
    spec:
      accessModes:
      resources:
        requests:
          storage: 25Gi
```

In the above YAML file, we have defined a simple StatefulSet to deploy a PostgreSQL database. We are creating a StatefulSet called postgresql-db with two pods (replicas: 2).

Additionally, we are creating a <u>Persistent Volume</u> using the volumeClaimTemplate and using it in the StatefulSet to store the PostgreSQL data. The default Persistent Volume provisioner will provision the volume, and we can deploy this by running the following command.

kubectl apply -f postgresql.yaml

Result:

```
> kubect1 apply -f postgresql.yaml
statefulset.apps/postgresql-db created
```

Now we have successfully

created a PostgreSQL StatefulSet yet need a service to expose it outside of the Kubernetes cluster. That can be done by creating a service that points to the StatefulSet.

#### postgresql-service.yaml

# PostgreSQL StatefulSet Service
apiVersion: v1

```
kind: Service
metadata:
   name: postgres-db-lb
spec:
   selector:
    app: postgresql-db
   type: LoadBalancer
   ports:
        - port: 5432
        targetPort: 5432
```

This will create a <u>load balancer</u>, and the service will expose our PostgreSQL database using the selector "app: postgresql-db." You can create the service by running the below command.

kubectl apply -f postgresql-service.yaml

Result:
> kubect1 apply -f postgresql-service.yaml
service/postgres-db-lb created

Now let's see if both the

StatefulSet and the service are successfully created in Kubernetes by running the following command.

kubectl get all

Result:

> kubectl get all							
NAME	READY	STATUS	RESTARTS	AGE			
pod/postgresql-db-0	1/1	Running	Θ	3m2	4s		
pod/postgresql-db-1	1/1	Running	0	2m5	Бs		
NAME	TYPE		CLUSTER-IP		EXTERNAL-IP	PORT(S)	AGE
service/kubernetes	Clus	terIP	10.96.0.1		<none></none>	443/TCP	21m
service/postgres-db-lb	b Load	Balancer	10.111.253	.4	10.111.253.4	5432:30975/TCP	3m7s
NAME		READY	AGE				
statefulset.apps/postgresql-db 2/2			3m24s				

The result indicates that two pods are created and running with a load balancer service exposing the StatefulSet via IP 10.111.253.4 using port 5432.

#### **Testing the connectivity**

Having deployed PostgreSQL, we need to verify that we can access it without any issues. We will use the <u>pgAdmin4 client</u> to initialize the connection:

- 1. Download the pgAdmin client in your environment.
- 2. Connect and try to initialize a connection.

With the above deployment, we will use the external IP of the postgres-db-lb service (10.111.253.4) with port 5432. Since we only defined a password in our environment variables for the PostgreSQL

StatefulSet, the configuration will have the default username "postgres" with the password we defined.

🖀 Create - Server		2 ×
General Connec	ction SSL SSH Tunnel Advanced	
Host name/address	10.111.253.4	
Port	5432	
Maintenance database	postgres	
Username	postgres	
Kerberos authentication?	False	
Password		
Save password?		
Role		
Service		
i ?	🗙 Cancel 🗳 Reset 🖬 S	ave

If all the details are correct, the

connection will be initiated when you click on "Save," and the user will be able to see the connected database.

ggAdmin 4     GgAdmin 4     File ∽ Object ∽	Tools∨ Hel <u>p</u> ∨							– 🗆 ×	
Browser 🗣 🎟 🖬 🔍 🔎	Dashboard Prop	erties SQL Statistics	Dependencies Depend	ents				×	
<ul> <li>Servers (1)</li> <li>W Kubernetes PostgreSQL</li> </ul>	Server sessions			Total 📕 Active 📕 Idle	Transactions per second		Transactions 📕 🤇	Commits <mark>–</mark> Rollbacks	
🗸 🛢 Databases (1)	6								
✓ ≡ postgres	5 4								
> 😵 Casis									
<ul> <li>&gt; Event Triggers</li> </ul>									
> 🛱 Extensions									
> 🥣 Foreign Data Wrapper	0				0				
Canguages	Tuples in	📕 Inserts 📕 U		out	Fetched Returned	Bloc	k I/O	Reads Hits	
> Schemas	1		120			12	0		
> Subscriptions			100			10	0		
> 📥 Login/Group Roles			80			8	0		
> 🔁 Tablespaces			60-			6	i0		
			40			4	0		
	0		20-			2	0		
	Server activity								
	Sessions Locks	s Prepared Transactions	Configuration				<b>Q</b> Search	9	
	PID	Database User	Application	Client E	ackend start S	State V	Vait event	Blocking PIDs	
	😣 🔳  69				021-06-23 19:11:12 UTC	ŀ	ctivity: CheckpointerMain		
	⑧ ■ ▶ 70			2	021-06-23 19:11:12 UTC	ļ	ctivity: BgWriterMain		
	⑧ ■ ▶ 71			2	021-06-23 19:11:12 UTC	ļ	ctivity: WalWriterMain		
	😢 🗖  72			2	2021-06-23 19:11:12 UTC	ļ	ctivity: AutoVacuumMain		
	8 ■ > 74	postares		2	021-06-23 19:11:12 UTC	ļ	ctivity: Logicall auncherMain		
			ngAdmin 4 - DB:nostgros	172 17 0 1	021-06-23 10:33:44 UTO	active			
	97	posigres posigres	pgAdmin 4 - DB.postgres	172.17.0.1 2	021 00 23 19.33.44 010	active			

## Using ConfigMap in StatefulSet

In our earlier example, we defined the environment variables with the StatefulSet YAML.

However, the best practice would be to separate the environment variables using <u>ConfigMaps</u> and call the ConfigMap from the StatefulSet deployment. This makes it easier to manage and maintain each component of the deployment.

So, let's create a ConfigMap and modify the StatefulSet YAML as shown below.

#### postgresql-configmap.yaml

```
# PostgreSQL StatefulSet ConfigMap
apiVersion: v1
kind: ConfigMap
metadata:
   name: postgres-db-config
   labels:
      app: postgresql-db
data:
   POSTGRES_DB: testdb
   POSTGRES_USER: testdbuser
   POSTGRES_PASSWORD: testdbuserpassword
   PGDATA: /data/pgdata
```

In the above ConfigMap, we have extended our environment variable to specify a PostgreSQL

database, user, password, and data store. Now let's create the configMap and view the configurations using this command:

#### kubectl apply -f .postgresql-configmap.yaml

Result:

```
> kubectl apply -f .\postgresql-configmap.yaml
configmap/postgres-db-config created
```

We can get the information of

the created ConfigMap using the describe function:

kubectl describe configmap postgres-db-config

Result:

> kubectl desc	cribe configmap postgres-db-config
Name:	postgres-db-config
Namespace:	default
Labels:	app=postgresql-db
Annotations:	<none></none>
Data	
====	
PGDATA:	
/data/pgdata	
POSTGRES_DB:	
testdb	
POSTGRES_PASS	WORD:
testdbuserpass	word
POSTGRES_USER:	
testdbuser	
Events: <none< td=""><td></td></none<>	

The next step is

to modify the StatefulSet to call the data from the ConfigMap. That can be done by using the envFrom field to point to the above ConfigMap. This will also enable us to create a StatefulSet using the data in the ConfigMap.

#### postgresql.yaml

```
# PostgreSQL StatefulSet - ConfigMap
apiVersion: apps/v1
kind: StatefulSet
metadata:
    name: postgresql-db
```

```
spec:
  serviceName: postgresql-db-service
  selector:
    matchLabels:
      app: postgresql-db
  replicas: 2
  template:
    metadata:
      labels:
        app: postgresql-db
    spec:
      containers:
        - name: postgresql-db
          image: postgres:latest
          volumeMounts:
            - name: postgresql-db-disk
              mountPath: /data
          # Config from ConfigMap
          envFrom:
          - configMapRef:
              name: postgres-db-config
 # Volume Claim
  volumeClaimTemplates:
    - metadata:
        name: postgresql-db-disk
      spec:
        accessModes:
        resources:
          requests:
            storage: 25Gi
```

With our Persistent Volumes, deletions in the underlying database will be preserved even in case of Pod errors. As a StatefulSet, the state of Pods will also be preserved, and they will get assigned to the desired volumes correctly when recreated.

We can identify the pods using the following command:

kubectl get pvc

Result:

> kubectl get p	pvc					
NAME			STATUS	VOLUME	CAPACITY	AC
CESS MODES STO	DRAGECLASS	AGE				
postgresql-db-di	isk-postgresq	1-db-0	Bound	pvc-16800078-b0e0-4911-a7c0-9e2ead637453	25Gi	RW
0 sta	andard 4	4m12s				
postgresql-db-di	isk-postgresq	1-db-1	Bound	pvc-165ed65c-a1f3-4a85-a3aa-bf214c53b767	25Gi	RW
0 sto	andard 4	4m5s				

## Managing state is crucial to app functionality

Kubernetes StatefulSets allows users to easily create and manage stateful applications or services within a Kubernetes cluster. However, these StatefulSets configurations involve some complexity—so you must carefully plan your deployments before them carrying out.

Additionally, StatefulSets are the ideal solution for dealing with database applications, payment services, etc., where managing state is a crucial part of the application functionality.

## **Related reading**

- BMC DevOps Blog
- Kubernetes Multi-Clusters: How & Why To Use Them
- <u>Kubernetes Best Practices for Enhanced Cluster Efficiency</u>
- <u>3 Kubernetes Patterns for Cloud Native Applications</u>
- <u>The State of Containers Today</u>
- <u>Containerized Machine Learning: An Intro to ML in Containers</u>