

OX PowerDNS Cloud Control

Overview

Aug 12, 2022

Release 2.1.0

Contents

1	Cloud Control	1
1.1	Simple deployment - Recursor	1
1.2	Simple deployment - Auth	1
1.3	Complex deployment	2
1.4	Rules & Actions	2
1.5	DNSdist with co-hosted Recursors	3
1.6	DNSdist with DoH and/or DoT listeners	3
1.7	ZoneControl deployment	4
2	Cloud Control on Kubernetes	5
2.1	Auth	5
2.1.1	Auth agent	6
2.2	DNSdist	6
2.2.1	DNSdist agent	7
2.3	Recursor	8
2.3.1	Recursor agent	9
2.4	Resolver	9
2.5	Ruleset	10
2.6	ZoneControl	10
2.6.1	ZoneControl Syncer	11
3	Helm Charts	12
3.1	Helm Chart: powerdns-crds	12
3.2	Helm Chart: powerdns	12
3.3	Helm Chart: powerdns-operators	12
4	Getting Started	13
4.1	Install Tools	13
4.2	Download Helm Charts	13
4.3	Install/Upgrade CloudControl CRDs	14
4.4	Install/Upgrade CloudControl	14
4.4.1	Registry Credentials	15
4.4.2	Cluster Networking	15
4.4.3	Deploying Recursor	17
4.4.4	Adding DNSdist	18
4.4.5	Adding an external Resolver	19
4.4.6	Adding a DNSdist rule	20
4.4.7	Using DNSdist rules to route traffic	22
4.4.8	Separating config into multiple files	23
4.4.9	Exposing dnsdist	25
4.4.10	Deploying ZoneControl	25

5	Advanced Examples	27
5.1	DNSdist: DoH	27
5.2	DNSdist: DoT	28
5.3	DNSdist: Co-hosted Recursor	29
6	Troubleshooting	30
6.1	Accessing DNSdist console	30
6.2	Pod Events	31


1 Cloud Control

Cloud Control facilitates orchestration, management & monitoring of OX PowerDNS products in Kubernetes deployments. OX PowerDNS products supported in this version are:

- OX PowerDNS DNSdist - A DNS, DoS and abuse-aware loadbalancer that brings out the best possible performance in any DNS deployment.
- OX PowerDNS Recursor - A high-performing, low latency DNS resolver.
- OX PowerDNS Authoritative Server - A versatile authoritative server for hosting domain names.
- OX PowerDNS ZoneControl - A graphical web-based interface for managing domains on the OX PowerDNS Authoritative Server.

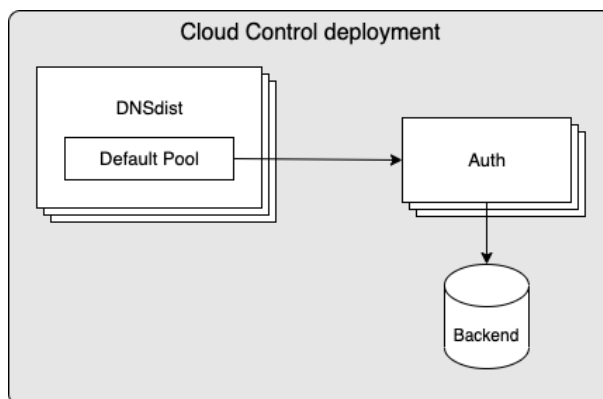
1.1 Simple deployment - Recursor

Cloud Control can be used to roll out a set of Recursor instances, with a set of DNSdists in front. In the below diagram you can see a set of DNSdist instances, with a default pool sending all traffic to a set of Recursor instances:



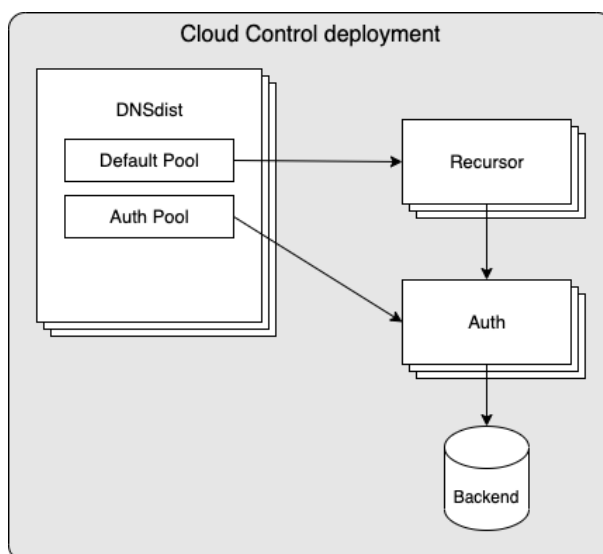
1.2 Simple deployment - Auth

Cloud Control can be used to roll out a set of Auth instances, with a set of DNSdists in front. In the below diagram you can see a set of Auth instances, with a default pool sending all traffic to a set of Auth instances:



1.3 Complex deployment

In a more complex deployment you can deploy both Recursor & Auth instances, having DNSdist using multiple pools to send traffic to the different instances based on the incoming queries/traffic. In the below example you see a setup where both Recursor & Auth are deployed, with DNSdist using rules to send some traffic to Auth, while defaulting to sending queries to Recursor. The Recursor > Auth arrow signifies the use of forward zones, which instructs the Recursor to forward queries for certain zones to Auth.



1.4 Rules & Actions


Deciding which traffic to send to each pool is handled by DNSdist's packet policies, which offers a mechanism to define rules and corresponding actions. In the context of the above diagram, such rules & actions could be:

Rule	Action
'QPS' of requests from the sender has exceeded a certain value	Answer request with 'REFUSED'
'Opcode' of request is 'Notify'	let Auth pool handle the request
'Qtype' of request is 'AXFR'	let Auth pool handle the request

Note: By default, all requests will be handled by the 'Default Pool'

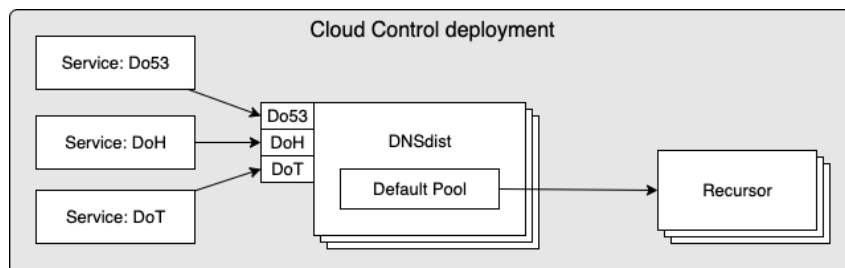
1.5 DNSdist with co-hosted Recursors

In a high load environment, the overhead on Kubernetes network components from the DNSdist to Recursor traffic can potentially become a bottleneck and/or lead to unacceptable latency. For these scenarios it is possible to have 1 or more Recursor instances running within the same Pod as DNSdist. Such a deployment would look as follows:




1.6 DNSdist with DoH and/or DoT listeners

Inbound traffic to DNSdist is supported not only via the standard UDP & TCP over port 53 (Do53), but also via DoH and DoT. When configured, you can have a deployment that looks as follows:



1.7 ZoneControl deployment

Cloud Control can be used to roll out a set of ZoneControl instances and configure the endpoints of Auth instances that it should be able to manage. In the below diagram you can see a set of ZoneControl instances, configured to manage 2 sets of Auth instances, one within the same Cloud Control deployment and another in a separate deployment:



2 Cloud Control on Kubernetes

Cloud Control provides a Helm Chart which allows for the definition & configuration of the following:

- **auth** - Definition of a set of OX PowerDNS Authoritative Server instances and corresponding configuration
- **dnsdist** - Definition of a set of OX PowerDNS DNSdist instances and corresponding configuration
- **recursor** - Definition of a set of OX PowerDNS Recursor instances and corresponding configuration
- **resolver** - Definition of a set of external resolver endpoints
- **ruleset** - Definition of a set of rules which can be applied to DNSdist instances
- **zonecontrol** - Definition of a set of OX PowerDNS ZoneControl instances and corresponding configuration


The following sections discuss each in more detail.

2.1 Auth

For each auth defined in the input to the Helm Chart, objects of the following types (aka kind in Kubernetes terminology) will be created in Kubernetes:

Kind	API Group	Description
Auth	cloudcontrol.powerdns.com	Object which holds configuration of the Auth instances
Deployment	core	Deployment of Auth pods (including ReplicaSet)
Service	core	Service which can be discovered by DNSdist & Recursor agents to direct traffic to the Auth pods

When an auth instance is configured using the Helm Chart, it will deploy the following to Kubernetes:



As the diagram shows an Auth pod will consist of 2 containers + 1 init container:

- **auth-init** - Prepares configuration for Auth.
- **auth** - Container running OX PowerDNS Authoritative Server.
- **agent** - Contains an agent that watches several kinds of objects in Kubernetes within the namespace. If any watched objects are created/updated/removed, the agent will sync any corresponding configuration items to the running Auth instance. The agent is described in detail in the next chapter.

2.1.1 Auth agent

The Auth agent is responsible for keeping the configuration of the running Auth process in sync with the desired configuration. If any configuration changes are needed, the agent will attempt to synchronize them without restarting the Auth process.

Items which are watched by the agent are:


Kind	Purpose
Auth	The object which contains the configuration details for an Auth deployment. If any updates are detected the agent will attempt to update the configuration of Auth without having to restart it.
Pod	The agent watches the pod which it is a part of. Particularly the statuses of each container inside the pod are observed, to ensure the agent can synchronize an Auth instance again if it's container was recycled for any reason.
GeoIP zone-files	The agent watches for changes in the GeoIP zonefiles that can be configured for the GeoIP backend using the <i>domains</i> attribute. If any changes are detected the agent will instruct Auth to reload the zonefiles.

2.2 DNSdist

For each dnsdist defined in the input to the Helm Chart, objects of the following types (kind in Kubernetes) will be created in Kubernetes:

Kind	API Group	Description
DNSDist	cloudcontrol. powerdns.com	Object which holds configuration of the DNSDist instances
Deployment	core	Deployment of DNSDist pods (including ReplicaSet)
Service	core	Service which can be used to direct traffic to the DNSDist pods

When a `dnsdist` instance is configured using the Helm Chart, it will deploy the following to Kubernetes:



As the diagram shows a DNSDist pod will consist of 3 containers + 1 init container:

- **dnsdist-init** - Prepares configuration for dnsdist.
- **dnsdist** - Container running OX PowerDNS DNSDist.
- **rpc-server** - Runs an API that is responsible for handling JSON messages over HTTP from the agent and forwarding them to dnsdist.
- **agent** - Contains an agent that watches several kinds of objects in Kubernetes within the namespace. If any watched objects are created/updated/removed, the agent will sync any corresponding configuration items to the running dnsdist instance. The agent is described in detail in the next chapter.

2.2.1 DNSDist agent

The DNSDist agent is responsible for keeping the configuration of the running DNSDist process in sync with the desired configuration. If any configuration changes are needed, the agent will attempt to synchronize them without restarting the DNSDist process. These configuration changes range from performance parameters defined in the DNSDist object to adjusting server pools according to changes observed in Recursor, Auth & Resolver deployments.

Items which are watched by the agent are:


Kind	Purpose
DNSDist	The object which contains the configuration details for a DNSdist deployment. If any updates are detected the agent will attempt to update the configuration of DNSdist without having to restart it.
Pod	The agent watches the pod which it is a part of. Particularly the statuses of each container inside the pod are observed, to ensure the agent can synchronize a DNSdist instance again if it's container was recycled for any reason.
DNSDistRule	Any rule objects which match the RuleSelector on the DNSDist object are watched and synchronized to the DNSdist process if needed. Any new rules that match the RuleSelector are also applied as soon as they are observed by the agent.
Service & Endpoints	The agent watches for changes in the Endpoints of any Service objects which match the ServiceSelector of the DNSDist object. This allows the agent to discover the servers that should be part of the pool(s) in DNSdist and works for Recursor, Auth & Resolver deployments.

2.3 Recursor

For each recursor defined in the input to the Helm Chart, objects of the following types (aka kind in Kubernetes terminology) will be created in Kubernetes:

Kind	API Group	Description
Recursor	ccloudcontrol.powerdns.com	Object which holds configuration of the Recursor instances
Deployment	core	Deployment of Recursor pods (including ReplicaSet)
Service	core	Service which can be discovered by DNSdist agents to direct traffic to the Recursor pods

When a recursor instance is configured using the Helm Chart, it will deploy the following to Kubernetes:



As the diagram shows a Recursor pod will consist of 2 containers + 1 init container:

- **recursor-init** - Prepares configuration for Recursor.
- **recursor** - Container running OX PowerDNS Recursor.

- **agent** - Contains an agent that watches several kinds of objects in Kubernetes within the namespace. If any watched objects are created/updated/removed, the agent will sync corresponding configuration items to the running Recursor instance. The agent is described in detail in the next chapter.

2.3.1 Recursor agent

The Recursor agent is responsible for keeping the configuration of the running Recursor process in sync with the desired configuration. If any configuration changes are needed, the agent will attempt to synchronize them without restarting the Recursor process.

Items which are watched by the agent are:


Kind	Purpose
Recursor	The object which contains the configuration details for a Recursor deployment. If any updates are detected the agent will attempt to update the configuration of Recursor without having to restart it.
Pod	The agent watches the pod which it is a part of. Particularly the statuses of each container inside the pod are observed, to ensure the agent can synchronize a Recursor instance again if it's container was recycled for any reason.
Service & Endpoints	The agent watches for changes in the Endpoints of any Service objects which match the ServiceSelector of the Recursor object. This allows the agent to discover the endpoints that should be part of the forward zones in Recursor.

2.4 Resolver

For each resolver defined in the input to the Helm Chart, objects of the following types (aka kind in Kubernetes terminology) will be created in Kubernetes:

Kind	API Group	Description
Endpoints	core	Object that holds each IP:port combination defined for the resolver
Service	core	Service which can be discovered by DNSdist & Recursor agents to direct traffic to the resolver's endpoints

When a resolver instance is configured using the Helm Chart, it will deploy the following to Kubernetes:



2.5 Ruleset

For each ruleset defined in the input to the Helm Chart, objects of the following types (aka kind in Kubernetes terminology) will be created in Kubernetes:


Kind	API Group	Description
DNSDistRule	cloudcontrol. powerdns.com	Object which holds configuration of a set of rules which can be discovered by DNSdist agents and applied to DNSdist without restarting

2.6 ZoneControl

For each zonecontrol defined in the input to the Helm Chart, objects of the following types (aka kind in Kubernetes terminology) will be created in Kubernetes:

Kind	API Group	Description
ZoneControl	cloudcontrol. powerdns.com	Object which holds configuration of the ZoneControl instances
Deployment	core	Deployment of ZoneControl pods (including ReplicaSet)
Service	core	Service which can be used to expose ZoneControl instances
Ingress	networking. k8s.io	Ingress which can be used to expose ZoneControl instances outside of the cluster via HTTP(S)

When a zonecontrol instance is configured using the Helm Chart, it will deploy the following to Kubernetes:



As the diagram shows a ZoneControl instance will consist of a ZoneControl deployment with 1 container + 1 init container and a ZoneControl Syncer deployment. The ZoneControl Deployment contains the GUI and can have multiple replicas, while the ZoneControl Syncer deployment has a single replica and is used to synchronise configuration changes to the ZoneControl instances.

- **zonecontrol-init** - Prepares configuration for ZoneControl.
- **zonecontrol** - Container running OX PowerDNS ZoneControl.
- **syncer** - Contains an operator that watches ZoneControl objects in Kubernetes within the namespace. If any watched objects are updated, the syncer will synchronise any corresponding configuration items to the running ZoneControl instances.

2.6.1 ZoneControl Syncer

The ZoneControl Syncer agent is responsible for keeping the configuration of the running ZoneControl processes in sync with the desired configuration. If any configuration changes are needed, the syncer will attempt to synchronize them without restarting the ZoneControl process.

Items which are watched by the syncer are:

Kind	Purpose
ZoneControl	The object which contains the configuration details for a ZoneControl deployment. If any updates are detected the syncer will attempt to update the configuration of ZoneControl without having to restart it.

3 Helm Charts

CloudControl has several Helm Charts available to manage & deploy PowerDNS environments to Kubernetes. The main charts are as follows:

- **powerdns-crds**: Chart to install/upgrade the CloudControl CRDs
- **powerdns**: Chart to install/upgrade CloudControl deployments
- **powerdns-operators**: Chart that allows for installation of optional operators

3.1 Helm Chart: powerdns-crds

This chart is used to deploy & upgrade the CRDs used by PowerDNS CloudControl deployments. Having these CRDs deployed to the cluster is a prerequisite to being able to install an environment using the *powerdns* Helm chart.

Scope of objects: cluster-scoped, requires cluster privileges on *CRD* objects.

3.2 Helm Chart: powerdns

This chart is used to deploy & upgrade the PowerDNS CloudControl deployments.

Scope of objects: namespace-scoped, does not require any cluster privileges.

3.3 Helm Chart: powerdns-operators

This optional chart is used to deploy auxiliary Kubernetes Operators that may be used to easily deploy additional components to support CloudControl PowerDNS deployments. Due to the complexity of persistent storage in a Kubernetes environment we recommend you leverage any existing facilities you may have to provide the services offered by this chart instead of using this chart to deploy them.

Currently contains Operators for:

- **Postgres**: Allows for automated creation of Postgres databases, potentially used by Auth & ZoneControl deployments.

Scope of objects: cluster-scoped & namespace-scoped, requires cluster privileges on *CRD*, *ClusterRole* and *ClusterRoleBinding* objects .

4 Getting Started

4.1 Install Tools

You will need the following software on the machine from which you want to deploy CloudControl:

- Kubectl (Configured for your target Kubernetes cluster)
- Helm v3 (<https://helm.sh/docs/intro/install/>)

4.2 Download Helm Charts

CloudControl Helm Charts are available on the Open-Xchange registry, located at: registry.open-xchange.com.

There are several methods for obtaining Helm Charts using Helm's CLI, in this chapter we are using a method that copies the chart locally to your filesystem prior to using it. Any Helm-supported method will work, but you will need to adjust the commands in this guide accordingly if you wish to utilise a different method.

First step will be to make Helm aware of the CloudControl repository (replace username & password with your OX registry credentials):

```
helm repo add cloudcontrol https://registry.open-xchange.com/chartrepo/cloudcontrol \
--username=REGISTRY_USERNAME_HERE --password=REGISTRY_PASSWORD_HERE
```

Once the repository has been added you can pull the CloudControl Helm Charts. To pull the powerdns Helm Chart and export it to your current working directory use the following commands:

```
# Ensure repo data is up-to-date
helm repo update

# Pull the Helm Charts & unpack
helm pull cloudcontrol/powerdns-crds -d . --version=2.1.0 --untar
helm pull cloudcontrol/powerdns -d . --version=2.1.0 --untar
helm pull cloudcontrol/powerdns-operators -d . --version=2.1.0 --untar
```


4.3 Install/Upgrade CloudControl CRDs

The CloudControl CRDs can be installed or upgraded using the *powerdns-crds* Helm Chart. While the chart only deploys cluster-scoped objects (CRDs), you need to provide a namespace to allow Helm to store the relevant information about this deployment. This ensures you can easily upgrade to a newer version in the future.

To install the CRDs with a Helm release name of 'pdns-crds' stored in a namespace 'pdns-crds':

```
helm install pdns-crds ./powerdns-crds --namespace pdns-crds
```

Note: you can add `--create-namespace` if the namespace does not exist yet and you have privileges to create it

Using `kubectl` you should now be able to see the corresponding Kubernetes objects created:

```
# Kubectl command to show CRD objects (filtered for 'cloudcontrol')
kubectl get crd | grep cloudcontrol

# Kubectl output
dnsdistrules.cloudcontrol.powerdns.com          <timestamp of creation>
zonecontrols.cloudcontrol.powerdns.com         <timestamp of creation>
auths.cloudcontrol.powerdns.com                <timestamp of creation>
dnsdistdynamicrules.cloudcontrol.powerdns.com  <timestamp of creation>
recursors.cloudcontrol.powerdns.com            <timestamp of creation>
dnsdists.cloudcontrol.powerdns.com             <timestamp of creation>
```

Result should be a list of CRDs within the *cloudcontrol.powerdns.com* group as shown above.

To upgrade the CRDs, you can use the *helm upgrade* command. For example:

```
helm upgrade pdns-crds ./powerdns-crds --namespace pdns-crds
```

Note: Since the Helm upgrade command needs to have awareness of the previous install/upgrade, it is crucial to specify the same *release* and *namespace* (both 'pdns-crds' in this example). If you try to upgrade but do not specify the existing *release* and *namespace*, the upgrade of the CRDs will fail (if it does fail, Helm will tell you and the old CRDs will remain untouched)

4.4 Install/Upgrade CloudControl

The CloudControl Helm Chart has a large amount of configurable options, which are detailed in the reference documentation. In the next few chapters the most important parts are discussed.

4.4.1 Registry Credentials

Since the CloudControl images are in a protected repository there is a requirement to configure credentials in the Helm Chart input YAML file. These need to be configured with the following block:

```
registrySecrets:
  registry: registry.open-xchange.com
  username: REGISTRY_USERNAME_HERE
  password: REGISTRY_PASSWORD_HERE
  email: admin@registry.open-xchange.com
```

Make sure the username & password match your credentials for the OX registry.

4.4.2 Cluster Networking

To be able to support Kubernetes clusters with IPv4, IPv6 or dual stack (IPv4 & IPv6) configurations, it is required to ensure the 'ipFamily' configuration in the helm values matches your cluster. The 'ipFamily' section contains the following parameters:

- **ipv4**: Whether or not your cluster has IPv4 enabled (Default: true)
- **ipv6**: Whether or not your cluster has IPv6 enabled (Default: false)
- **families**: Preference of IP families on your cluster, if it is a dualstack cluster

To ensure your deployment is correctly configured, you need to provide one of the 4 possible variations:

IPv4 only (default)

```
# Networking configuration
ipFamily:
  ipv4: true
  ipv6: false
  families:
    - "IPv4"
    - "IPv6"
```

Note: 'families' is ignored in this configuration. It is only used in a dualstack setup.

IPv6 only

```
# Networking configuration
ipFamily:
  ipv4: false
  ipv6: true
  families:
    - "IPv4"
    - "IPv6"
```

Note: 'families' is ignored in this configuration. It is only used in a dualstack setup.

Dualstack - IPv4 primary

If you are running a dualstack cluster, you can check any Pod to see if your cluster has a preference for IPv4 or IPv6. Your pods will have a 'podIP' and 2 values for 'podIPs'. If the 'podIP' is an IPv4 address as shown in the example below, then you are running a cluster with IPv4 as primary:

```
podIP: 172.17.183.4 # IPv4
podIPs:
- ip: 172.17.183.4 # IPv4
- ip: fd43:128b:8658:b73b:3eb7:2e30:8815:3f6 # IPv6
```

Configuration for dualstack with IPv4 primary:

```
# Networking configuration
ipFamily:
  ipv4: true
  ipv6: true
families:
- "IPv4" # IPv4 is primary
- "IPv6"
```

Dualstack - IPv6 primary

If you are running a dualstack cluster, you can check any Pod to see if your cluster has a preference for IPv4 or IPv6. Your pods will have a 'podIP' and 2 values for 'podIPs'. If the 'podIP' is an IPv6 address as shown in the example below, then you are running a cluster with IPv6 as primary:

```
podIP: fd43:128b:8658:b73b:3eb7:2e30:8815:3f6 # IPv6
podIPs:
- ip: fd43:128b:8658:b73b:3eb7:2e30:8815:3f6 # IPv6
- ip: 172.17.183.4 # IPv4
```

Configuration for dualstack with IPv6 primary:

```
# Networking configuration
ipFamily:
  ipv4: true
  ipv6: true
families:
- "IPv6" # IPv6 is primary
- "IPv4"
```

For the remainder of the guide we will assume the cluster is running on the 'IPv4 only' scenario. If your cluster has a different setup please make sure you substitute accordingly.

4.4.3 Deploying Recursor

To deploy a set of Recursor instances, include an entry in the YAML file under the 'recursors' parent, such as:

```
recursors:
  myrecursor:
    replicas: 3
registrySecrets:
  registry: registry.open-xchange.com
  username: REGISTRY_USERNAME_HERE
  password: REGISTRY_PASSWORD_HERE
  email: admin@registry.open-xchange.com
ipFamily:
  ipv4: true
  ipv6: false
families:
  - "IPv4"
  - "IPv6"
```

The above file will create a set of Recursor instances named 'myrecursor' and the Deployment in Kubernetes will have a ReplicaSet with replicas=3. If you save this file as 'values.yaml' in your current working directory you should be able to use the Helm Chart to create the Recursor instances:

```
# The namespace
CC_NAMESPACE=my-namespace
HELM_RELEASE=ccdemo

helm install $HELM_RELEASE ./powerdns --namespace $CC_NAMESPACE --create-namespace \
--values ./values.yaml
```

Note: you can remove --create-namespace if you have an existing namespace to deploy into

Using kubectl you should now be able to see the corresponding Kubernetes objects created:

```
# Kubectl command to show all objects in a namespace
kubectl get all --namespace=$CC_NAMESPACE

# Kubectl output
NAME                                READY   STATUS    RESTARTS   AGE
pod/myrecursor-589559675d-d57jk     1/1    Running   0          3m12s
pod/myrecursor-589559675d-m779s     1/1    Running   0          3m12s
pod/myrecursor-589559675d-xxrvc     1/1    Running   0          3m12s

NAME                                TYPE          CLUSTER-IP   EXTERNAL-IP   PORT(S)    AGE
service/recursor-myrecursor         ClusterIP    None         <none>        5353/TCP   3m12s

NAME                                READY   UP-TO-DATE   AVAILABLE   AGE
deployment.apps/myrecursor          3/3     3            3           3m12s

NAME                                DESIRED   CURRENT   READY   AGE
replicaset.apps/myrecursor-589559675d  3         3         3       3m12
```

Result should be a deployment + replicaset + service + a number of pods equal to the 'replicas'

value from the values.yaml file.

4.4.4 Adding DNSdist

To add a set of DNSdist instances to our deployment, include an entry in the YAML file under the 'dnsdists' parent, such as:

```
dnsdists:
  mydnsdist:
    replicas: 2
    pools:
      default:
        serverGroups:
          - group: myrecursor
        packetcache:
          maxEntries: 200000
  recursors:
    myrecursor:
      replicas: 3
  registrySecrets:
    registry: registry.open-xchange.com
    username: REGISTRY_USERNAME_HERE
    password: REGISTRY_PASSWORD_HERE
    email: admin@registry.open-xchange.com
  ipFamily:
    ipv4: true
    ipv6: false
    families:
      - "IPv4"
      - "IPv6"
```

The above will add a set of DNSdist instances named 'mydnsdist' and the Deployment in Kubernetes will have a ReplicaSet with replicas=2. The 'pools' configuration instructs DNSdist's agent to make sure all instances of 'myrecursor' are added to the default pool in DNSdist. The 'packetcache' with 'maxEntries' configuration ensures the cache for this pool will be able to hold 200000 entries.

Save the values.yaml file and upgrade the environment using the Helm Chart:

```
# The namespace
CC_NAMESPACE=my-namespace

# Helm release name
HELM_RELEASE=ccdemo

helm upgrade $HELM_RELEASE ./powerdns --namespace $CC_NAMESPACE --values=./values.yaml
```

Using kubectl you should now be able to see the corresponding Kubernetes objects created for DNSdist:

```
# Kubectl command to show all objects in a namespace
kubectl get all --namespace=$CC_NAMESPACE

# Kubectl output
NAME                                READY   STATUS    RESTARTS   AGE
```

(continues on next page)

(continued from previous page)

pod/mydnsdist-775cbf55d9-qjtk	3/3	Running	1	15m	
pod/mydnsdist-775cbf55d9-t8fbk	3/3	Running	1	15m	
pod/myrecursor-589559675d-d57jk	1/1	Running	0	27m	
pod/myrecursor-589559675d-m779s	1/1	Running	0	27m	
pod/myrecursor-589559675d-xrvc	1/1	Running	0	27m	
NAME	TYPE	CLUSTER-IP	EXTERNAL-IP	PORT(S)	AGE
service/recursor-myrecursor	ClusterIP	None	<none>	5353/TCP	27m
NAME	READY	UP-TO-DATE	AVAILABLE	AGE	
deployment.apps/mydnsdist	2/2	2	2	15m	
deployment.apps/myrecursor	3/3	3	3	27m	
NAME	DESIRED	CURRENT	READY	AGE	
replicaset.apps/mydnsdist-775cbf55d9	2	2	2	15m	
replicaset.apps/myrecursor-589559675d	3	3	3	27m	

4.4.5 Adding an external Resolver

To add a set of external resolvers to our deployment, include an entry in the YAML file under the 'resolvers' parent, such as:

```
dnsdists:
  mydnsdist:
    replicas: 2
    pools:
      default:
        serverGroups:
          - group: myrecursor
          - group: myresolver
        packetcache:
          maxEntries: 200000
recursors:
  myrecursor:
    replicas: 3
resolvers:
  myresolver:
    ips:
      - 9.9.9.9
      - 149.112.112.112
registrySecrets:
  registry: registry.open-xchange.com
  username: REGISTRY_USERNAME_HERE
  password: REGISTRY_PASSWORD_HERE
  email: admin@registry.open-xchange.com
ipFamily:
  ipv4: true
  ipv6: false
  families:
    - "IPv4"
    - "IPv6"
```

The above will add a Service named 'myresolver' in Kubernetes which will have an Endpoints object containing the IP addresses (in this example the Quad9 IPs). By adding 'myresolver'

to the 'default' pool in DNSdist, traffic will be loadbalanced between the Recursor & resolver endpoints (not a realistic scenario, which will be tackled in the next chapter).

Save the values.yaml file and upgrade the environment using the Helm Chart:

```
# The namespace
CC_NAMESPACE=my-namespace

# Helm release name
HELM_RELEASE=ccdemo

helm upgrade $HELM_RELEASE ./powerdns --namespace $CC_NAMESPACE --values=./values.yaml
```

Using kubectl you should now be able to see the corresponding Kubernetes objects created for resolver (the service object named 'myresolver'):

```
# Kubectl command to show all objects in a namespace
kubectl get all --namespace=$CC_NAMESPACE

# Kubectl output
```

NAME	READY	STATUS	RESTARTS	AGE
pod/mydnsdist-775cbf55d9-qwvrq	3/3	Running	0	22s
pod/mydnsdist-775cbf55d9-swz2w	3/3	Running	0	22s
pod/myrecursor-589559675d-5sqmg	1/1	Running	0	22s
pod/myrecursor-589559675d-cv6bl	1/1	Running	0	22s
pod/myrecursor-589559675d-sptfh	1/1	Running	0	22s

NAME	TYPE	CLUSTER-IP	EXTERNAL-IP	PORT(S)	AGE
service/recursor-myrecursor	ClusterIP	None	<none>	5353/TCP	22s
service/resolver-myresolver	ClusterIP	None	<none>	53/TCP	22s

NAME	READY	UP-TO-DATE	AVAILABLE	AGE
deployment.apps/mydnsdist	2/2	2	2	22s
deployment.apps/myrecursor	3/3	3	3	22s

NAME	DESIRED	CURRENT	READY	AGE
replicaset.apps/mydnsdist-775cbf55d9	2	2	2	22s
replicaset.apps/myrecursor-589559675d	3	3	3	22s

4.4.6 Adding a DNSdist rule

To add more logic to DNSdist instances you can create rules under the 'rulesets' parent and assigning them to DNSdist objects, such as:

```
dnsdists:
  mydnsdist:
    replicas: 2
    pools:
      default:
        serverGroups:
          - group: myrecursor
          - group: myresolver
        packetcache:
          maxEntries: 200000
    rulegroups:
```

(continues on next page)

(continued from previous page)

```

- traffic-filters
recursors:
  myrecursor:
    replicas: 3
resolvers:
  myresolver:
    ips:
      - 9.9.9.9
      - 149.112.112.112
rulesets:
  block-traffic-ruleset:
    group: traffic-filters
    type: DNSDistRule
    priority: 100
    rules:
      - name: Block ANY
        combinator: AND
        selectors:
          - QType: ANY
        action:
          RCode:
            rcode: "REFUSED"
registrySecrets:
  registry: registry.open-xchange.com
  username: REGISTRY_USERNAME_HERE
  password: REGISTRY_PASSWORD_HERE
  email: admin@registry.open-xchange.com
ipFamily:
  ipv4: true
  ipv6: false
  families:
    - "IPv4"
    - "IPv6"

```

The above will add a DNSDistRule object named 'block-traffic-ruleset' in Kubernetes. This rule will select incoming queries with QType='ANY' and send a response 'REFUSED'. This rule is tagged with 'group' = 'traffic-filters', which is also added to the 'mydnsdist' rulegroups list, associating this rule to the DNSdist instances. More details on the specification of rules can be found in the reference guide.

Save the values.yaml file and upgrade the environment using the Helm Chart:

```

# The namespace
CC_NAMESPACE=my-namespace

# Helm release name
HELM_RELEASE=ccdemo

helm upgrade $HELM_RELEASE ./powerdns --namespace $CC_NAMESPACE --values=./values.yaml

```

Using kubectl you should now be able to see the corresponding Kubernetes objects if you specifically request them (since kubectl will not show any custom object types with 'get all'):

```

# Kubectl command to show all DNSDistRule objects in a namespace
kubectl get dnsdistrule --namespace=$CC_NAMESPACE

```

(continues on next page)

(continued from previous page)

```
# Kubectl output
NAME          AGE
block-traffic-ruleset  6s
```

4.4.7 Using DNSdist rules to route traffic

In a previous step we added recursors & resolvers to the default pool, but it would make more sense to have them in separate pools so they can serve different purposes. Rules allow this behaviour to be configured, such as:

```
dnsdists:
  mydnsdist:
    replicas: 2
    pools:
      default:
        serverGroups:
          - group: myrecursor
        packetcache:
          maxEntries: 200000
      external:
        serverGroups:
          - group: myresolver
        packetcache:
          maxEntries: 200000
    rulegroups:
      - traffic-filters
      - traffic-routers
recursors:
  myrecursor:
    replicas: 3
resolvers:
  myresolver:
    ips:
      - 9.9.9.9
      - 149.112.112.112
rulesets:
  route-traffic-ruleset:
    group: traffic-routers
    type: DNSDistRule
    priority: 200
    rules:
      - name: External IPv6 resolution
        combinator: AND
        selectors:
          - QType: AAAA
        action:
          Pool:
            poolname: "external"
  block-traffic-ruleset:
    group: traffic-filters
    type: DNSDistRule
    priority: 100
    rules:
```

(continues on next page)

(continued from previous page)

```

- name: Block ANY
  combinator: AND
  selectors:
    - QType: ANY
  action:
    RCode:
      rcode: "REFUSED"
registrySecrets:
  registry: registry.open-xchange.com
  username: REGISTRY_USERNAME_HERE
  password: REGISTRY_PASSWORD_HERE
  email: admin@registry.open-xchange.com
ipFamily:
  ipv4: true
  ipv6: false
  families:
    - "IPv4"
    - "IPv6"

```

In the above example we moved the 'myresolver' group to a new pool named 'external'. Also, a new ruleset 'route-traffic-ruleset' was added which will match any queries with 'QType' = 'AAAA' and assign the pool named 'external' to handle those queries.

Save the values.yaml file and upgrade the environment using the Helm Chart:

```

# The namespace
CC_NAMESPACE=my-namespace

# Helm release name
HELM_RELEASE=ccdemo

helm upgrade $HELM_RELEASE ./powerdns --namespace $CC_NAMESPACE --values=./values.yaml

```

Using kubectl you should now be able to see the new Kubernetes objects if you specifically request them (since kubectl will not show any custom object types with 'get all'):

```

# Kubectl command to show all DNSDistRule objects in a namespace
kubectl get dnsdistrule --namespace=$CC_NAMESPACE

# Kubectl output
NAME                                AGE
block-traffic-ruleset              33m
route-traffic-ruleset              2s

```

4.4.8 Separating config into multiple files

As you start adding more instances & configuration options to the Helm Chart input file it becomes harder to make sense of the config. A recommended approach to improving this is to make use of Helm's ability to add multiple values files to the arguments of the helm command line. For example:

generic.yaml:

```
registrySecrets:
  registry: registry.open-xchange.com
  username: REGISTRY_USERNAME_HERE
  password: REGISTRY_PASSWORD_HERE
  email: admin@registry.open-xchange.com
ipFamily:
  ipv4: true
  ipv6: false
  families:
    - "IPv4"
    - "IPv6"
```

rulesets.yaml:

```
rulesets:
  block-traffic-ruleset:
    group: traffic-filters
    type: DNSDistRule
    priority: 100
    rules:
      - name: Block ANY
        combinator: AND
        selectors:
          - QType: ANY
        action:
          RCode:
            rcode: "REFUSED"
```

instances.yaml:

```
dnsdists:
  mydnsdist:
    replicas: 2
    pools:
      default:
        serverGroups:
          - group: myrecursor
        packetcache:
          maxEntries: 200000
    rulegroups:
      - traffic-filters
recursors:
  myrecursor:
    replicas: 3
```

You can then run helm as follows:

```
# The namespace
CC_NAMESPACE=my-namespace

# Helm release name
HELM_RELEASE=ccdemo

helm upgrade $HELM_RELEASE ./powerdns --namespace $CC_NAMESPACE \
--values=./generic.yaml --values=./rulesets.yaml --values=./instances.yaml
```

4.4.9 Exposing dnssdist

We now have a set of dnssdist instances running, but to complete the setup we need to make sure we have a method to direct traffic to the dnssdist instances. You can find out the different methods to expose dnssdist instances by reading the chapter 'Exposing dnssdist' in the reference guide.

4.4.10 Deploying ZoneControl

If you have one or more deployments of Auth running, you can deploy ZoneControl to manage the zones and records using a graphical user interface. This can be done by including an entry under the 'zonecontrols' parent.

Since this will require a Postgres database, we either need to have an existing database available for usage, or the extra Helm chart named *powerdns-operators* can be used to provision an Operators that creates Postgres databases for us. In the below example we will make use of the operator approach. To do so, we need to make sure the operator is installed, which can be done as follows:

```
# The release we're working with
CCTAG=2.1.0

# The namespace
CCOPS_NAMESPACE=ccops

# Helm release name
HELM_RELEASE=ccops

# Ensure repo data is up-to-date
helm repo update

# Pull the Helm Chart & unpack
helm pull cloudcontrol/powerdns-operators -d . --version=$CCTAG --untar

# Deploy the operator
helm install $HELM_RELEASE ./powerdns-operators --namespace $CCOPS_NAMESPACE
```

As a result there should be a Postgres Operator running in the 'ccops' namespace. We can then deploy ZoneControl:

generic.yaml:

```
registrySecrets:
  registry: registry.open-xchange.com
  username: REGISTRY_USERNAME_HERE
  password: REGISTRY_PASSWORD_HERE
  email: admin@registry.open-xchange.com
ipFamily:
  ipv4: true
  ipv6: false
families:
  - "IPv4"
  - "IPv6"
```

zonecontrols.yaml:

```
zonecontrols:
  myzonecontrol:
    replicas: 2
  postgres:
    operator: true
  authEndpoints:
    - name: auth1
      url: https://auth1.example.com
      key: "apiKeyForAuth1"
    - name: auth2
      url: https://auth1.example.com
      key: "apiKeyForAuth2"
```

The above example assumes there are 2 deployments of Auth, named 'auth1' and 'auth2', with the Auth API endpoints accessible via the corresponding url and key. For more configuration options you can refer to the reference guide.

You can deploy these as follows:

```
# The namespace
ZC_NAMESPACE=zonecontrol

# Helm release name
HELM_RELEASE=ccdemo

helm install $HELM_RELEASE ./powerdns --namespace $ZC_NAMESPACE \
--values=./generic.yaml --values=./zonecontrols.yaml
```

Note: In the above example we deploy ZoneControl in a dedicated namespace 'zonecontrol'. Whilst not strictly necessary, it is generally advisable to deploy ZoneControl in a dedicated namespace to keep the management & delivery functions of CloudControl separated.

5 Advanced Examples

5.1 DNSdist: DoH

To deploy a set of DNSdist instances with DoH enabled, include a 'doh' configuration node in the dnsdist instance. The example below shows a basic DoH-enabled deployment of a set of DNSdist instances with Recursors:

```
dnsdists:
  mydohdist:
    replicas: 2
    pools:
      default:
        serverGroups:
          - group: myrecursor
        packetcache:
          maxEntries: 200000
    doh:
      - name: mydoh
        certificates:
          - key: |-
              -----BEGIN RSA PRIVATE KEY-----
              << CONTENTS OF PRIVATE KEY HERE>>
              -----END RSA PRIVATE KEY-----
            cert: |-
              -----BEGIN CERTIFICATE-----
              << CONTENTS OF CERTIFICATE HERE>>
              -----END CERTIFICATE-----
              -----BEGIN CERTIFICATE-----
              << CONTENTS OF ANY INTERMEDIATE CERTIFICATE(S) HERE>>
              -----END CERTIFICATE-----
recursors:
  myrecursor:
    replicas: 2
registrySecrets:
  registry: registry.open-xchange.com
  username: REGISTRY_USERNAME_HERE
  password: REGISTRY_PASSWORD_HERE
  email: admin@registry.open-xchange.com
ipFamily:
  ipv4: true
  ipv6: false
  families:
    - "IPv4"
    - "IPv6"
```

Note: Make sure to replace the contents of the 'key' and 'cert' with the data of a valid pair.

The above will result in a DNSdist deployment with the regular 'dnsdist-mydohdist' Service created, plus an additional Service named 'dnsdist-mydohdist-doh-mydoh'. This additional Service will have (by default) an inbound listener for traffic over port '443'.

You can refer to the 'Reference' guide for all available options to configure DoH. Options available include the configuration of STEK tickets (enabled & rotated by default) and loading certificates from pre-existing TLS Secrets to leverage a certificate manager such as certmanager.

5.2 DNSdist: DoT

To deploy a set of DNSdist instances with DoT enabled, include a 'dot' configuration node in the dnsdist instance. The example below shows a basic DoT-enabled deployment of a set of DNSdist instances with Recursors:

```
dnsdists:
  mydotdist:
    replicas: 2
    pools:
      default:
        serverGroups:
          - group: myrecursor
        packetcache:
          maxEntries: 200000
  dot:
    - name: mydot
      certificates:
        - key: |-
            -----BEGIN RSA PRIVATE KEY-----
            << CONTENTS OF PRIVATE KEY HERE>>
            -----END RSA PRIVATE KEY-----
          cert: |-
            -----BEGIN CERTIFICATE-----
            << CONTENTS OF CERTIFICATE HERE>>
            -----END CERTIFICATE-----
            -----BEGIN CERTIFICATE-----
            << CONTENTS OF ANY INTERMEDIATE CERTIFICATE(S) HERE>>
            -----END CERTIFICATE-----
recursors:
  myrecursor:
    replicas: 2
registrySecrets:
  registry: registry.open-xchange.com
  username: REGISTRY_USERNAME_HERE
  password: REGISTRY_PASSWORD_HERE
  email: admin@registry.open-xchange.com
ipFamily:
  ipv4: true
  ipv6: false
  families:
    - "IPv4"
    - "IPv6"
```

Note: Make sure to replace the contents of the 'key' and 'cert' with the data of a valid pair.

The above will result in a DNSdist deployment with the regular 'dnsdist-mydotdist' Service created, plus an additional Service named 'dnsdist-mydotdist-dot-mydot'. This additional Service

will have (by default) an inbound listener for traffic over port '853'.

You can refer to the 'Reference' guide for all available options to configure DoT. Options available include the configuration of STEK tickets (enabled & rotated by default) and loading certificates from pre-existing TLS Secrets to leverage a certificate manager such as certmanager.

5.3 DNSdist: Co-hosted Recursor

To deploy a set of DNSdist instances with co-hosted Recursor instances, include a 'recursor' configuration node in the dnsdist instance. The example below shows an example of DNSdist instances with co-hosted Recursors:

```
dnsdists:
  mydnsdist:
    replicas: 2
    pools:
      default:
        packetcache:
          maxEntries: 200000
    recursor:
      replicas: 2
registrySecrets:
  registry: registry.open-xchange.com
  username: REGISTRY_USERNAME_HERE
  password: REGISTRY_PASSWORD_HERE
  email: admin@registry.open-xchange.com
ipFamily:
  ipv4: true
  ipv6: false
families:
  - "IPv4"
  - "IPv6"
```

The above will result in a DNSdist deployment where each DNSdist Pod also contains 2 Recursor containers (+ a Recursor agent container to keep the Recursors synchronised)

6 Troubleshooting

6.1 Accessing DNSdist console

DNSdist offers a commandline console which allows for debugging of issues and retrieving statistics. In Cloud Control deployments this is enabled by default and can be accessed via kubectl's exec command. This chapter will show how to gain access to the console and a few sample commands. For full documentation on the DNSdist console you can refer to: [DNSdist reference guide](#)

Note: While DNSdist's console exposes methods to modify a running instance we highly encourage users NOT to do this. Any change made to a running instance using the console will not persist and will not be synchronized to other DNSdist instances.

The following command can be used to gain access to the console:

```
# Pod name (make sure to replace with an existing DNSdist pod's name)
POD=mydnsdist-775cbf55d9-qjtk5

# The namespace
CC_NAMESPACE=my-namespace

# Kubectl command to access the DNSdist console
kubectl exec -it $POD --namespace=$CC_NAMESPACE -c dnsdist -- dnsdist -c \
--config=/config/dnsdist.conf
```

You should then be presented with a console session as follows:

```
* dnsdist-state loaded
* Control socket set to 127.0.0.1:5199 with provided key
>
```

To see the status of the recursor and/or resolver instances that DNSdist will send queries to use showServers():

```
> showServers()
# Name Address State Qps Ord Wt Queries Pools
0 Endpoints/my-namespa 10.244.1.7:5353 up 0.0 1 1 546
1 Endpoints/my-namespa 10.244.1.8:5353 up 0.0 1 1 0
2 Endpoints/my-namespa 10.244.1.9:5353 up 0.0 1 1 0
3 Endpoints/my-namespa 149.112.112.112:53 up 0.0 1 1 0 external
4 Endpoints/my-namespa 9.9.9.9:53 up 0.0 1 1 0 external
All 0.0 546
```

Show the pools using showPools():

```
> showPools()
Name          Cache      ServerPolicy      Servers
external      leastOutstanding  leastOutstanding  10.244.1.7:5353, 10.244.1.8:5353, 10.244.1.9:5353
external      leastOutstanding  leastOutstanding  149.112.112.112:53, 9.9.9.9:53
```

List all rules with showRules():

```
> showRules()
#   Name      Matches   Rule              Action
0   Name      0         qtype==ANY       set rcode 5
1   Name      0         qtype==AAAA      to pool external
```

6.2 Pod Events

Cloud Control pods, primarily DNSdist, emit events to indicate potential problematic behaviour and provide tracability into the synchronisation processes.

There are many ways to list events in a namespace, for a pod, etc.. In the below example we'll use kubectl's get event to show the events for a specific pod, but in a production setting we recommend capturing these in your logging/monitoring infrastructure.

```
# Pod name (make sure to replace with an existing DNSdist pod's name)
POD=mydnsdist-775cbf55d9-qjtk

# The namespace
CC_NAMESPACE=my-namespace

# Kubectl command to list recent events emitted by a pod in a given namespace
kubectl get event --namespace=$CC_NAMESPACE --field-selector involvedObject.name=$POD
```

Examples of events generated by DNSdist pods (reformatted to fit):

```
# Event emitted by agent when a rule is updated
Type: Normal
Reason: DNSDistRuleUpdated
Object: pod/mydnsdist-775cbf55d9-gvjwk
Message: DNSDistRule 'my-namespace/block-traffic-ruleset' has been synchronised

# Event emitted by agent when a recursor/resolver endpoint changes
Type: Normal
Reason: EndpointsUpdated
Object: pod/mydnsdist-775cbf55d9-gvjwk
Message: Endpoints 'my-namespace/recursor-myrecursor' has been synchronised

# Event emitted by Kubernetes when a readiness probe fails
Type: Warning
Reason: Unhealthy
Object: pod/mydnsdist-775cbf55d9-gvjwk
Message: Readiness probe failed: HTTP probe failed with statuscode: 500
```