

---

# **cert-manager Documentation**

**Jetstack Ltd**

**Apr 12, 2019**



---

## Contents:

---

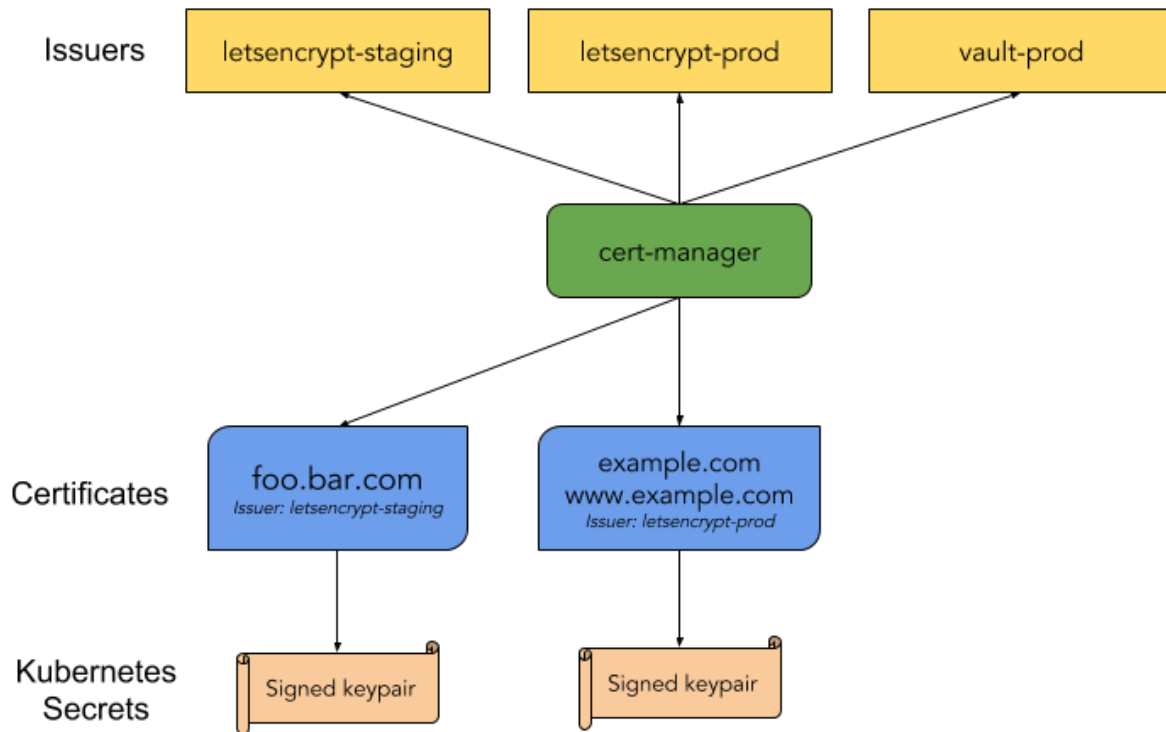
<b>1</b>	<b>Get started</b>	<b>3</b>
1.1	Installing cert-manager . . . . .	3
1.2	Webhook component . . . . .	7
1.3	Troubleshooting installation . . . . .	10
<b>2</b>	<b>Tutorials</b>	<b>13</b>
2.1	ACME Issuer Tutorials . . . . .	13
<b>3</b>	<b>Tasks</b>	<b>39</b>
3.1	Setting up Issuers . . . . .	39
3.2	Issuing Certificates . . . . .	50
3.3	ACME specific tasks . . . . .	53
3.4	Backing up and restoring . . . . .	67
3.5	Upgrading cert-manager . . . . .	67
<b>4</b>	<b>Reference documentation</b>	<b>75</b>
4.1	Certificates . . . . .	75
4.2	Orders . . . . .	77
4.3	Challenges . . . . .	78
4.4	Issuers . . . . .	80
4.5	ClusterIssuers . . . . .	81
4.6	API documentation . . . . .	82
<b>5</b>	<b>Development documentation</b>	<b>83</b>
5.1	Develop with minikube . . . . .	83
5.2	Running end-to-end tests . . . . .	85
5.3	Contributing DNS01 providers . . . . .	85
5.4	DCO Sign off . . . . .	86
5.5	Release process . . . . .	87
5.6	Generating Documentation . . . . .	89



cert-manager is a native [Kubernetes](#) certificate management controller. It can help with issuing certificates from a variety of sources, such as [Let's Encrypt](#), [HashiCorp Vault](#), **'Venafi'**, a simple signing keypair, or self signed.

It will ensure certificates are valid and up to date, and attempt to renew certificates at a configured time before expiry.

It is loosely based upon the work of [kube-lego](#) and has borrowed some wisdom from other similar projects e.g. [kube-cert-manager](#).



This is the full technical documentation for the project, and should be used as a source of references when seeking help with the project.



The guides in this section will explain how to install and set up cert-manager.

If you run into issues deploying, upgrading or running cert-manager please check the *troubleshooting* document.

## 1.1 Installing cert-manager

cert-manager runs within your Kubernetes cluster as a series of deployment resources. It utilises `CustomResourceDefinitions` to configure Certificate Authorities and request certificates.

It is deployed using regular YAML manifests, like any other applications on Kubernetes.

Once cert-manager has been deployed, you must configure Issuer or ClusterIssuer resources which represent certificate authorities. More information on configuring different Issuer types can be found in the *respective setup guides*.

### 1.1.1 Installing with regular manifests

In order to install cert-manager, we must first create a namespace to run it within. This guide will install cert-manager into the `cert-manager` namespace. It is possible to run cert-manager in a different namespace, although you will need to make modifications to the deployment manifests.

```
# Create a namespace to run cert-manager in
kubectl create namespace cert-manager
```

As part of the installation, cert-manager also deploys a `ValidatingWebhookConfiguration` resource in order to validate that the Issuer, ClusterIssuer and Certificate resources we will create after installation are valid.

In order to deploy the `ValidatingWebhookConfiguration`, cert-manager creates a number of ‘internal’ Issuer and Certificate resources in its own namespace.

This creates a chicken-and-egg problem, where cert-manager requires the webhook in order to create the resources, and the webhook requires cert-manager in order to run.

We avoid this problem by disabling resource validation on the namespace that cert-manager runs in:

```
# Disable resource validation on the cert-manager namespace
kubectl label namespace cert-manager certmanager.k8s.io/disable-validation=true
```

You can read more about the webhook on the [webhook document](#).

We can now go ahead and install cert-manager. All resources (the CustomResourceDefinitions, cert-manager, and the webhook component) are included in a single YAML manifest file:

```
# Install the CustomResourceDefinitions and cert-manager itself
kubectl apply -f https://raw.githubusercontent.com/jetstack/cert-manager/release-0.7/
↳deploy/manifests/cert-manager.yaml
```

---

**Note:** If you are running kubectl v1.12 or below, you will need to add the `--validate=false` flag to your `kubectl apply` command above else you will receive a validation error relating to the `caBundle` field of the `ValidatingWebhookConfiguration` resource. This issue is resolved in Kubernetes 1.13 onwards. More details can be found in [kubernetes/kubernetes#69590](#).

---

**Note:** When running on GKE (Google Kubernetes Engine), you may encounter a ‘permission denied’ error when creating some of these resources. This is a nuance of the way GKE handles RBAC and IAM permissions, and as such you should ‘elevate’ your own privileges to that of a ‘cluster-admin’ **before** running the above command. If you have already run the above command, you should run them again after elevating your permissions:

```
kubectl create clusterrolebinding cluster-admin-binding \
  --clusterrole=cluster-admin \
  --user=$(gcloud config get-value core/account)
```

---

### 1.1.2 Installing with Helm

As an alternative to the YAML manifests referenced above, we also provide an official Helm chart for installing cert-manager.

#### Pre-requisites

- Helm and Tiller installed (or alternatively, use [Tillerless Helm v2](#))
- cluster-admin privileges bound to the Tiller pod

#### Foreword

Before deploying cert-manager with Helm, you must ensure [Tiller](#) is up and running in your cluster. Tiller is the server side component to Helm.

Your cluster administrator may have already setup and configured Helm for you, in which case you can skip this step.

Full documentation on installing Helm can be found in the [Installing helm docs](#).

If your cluster has RBAC (Role Based Access Control) enabled (default in GKE v1.7+), you will need to take special care when deploying Tiller, to ensure Tiller has permission to create resources as a cluster administrator. More information on deploying Helm with RBAC can be found in the [Helm RBAC docs](#).



## Steps

In order to install the Helm chart, you must run:

```
# Install the CustomResourceDefinition resources separately
kubectl apply -f https://raw.githubusercontent.com/jetstack/cert-manager/release-0.7/
↳deploy/manifests/00-crds.yaml

# Create the namespace for cert-manager
kubectl create namespace cert-manager

# Label the cert-manager namespace to disable resource validation
kubectl label namespace cert-manager certmanager.k8s.io/disable-validation=true

# Add the Jetstack Helm repository
helm repo add jetstack https://charts.jetstack.io

# Update your local Helm chart repository cache
helm repo update

# Install the cert-manager Helm chart
helm install \
  --name cert-manager \
  --namespace cert-manager \
  --version v0.7.0 \
  jetstack/cert-manager
```

The default cert-manager configuration is good for the majority of users, but a full list of the available options can be found in the [Helm chart README](#).

### 1.1.3 Verifying the installation

Once you've installed cert-manager, you can verify it is deployed correctly by checking the cert-manager namespace for running pods:

```
kubectl get pods --namespace cert-manager
```

NAME	READY	STATUS	RESTARTS	AGE
cert-manager-5c6866597-zw7kh	1/1	Running	0	2m
webhook-78fb756679-9bsmf	1/1	Running	0	2m
webhook-ca-sync-1543708620-n82gj	0/1	Completed	0	1m

You should see both the cert-manager and webhook component in a Running state, and the ca-sync pod is Completed. If the webhook has not Completed but the cert-manager pod has recently started, wait a few minutes for the ca-sync pod to be retried. If you experience problems, please check the [troubleshooting guide](#).

The following steps will confirm that cert-manager is set up correctly and able to issue basic certificate types:

```
# Create a ClusterIssuer to test the webhook works okay
cat <<EOF > test-resources.yaml
apiVersion: v1
kind: Namespace
metadata:
  name: cert-manager-test
---
apiVersion: certmanager.k8s.io/v1alpha1
```

(continues on next page)

```

kind: Issuer
metadata:
  name: test-selfsigned
  namespace: cert-manager-test
spec:
  selfSigned: {}
---
apiVersion: certmanager.k8s.io/v1alpha1
kind: Certificate
metadata:
  name: selfsigned-cert
  namespace: cert-manager-test
spec:
  commonName: example.com
  secretName: selfsigned-cert-tls
  issuerRef:
    name: test-selfsigned
EOF

# Create the test resources
kubectl apply -f test-resources.yaml

# Check the status of the newly created certificate
# You may need to wait a few seconds before cert-manager processes the
# certificate request
kubectl describe certificate -n cert-manager-test
...
Spec:
  Common Name:  example.com
  Issuer Ref:
    Name:      test-selfsigned
  Secret Name: selfsigned-cert-tls
Status:
  Conditions:
    Last Transition Time:  2019-01-29T17:34:30Z
    Message:              Certificate is up to date and has not expired
    Reason:               Ready
    Status:               True
    Type:                 Ready
  Not After:             2019-04-29T17:34:29Z
Events:
  Type      Reason      Age   From          Message
  ----      -
  Normal    CertIssued  4s    cert-manager  Certificate issued successfully

# Clean up the test resources
kubectl delete -f test-resources.yaml

```

If all the above steps have completed without error, you are good to go!

If you experience problems, please check the [troubleshooting guide](#).

## 1.1.4 Configuring your first Issuer

Before you can begin issuing certificates, you must configure at least one Issuer or ClusterIssuer resource in your cluster.

You should read the [Setting up Issuers](#) guide to learn how to configure cert-manager to issue certificates from one of the supported backends.

## 1.1.5 Alternative installation methods

### kubeprod

[Bitnami Kubernetes Production Runtime](#) (BKPR, kubeprod) is a curated collection of the services you would need to deploy on top of your Kubernetes cluster to enable logging, monitoring, certificate management, automatic discovery of Kubernetes resources via public DNS servers and other common infrastructure needs.

It depends on cert-manager for certificate management, and it is [regularly tested](#) so the components are known to work together for GKE and AKS clusters (EKS to be added soon). For its ingress stack it creates a DNS entry in the configured DNS zone and requests a TLS certificate from the Let's Encrypt staging server.

BKPR can be deployed using the `kubeprod install` command, which will deploy cert-manager as part of it. Details available in the [BKPR installation guide](#).

## 1.1.6 Debugging installation issues

If you have any issues with your installation, please refer to the [troubleshooting guide](#).

## 1.2 Webhook component

In order to provide advanced resource validation, cert-manager includes a [ValidatingWebhookConfiguration](#) resource which is deployed into the cluster.

This allows cert-manager to validate that Issuer, ClusterIssuer and Certificate resources that are submitted to the apiserver are syntactically valid, and catch issues with your resources early on.

If you disable the webhook component, cert-manager will still perform the same resource validation however it will not reject 'create' events when the resources are submitted to the apiserver if they are invalid. This means it may be possible for a user to submit a resource that renders the controller inoperable. For this reason, it is strongly advised to keep the webhook **enabled**.

---

**Note:** This feature requires Kubernetes v1.9 or greater.

---

### 1.2.1 How it works

This section walks through how the resource validation webhook is configured and explains the process required for it to provision.

The webhook is a [ValidatingWebhookConfiguration](#) resource combined with an extra pod that is deployed alongside the cert-manager-controller.

The [ValidatingWebhookConfiguration](#) instructs the Kubernetes apiserver to POST the contents of any Create or Update operations performed on cert-manager resource types in order to validate that they are setting valid configurations.

This allows us to ensure mis-configurations are caught early on and communicated to you.

In order for this to work, the webhook requires a TLS certificate that the apiserver is configured to trust.

The cert-manager deployment manifests define two Issuer resources, and two Certificate resources:

- issuer/cert-manager-webhook-selfsign - A self signing Issuer that is used to issue a self signed root CA certificate.
- certificate/cert-manager-webhook-ca - A self-signed root CA certificate which is used to sign certificates for the webhook pod.
- issuer/cert-manager-webhook-ca - A CA Issuer that is used to issue certificates used by the webhook pod to serve with.
- certificate/cert-manager-webhook-webhook-tls - A TLS certificate issued by the root CA above, served by the webhook.

You can check the status of these resources to ensure they're functioning correctly by running:

```
kubectl get issuer --namespace cert-manager
NAME                                AGE
cert-manager-webhook-ca             10m
cert-manager-webhook-selfsign       10m

kubectl get certificate -o wide --namespace cert-manager
NAME                                READY  SECRET  ISSUER  AGE
↔ STATUS
cert-manager-webhook-ca             True   cert-manager-webhook-ca  cert-
↔manager-webhook-selfsign          Certificate is up to date and has not expired  10m
cert-manager-webhook-webhook-tls    True   cert-manager-webhook-webhook-tls  cert-
↔manager-webhook-ca                Certificate is up to date and has not expired  10m
```

If the certificates or issuer are not Ready or you cannot see them, you should check the [troubleshooting](#) guide for help.

---

**Note:** If you are running Kubernetes v1.10 or earlier, you may need to run `kubectl describe` instead of `kubectl get` as the 'additionalPrinterColumns' functionality only moved to beta in v1.11.

---

### cainjector

The cert-manager CA injector is responsible for injecting the two CA bundles above into the webhook's Validating-WebhookConfiguration and APIService resource in order to allow the Kubernetes apiserver to 'trust' the webhook apiserver.

This component is configured using the `certmanager.k8s.io/inject-apiserver-ca: "true"` and `certmanager.k8s.io/inject-apiserver-ca: "true"` annotations on the APIService and Validating-WebhookConfiguration resources.

It copies across the CA defined in the 'cert-manager-webhook-ca' Secret generated above to the `caBundle` field on the APIService resource. It also sets the webhook's `clientConfig.caBundle` field on the `cert-manager-webhook` ValidatingWebhookConfiguration resource to that of your Kubernetes API server in order to support Kubernetes versions earlier than v1.11.

### Known issues

This section contains known issues with the webhook component.

If you're having problems, or receiving errors when creating cert-manager resources, please read through this section for help.

## Disabling validation on the cert-manager namespace

If you've installed cert-manager with custom manifests, or have performed an upgrade from an earlier version, it's important to make sure that the namespace that the webhook is running in has an additional label applied to it in order to disable resource validation on the namespace that the webhook runs in.

If this step is not completed, cert-manager will not be able to provision certificates for the webhook correctly, causing a chicken-egg situation.

To apply the label, run:

```
kubectl label namespace cert-manager certmanager.k8s.io/disable-validation=true
```

You may need to wait a little while before cert-manager retries issuing the certificates if they have been failing for a while due to cert-manager's built in back-offs.

## Running on private GKE clusters

When Google configure the control plane for private clusters, they automatically configure VPC peering between your Kubernetes cluster's network and a separate Google managed project.

In order to restrict what Google are able to access within your cluster, the firewall rules configured restrict access to your Kubernetes pods.

This means that in order to use the webhook component with a GKE private cluster, you must configure an additional firewall rule to allow the GKE control plane access to your webhook pod.

You can read more information on how to add firewall rules for the GKE control plane nodes in the [GKE docs](#).

Alternatively, you can read how to *disable the webhook component* below.

---

**Todo:** add an example command for how to do this here & explain any security implications

---

## 1.2.2 Disable the webhook component

If you are having issues with the webhook and cannot use it at this time, you can optionally disable the webhook altogether.

Doing this may expose your cluster to mis-configuration problems that in some cases could cause cert-manager to stop working altogether (i.e. if invalid types are set for fields on cert-manager resources).

How you disable the webhook depends on your deployment method.

### With Helm

The Helm chart exposes an option that can be used to disable the webhook.

To do so with an existing installation, you can run:

```
helm upgrade cert-manager \
  --reuse-values \
  --set webhook.enabled=false
```

If you have not installed cert-manager yet, you can add the `--set webhook.enabled=false` to the `helm install` command used to install cert-manager.

### With static manifests

Because we cannot specify options when installing the static manifests to conditionally disable different components, we also ship a copy of the deployment files that do not include the webhook.

Instead of installing with `cert-manager.yaml` file, you should instead use the `cert-manager-no-webhook.yaml` file located in the `deploy` directory.

This is a destructive operation, as it will remove the CustomResourceDefinition resources, causing your configured Issuers, Certificates etc to be deleted.

You should first *backup your configuration* before running the following commands.

To re-install cert-manager without the webhook, run:

```
kubectl delete -f https://raw.githubusercontent.com/jetstack/cert-manager/release-0.7/
↳deploy/manifests/cert-manager.yaml

kubectl apply -f https://raw.githubusercontent.com/jetstack/cert-manager/release-0.7/
↳deploy/manifests/cert-manager-no-webhook.yaml
```

Once you have re-installed cert-manager, you should then *restore your configuration*.

## 1.3 Troubleshooting installation

### 1.3.1 Internal error occurred: failed calling admission webhook ... the server is currently unable to handle the request

When installing or upgrading cert-manager, you may run into issues when going through the Validation Steps in the install guide which relate to the admission webhook.

If you see an error like the above, this guide will talk you through a few checks that can pick up common installation problems.

#### 1. Check the namespace cert-manager is running in

As described in the *Webhook component* documentation, the webhook component requires TLS certificates in order to start and communicate securely with the Kubernetes API server.

In order for cert-manager to be able to issue certificates for the webhook before it has started, we must **disable** resource validation on the namespace that cert-manager is running in.

Assuming you have deployed into the `cert-manager` namespace, run the following command to verify that your cert-manager namespace has the necessary label:

```
kubectl get namespace

Name:          cert-manager
Labels:        certmanager.k8s.io/disable-validation=true
...
```

If you cannot see the `certmanager.k8s.io/disable-validation=true` label on your namespace, you should add it with:

```
kubectl label namespace cert-manager certmanager.k8s.io/disable-validation=true
```

Please continue reading this guide once you have added the label.

## 2. Verify that the webhook Issuer and Certificate resources exist

If you had any issues upgrading, especially if you install cert-manager using Helm, you may run into an issue where either:

- the CustomResourceDefinition resources do not exist
- the webhook's Issuer and Certificate resources do not exist

We can first check for the existence of the CustomResourceDefinition resources:

```
kubectl get crd | grep certmanager
```

NAME	CREATED AT
certificates.certmanager.k8s.io	2018-08-17T20:12:26Z
challenges.certmanager.k8s.io	2018-08-02T15:33:02Z
clusterissuers.certmanager.k8s.io	2018-08-17T20:12:26Z
issuers.certmanager.k8s.io	2018-08-17T20:12:26Z
orders.certmanager.k8s.io	2018-08-02T14:40:11Z

We should then also check for that the webhook's Issuer and Certificate resources exist and have been issued correctly:

```
kubectl get issuer,certificate
```

NAME	AGE
issuer.certmanager.k8s.io/cert-manager-webhook-ca	22d
issuer.certmanager.k8s.io/cert-manager-webhook-selfsign	22d

NAME	AGE	READY	SECRET
↔ certificate.certmanager.k8s.io/cert-manager-webhook-ca		True	cert-
↔ manager-webhook-ca	22d		
↔ certificate.certmanager.k8s.io/cert-manager-webhook-webhook-tls		True	cert-
↔ manager-webhook-webhook-tls	22d		

If you do not see the CustomResourceDefinitions installed, or cannot see the webhook's Issuer and Certificate resources, please go back to the install guide and ensure you've followed every step closely.

Take particular care to install the CRD manifest **before** installing cert-manager itself.

## 3. Verify all cert-manager pods are running successfully

You can verify that cert-manager has managed to start successfully by checking the state of the pods that have been deployed:

```
# Get all pods, including Completed and Errored pods
kubectl get pods --show-all --namespace cert-manager
```

NAME	READY	STATUS	RESTARTS	AGE
cert-manager-7cbdc48784-rpgnt	1/1	Running	0	3m
cert-manager-webhook-5b5dd6999-kst4x	1/1	Running	0	3m
cert-manager-cainjector-3ba5cd2bcd-de332x	1/1	Running	0	3m

If the 'webhook' pod (2nd line) is in a ContainerCreating state, it may still be waiting for the Secret in step 2 to be mounted into the pod.

Provided the Secret resource **does** now exist, Waiting a few minutes, or deleting the pod and allowing it to be recreated should get things moving again.

---

**Note:** Check if the Secret exists by running:

```
kubectl get secret cert-manager-webhook-webhook-tls
```

---



This section contains guides that help you get started using cert-manager for more specific use cases. For more information on performing individual tasks, read the *tasks section*.

## 2.1 ACME Issuer Tutorials

This sections contains tutorials relating to the ACME issuer.

### 2.1.1 Quick-Start using Cert-Manager with NGINX Ingress

#### Step 0 - Install Helm Client

**Skip this section if you have helm installed.**

The easiest way to install *cert-manager* is to use [Helm](#), a templating and deployment tool for Kubernetes resources. First, ensure the Helm client is installed following the [Helm installation instructions](#).

For example, on macOS:

```
$ brew install kubernetes-helm
```

#### Step 1 - Installer Tiller

**Skip this section if you have Tiller set-up.**

Tiller is Helm's server-side component, which the `helm` client uses to deploy resources.

Deploying resources is a privileged operation; in the general case requiring arbitrary privileges. With this example, we give Tiller complete control of the cluster. View the documentation on [securing helm](#) for details on setting up appropriate permissions for your environment.

Create the a ServiceAccount for tiller:

```
$ kubectl create serviceaccount tiller --namespace=kube-system
serviceaccount "tiller" created
```

Grant the tiller service account cluster admin privileges:

```
$ kubectl create clusterrolebinding tiller-admin --serviceaccount=kube-system:tiller -
↳-clusterrole=cluster-admin
clusterrolebinding.rbac.authorization.k8s.io "tiller-admin" created
```

Install tiller with the tiller service account:

```
$ helm init --service-account=tiller
$HELM_HOME has been configured at /Users/myaccount/.helm.

Tiller (the Helm server-side component) has been installed into your Kubernetes_
↳Cluster.

Please note: by default, Tiller is deployed with an insecure 'allow unauthenticated_
↳users' policy.
To prevent this, run `helm init` with the --tiller-tls-verify flag.
For more information on securing your installation see: https://docs.helm.sh/using_
↳helm/#securing-your-helm-installation
Happy Helming!
```

Update the helm repository with the latest charts:

```
$ helm repo update
Hang tight while we grab the latest from your chart repositories...
...Skip local chart repository
...Successfully got an update from the "stable" chart repository
...Successfully got an update from the "coreos" chart repository
Update Complete. Happy Helming!
```

## Step 2 - Deploy the NGINX Ingress Controller

A [kubernetes ingress controller](#) is designed to be the access point for HTTP and HTTPS traffic to the software running within your cluster. The nginx-ingress controller does this by providing an HTTP proxy service supported by your cloud provider's load balancer.

You can get more details about nginx-ingress and how it works from the [documentation for nginx-ingress](#).

Use helm to install an Nginx Ingress controller:

```
$ helm install stable/nginx-ingress --name quickstart

NAME: quickstart
LAST DEPLOYED: Sat Nov 10 10:25:06 2018
NAMESPACE: default
STATUS: DEPLOYED

RESOURCES:
==> v1/ConfigMap
NAME                                AGE
quickstart-nginx-ingress-controller 0s
```

(continues on next page)

(continued from previous page)

```

==> v1beta1/ClusterRole
quickstart-nginx-ingress 0s

==> v1beta1/Deployment
quickstart-nginx-ingress-controller 0s
quickstart-nginx-ingress-default-backend 0s

==> v1/Pod(related)

NAME                                READY  STATUS
↔RESTARTS  AGE
quickstart-nginx-ingress-controller-6cfc45747-wcxrg 0/1    ContainerCreating 0
↔          0s
quickstart-nginx-ingress-default-backend-bf9db5c67-dkg4l 0/1    ContainerCreating 0
↔          0s

==> v1/ServiceAccount

NAME                                AGE
quickstart-nginx-ingress 0s

==> v1beta1/ClusterRoleBinding
quickstart-nginx-ingress 0s

==> v1beta1/Role
quickstart-nginx-ingress 0s

==> v1beta1/RoleBinding
quickstart-nginx-ingress 0s

==> v1/Service
quickstart-nginx-ingress-controller 0s
quickstart-nginx-ingress-default-backend 0s

NOTES:
The nginx-ingress controller has been installed.
It may take a few minutes for the LoadBalancer IP to be available.
You can watch the status by running 'kubectl --namespace default get services -o wide
↔-w quickstart-nginx-ingress-controller'

An example Ingress that makes use of the controller:

  apiVersion: extensions/v1beta1
  kind: Ingress
  metadata:
    annotations:
      kubernetes.io/ingress.class: nginx
    name: example
    namespace: foo
  spec:
    rules:
      - host: www.example.com
        http:
          paths:
            - backend:
                serviceName: exampleService

```

(continues on next page)

(continued from previous page)

```

        servicePort: 80
        path: /
        # This section is only required if TLS is to be enabled for the Ingress
        tls:
          - hosts:
              - www.example.com
            secretName: example-tls

```

If TLS is enabled **for** the Ingress, a Secret containing the certificate and key must also be provided:

```

apiVersion: v1
kind: Secret
metadata:
  name: example-tls
  namespace: foo
data:
  tls.crt: <base64 encoded cert>
  tls.key: <base64 encoded key>
type: kubernetes.io/tls

```

It can take a minute or two for the cloud provider to provide and link a public IP address. When it is complete, you can see the external IP address using the `kubectl` command:

```

$ kubectl get svc

```

NAME	PORT(S)	AGE	TYPE	CLUSTER-IP	EXTERNAL-IP
kubernetes	443/TCP	23m	ClusterIP	10.63.240.1	<none>
quickstart-nginx-ingress-controller	161 80:31345/TCP, 443:31376/TCP	16m	LoadBalancer	10.63.248.177	35.233.154.
quickstart-nginx-ingress-default-backend	80/TCP	16m	ClusterIP	10.63.250.234	<none>

This command shows you all the services in your cluster (in the `default` namespace), and any external IP addresses they have. When you first create the controller, your cloud provider won't have assigned and allocated an IP address through the LoadBalancer yet. Until it does, the external IP address for the service will be listed as `<pending>`.

Your cloud provider may have options for reserving an IP address prior to creating the ingress controller and using that IP address rather than assigning an IP address from a pool. Read through the documentation from your cloud provider on how to arrange that.

### Step 3 - Assign a DNS name

The external IP that is allocated to the ingress-controller is the IP to which all incoming traffic should be routed. To enable this, add it to a DNS zone you control, for example as `example.your-domain.com`.

This quickstart assumes you know how to assign a DNS entry to an IP address and will do so.

### Step 4 - Deploy an Example Service

Your service may have its own chart, or you may be deploying it directly with manifests. This quickstart uses manifests to create and expose a sample service. The example service uses `kuard`, a demo application which makes an excellent back-end for examples.

The quickstart example uses three manifests for the sample. The first two are a sample deployment and an associated service:

- deployment manifest: [deployment.yaml](#)

```
apiVersion: extensions/v1beta1
kind: Deployment
metadata:
  name: kuard
spec:
  replicas: 1
  template:
    metadata:
      labels:
        app: kuard
    spec:
      containers:
      - image: gcr.io/kuar-demo/kuard-amd64:1
        imagePullPolicy: Always
        name: kuard
        ports:
        - containerPort: 8080
```

- service manifest: [service.yaml](#)

```
apiVersion: v1
kind: Service
metadata:
  name: kuard
spec:
  ports:
  - port: 80
    targetPort: 8080
    protocol: TCP
  selector:
    app: kuard
```

You can create download and reference these files locally, or you can reference them from the GitHub source repository for this documentation. To install the example service from the tutorial files straight from GitHub, you may use the commands:

```
$ kubectl apply -f https://raw.githubusercontent.com/jetstack/cert-manager/release-0.
↪7/docs/tutorials/acme/quick-start/example/deployment.yaml
deployment.extensions "kuard" created

$ kubectl apply -f https://raw.githubusercontent.com/jetstack/cert-manager/release-0.
↪7/docs/tutorials/acme/quick-start/example/service.yaml
service "kuard" created
```

An [ingress resource](#) is what Kubernetes uses to expose this example service outside the cluster. You will need to download and modify the example manifest to reflect the domain that you own or control to complete this example.

A sample ingress you can start with is:

- ingress manifest: [ingress.yaml](#)

```
apiVersion: extensions/v1beta1
kind: Ingress
metadata:
```

(continues on next page)

(continued from previous page)

```

name: kuard
annotations:
  kubernetes.io/ingress.class: "nginx"
  #certmanager.k8s.io/issuer: "letsencrypt-staging"
  #certmanager.k8s.io/acme-challenge-type: http01

spec:
  tls:
    - hosts:
        - example.example.com
      secretName: quickstart-example-tls
  rules:
    - host: example.example.com
      http:
        paths:
          - path: /
            backend:
              serviceName: kuard
              servicePort: 80

```

You can download the sample manifest from github, edit it, and submit the manifest to Kubernetes with the command:

```

$ kubectl create --edit -f https://raw.githubusercontent.com/jetstack/cert-manager/
→release-0.7/docs/tutorials/acme/quick-start/example/ingress.yaml

# edit the file in your editor, and once it is saved:
ingress.extensions "kuard" created

```

**Note:** The ingress example we show above has a *host* definition within it. The nginx-ingress-controller will route traffic when the hostname requested matches the definition in the ingress. You *can* deploy an ingress without a *host* definition in the rule, but that pattern isn't usable with a TLS certificate, which expects a fully qualified domain name.

Once it is deployed, you can use the command *kubectl get ingress* to see the status of the ingress:

NAME	HOSTS	ADDRESS	PORTS	AGE
kuard	*		80, 443	17s

It may take a few minutes, depending on your service provider, for the ingress to be fully created. When it has been created and linked into place, the ingress will show an address as well:

NAME	HOSTS	ADDRESS	PORTS	AGE
kuard	*	35.199.170.62	80	9m

**Note:** The IP address on the ingress *may not* match the IP address that the nginx-ingress-controller. This is fine, and is a quirk/implementation detail of the service provider hosting your Kubernetes cluster. Since we are using the nginx-ingress-controller instead of any cloud-provider specific ingress backend, use the IP address that was defined and allocated for the nginx-ingress-service LoadBalancer resource as the primary access point for your service.

Make sure the service is reachable at the domain name you added above, for example *http://example.your-domain.com*. The simplest way is to open a browser and enter the name that you set up in DNS, and for which we just added the ingress.

You may also use a command line tool like *curl* to check the ingress.

```
$ curl -kivL -H 'Host: example.your-domain.com' 'http://35.199.164.14'
```

The options on this curl command will provide verbose output, following any redirects, show the TLS headers in the output, and not error on insecure certificates. With nginx-ingress-controller, the service will be available with a TLS certificate, but it will be using a self-signed certificate provided as a default from the nginx-ingress-controller. Browsers will show a warning that this is an invalid certificate. This is expected and normal, as we have not yet used cert-manager to get a fully trusted certificate for our site.

**Warning:** It is critical to make sure that your ingress is available and responding correctly on the internet. This quickstart example uses Let's Encrypt to provide the certificates, which expects and validates both that the service is available and that during the process of issuing a certificate uses that validation as proof that the request for the domain belongs to someone with sufficient control over the domain.

## Step 5 - Deploy Cert Manager

We need to install cert-manager to do the work with kubernetes to request a certificate and respond to the challenge to validate it. We can use helm to install cert-manager. This example installed cert-manager into the *kube-system* namespace from the public helm charts.

```
# Install the cert-manager CRDs. We must do this before installing the Helm
# chart in the next step for `release-0.7` of cert-manager:
$ kubectl apply -f https://raw.githubusercontent.com/jetstack/cert-manager/release-0.
↪7/deploy/manifests/00-crds.yaml

## IMPORTANT: you MUST install the cert-manager CRDs before installing the
## cert-manager Helm chart
$ kubectl apply \
  -f https://raw.githubusercontent.com/jetstack/cert-manager/release-0.7/deploy/
↪manifests/00-crds.yaml

## IMPORTANT: if the cert-manager namespace already exists, you MUST ensure
## it has an additional label on it in order for the deployment to succeed
$ kubectl label namespace cert-manager certmanager.k8s.io/disable-validation="true"

## Add the Jetstack Helm repository
$ helm repo add jetstack https://charts.jetstack.io
## Updating the repo just incase it already existed
$ helm repo update

## Install the cert-manager helm chart
$ helm install --name cert-manager --namespace cert-manager jetstack/cert-manager

NAME:      cert-manager
LAST DEPLOYED: Wed Jan  9 13:36:13 2019
NAMESPACE: cert-manager
STATUS:    DEPLOYED

RESOURCES:
==> v1beta1/ClusterRoleBinding
NAME                                AGE
cert-manager-webhook-ca-sync        2s
cert-manager-webhook:auth-delegator 2s
cert-manager                        2s
```

(continues on next page)

(continued from previous page)

```

==> v1beta1/APIService
NAME                                     AGE
v1beta1.admission.certmanager.k8s.io  2s

==> v1alpha1/Certificate
cert-manager-webhook-webhook-tls  1s
cert-manager-webhook-ca           1s

==> v1beta1/ValidatingWebhookConfiguration
cert-manager-webhook  1s

==> v1/ServiceAccount
NAME                SECRETS  AGE
cert-manager-webhook-ca-sync  1        2s
cert-manager-webhook          1        2s
cert-manager                 1        2s

==> v1beta1/RoleBinding
NAME                                     AGE
cert-manager-webhook:webhook-authentication-reader  2s

==> v1beta1/Deployment
NAME                DESIRED  CURRENT  UP-TO-DATE  AVAILABLE  AGE
cert-manager-webhook  1        1        1            0          2s
cert-manager          1        1        1            0          2s

==> v1/Job
NAME                DESIRED  SUCCESSFUL  AGE
cert-manager-webhook-ca-sync  1        0          2s

==> v1beta1/CronJob
NAME                SCHEDULE      SUSPEND  ACTIVE  LAST SCHEDULE  AGE
cert-manager-webhook-ca-sync  * * */24 * *  False   0        <none>        2s

==> v1beta1/ClusterRole
NAME                AGE
cert-manager-webhook-ca-sync  2s
cert-manager          2s

==> v1/ClusterRole
cert-manager-webhook:webhook-requester  2s
cert-manager-view                        2s
cert-manager-edit                        2s

==> v1/Service
NAME                TYPE          CLUSTER-IP    EXTERNAL-IP  PORT(S)  AGE
cert-manager-webhook  ClusterIP    10.3.244.237 <none>       443/TCP  2s

==> v1/ConfigMap
NAME                DATA  AGE
cert-manager-webhook-ca-sync  1      2s

==> v1alpha1/Issuer
NAME                AGE
cert-manager-webhook-ca  1s
cert-manager-webhook-selfsign  1s

```

(continues on next page)



(continued from previous page)

```

==> v1/Pod(related)
NAME                                READY  STATUS             RESTARTS  AGE
cert-manager-webhook-745b49d445-rnxm2  0/1   ContainerCreating  0          2s
cert-manager-9cdd9f774-t856z          0/1   ContainerCreating  0          2s
cert-manager-webhook-ca-sync-ddf4b    0/1   ContainerCreating  0          2s

NOTES:
cert-manager has been deployed successfully!

In order to begin issuing certificates, you will need to set up a ClusterIssuer
or Issuer resource (for example, by creating a 'letsencrypt-staging' issuer).

More information on the different types of issuers and how to configure them
can be found in our documentation:

https://docs.cert-manager.io/en/latest/reference/issuers.html

For information on how to configure cert-manager to automatically provision
Certificates for Ingress resources, take a look at the `ingress-shim`
documentation:

https://docs.cert-manager.io/en/latest/reference/ingress-shim.html

```

Cert-manager uses two different custom resources, also known as **CRD's**, to configure and control how it operates, as well as share status of its operation. These two resources are:

#### *Issuers (or ClusterIssuers)*

An Issuer is the definition for where cert-manager will get request TLS certificates. An Issuer is specific to a single namespace in Kubernetes, and a ClusterIssuer is meant to be a cluster-wide definition for the same purpose.

#### *Certificate*

A certificate is the resource that cert-manager uses to expose the state of a request as well as track upcoming expirations.

## Step 6 - Configure Let's Encrypt Issuer

We will set up two issuers for Let's Encrypt in this example. The Let's Encrypt production issuer has **very strict rate limits**. When you are experimenting and learning, it is very easy to hit those limits, and confuse rate limiting with errors in configuration or operation.

Because of this, we will start with the Let's Encrypt staging issuer, and once that is working switch to a production issuer.

Create this definition locally and update the email address to your own. This email required by Let's Encrypt and used to notify you of certificate expirations and updates.

- staging issuer: `staging-issuer.yaml`

```

apiVersion: certmanager.k8s.io/v1alpha1
kind: Issuer
metadata:
  name: letsencrypt-staging
spec:
  acme:

```

(continues on next page)

(continued from previous page)

```

# The ACME server URL
server: https://acme-staging-v02.api.letsencrypt.org/directory
# Email address used for ACME registration
email: user@example.com
# Name of a secret used to store the ACME account private key
privateKeySecretRef:
  name: letsencrypt-staging
# Enable the HTTP-01 challenge provider
http01: {}

```

Once edited, apply the custom resource:

```

$ kubectl create --edit -f https://raw.githubusercontent.com/jetstack/cert-manager/
↪release-0.7/docs/tutorials/acme/quick-start/example/staging-issuer.yaml
issuer.certmanager.k8s.io "letsencrypt-staging" created

```

Also create a production issuer and deploy it. As with the staging issuer, you will need to update this example and add in your own email address.

- production issuer: `production-issuer.yaml`

```

apiVersion: certmanager.k8s.io/v1alpha1
kind: Issuer
metadata:
  name: letsencrypt-prod
spec:
  acme:
    # The ACME server URL
    server: https://acme-v02.api.letsencrypt.org/directory
    # Email address used for ACME registration
    email: user@example.com
    # Name of a secret used to store the ACME account private key
    privateKeySecretRef:
      name: letsencrypt-prod
    # Enable the HTTP-01 challenge provider
    http01: {}

```

```

$ kubectl create --edit -f https://raw.githubusercontent.com/jetstack/cert-manager/
↪release-0.7/docs/tutorials/acme/quick-start/example/production-issuer.yaml
issuer.certmanager.k8s.io "letsencrypt-prod" created

```

Both of these issuers are configured to use the *HTTP01* challenge provider.

Check on the status of the issuer after you create it:

You should see the issuer listed with a registered account.

## Step 7 - Deploy a TLS Ingress Resource

With all the pre-requisite configuration in place, we can now do the pieces to request the TLS certificate. There are two primary ways to do this: using annotations on the ingress with *ingress-shim* or directly creating a certificate resource.

In this example, we will add annotations to the ingress, and take advantage of *ingress-shim* to have it create the certificate resource on our behalf. After creating a certificate, the cert-manager will update or create an ingress resource and use that to validate the domain. Once verified and issued, cert-manager will create or update the secret defined in the certificate.

**Note:** The secret that is used in the ingress should match the secret defined in the certificate. There isn't any explicit checking, so a typo will result in the nginx-ingress-controller falling back to its self-signed certificate. In our example, we are using annotations on the ingress (and ingress-shim) which will create the correct secrets on your behalf.

Edit the ingress add the annotations that were commented out in our earlier example:

- ingress tls: `ingress-tls.yaml`

```
apiVersion: extensions/v1beta1
kind: Ingress
metadata:
  name: kuard
  annotations:
    kubernetes.io/ingress.class: "nginx"
    certmanager.k8s.io/issuer: "letsencrypt-staging"
    certmanager.k8s.io/acme-challenge-type: http01
spec:
  tls:
  - hosts:
    - example.example.com
    secretName: quickstart-example-tls
  rules:
  - host: example.example.com
    http:
      paths:
      - path: /
        backend:
          serviceName: kuard
          servicePort: 80
```

and apply it:

```
$ kubectl create --edit -f https://raw.githubusercontent.com/jetstack/cert-manager/
↪release-0.7/docs/tutorials/acme/quick-start/example/ingress-tls.yaml
ingress.extensions "kuard" configured
```

Cert-manager will read these annotations and use them to create a certificate, which you can request and see:

```
$ kubectl get certificate
NAME                                AGE
quickstart-example-tls             38s
```

Cert-manager reflects the state of the process for every request in the certificate object. You can view this information using the `kubectl describe` command:

```
$ kubectl describe certificate quickstart-example-tls

Name:                                quickstart-example-tls
Namespace:                            default
Labels:                                <none>
Annotations:                            <none>
API Version:                          certmanager.k8s.io/v1alpha1
Kind:                                   Certificate
Metadata:
  Cluster Name:
```

(continues on next page)

```

Creation Timestamp: 2018-11-17T17:58:37Z
Generation: 0
Owner References:
  API Version: extensions/v1beta1
  Block Owner Deletion: true
  Controller: true
  Kind: Ingress
  Name: kuard
  UID: a3e9f935-ea87-11e8-82f8-42010a8a00b5
Resource Version: 9295
Self Link: /apis/certmanager.k8s.io/v1alpha1/namespaces/default/
↪certificates/quickstart-example-tls
  UID: 68d43400-ea92-11e8-82f8-42010a8a00b5
Spec:
  Acme:
    Config:
      Domains:
        example.your-domain.com
      Http 01:
        Ingress:
          Ingress Class: nginx
  Dns Names:
    example.your-domain.com
  Issuer Ref:
    Kind: Issuer
    Name: letsencrypt-staging
  Secret Name: quickstart-example-tls
Status:
  Acme:
    Order:
      URL: https://acme-staging-v02.api.letsencrypt.org/acme/order/7374163/13665676
  Conditions:
    Last Transition Time: 2018-11-17T18:05:57Z
    Message: Certificate issued successfully
    Reason: CertIssued
    Status: True
    Type: Ready
Events:
  Type      Reason      Age      From      Message
  ----      -
  Normal    CreateOrder  9m      cert-manager    Created new ACME order,
↪attempting validation...
  Normal    DomainVerified  8m      cert-manager    Domain "example.your-
↪domain.com" verified with "http-01" validation
  Normal    IssueCert    8m      cert-manager    Issuing certificate...
  Normal    CertObtained  7m      cert-manager    Obtained certificate,
↪from ACME server
  Normal    CertIssued   7m      cert-manager    Certificate issued,
↪Successfully

```

The events associated with this resource and listed at the bottom of the *describe* results show the state of the request. In the above example the certificate was validated and issued within a couple of minutes.

Once complete, cert-manager will have created a secret with the details of the certificate based on the secret used in the ingress resource. You can use the describe command as well to see some details:

```
$ kubectl describe secret quickstart-example-tls

Name:          quickstart-example-tls
Namespace:    default
Labels:       certmanager.k8s.io/certificate-name=quickstart-example-tls
Annotations:  certmanager.k8s.io/alt-names=example.your-domain.com
              certmanager.k8s.io/common-name=example.your-domain.com
              certmanager.k8s.io/issuer-kind=Issuer
              certmanager.k8s.io/issuer-name=letsencrypt-staging

Type: kubernetes.io/tls

Data
====
tls.crt:  3566 bytes
tls.key:  1675 bytes
```

Now that we have confidence that everything is configured correctly, you can update the annotations in the ingress to specify the production issuer:

- ingress tls final: [ingress-tls-final.yaml](#)

```
apiVersion: extensions/v1beta1
kind: Ingress
metadata:
  name: kuard
  annotations:
    kubernetes.io/ingress.class: "nginx"
    certmanager.k8s.io/issuer: "letsencrypt-prod"
    certmanager.k8s.io/acme-challenge-type: http01
spec:
  tls:
  - hosts:
    - example.example.com
    secretName: quickstart-example-tls
  rules:
  - host: example.example.com
    http:
      paths:
      - path: /
        backend:
          serviceName: kuard
          servicePort: 80
```

```
$ kubectl create --edit -f https://raw.githubusercontent.com/jetstack/cert-manager/
↪release-0.7/docs/tutorials/acme/quick-start/example/ingress-tls-final.yaml

ingress.extensions "kuard" configured
```

You will also need to delete the existing secret, which cert-manager is watching and will cause it to reprocess the request with the updated issuer.

```
$ kubectl delete secret quickstart-example-tls

secret "quickstart-example-tls" deleted
```

This will start the process to get a new certificate, and using describe you can see the status. Once the production cer-

tificate has been updated, you should see the example KUARD running at your domain with a signed TLS certificate.

```
$ kubectl describe certificate
Name:          quickstart-example-tls
Namespace:     default
Labels:        <none>
Annotations:   <none>
API Version:   certmanager.k8s.io/v1alpha1
Kind:          Certificate
Metadata:
  Cluster Name:
  Creation Timestamp: 2018-11-17T18:36:48Z
  Generation:        0
  Owner References:
    API Version:      extensions/v1beta1
    Block Owner Deletion: true
    Controller:       true
    Kind:              Ingress
    Name:              kuard
    UID:               a3e9f935-ea87-11e8-82f8-42010a8a00b5
  Resource Version:  283686
  Self Link:         /apis/certmanager.k8s.io/v1alpha1/namespaces/default/
↪certificates/quickstart-example-tls
  UID:               bdd93b32-ea97-11e8-82f8-42010a8a00b5
Spec:
  Acme:
    Config:
      Domains:
        example.your-domain.com
      Http 01:
        Ingress:
          Ingress Class: nginx
    Dns Names:
      example.your-domain.com
    Issuer Ref:
      Kind:      Issuer
      Name:      letsencrypt-prod
      Secret Name: quickstart-example-tls
Status:
  Conditions:
    Last Transition Time: 2019-01-09T13:52:05Z
    Message:              Certificate does not exist
    Reason:               NotFound
    Status:               False
    Type:                 Ready
Events:
  Type    Reason      Age   From          Message
  ----    -
  Normal  Generated   18s   cert-manager  Generated new private key
  Normal  OrderCreated 18s   cert-manager  Created Order resource "quickstart-
↪example-tls-889745041"
```

You can see the current state of the ACME Order by running `kubectl describe` on the Order resource that cert-manager has created for your Certificate:

```
$ kubectl describe order quickstart-example-tls-889745041
...
```

(continues on next page)

(continued from previous page)

```

Events:
  Type      Reason      Age   From           Message
  ----      -
  Normal    Created     90s   cert-manager   Created Challenge resource "quickstart-
↳example-tls-889745041-0" for domain "example.your-domain.com"

```

Here, we can see that cert-manager has created 1 ‘Challenge’ resource to fulfil the Order. You can dig into the state of the current ACME challenge by running `kubectl describe` on the automatically created Challenge resource:

```

$ kubectl describe challenge quickstart-example-tls-889745041-0
...

Status:
  Presented:   true
  Processing:  true
  Reason:      Waiting for http-01 challenge propagation
  State:       pending
Events:
  Type      Reason      Age   From           Message
  ----      -
  Normal    Started     15s   cert-manager   Challenge scheduled for processing
  Normal    Presented   14s   cert-manager   Presented challenge using http-01 challenge_
↳mechanism

```

From above, we can see that the challenge has been ‘presented’ and cert-manager is waiting for the challenge record to propagate to the ingress controller. You should keep an eye out for new events on the challenge resource, as a ‘success’ event should be printed after a minute or so (depending on how fast your ingress controller is at updating rules):

```

$ kubectl describe challenge quickstart-example-tls-889745041-0
...

Status:
  Presented:   false
  Processing:  false
  Reason:      Successfully authorized domain
  State:       valid
Events:
  Type      Reason      Age   From           Message
  ----      -
  Normal    Started     71s   cert-manager   Challenge scheduled for processing
  Normal    Presented   70s   cert-manager   Presented challenge using http-01_
↳challenge mechanism
  Normal    DomainVerified 2s    cert-manager   Domain "example.your-domain.com"
↳verified with "http-01" validation

```

**Note:** If your challenges are not becoming ‘valid’ and remain in the ‘pending’ state (or enter into a ‘failed’ state), it is likely there is some kind of configuration error. Read the [Challenge resource reference docs](#) for more information on debugging failing challenges.

Once the challenge(s) have been completed, their corresponding challenge resources will be *deleted*, and the ‘Order’ will be updated to reflect the new state of the Order:

```

$ kubectl describe order quickstart-example-tls-889745041
...

```

(continues on next page)

(continued from previous page)

```

Events:
  Type          Reason          Age   From          Message
  ----          -
  Normal       Created         90s   cert-manager  Created Challenge resource "quickstart-
↪example-tls-889745041-0" for domain "example.your-domain.com"
  Normal       OrderValid     16s   cert-manager  Order completed successfully

```

Finally, the 'Certificate' resource will be updated to reflect the state of the issuance process. If all is well, you should be able to 'describe' the Certificate and see something like the below:

```

$ kubectl describe certificate quickstart-example-tls

Status:
  Conditions:
    Last Transition Time: 2019-01-09T13:57:52Z
    Message:             Certificate is up to date and has not expired
    Reason:              Ready
    Status:              True
    Type:                Ready
    Not After:           2019-04-09T12:57:50Z
  Events:
    Type          Reason          Age   From          Message
    ----          -
    Normal       Generated       11m   cert-manager  Generated new private key
    Normal       OrderCreated    11m   cert-manager  Created Order resource
↪"quickstart-example-tls-889745041"
    Normal       OrderComplete  10m   cert-manager  Order "quickstart-example-
↪tls-889745041" completed successfully

```

## 2.1.2 Issuing an ACME certificate using DNS validation

**Todo:** This guide needs rewriting to be clearer, splitting into sections and potentially rewriting altogether.

cert-manager can be used to obtain certificates from a CA using the [ACME](#) protocol. The ACME protocol supports various challenge mechanisms which are used to prove ownership of a domain so that a valid certificate can be issued for that domain.

One such challenge mechanism is DNS-01. With a DNS-01 challenge, you prove ownership of a domain by proving you control its DNS records. This is done by creating a TXT record with specific content that proves you have control of the domains DNS records.

The following Issuer defines the necessary information to enable DNS validation. You can read more about the Issuer resource in the [Issuer reference docs](#).

```

1 apiVersion: certmanager.k8s.io/v1alpha1
2 kind: Issuer
3 metadata:
4   name: letsencrypt-staging
5   namespace: default
6 spec:
7   acme:
8     server: https://acme-staging-v02.api.letsencrypt.org/directory
9     email: user@example.com

```

(continues on next page)



(continued from previous page)

```

10
11 # Name of a secret used to store the ACME account private key
12 privateKeySecretRef:
13   name: letsencrypt-staging
14
15 # ACME DNS-01 provider configurations
16 dns01:
17
18 # Here we define a list of DNS-01 providers that can solve DNS challenges
19 providers:
20
21 - name: prod-dns
22   clouddns:
23     # A secretKeyRef to a google cloud json service account
24     serviceAccountSecretRef:
25       name: clouddns-service-account
26       key: service-account.json
27     # The project in which to update the DNS zone
28     project: gcloud-prod-project
29
30 - name: cf-dns
31   cloudflare:
32     email: user@example.com
33     # A secretKeyRef to a cloudflare api key
34     apiKeySecretRef:
35       name: cloudflare-api-key
36     key: api-key.txt

```

We have specified the ACME server URL for Let's Encrypt's [staging environment](#). The staging environment will not issue trusted certificates but is used to ensure that the verification process is working properly before moving to production. Let's Encrypt's production environment imposes much stricter [rate limits](#), so to reduce the chance of you hitting those limits it is highly recommended to start by using the staging environment. To move to production, simply create a new Issuer with the URL set to `https://acme-v02.api.letsencrypt.org/directory`.

The first stage of the ACME protocol is for the client to register with the ACME server. This phase includes generating an asymmetric key pair which is then associated with the email address specified in the Issuer. Make sure to change this email address to a valid one that you own. It is commonly used to send expiry notices when your certificates are coming up for renewal. The generated private key is stored in a Secret named `letsencrypt-staging`.

The `dns01` stanza contains a list of DNS-01 providers that can be used to solve DNS challenges. Our Issuer defines two providers. This gives us a choice of which one to use when obtaining certificates.

More information about the DNS provider configuration, including a list of supported providers, can be found [in the dns01 reference docs](#).

Once we have created the above Issuer we can use it to obtain a certificate.

```

1 apiVersion: certmanager.k8s.io/v1alpha1
2 kind: Certificate
3 metadata:
4   name: example-com
5   namespace: default
6 spec:
7   secretName: example-com-tls
8   issuerRef:
9     name: letsencrypt-staging
10  commonName: '*.example.com'

```

(continues on next page)

(continued from previous page)

```

11  dnsNames:
12  - example.com
13  - foo.com
14  acme:
15    config:
16    - dns01:
17      provider: prod-dns
18      domains:
19      - '*.example.com'
20      - example.com
21    - dns01:
22      provider: cf-dns
23      domains:
24      - foo.com

```

The Certificate resource describes our desired certificate and the possible methods that can be used to obtain it. You can obtain certificates for wildcard domains just like any other. Make sure to wrap wildcard domains with asterisks in your YAML resources, to avoid formatting issues. If you specify both `example.com` and `*.example.com` on the same Certificate, it will take slightly longer to perform validation as each domain will have to be validated one after the other. You can learn more about the Certificate resource in the [reference docs](#). If the certificate is obtained successfully, the resulting key pair will be stored in a secret called `example-com-tls` in the same namespace as the Certificate.

The certificate will have a common name of `*.example.com` and the [Subject Alternative Names \(SANs\)](#) will be `*.example.com`, `example.com` and `foo.com`.

In our Certificate we have referenced the `letsencrypt-staging` Issuer above. The Issuer must be in the same namespace as the Certificate. If you want to reference a `ClusterIssuer`, which is a cluster-scoped version of an Issuer, you must add `kind: ClusterIssuer` to the `issuerRef` stanza.

For more information on `ClusterIssuers`, read the [ClusterIssuer reference docs](#).

The `acme` stanza defines the configuration for our ACME challenges. Here we have defined the configuration for our DNS challenges which will be used to verify domain ownership. For each domain mentioned in a `dns01` stanza, cert-manager will use the provider's credentials from the referenced Issuer to create a TXT record called `_acme-challenge`. This record will then be verified by the ACME server in order to issue the certificate. Once domain ownership has been verified, any cert-manager affected records will be cleaned up.

---

**Note:** It is your responsibility to ensure the selected provider is authoritative for your domain.

---

After creating the above Certificate, we can check whether it has been obtained successfully using `kubectl describe`:

```

$ kubectl describe certificate example-com
Events:
  Type      Reason          Age          From          Message
  ----      -
  Normal    CreateOrder     57m         cert-manager  Created new ACME order, attempting_
↪validation...
  Normal    DomainVerified  55m         cert-manager  Domain "*.example.com" verified with
↪"dns-01" validation
  Normal    DomainVerified  55m         cert-manager  Domain "example.com" verified with
↪"dns-01" validation
  Normal    DomainVerified  55m         cert-manager  Domain "foo.com" verified with "dns-
↪01" validation

```

(continues on next page)

(continued from previous page)

Normal	IssueCert	55m	cert-manager	Issuing certificate...
Normal	CertObtained	55m	cert-manager	Obtained certificate from ACME server
Normal	CertIssued	55m	cert-manager	Certificate issued successfully

You can also check whether issuance was successful with `kubectl get secret example-com-tls -o yaml`. You should see a base64 encoded signed TLS key pair.

Once our certificate has been obtained, cert-manager will periodically check its validity and attempt to renew it if it gets close to expiry. cert-manager considers certificates to be close to expiry when the 'Not After' field on the certificate is less than the current time plus 30 days.

### 2.1.3 Issuing an ACME certificate using HTTP validation

cert-manager can be used to obtain certificates from a CA using the [ACME](#) protocol. The ACME protocol supports various challenge mechanisms which are used to prove ownership of a domain so that a valid certificate can be issued for that domain.

One such challenge mechanism is the HTTP-01 challenge. With a HTTP-01 challenge, you prove ownership of a domain by ensuring that a particular file is present at the domain. It is assumed that you control the domain if you are able to publish the given file under a given path.

The following Issuer defines the necessary information to enable HTTP validation. You can read more about the Issuer resource in the [Issuer reference docs](#).

```

1 apiVersion: certmanager.k8s.io/v1alpha1
2 kind: Issuer
3 metadata:
4   name: letsencrypt-staging
5   namespace: default
6 spec:
7   acme:
8     # The ACME server URL
9     server: https://acme-staging-v02.api.letsencrypt.org/directory
10    # Email address used for ACME registration
11    email: user@example.com
12    # Name of a secret used to store the ACME account private key
13    privateKeySecretRef:
14      name: letsencrypt-staging
15    # Enable the HTTP-01 challenge provider
16    http01: {}

```

We have specified the ACME server URL for Let's Encrypt's [staging environment](#). The staging environment will not issue trusted certificates but is used to ensure that the verification process is working properly before moving to production. Let's Encrypt's production environment imposes much stricter [rate limits](#), so to reduce the chance of you hitting those limits it is highly recommended to start by using the staging environment. To move to production, simply create a new Issuer with the URL set to `https://acme-v02.api.letsencrypt.org/directory`.

The first stage of the ACME protocol is for the client to register with the ACME server. This phase includes generating an asymmetric key pair which is then associated with the email address specified in the Issuer. Make sure to change this email address to a valid one that you own. It is commonly used to send expiry notices when your certificates are coming up for renewal. The generated private key is stored in a Secret named `letsencrypt-staging`.

The presence of the `http01` field simply enables the HTTP-01 challenge for this Issuer. No further configuration is necessary or currently possible.

Once we have created the above Issuer we can use it to obtain a certificate.

```
1 apiVersion: certmanager.k8s.io/v1alpha1
2 kind: Certificate
3 metadata:
4   name: example-com
5   namespace: default
6 spec:
7   secretName: example-com-tls
8   issuerRef:
9     name: letsencrypt-staging
10  commonName: example.com
11  dnsNames:
12    - www.example.com
13  acme:
14    config:
15      - http01:
16          ingressClass: nginx
17          domains:
18            - example.com
19      - http01:
20          ingress: my-ingress
21          domains:
22            - www.example.com
```

The Certificate resource describes our desired certificate and the possible methods that can be used to obtain it. You can learn more about the Certificate resource in the [reference docs](#). If the certificate is obtained successfully, the resulting key pair will be stored in a secret called `example-com-tls` in the same namespace as the Certificate.

The certificate will have a common name of `example.com` and the [Subject Alternative Names \(SANs\)](#) will be `example.com` and `www.example.com`.

In our Certificate we have referenced the `letsencrypt-staging` Issuer above. The Issuer must be in the same namespace as the Certificate. If you want to reference a `ClusterIssuer`, which is a cluster-scoped version of an Issuer, you must add `kind: ClusterIssuer` to the `issuerRef` stanza.

For more information on ClusterIssuers, read the [ClusterIssuer reference docs](#).

The `acme` stanza defines the configuration for our ACME challenges. Here we have defined the configuration for our HTTP-01 challenges which will be used to verify domain ownership. To verify ownership of each domain mentioned in an `http01` stanza, cert-manager will create a Pod, Service and Ingress that exposes an HTTP endpoint that satisfies the HTTP-01 challenge.

The fields `ingress` and `ingressClass` in the `http01` stanza can be used to control how cert-manager interacts with Ingress resources:

- If the `ingress` field is specified, then an Ingress resource with the same name in the same namespace as the Certificate must already exist and it will be modified only to add the appropriate rules to solve the challenge. This field is useful for the GCLB ingress controller, as well as a number of others, that assign a single public IP address for each ingress resource. Without manual intervention, creating a new ingress resource would cause any challenges to fail.
- If the `ingressClass` field is specified, a new ingress resource with a randomly generated name will be created in order to solve the challenge. This new resource will have an annotation with key `kubernetes.io/ingress.class` and value set to the value of the `ingressClass` field. This works for the likes of the NGINX ingress controller.
- If neither are specified, new ingress resources will be created with a randomly generated name, but they will not have the ingress class annotation set.
- If both are specified, then the `ingress` field will take precedence.

Once domain ownership has been verified, any cert-manager affected resources will be cleaned up or deleted.

**Note:** It is your responsibility to point each domain name at the correct IP address for your ingress controller.

After creating the above Certificate, we can check whether it has been obtained successfully using `kubectl describe`:

```
$ kubectl describe certificate example-com
Events:
  Type      Reason          Age          From          Message
  ----      -
  Normal    CreateOrder     57m         cert-manager  Created new ACME order, attempting_
↪validation...
  Normal    DomainVerified  55m         cert-manager  Domain "example.com" verified with
↪"http-01" validation
  Normal    DomainVerified  55m         cert-manager  Domain "www.example.com" verified_
↪with "http-01" validation
  Normal    IssueCert       55m         cert-manager  Issuing certificate...
  Normal    CertObtained    55m         cert-manager  Obtained certificate from ACME server
  Normal    CertIssued      55m         cert-manager  Certificate issued successfully
```

You can also check whether issuance was successful with `kubectl get secret example-com-tls -o yaml`. You should see a base64 encoded signed TLS key pair.

Once our certificate has been obtained, cert-manager will periodically check its validity and attempt to renew it if it gets close to expiry. cert-manager considers certificates to be close to expiry when the ‘Not After’ field on the certificate is less than the current time plus 30 days.

## 2.1.4 Migrating from kube-lego

`kube-lego` is an older Jetstack project for obtaining TLS certificates from Let’s Encrypt (or another ACME server).

Since cert-managers release, kube-lego has been gradually deprecated in favour of this project. There are a number of key differences between the two:

Feature	kube-lego	cert-manager
Configuration	Annotations on Ingress resources	CRDs
CAs	ACME	ACME, signing keypair
Kubernetes	v1.2 - v1.8	v1.7+
Debugging	Look at logs	Kubernetes Events API
Multi-tenancy	Not supported	Supported
Distinct issuance sources per Certificate	Not supported	Supported
Ingress controller support (ACME)	GCE, nginx	All

This guide will walk through how you can safely migrate your kube-lego installation to cert-manager, without service interruption.

By the end of the guide, we should have:

1. Scaled down and removed kube-lego
2. Installed cert-manager
3. Migrated ACME private key to cert-manager
4. Created an ACME ClusterIssuer using this private key, to issue certificates throughout your cluster

5. Configured cert-manager's *ingress-shim* to automatically provision Certificate resources for all Ingress resources with the `kubernetes.io/tls-acme: "true"` annotation, using the ClusterIssuer we have created
6. Verified that the cert-manager installation is working

### 1. Scale down kube-lego

Before we begin deploying cert-manager, it is best we scale our kube-lego deployment down to 0 replicas. This will prevent the two controllers potentially 'fighting' each other. If you deployed kube-lego using the official deployment YAMLs, a command like so should do:

```
$ kubectl scale deployment kube-lego \
  --namespace kube-lego \
  --replicas=0
```

You can then verify your kube-lego pod is no longer running with:

```
$ kubectl get pods --namespace kube-lego
```

### 2. Deploy cert-manager

cert-manager should be deployed using Helm, according to our official *Get started* guide. No special steps are required here. We will return to this deployment at the end of this guide and perform an upgrade of some of the CLI flags we deploy cert-manager with however.

Please take extra care to ensure you have configured RBAC correctly when deploying Helm and cert-manager - there are some nuances described in our deploying document!

### 3. Obtaining your ACME account private key

In order to continue issuing and renewing certificates on your behalf, we need to migrate the user account private key that kube-lego has created for you over to cert-manager.

Your ACME user account identity is a private key, stored in a secret resource. By default, kube-lego will store this key in a secret named `kube-lego-account` in the same namespace as your kube-lego Deployment. You may have overridden this value when you deploy kube-lego, in which case the secret name to use will be the value of the `LEGO_SECRET_NAME` environment variable.

You should download a copy of this secret resource and save it in your local directory:

```
$ kubectl get secret kube-lego-account -o yaml \
  --namespace kube-lego \
  --export > kube-lego-account.yaml
```

Once saved, open up this file and change the `metadata.name` field to something more relevant to cert-manager. For the rest of this guide, we'll assume you chose `letsencrypt-private-key`.

Once done, we need to create this new resource in the `kube-system` namespace. By default, cert-manager stores supporting resources for ClusterIssuers in the namespace that it is running in, and we used `kube-system` when deploying cert-manager above. You should change this if you have deployed cert-manager into a different namespace.

```
$ kubectl create -f kube-lego-account.yaml \
  --namespace kube-system
```

## 4. Creating an ACME ClusterIssuer using your old ACME account

We need to create a ClusterIssuer which will hold information about the ACME account previously registered via kube-lego. In order to do so, we need two more pieces of information from our old kube-lego deployment: the server URL of the ACME server, and the email address used to register the account.

Both of these bits of information are stored within the kube-lego ConfigMap.

To retrieve them, you should be able to get the ConfigMap using `kubectl`:

```
$ kubectl get configmap kube-lego -o yaml \
  --namespace kube-lego \
  --export
```

Your email address should be shown under the `.data.lego.email` field, and the ACME server URL under `.data.lego.url`.

For the purposes of this guide, we will assume the lego email is `user@example.com` and the URL `https://acme-staging-v02.api.letsencrypt.org/directory`.

Now that we have migrated our private key to the new Secret resource, as well as obtaining our ACME email address and URL, we can create a ClusterIssuer resource!

Create a file named `cluster-issuer.yaml`:

```
1 apiVersion: certmanager.k8s.io/v1alpha1
2 kind: ClusterIssuer
3 metadata:
4   # Adjust the name here accordingly
5   name: letsencrypt-staging
6 spec:
7   acme:
8     # The ACME server URL
9     server: https://acme-staging-v02.api.letsencrypt.org/directory
10    # Email address used for ACME registration
11    email: user@example.com
12    # Name of a secret used to store the ACME account private key from step 3
13    privateKeySecretRef:
14      name: letsencrypt-private-key
15    # Enable the HTTP-01 challenge provider
16    http01: {}
```

We then submit this file to our Kubernetes cluster:

```
$ kubectl create -f cluster-issuer.yaml
```

You should be able to verify the ACME account has been verified successfully:

```
$ kubectl describe clusterissuer letsencrypt-staging
Name:          letsencrypt-staging
Namespace:
Labels:        <none>
Annotations:   <none>
API Version:   certmanager.k8s.io/v1alpha1
Kind:          ClusterIssuer
Metadata:
  Cluster Name:
  Creation Timestamp:  2017-11-30T22:33:40Z
  Generation:         0
```

(continues on next page)

(continued from previous page)

```

Resource Version: 4450170
Self Link: /apis/certmanager.k8s.io/v1alpha1/letsencrypt-staging
UID: 83d04e6b-d61e-11e7-ac26-42010a840044
Spec:
  Acme:
    Email: user@example.com
    Http 01:
    Private Key Secret Ref:
      Key:
        Name: letsencrypt-private-key
        Server: https://acme-staging-v02.api.letsencrypt.org/directory
  Status:
    Acme:
      Uri: https://acme-staging-v02.api.letsencrypt.org/acme/acct/11217539
    Conditions:
      Last Transition Time: 2018-04-12T17:32:30Z
      Message: The ACME account was registered with the ACME server
      Reason: ACMEAccountRegistered
      Status: True
      Type: Ready

```

## 5. Configuring ingress-shim to use our new ClusterIssuer by default

Now that our ClusterIssuer is ready to issue certificates, we have one last thing to do: we must reconfigure ingress-shim (deployed as part of cert-manager) to automatically create Certificate resources for all Ingress resources it finds with appropriate annotations.

More information on the role of ingress-shim can be found *in the docs*, but for now we can just run a `helm upgrade` in order to add a few additional flags. Assuming you've named your ClusterIssuer `letsencrypt-staging` (as above), run:

```

helm upgrade cert-manager \
  stable/cert-manager \
  --namespace kube-system \
  --set ingressShim.defaultIssuerName=letsencrypt-staging \
  --set ingressShim.defaultIssuerKind=ClusterIssuer

```

You should see the cert-manager pod be re-created, and once started it should automatically create Certificate resources for all of your ingresses that previously had kube-lego enabled.

## 6. Verify each ingress now has a corresponding Certificate

Before we finish, we should make sure there is now a Certificate resource for each ingress resource you previously enabled kube-lego on.

You should be able to check this by running:

```
$ kubectl get certificates --all-namespaces
```

There should be an entry for each ingress in your cluster with the kube-lego annotation.

We can also verify that cert-manager has 'adopted' the old TLS certificates by viewing the logs for cert-manager:



```
$ kubectl logs -n kube-system -l app=cert-manager -c cert-manager
...
I1025 21:54:02.869269      1 sync.go:206] Certificate my-example-certificate_
↳scheduled for renewal in 292 hours
```

Here we can see cert-manager has verified the existing TLS certificate and scheduled it to be renewed in 292h time.



This section contains guides on using specific features of cert-manager, such as configuring different Issuer types and any special settings that you may want to configure.

### 3.1 Setting up Issuers

Before you can begin issuing certificates, you must configure at least one Issuer or ClusterIssuer resource in your cluster.

These represent a certificate authority from which signed x509 certificates can be obtained, such as Let's Encrypt, or your own signing key pair stored in a Kubernetes Secret resource. They are referenced by Certificate resources in order to request certificates from them.

An *Issuer* is scoped to a single namespace, and can only fulfill *Certificate* resources within its own namespace. This is useful in a multi-tenant environment where multiple teams or independent parties operate within a single cluster.

On the other hand, a *ClusterIssuer* is a cluster wide version of an *Issuer*. It is able to be referenced by *Certificate* resources in any namespace.

Users often create `letsencrypt-staging` and `letsencrypt-prod` *ClusterIssuers* if they operate a single-tenant environment and want to expose a cluster-wide mechanism for obtaining TLS certificates from Let's Encrypt.

#### 3.1.1 Supported issuer types

cert-manager supports a number of different issuer backends, each with their own different types of configuration.

Please follow one of the below linked guides to learn how to set up the issuer types you require:

- *CA* - issue certificates signed by a X509 signing keypair, stored in a Secret in the Kubernetes API server.
- *Self signed* - issue self signed certificates.
- *ACME* - issue certificates obtained by performing challenge validations against an ACME server such as Let's Encrypt.

- *Vault* - issue certificates from a Vault instance configured with the [Vault PKI backend](#).
- *Venafi* - issue certificates from a [Venafi Cloud](#) or [Trust Protection Platform](#) instance.

### 3.1.2 Additional information

There are a few key things to know about Issuers, but for full information you can refer to the [Issuer reference docs](#).

#### Difference between Issuers and ClusterIssuers

ClusterIssuers are a resource type similar to *Issuers*. They are specified in exactly the same way, but they do not belong to a single namespace and can be referenced by Certificate resources from multiple different namespaces.

They are particularly useful when you want to provide the ability to obtain certificates from a central authority (e.g. Letsencrypt, or your internal CA) and you run single-tenant clusters.

The resource spec is identical, and you should set the `certificate.spec.issuerRef.kind` field to `ClusterIssuer` when creating your Certificate resources.

#### Setting up ACME Issuers

The ACME Issuer type represents a single Account registered with the ACME server.

When you create a new ACME Issuer, cert-manager will generate a private key which is used to identify you with the ACME server.

To set up a basic ACME issuer, you should create a new Issuer or ClusterIssuer resource.

In this example, we will create a non-namespaced ClusterIssuer resource for the [Let's Encrypt staging endpoint](#) that has only the [HTTP01 Challenge Provider](#) enabled.

You should read the guides linked at the bottom of this page to learn more about the ACME challenge validation mechanisms that cert-manager supports and how to configure the various DNS01 provider implementations.

#### Creating a basic ACME Issuer

The below example configures a ClusterIssuer named `letsencrypt-staging` that is configured to enable the HTTP01 challenge validation mechanism **only**.

You should copy and paste this example into a new file named `letsencrypt-staging.yaml` and update the `spec.acme.email` field to be your own email address.

```
1 apiVersion: certmanager.k8s.io/v1alpha1
2 kind: ClusterIssuer
3 metadata:
4   name: letsencrypt-staging
5 spec:
6   acme:
7     # You must replace this email address with your own.
8     # Let's Encrypt will use this to contact you about expiring
9     # certificates, and issues related to your account.
10    email: user@example.com
11    server: https://acme-staging-v02.api.letsencrypt.org/directory
12    privateKeySecretRef:
13    # Secret resource used to store the account's private key.
```

(continues on next page)

(continued from previous page)

```
14   name: example-issuer-account-key
15   # Enable the HTTP01 challenge mechanism for this Issuer
16   http01: {}
```

You can then create this resource:

```
kubectl apply -f letsencrypt-staging.yaml
```

To verify that the account has been registered successfully, you can run `kubectl describe` and check the 'Ready' condition:

```
kubectl describe clusterissuer letsencrypt-staging
...
Status:
  Acme:
    Uri: https://acme-staging-v02.api.letsencrypt.org/acme/acct/7571319
  Conditions:
    Last Transition Time: 2019-01-30T14:52:03Z
    Message:              The ACME account was registered with the ACME server
    Reason:               ACMEAccountRegistered
    Status:               True
    Type:                 Ready
```

## Notes on issuing ACME certificates

Currently, there is some additional configuration needed on Certificate resources when issuing certificates from ACME issuers.

You should read the *Issuing Certificates using ACME* documentation for more information on how to configure these additional fields.

## Advanced HTTP01 configuration

There are a few additional options that can be set on the Issuer resource to alter the behaviour of the HTTP01 solver. For full details, read the *HTTP01 Challenge Provider* documentation to learn about these options.

## Configuring DNS01 providers

It is also possible to validate domain ownership using DNS01 validation.

In order to do this, your Issuer resource must be configured with credentials for a supported DNS provider's account. The full list of support DNS providers, and information on how to configure them can be found in the *DNS01 Challenge Provider* documentation.

## Setting up CA Issuers

cert-manager can be used to obtain certificates using an arbitrary signing key pair stored in a Kubernetes Secret resource.

This guide will show you how to configure and create a CA based issuer, backed by a signing key pair stored in a Secret resource.

### 1. (Optional) Generate a signing key pair

The CA Issuer does not automatically create and manage a signing key pair for you. As a result, you will need to either supply your own or generate a self signed CA using a tool such as [openssl](#) or [cfssl](#).

This guide will explain how to generate a new signing key pair, however you can substitute it for your own so long as it has the `CA` flag set.

```
# Generate a CA private key
$ openssl genrsa -out ca.key 2048

# Create a self signed Certificate, valid for 10yrs with the 'signing' option set
$ openssl req -x509 -new -nodes -key ca.key -subj "/CN=${COMMON_NAME}" -days 3650 -
↪reqexts v3_req -extensions v3_ca -out ca.crt
```

The output of these commands will be two files, `ca.key` and `ca.crt`, the key and certificate for your signing key pair. If you already have your own key pair, you should name the private key and certificate `ca.key` and `ca.crt` respectively.

### 2. Save the signing key pair as a Secret

We are going to create an Issuer that will use this key pair to generate signed certificates. You can read more about the Issuer resource in [the Issuer reference docs](#). To allow the Issuer to reference our key pair we will store it in a Kubernetes Secret resource.

Issuers are namespaced resources and so they can only reference Secrets in their own namespace. We will therefore put the key pair into the same namespace as the Issuer. We could alternatively create a [ClusterIssuer](#), a cluster-scoped version of an Issuer. For more information on ClusterIssuers, read the [ClusterIssuer reference documentation](#).

The following command will create a Secret containing a signing key pair in the default namespace:

```
kubectl create secret tls ca-key-pair \
  --cert=ca.crt \
  --key=ca.key \
  --namespace=default
```

### 3. Creating an Issuer referencing the Secret

We can now create an Issuer referencing the Secret resource we just created:

```
1 apiVersion: certmanager.k8s.io/v1alpha1
2 kind: Issuer
3 metadata:
4   name: ca-issuer
5   namespace: default
6 spec:
7   ca:
8     secretName: ca-key-pair
```

We are now ready to obtain certificates!

## 4. Obtain a signed Certificate

We can now create the following Certificate resource which specifies the desired certificate. You can read more about the Certificate resource in *the reference docs*.

```

1 apiVersion: certmanager.k8s.io/v1alpha1
2 kind: Certificate
3 metadata:
4   name: example-com
5   namespace: default
6 spec:
7   secretName: example-com-tls
8   issuerRef:
9     name: ca-issuer
10    # We can reference ClusterIssuers by changing the kind here.
11    # The default value is Issuer (i.e. a locally namespaced Issuer)
12    kind: Issuer
13   commonName: example.com
14   organization:
15     - Example CA
16   dnsNames:
17     - example.com
18     - www.example.com

```

In order to use the Issuer to obtain a Certificate, we must create a Certificate resource in the **same namespace as the Issuer**, as an Issuer is a namespaced resource. We could alternatively create a *ClusterIssuer* if we wanted to reuse the signing key pair across multiple namespaces.

Once we have created the Certificate resource, cert-manager will attempt to use the Issuer `ca-issuer` to obtain a certificate. If successful, the certificate will be stored in a Secret resource named `example-com-tls` in the same namespace as the Certificate resource (default).

The example above explicitly sets the `commonName` field to `example.com`. cert-manager automatically adds the `commonName` field as a **DNS SAN** if it is not already contained in the `dnsNames` field.

If we had **not** specified the `commonName` field, then the **first** DNS SAN that is specified (under `dnsNames`) would be used as the certificate's common name.

After creating the above Certificate, we can check whether it has been obtained successfully like so:

```

$ kubectl describe certificate example-com
Events:
  Type          Reason              Age             From              Message
  ----          -
  Warning       ErrorCheckCertificate 26s            cert-manager-controller Error_
  ↪checking existing TLS certificate: secret "example-com-tls" not found
  Normal        PrepareCertificate    26s            cert-manager-controller Preparing_
  ↪certificate with issuer
  Normal        IssueCertificate     26s            cert-manager-controller Issuing_
  ↪certificate...
  Normal        CertificateIssued    25s            cert-manager-controller _
  ↪Certificate issued successfully

```

You can also check whether issuance was successful with `kubectl get secret example-com-tls -o yaml`. You should see a base64 encoded signed TLS key pair.

Once the certificate has been obtained, cert-manager will keep checking its validity and attempt to renew it if it gets close to expiry. cert-manager considers certificates to be close to expiry when the 'Not After' field on the certificate is

less than the current time plus 30 days. For CA based Issuers, cert-manager will issue certificates with the ‘Not After’ field set to the current time plus 365 days.

### Setting up self signing Issuers

Self signed Issuers will issue self signed certificates.

This is useful when building PKI within Kubernetes, or as a means to generate a root CA for use with the *CA Issuer*.

A self-signed Issuer contains no additional configuration fields, and can be created with a resource like so:

```
apiVersion: certmanager.k8s.io/v1alpha1
kind: ClusterIssuer
metadata:
  name: selfsigning-issuer
spec:
  selfSigned: {}
```

---

**Note:** The presence of the `selfSigned: {}` line is enough to indicate that this Issuer is of type ‘self signed’.

---

Once created, you should be able to issue certificates like usual by referencing the newly created Issuer in your `issuerRef`:

```
apiVersion: certmanager.k8s.io/v1alpha1
kind: Certificate
metadata:
  name: example-crt
spec:
  secretName: my-selfsigned-cert
  commonName: "my-selfsigned-root-ca"
  isCA: true
  issuerRef:
    name: selfsigning-issuer
    kind: ClusterIssuer
```

### Setting up Vault Issuers

#### Installing Vault

Vault installation is a complex subject. For a thorough tour of the subject you can read the official HashiCorp Vault [documentation](#).

#### Vault PKI Backend

The PKI Secrets Engine needs to be initialized for cert-manager to be able to generate certificate. The official Vault documentation can be found [here](#).

#### Vault Authentication with a AppRole

This Vault authentication method uses a [Vault AppRole](#).



The secret ID of the AppRole is stored in a secret.

Here an example of a secret containing the `secretId` of the AppRole:

```
apiVersion: v1
kind: Secret
type: Opaque
metadata:
  name: cert-manager-vault-approle
  namespace: default
data:
  secretId: "MDI..."
```

Where the `secretId` is the base 64 encoded value of the appRole `secretId` giving access to the pki backend in Vault.

We can now create a cluster issuer referencing this secret:

```
apiVersion: certmanager.k8s.io/v1alpha1
kind: Issuer
metadata:
  name: vault-issuer
  namespace: default
spec:
  vault:
    path: pki_int/sign/example-dot-com
    server: https://vault
    caBundle: <base64 encoded caBundle PEM file>
    auth:
      appRole:
        path: approle
        roleId: "291b9d21-8ff5-..."
        secretRef:
          name: cert-manager-vault-approle
          key: secretId
```

Where `path` is the Vault role path of the PKI backend and `server` is the Vault server base URL. The `path` MUST USE the vault `sign` endpoint. The Vault appRole credentials are supplied as the Vault authentication method using the appRole created in Vault. The `secretRef` references the Kubernetes secret created previously. More specifically, the field `name` is the Kubernetes secret name and `key` is the name given as the key value that store the `secretId`. The optional attribute `path` specifies where the AppRole authentication is mounted in Vault. The attribute `path` default value is `approle`.

An optional base64 encoded `caBundle` in PEM format can be provided to validate the TLS connection to the Vault Server. When `caBundle` is set it replaces the CA bundle inside the container running cert-manager. This parameter has no effect if the connection used is in plain HTTP.

Once we have created the above Issuer we can use it to obtain a certificate.

```
apiVersion: certmanager.k8s.io/v1alpha1
kind: Certificate
metadata:
  name: example-com
  namespace: default
spec:
  secretName: example-com-tls
  issuerRef:
    name: vault-issuer
  commonName: example.com
```

(continues on next page)

```
dnsNames:
- www.example.com
```

The Certificate resource describes our desired certificate and the possible methods that can be used to obtain it. You can learn more about the Certificate resource in the [reference docs](#). If the certificate is obtained successfully, the resulting key pair will be stored in a secret called `example-com-tls` in the same namespace as the Certificate.

The certificate will have a common name of `example.com` and the [Subject Alternative Names \(SANs\)](#) will be `example.com` and `www.example.com`.

In our Certificate we have referenced the `vault-issuer` Issuer above. The Issuer must be in the same namespace as the Certificate. If you want to reference a `ClusterIssuer`, which is a cluster-scoped version of an Issuer, you must add `kind: ClusterIssuer` to the `issuerRef` stanza.

For more information on ClusterIssuers, read the [ClusterIssuer reference docs](#).

## Vault Authentication with a Token

This Vault authentication method uses a plain token. A Vault token is generated by one of the many authentication backends supported by Vault. Tokens in Vault have expiration and need to be refreshed. You need to be aware that cert-manager does not refresh these tokens. Another process must be put in place to keep them from expiring.

For testing purposes a root token is generated at Vault installation time. **WARNING: Root tokens do not expire, so should only be used for testing purposes.**

Please refer to the official token [documentation](#) for all the details.

Here an example of a secret Kubernetes resource containing the Vault token:

```
apiVersion: v1
kind: Secret
type: Opaque
metadata:
  name: cert-manager-vault-token
  namespace: kube-system
data:
  token: "MjI..."
```

Where the token value is the base 64 encoded value of the token giving access to the PKI backend in Vault.

We can now create an issuer referencing this secret:

```
apiVersion: certmanager.k8s.io/v1alpha1
kind: Issuer
metadata:
  name: vault-issuer
  namespace: default
spec:
  vault:
    auth:
      tokenSecretRef:
        name: cert-manager-vault-token
        key: token
    path: pki_int/sign/example-dot-com
    server: https://vault
    caBundle: <base64 encoded caBundle PEM file>
```

Where *path* is the Vault role path of the PKI backend and *server* is the Vault server base URL. The secret created previously is referenced in the issuer with its *name* and *key* corresponding to the name of the Kubernetes secret and the property name containing the token value respectively.

An optional base64 encoded *caBundle* in PEM format can be provided to validate the TLS connection to the Vault Server. When *caBundle* is set it replaces the CA bundle inside the container running cert-manager. This parameter has no effect if the connection used is in plain HTTP.

Once we have created the above Issuer we can use it to obtain a certificate.

```
apiVersion: certmanager.k8s.io/v1alpha1
kind: Certificate
metadata:
  name: example-com
  namespace: default
spec:
  secretName: example-com-tls
  issuerRef:
    name: vault-issuer
  commonName: example.com
  dnsNames:
  - www.example.com
```

The Certificate resource describes our desired certificate and the possible methods that can be used to obtain it. You can learn more about the Certificate resource in the *reference docs*. If the certificate is obtained successfully, the resulting key pair will be stored in a secret called `example-com-tls` in the same namespace as the Certificate.

The certificate will have a common name of `example.com` and the Subject Alternative Names (SANs) will be `example.com` and `www.example.com`.

In our Certificate we have referenced the `vault-issuer` Issuer above. The Issuer must be in the same namespace as the Certificate. If you want to reference a ClusterIssuer, which is a cluster-scoped version of an Issuer, you must add `kind: ClusterIssuer` to the `issuerRef` stanza.

For more information on ClusterIssuers, read the *ClusterIssuer reference docs*.

## Setting up Venafi Issuers

The Venafi Issuer types allows you to obtain certificates from Venafi Cloud and ‘Venafi Trust Protection Platform’ instances.

Register your account at <https://api.venafi.cloud/login> and get an API key from your dashboard.

You can have multiple different Venafi Issuer types installed within the same cluster, including mixtures of Cloud and TPP issuer types. This allows you to be flexible with the types of Venafi account you use.

Automated certificate renewal and management are provided for Certificates using the Venafi issuer.

---

**Note:** The Venafi Issuer has been recently added, and the exact structure of the Issuer resource is subject to change. Such changes will be clearly documented, and migration steps will be provided.

---

## Creating an Issuer resource

A single Venafi Issuer represents a single ‘zone’ within the Venafi API, therefore you must create an Issuer resource for each Venafi Zone you want to obtain certificates from.

You can configure your Issuer resource to either issue certificates only within a single namespace, or cluster-wide (using a ClusterIssuer resource). For more information on the distinction between Issuer and ClusterIssuer resources, read the [issuer\\_vs\\_clusterissuer](#) section.

### Creating a Venafi Cloud Issuer

In order to set up a Venafi Cloud Issuer, you must first create a Kubernetes Secret resource containing your Venafi Cloud API credentials:

```
kubectl create secret generic \
  cloud-secret \
  --namespace='NAMESPACE OF YOUR ISSUER RESOURCE' \
  --from-literal=apikey='YOUR_CLOUD_API_KEY_HERE'
```

**Note:** If you are configuring your Issuer as a ClusterIssuer resource in order to issue Certificates across your whole cluster, you must set the `--namespace` parameter to `cert-manager`, which is the default 'cluster resource namespace'.

This API key will be used by cert-manager to interact with the Venafi Cloud service on your behalf.

Once the API key Secret has been created, you can create your Issuer or ClusterIssuer resource. If you are creating a ClusterIssuer resource, you must change the `kind` field to `ClusterIssuer` and remove the `metadata.namespace` field.

Save the below content after making your amendments to a file named `venafi-cloud-issuer.yaml`:

```
apiVersion: certmanager.k8s.io/v1alpha1
kind: Issuer
metadata:
  name: cloud-venafi-issuer
  namespace: <NAMESPACE YOU WANT TO ISSUE CERTIFICATES IN>
spec:
  venafi:
    zone: "DevOps" # Set this to the Venafi policy zone you want to use
  cloud:
    apiTokenSecretRef:
      name: cloud-secret
      key: apikey
```

You can then create the Issuer using `kubectl create -f`:

```
kubectl create -f venafi-cloud-issuer.yaml
```

Verify the Issuer has been initialised correctly using `kubectl describe`:

```
kubectl describe issuer cloud-venafi-issuer --namespace='NAMESPACE OF YOUR ISSUER_
↳RESOURCE'

(TODO) include sample output
```

You are now ready to issue certificates using the newly provisioned Venafi Issuer.

Read the [Issuing Certificates](#) document for more information on how to create Certificate resources.

## Creating a Venafi Trust Protection Platform Issuer

The Venafi Trust Protection integration allows you to obtain certificates from a properly configured Venafi TPP instance.

The setup is similar to the Venafi Cloud configuration above, however some of the connection parameters are slightly different.

**Note:** You **must** allow “User Provided CSRs” as part of your TPP policy, as this is the only type supported by cert-manager at this time.

In order to set up a Venafi Trust Protection Platform Issuer, you must first create a Kubernetes Secret resource containing your Venafi TPP API credentials:

```
kubectl create secret generic \
  tpp-secret \
  --namespace=<NAMESPACE OF YOUR ISSUER RESOURCE> \
  --from-literal=user='YOUR_TPP_USERNAME_HERE' \
  --from-literal=password='YOUR_TPP_PASSWORD_HERE'
```

**Note:** If you are configuring your Issuer as a ClusterIssuer resource in order to issue Certificates across your whole cluster, you must set the `--namespace` parameter to `cert-manager`, which is the default ‘cluster resource namespace’.

These credentials will be used by cert-manager to interact with your Venafi TPP instance.

Once the Secret containing credentials has been created, you can create your Issuer or ClusterIssuer resource. If you are creating a ClusterIssuer resource, you must change the `kind` field to `ClusterIssuer` and remove the `metadata.namespace` field.

Save the below content after making your amendments to a file named `venafi-tpp-issuer.yaml`:

```
apiVersion: certmanager.k8s.io/v1alpha1
kind: Issuer
metadata:
  name: tpp-venafi-issuer
  namespace: <NAMESPACE YOU WANT TO ISSUE CERTIFICATES IN>
spec:
  venafi:
    zone: devops\cert-manager # Set this to the Venafi policy zone you want to use
  tpp:
    url: https://tpp.venafi.example/vedsdk # Change this to the URL of your TPP_
    ↪instance
    caBundle: <base64 encoded string of caBundle PEM file>
    credentialsRef:
      name: tpp-secret
```

You can then create the Issuer using `kubectl create -f`:

```
kubectl create -f venafi-tpp-issuer.yaml
```

Verify the Issuer has been initialised correctly using `kubectl describe`:

```
kubectl describe issuer tpp-venafi-issuer --namespace='NAMESPACE OF YOUR ISSUER_  
↳RESOURCE'  
  
(TODO) include sample output
```

You are now ready to issue certificates using the newly provisioned Venafi Issuer.

Read the *Issuing Certificates* document for more information on how to create Certificate resources.

## 3.2 Issuing Certificates

The Certificate resource type is used to request certificates from different Issuers.

In order to issue any certificates, you'll need to configure an Issuer resource first.

If you have not configured any issuers yet, you should read the *Setting up Issuers* guide.

### 3.2.1 Creating Certificate resources

A Certificate resource specifies fields that are used to generate certificate signing requests which are then fulfilled by the issuer type you have referenced.

Certificates specify which issuer they want to obtain the certificate from by specifying the `certificate.spec.issuerRef` field.

A basic Certificate resource, for the `example.com` and `www.example.com` DNS names that is valid for 90d and renews 15d before expiry is below:

```
1 apiVersion: certmanager.k8s.io/v1alpha1  
2 kind: Certificate  
3 metadata:  
4   name: example-com  
5   namespace: default  
6 spec:  
7   secretName: example-com-tls  
8   duration: 2160h # 90d  
9   renewBefore: 360h # 15d  
10  commonName: example.com  
11  dnsNames:  
12  - example.com  
13  - www.example.com  
14  issuerRef:  
15    name: ca-issuer  
16    # We can reference ClusterIssuers by changing the kind here.  
17    # The default value is Issuer (i.e. a locally namespaced Issuer)  
18  kind: Issuer
```

The signed certificate will be stored in a Secret resource named `example-com-tls` once the issuer has successfully issued the requested certificate.

The Certificate will be issued using the issuer named `ca-issuer` in the default namespace (the same namespace as the Certificate resource).

---

**Note:** If you want to create an Issuer that can be referenced by Certificate resources in **all** namespaces, you should create a *ClusterIssuer* resource and set the `certificate.spec.issuerRef.kind` field to `ClusterIssuer`.

---

**Note:** The `renewBefore` and `duration` fields must be specified using Golang's `time.Time` string format, which does not allow the `d` (days) suffix. You must specify these values using `s`, `m` and `h` suffixes instead. Failing to do so without installing the *webhook* component can prevent cert-manager from functioning correctly (#1269).

---

A full list of the fields supported on the Certificate resource can be found in the [API reference documentation](#).

### 3.2.2 Temporary certificates whilst issuing

With some Issuer types, certificates can take a few minutes to be issued.

A temporary untrusted certificate will be issued whilst this process takes place if another certificate does not already exist in the target Secret resource.

This helps to improve compatibility with certain ingress controllers (e.g. *ingress-gce*) which require a TLS certificate to be present at all times in order to function.

After the real, valid certificate has been obtained, cert-manager will replace the temporary self signed certificate with the valid one, **but will retain the same private key**.

### 3.2.3 Special fields on Certificate resources for ACME Issuers

When creating Certificate resources that reference ACME Issuers, you must set an additional `certificate.spec.acme` stanza on the resource to configure what challenge mechanism to use for each DNS name specified on the certificate.

More information on setting these fields can be found in the *Issuing Certificates using ACME* guide.

#### Automatically creating Certificates for Ingress resources

cert-manager can be configured to automatically provision TLS certificates for Ingress resources via annotations on your Ingresses.

A small sub-component of cert-manager, *ingress-shim*, is responsible for this.

#### How it works

*ingress-shim* watches Ingress resources across your cluster. If it observes an Ingress with *any* of the annotations described in the 'Usage' section, it will ensure a Certificate resource with the same name as the Ingress, and configured as described on the Ingress exists. For example:

```
apiVersion: extensions/v1beta1
kind: Ingress
metadata:
  annotations:
    # add an annotation indicating the issuer to use.
    certmanager.k8s.io/cluster-issuer: nameOfClusterIssuer
  name: myIngress
```

(continues on next page)

```
namespace: myIngress
spec:
  rules:
  - host: myingress.com
    http:
      paths:
      - backend:
          serviceName: myservice
          servicePort: 80
        path: /
  tls: # < placing a host in the TLS config will indicate a cert should be created
  - hosts:
    - myingress.com
      secretName: myingress-cert # < cert-manager will store the created certificate in
      ↪ this secret.
```

## Configuration

Since cert-manager v0.2.2, ingress-shim is deployed automatically as part of a Helm chart installation.

If you would also like to use the old `kube-lego` `kubernetes.io/tls-acme: "true"` annotation for fully automated TLS, you will need to configure a default Issuer when deploying cert-manager. This can be done by adding the following `--set` when deploying using Helm:

```
--set ingressShim.defaultIssuerName=letsencrypt-prod \
--set ingressShim.defaultIssuerKind=ClusterIssuer
```

In the above example, cert-manager will create Certificate resources that reference the ClusterIssuer `letsencrypt-prod` for all Ingresses that have a `kubernetes.io/tls-acme: "true"` annotation.

For more information on deploying cert-manager, read the [deployment guide](#).

## Supported annotations

You can specify the following annotations on ingresses in order to trigger Certificate resources to be automatically created:

- `certmanager.k8s.io/issuer` - the name of an Issuer to acquire the certificate required for this ingress from. The Issuer **must** be in the same namespace as the Ingress resource.
- `certmanager.k8s.io/cluster-issuer` - the name of a ClusterIssuer to acquire the certificate required for this ingress from. It does not matter which namespace your Ingress resides, as ClusterIssuers are non-namespaced resources.
- `certmanager.k8s.io/acme-challenge-type` - by default, if the Issuer specified is an ACME issuer (either through ingress-shim's defaults, or with one of the above annotations), the ingress-shim will set the ACME challenge mechanism on the Certificate resource it creates to 'http01'. This annotation can be used to alter this behaviour. Must be one of 'http01' or 'dns01'.
- `certmanager.k8s.io/acme-dns01-provider` - if the ACME challenge type has been set to dns01, this annotation **must** be specified to instruct cert-manager which DNS provider (as configured on the specified Issuer resource) should be used. This field is required if the challenge type is set to DNS01.
- `certmanager.k8s.io/acme-http01-ingress-class` - if the ACME challenge type has been set to http01, this annotation allows you to configure ingress class that will be used to solve challenges for this ingress.



Customising this is useful when you are trying to secure internal services, and need to solve challenges using different ingress class to that of the ingress. If not specified and the 'acme-http01-edit-in-place' annotation is not set, this defaults to the ingress class of the ingress resource.

- `kubernetes.io/tls-acme: "true"` - this annotation requires additional configuration of the ingress-shim (see above). Namely, a default issuer must be specified as arguments to the ingress-shim container.
- `certmanager.k8s.io/acme-http01-edit-in-place: "true"` - if the ACME challenge type has been set to http01, and the ingress has the 'kubernetes.io/tls-acme: true' annotation, this controls whether the ingress is modified 'in-place', or a new one created specifically for the http01 challenge. If present, and set to "true" the existing ingress will be modified. Any other value, or the absence of the annotation assumes "false".

## 3.3 ACME specific tasks

In order to use the ACME provider, there are a number of required fields. For your ACME issuer to support the various ACME challenge mechanisms, you may need to provide some additional configuration on your resource, such as configuring credentials for a DNS provider or enabling HTTP01 validation.

### 3.3.1 Issuing Certificates using ACME

ACME certificates currently require additional configuration on the Certificate resource that you create in order to determine how to solve the [ACME challenges](#) that the ACME protocol requires.

In future releases of cert-manager, this configuration is likely to move off of the Certificate resource and onto the Issuer resource in order to create a better separation of concerns. More info can be found on [issue #1450](#).

#### Configuring Certificates for ACME issuance

In order to issue certificates using the ACME issuer type, you must configure which ACME challenge provider is used for each domain name you are requesting a Certificate for.

This is done by configuring a mapping between domain names and the solver types that have been configured on the corresponding Issuer resource.

#### Using HTTP01 challenges

In order to use the HTTP01 challenge provider, you must first configure your Issuer with the appropriate settings described in the [HTTP01 Challenge Provider](#) documentation.

Assuming you've created the same example ACME Issuer with http01 enabled as in the [Setting up ACME Issuers](#) guide:

```

1 apiVersion: certmanager.k8s.io/v1alpha1
2 kind: ClusterIssuer
3 metadata:
4   name: letsencrypt-staging
5 spec:
6   acme:
7     # You must replace this email address with your own.
8     # Let's Encrypt will use this to contact you about expiring
9     # certificates, and issues related to your account.

```

(continues on next page)

(continued from previous page)

```
10 email: user@example.com
11 server: https://acme-staging-v02.api.letsencrypt.org/directory
12 privateKeySecretRef:
13   # Secret resource used to store the account's private key.
14   name: example-issuer-account-key
15   # Enable the HTTP01 challenge mechanism for this Issuer
16 http01: {}
```

We must configure our Certificate resource with the ‘ingress class’ that will be used to solve the ACME HTTP01 challenges:

```
1 apiVersion: certmanager.k8s.io/v1alpha1
2 kind: Certificate
3 metadata:
4   name: example-com
5   namespace: default
6 spec:
7   secretName: example-com-tls
8   issuerRef:
9     name: letsencrypt-staging
10  commonName: example.com
11  dnsNames:
12  - example.com
13  - www.example.com
14  acme:
15    config:
16    - http01:
17      ingressClass: nginx
18    domains:
19    - example.com
20    - www.example.com
```

**Note:** If you use ‘ingress-gce’, aka the GCLB ingress controller, you will need to modify your Certificate definition to specify the `certificate.spec.acme.config.http01.ingress` field instead of `ingressClass`, like so:

```
...
acme:
  config:
  - http01:
    ingress: name-of-gce-ingress-resource
  domains:
  - example.com
  - www.example.com
```

### Using DNS01 challenges

In order to use DNS01 validation, you must first configure your Issuer resource with credentials and connection information needed to access your DNS provider’s administrative console.

You can find more information on the different supported DNS providers and how to configure them in the *DNS01 Challenge Provider* documentation.

The example Issuer on the [DNS01 Challenge Provider](#) page is configured with credentials for a Google Cloud DNS account:

```

1 apiVersion: certmanager.k8s.io/v1alpha1
2 kind: ClusterIssuer
3 metadata:
4   name: letsencrypt-staging
5 spec:
6   acme:
7     email: user@example.com
8     server: https://acme-staging-v02.api.letsencrypt.org/directory
9     privateKeySecretRef:
10    name: example-issuer-account-key
11    dns01:
12     providers:
13     - name: prod-clouddns
14       clouddns:
15         project: my-project
16         serviceAccountSecretRef:
17         name: prod-clouddns-svc-acct-secret
18         key: service-account.json

```

In the above example on line 13, you can see we have named this DNS provider `prod-clouddns`.

When creating Certificates that intend to utilise this DNS01 provider for validations, we must remember to include this “provider name” in our Certificate’s spec:

```

1 apiVersion: certmanager.k8s.io/v1alpha1
2 kind: Certificate
3 metadata:
4   name: example-com
5   namespace: default
6 spec:
7   secretName: example-com-tls
8   issuerRef:
9     name: letsencrypt-staging
10  commonName: example.com
11  dnsNames:
12  - example.com
13  - www.example.com
14  acme:
15    config:
16    - dns01:
17      provider: prod-clouddns
18    domains:
19    - example.com
20    - www.example.com

```

If you do not specify a provider name, cert-manager will not know how to solve challenges for your domains and the issuance process **will not succeed**.

### 3.3.2 DNS01 Challenge Provider

The ACME issuer can also contain DNS provider configuration, which can be used by Certificates using this Issuer in order to validate DNS01 challenge requests:

You can read about how the DNS01 challenge type works on the [Let’s Encrypt challenge types page](#).

```
1 apiVersion: certmanager.k8s.io/v1alpha1
2 kind: Issuer
3 metadata:
4   name: example-issuer
5 spec:
6   acme:
7     email: user@example.com
8     server: https://acme-staging-v02.api.letsencrypt.org/directory
9     privateKeySecretRef:
10    name: example-issuer-account-key
11   dns01:
12     providers:
13     - name: prod-clouddns
14       clouddns:
15         project: my-project
16         serviceAccountSecretRef:
17         name: prod-clouddns-svc-acct-secret
18         key: service-account.json
```

Each issuer can specify multiple different DNS01 challenge providers, and it is also possible to have multiple instances of the same DNS provider on a single Issuer (e.g. two clouddns accounts could be set, each with their own name).

### Setting nameservers for DNS01 self check

Cert-manager will check the correct DNS records exist before attempting a DNS01 challenge. By default, the DNS servers for this check will be taken from `/etc/resolv.conf`. If this is not desired (for example with multiple authoritative nameservers or split-horizon DNS), the cert-manager controller provides the `--dns01-self-check-nameservers` flag, which allows overriding the default nameservers with a comma separated list of custom nameservers.

Example usage:

```
--dns01-self-check-nameservers "8.8.8.8:53,1.1.1.1:53"
```

### Supported DNS01 providers

A number of different DNS providers are supported for the ACME issuer. Below is a listing of available providers, their `.yaml` configurations, along with additional Kubernetes and provider specific notes regarding their usage.

#### ACME-DNS

```
acmedns:
  host: https://acme.example.com
  accountSecretRef:
    name: acme-dns
    key: acmedns.json
```

In general, clients to acme-dns perform registration on the users behalf and inform them of the CNAME entries they must create. This is not possible in cert-manager, it is a non-interactive system. Registration must be carried out beforehand and the resulting credentials JSON uploaded to the cluster as a secret. In this example, we use `curl` and the API endpoints directly. Information about setting up and configuring acme-dns is available on the [acme-dns project page](#).

1. First, register with the acme-dns server, in this example, there is one running at “auth.example.com”

`curl -X POST http://auth.example.com/register` will return a JSON with credentials for your registration:

```
{
  "username": "eabcb41-d89f-4580-826f-3e62e9755ef2",
  "password": "pbAXVj1IOE01xbut7YnAbkhMQIkwoHO0ek2j4Q0",
  "fulldomain": "d420c923-bbd7-4056-ab64-c3ca54c9b3cf.auth.example.com",
  "subdomain": "d420c923-bbd7-4056-ab64-c3ca54c9b3cf",
  "allowfrom": []
}
```

It is strongly recommended to restrict the update endpoint to the IP range of your pods. This is done at registration time as follows:

```
curl -X POST http://auth.example.com/register -H "Content-Type: application/json" --data '{"allowfrom": ["10.244.0.0/16"]}'
```

Make sure to update the `allowfrom` field to match your cluster configuration. The JSON will now look like

```
{
  "username": "eabcb41-d89f-4580-826f-3e62e9755ef2",
  "password": "pbAXVj1IOE01xbut7YnAbkhMQIkwoHO0ek2j4Q0",
  "fulldomain": "d420c923-bbd7-4056-ab64-c3ca54c9b3cf.auth.example.com",
  "subdomain": "d420c923-bbd7-4056-ab64-c3ca54c9b3cf",
  "allowfrom": ["10.244.0.0/16"]
}
```

2. Save this JSON to a file with the key as your domain. You can specify multiple domains with the same credentials if you like. In our example, the returned credentials can be used to verify ownership of “example.com” and “example.org”.

```
{
  "example.com": {
    "username": "eabcb41-d89f-4580-826f-3e62e9755ef2",
    "password": "pbAXVj1IOE01xbut7YnAbkhMQIkwoHO0ek2j4Q0",
    "fulldomain": "d420c923-bbd7-4056-ab64-c3ca54c9b3cf.auth.example.com",
    "subdomain": "d420c923-bbd7-4056-ab64-c3ca54c9b3cf",
    "allowfrom": ["10.244.0.0/16"]
  },
  "example.org": {
    "username": "eabcb41-d89f-4580-826f-3e62e9755ef2",
    "password": "pbAXVj1IOE01xbut7YnAbkhMQIkwoHO0ek2j4Q0",
    "fulldomain": "d420c923-bbd7-4056-ab64-c3ca54c9b3cf.auth.example.com",
    "subdomain": "d420c923-bbd7-4056-ab64-c3ca54c9b3cf",
    "allowfrom": ["10.244.0.0/16"]
  }
}
```

3. Next update your primary DNS server with CNAME record that will tell the verifier how to locate the challenge TXT record. This is obtained from the “fulldomain” field in the registration:

```
_acme-challenge.example.com CNAME d420c923-bbd7-4056-ab64-c3ca54c9b3cf.auth.example.com
_acme-challenge.example.org CNAME d420c923-bbd7-4056-ab64-c3ca54c9b3cf.auth.example.com
```

Note that the “name” of the record is always the “\_acme-challenge” subdomain, and the “value” of the record matches exactly the “fulldomain” field from registration.

At verification time, the domain name `d420c923-bbd7-4056-ab64-c3ca54c9b3cf.auth.example.com` will be a TXT record that is set to your validation token. When the verifier queries `_acme-challenge.example.com`, it will be directed to the correct location by this CNAME record. This proves that you control “example.com”

4. Create a secret from the credentials json that was saved in step 2, this secret is referenced in the `accountSecretRef` field of your `dns01` issuer settings.

```
kubectl create secret generic acme-dns --from-file acmedns.json
```

### Akamai FastDNS

```
akamai:
  serviceConsumerDomain: akab-tho6xie2aiteip8p-poith5aej0ughaba.luna.akamaiapis.net
  clientTokenSecretRef:
    name: akamai-dns
    key: clientToken
  clientSecretSecretRef:
    name: akamai-dns
    key: clientSecret
  accessTokenSecretRef:
    name: akamai-dns
    key: accessToken
```

### AzureDNS

Configuring the AzureDNS DNS-01 Challenge for a Kubernetes cluster requires creating a service principal in Azure.

For security purposes, it is appropriate to utilize RBAC to ensure that you properly maintain access control to your resources in Azure. The service principal that is generated by this tutorial has fine grained access to ONLY the DNS Zone in the specific resource group specified. It requires this permission so that it can read/write the `_acme_challenge` TXT records to the zone.

To create the service principal:

```
1 AZURE_CERT_MANAGER_SP_NAME=SOME_SERVICE_PRINCIPAL_NAME
2 AZURE_CERT_MANAGER_SP_PASSWORD=SOME_PASSWORD
3 AZURE_CERT_MANAGER_DNS_RESOURCE_GROUP=SOME_RESOURCE_GROUP
4 AZURE_CERT_MANAGER_DNS_NAME=SOME_DNS_ZONE
5
6 AZURE_CERT_MANAGER_SP_APP_ID=$(az ad sp create-for-rbac --name $AZURE_CERT_MANAGER_SP_
  ↳NAME --password $AZURE_CERT_MANAGER_SP_PASSWORD --query "appId" --output tsv)
7
8 # Lower the Permissions of the SP
9 az role assignment delete --assignee $AZURE_CERT_MANAGER_SP_APP_ID --role Contributor
10
11 # Give Access to DNS Zone
12 DNS_ID=$(az network dns zone show --name $AZURE_CERT_MANAGER_DNS_NAME --resource-
  ↳group $AZURE_CERT_MANAGER_DNS_RESOURCE_GROUP --query "id" --output tsv)
13
14 az role assignment create --assignee $AZURE_CERT_MANAGER_SP_APP_ID --role "DNS Zone_
  ↳Contributor" --scope $DNS_ID
15
16 # Check Permissions
17 az role assignment list --assignee $AZURE_CERT_MANAGER_SP_APP_ID
```

(continues on next page)

(continued from previous page)

```

18
19 # Create Secret
20 kubectl create secret generic azuredns-config \
21   --from-literal=CLIENT_SECRET=$AZURE_CERT_MANAGER_SP_PASSWORD
22
23 # Get the Service Principal App ID for configuration
24 echo $AZURE_CERT_MANAGER_SP_APP_ID

```

You can configure the issuer like so:

```

apiVersion: certmanager.k8s.io/v1alpha1
kind: ClusterIssuer
metadata:
  name: letsencrypt-prod
spec:
  acme:
    server: https://acme-v02.api.letsencrypt.org/directory
    email: example@example.com
    privateKeySecretRef:
      name: letsencrypt-prod
    dns01:
      providers:
      - name: azure
        azuredns:
          # Service principal clientId (also called appId)
          clientID: AZURE_SERVICE_PRINCIPAL_ID
          # A secretKeyRef to a service principal ClientSecret (password)
          # ref: https://docs.microsoft.com/en-us/azure/container-service/
          ↪ kubernetes/container-service-kubernetes-service-principal
          clientSecretSecretRef:
            name: AZUREDNS_SECRET_KEY_NAME
            key: CLIENT_SECRET
          # Azure subscription Id
          subscriptionID: AZURE_SUBSCRIPTION_ID
          # Azure AD tenant Id
          tenantID: AZURE_TENANT_ID
          # ResourceGroup name where dns zone is provisioned
          resourceGroupName: AZURE_RESOURCE_GROUP
          hostedZoneName: AZURE_DNS_ZONE_NAME

```

## Cloudflare

```

cloudflare:
  email: my-cloudflare-acc@example.com
  apiKeySecretRef:
    name: cloudflare-api-key-secret
    key: api-key

```

## Google CloudDNS

This guide explains how to set up an Issuer, or ClusterIssuer, to use Google CloudDNS to solve DNS01 ACME challenges. It's advised you read the *DNS01 Challenge Provider* page first for a more general understanding of how cert-manager handles DNS01 challenges.

**Note:** This guide assumes that your cluster is hosted on Google Cloud Platform (GCP) and that you already have a domain set up with CloudDNS.

---

### Set up a Service Account

Cert-manager needs to be able to add records to CloudDNS in order to solve the DNS01 challenge. To enable this, a GCP service account must be created with the `dns.admin` role.

---

**Note:** For this guide the `gcloud` command will be used to set up the service account. Ensure that `gcloud` is in using the correct project and zone before entering the commands. These steps could also be completed using the Cloud Console.

---

```
gcloud iam service-accounts create dns01-solver \
  --display-name "dns01-solver"
# Replace both uses of project-id with the id of your project
gcloud projects add-iam-policy-binding project-id \
  --member serviceAccount:dns01-solver@project-id.iam.gserviceaccount.com \
  --role roles/dns.admin
```

### Create a Service Account Secret

To access this service account cert-manager uses a key stored in a Kubernetes Secret. First, create a key for the service account and download it as JSON file, then create a Secret from this file.

```
# Replace use of project-id with the id of your project
gcloud iam service-accounts keys create key.json \
  --iam-account dns01-solver@project-id.iam.gserviceaccount.com
kubectl create secret generic clouddns-dns01-solver-svc-acct \
  --from-file=key.json
```

**Note:** Keep the key file safe and do not share it, as it could be used to gain access to your cloud resources. The key file can be deleted once it has been used to generate the Secret.

---

### Create an Issuer That Uses CloudDNS

Next, create an Issuer (or ClusterIssuer) with a `clouddns` provider. An example Issuer manifest can be seen below with annotations.

```
1 apiVersion: certmanager.k8s.io/v1alpha1
2 kind: Issuer
3 metadata:
4   name: letsencrypt-staging
5 spec:
6   acme:
7     # Replace with your email address so you can be notified of expiring certificates
8     email: user@example.com
```

(continues on next page)



(continued from previous page)

```

9     # Use Let's Encrypt staging for testing as production enforces stricter usage_
↪limits
10    server: https://acme-staging-v02.api.letsencrypt.org/directory
11    privateKeySecretRef:
12      # The secret that holds the generated private key used to communicate with Let
↪'s Encrypt
13      name: letsencrypt-staging-account-key
14      dns01:
15        providers:
16          # The name given to this CloudDNS provider, multiple CloudDNS providers can be_
↪added with different names
17          - name: my-clouddns-provider
18            clouddns:
19              # The ID of the GCP project
20              project: my-project-id
21              # This is the secret used to access the service account
22              serviceAccountSecretRef:
23                name: clouddns-dns01-solver-svc-acct
24                key: key.json

```

For more information about Issuers and ClusterIssuers, see [Setting Up Issuers](#).

Once an Issuer (or ClusterIssuer) has been created successfully a Certificate can then be added to verify that everything works.

```

1  apiVersion: certmanager.k8s.io/v1alpha1
2  kind: Certificate
3  metadata:
4    name: example-com
5    namespace: default
6  spec:
7    secretName: example-com-tls
8    issuerRef:
9      # The issuer created previously
10     name: letsencrypt-staging
11     commonName: example.com
12     dnsNames:
13     - example.com
14     - www.example.com
15     acme:
16       config:
17       - dns01:
18         # The provider in the previously created issuer
19         provider: my-clouddns-provider
20       domains:
21       - example.com
22       - www.example.com

```

For more details about Certificates, see [Issuing Certificates](#).

### Amazon Route53

```

route53:
  region: eu-west-1

```

(continues on next page)

(continued from previous page)

```
# optional if ambient credentials are available; see ambient credentials_
↪documentation
accessKeyID: AKIAIOSFODNN7EXAMPLE
secretAccessKeySecretRef:
  name: prod-route53-credentials-secret
  key: secret-access-key
```

Cert-manager requires the following IAM policy.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "route53:GetChange",
      "Resource": "arn:aws:route53::change/*"
    },
    {
      "Effect": "Allow",
      "Action": "route53:ChangeResourceRecordSets",
      "Resource": "arn:aws:route53::hostedzone/*"
    },
    {
      "Effect": "Allow",
      "Action": "route53:ListHostedZonesByName",
      "Resource": "*"
    }
  ]
}
```

The `route53:ListHostedZonesByName` statement can be removed if you specify the optional hosted zone ID (`spec.acme.dns01.providers[].hostedZoneID`) on the Issuer resource. You can further tighten this policy by limiting the hosted zone that cert-manager has access to (replace `arn:aws:route53::hostedzone/*` with `arn:aws:route53::hostedzone/DIKER8JPL21PSA`, for instance).

## DigitalOcean

This provider uses a Kubernetes `Secret` Resource to work. In the following example, the secret will have to be named `digitalocean-dns` and have a subkey `access-token` with the token in it.

To create a Personal Access Token, see [DigitalOcean documentation](https://cloud.digitalocean.com/account/api/tokens/new). Handy direct link: <https://cloud.digitalocean.com/account/api/tokens/new>

```
digitalocean:
  tokenSecretRef:
    name: digitalocean-dns
    key: access-token
```

## RFC-2136

The goal of this document is to provide a configuration overview of the various facilities required to deploy cert-manager against a RFC-2136 compliant DNS server such as BIND named. This capability is also commonly known as “dynamic DNS”.

Unlike the peer of other cert-manager DNS integrations, `named` is a bit of a “Swiss Army Knife” of domain name servers. Over the years, it has been highly optimized to provide maximal vertical scalability for a single node, as well as horizontal scalability with service provider interfaces. This flexibility makes it impossible to go into every possible `named` deployment that a user may run in to though. Instead, this document will try to make sure your server is ready to accept requests from cert-manager using command line tools, then get on to the making the two work together.

## Transaction Signatures TSIG

Dynamic DNS updates are essentially server queries which otherwise might return resource records (RRs). Since DNS servers are commonly exposed to the public internet, being able to push an unauthenticated update to any server that responds to queries would be immediately untenable.

In the eyes of the `named` architects, the generic solution to this problem space was twofold. The first is to require manual enablement of updates at a zone level, such as `example.com`. In a naive network, there is no requirement that zone updates have any security to them, and clients can be configured such that they can provide updates without any authentication. An example of where this is useful is for machines booting using DHCP, in this case the machines know about themselves and the DNS server can be configured to accept updates when they come from the address being configured.

This clearly has limitations in situations such as cert-manager and the DNS-01 challenge. In this environment, a TXT RR must be created after coordination with the ACME server. After negotiating with the ACME server, a the TXT RR that is published on the domain validates that the domain is legitimately engaged with the process of creating a certificate for it. In the bigger picture of DNS, this means that an arbitrary actor (cert-manager, in this case) must be able to add one of these KV mappings to the domain and delete it after the certificate has been issued. `cert-manager` does not have a convenient physical characteristic such as a DHCP allocation to validate it's requests.

For cases like this, we need to be able to sign a request that is being sent to the DNS server. We do that through TSIGs, or Transaction SIGNatures.

## Configuration Step 1 - Set up your DNS server for secure dynamic updates

There are many excellent tutorials on the net that walk through preparing a basic `named` server for dynamic updates:

- <https://www.cyberciti.biz/faq/unix-linux-bind-named-configuring-tsig/>
- <https://tomthorp.me/blog/using-tsig-enable-secure-zone-transfers-between-bind-9x-servers>

More complex `named` deployments will not use text files, but rather may use LDAP or SQL for a database for resource records. An additional wrinkle is metadata configuration, such as for zone metadata like enabling dynamic updates or access control lists (ACLs) for a zone. There are too many configurations to go into here, but you should be able to find the documentation to do so.

Whatever your deployment is, the goal at this stage has nothing to do with cert-manager and everything to do with a tool called `nsupdate` generating updates signed with TSIG. Once this is out of the way, you can attack the cert-manager configuration with far greater confidence.

## Using `nsupdate`

Most paths to configuring BIND `named` will go through using `dnssec-keygen`. This command-line tool generates a named private key that is used for signing TSIG requests. When a request is signed, both the signature and the name of the private key are attached to the request in an unencrypted form. In this manner, when the request is received, the name of the private key can be used to by the recipient to find the private key itself, build a new signature with it, and compare the two for acceptance.

Since there are dozens of ways to have your `named` server misconfigured, we'll use `nsupdate` to test that the server behaves as expected before we get there. [https://debian-administration.org/article/591/Using\\_the\\_dynamic\\_DNS\\_editor\\_nsupdate](https://debian-administration.org/article/591/Using_the_dynamic_DNS_editor_nsupdate) is a solid breakdown of how to use the tool.

To get started, we'll simply run `nsupdate -k <keyID>` where `keyID` is the value returned from `dnssec-keygen`. This will read the key from disk and provide a command prompt to issue commands. In general, we want to write a simple TXT RR and make sure we can delete it.

```
$ nsupdate -k <keyID>
> update add www1.example.com txt testing
> send
> ... test here with ``nslookup``
> update delete www1.example.com txt
> send
> ... test here with ``nslookup``
```

Any failures to write, read or delete the record will mean that `cert-manager` will not be able to do so either, no matter how well it is configured.

### Configuration Step 2 - Set up cert-manager

Now we get to the fun stuff, seeing everything work. Remember that we need to set up the ACME DNS-01 issuer and challenge mechanism as well as the `rfc2136` provider. Since the documentation covers the other parts sufficiently, let's focus on the provider here.

Example:

```
rfc2136:
  nameserver: 1.2.3.4:53
  tsigKeyName: example-com-secret
  tsigAlgorithm: HMACSHA512
  tsigSecretSecretRef:
    name: tsig-secret
    namespace: cert-manager
  key: tsig-secret-key
```

For this example configuration, we'll need the following two commands. The first, on your `named` server generates the key. Note how `example-com-secret` is both in the `tsigKeyName` above and the `dnssec-keygen` command that follows.

```
dnssec-keygen -r /dev/urandom -a HMAC-SHA512 -b 512 -n HOST example-com-secret
```

Also note how the `tsigAlgorithm` is provided in both the configuration and the keygen command. They are listed at <https://github.com/miekg/dns/blob/v1.0.12/tsig.go#L18-L23>.

The second bit of configuration you need on the kubernetes side is to create a secret. Pulling the secret key string from the `<key>.private` file generated above, use the secret in the placeholder below:

```
kubectl -n cert-manager create secret generic tsig-secret --from-literal=tsig-secret-
↪key=<somesecret>
```

Note how the `tsig-secret` and `tsig-secret-key` match the configuration in the `tsigSecretSecretRef` above.

## Rate Limits

The `rfc2136` provider waits until *all* nameservers to in your domain's SOA RR respond with the same result before it contacts Let's Encrypt to complete the challenge process. This is because the challenge server contacts a non-authoritative DNS server that does a recursive query (a query for records it does not maintain locally). If the servers in the SOA do not contain the correct values, it's likely that the non-authoritative server will have bad information as well, causing the request to go against rate limits and eventually locking the process out.

This process is in place to protect users from server misconfigurations creating a more subtle lockout that persists after the server configuration has been repaired.

As documented elsewhere, it is prudent to fully debug configurations using the ACME staging servers before using the production servers. The staging servers have less aggressive rate limits, but the certificates they issue are not signed with a root certificate trusted by browsers.

## What's next?

This configuration so far will actually do nothing. You still have to request a certificate as in `dns-validation`. Once a certificate is requested, the provider will begin processing the request.

## Troubleshooting

- Be sure that you have fully tested the DNS server updates using `nsupdate` first. Ideally, this is done from a pod in the same namespace as the `rfc2136` provider to ensure there are no firewall issues.
- The logs for the `cert-manager` pod are your friend. Additional logs can be generated by adding the `--v=5` argument to the container launch.
- The TSIG key is encoded with `base64`, but the Kubernetes API server also expects that key literals will be decoded before they are stored. In some cases, a key must be double-encoded. (If you've tested using `nsupdate`, it's pretty easy to spot when you are running into this.)
- Pay attention to the refresh time of the zone you are working with. For zones with low traffic, it will not make a significant difference to reduce the refresh time down to about five minutes while getting initial certificates. Once the process is working, the beauty of `cert-manager` is it doesn't matter if a renewal takes hours due to refresh times, it's all automated!
- Compared to the other providers that often use REST APIs to modify DNS RRs, this provider can take a little longer. You can watch `kubectl certificate yourcert` to get a display of what's going on. It's not uncommon for the process to take five minutes in total.

### 3.3.3 HTTP01 Challenge Provider

In order to allow HTTP01 challenges to be solved, we must enable the HTTP01 challenge provider on our Issuer resource.

This is done through setting the `http01` field on the `issuer.spec.acme` stanza. Cert-manager will then attempt to solve ACME HTTP-01 challenges by using Ingress resources

```

1 apiVersion: certmanager.k8s.io/v1alpha1
2 kind: Issuer
3 metadata:
4   name: example-issuer
5 spec:
```

(continues on next page)

(continued from previous page)

```
6  acme:
7    email: user@example.com
8    server: https://acme-staging-v02.api.letsencrypt.org/directory
9    privateKeySecretRef:
10     name: example-issuer-account-key
11   http01: {}
```

---

**Note:** Let's Encrypt does not support issuing wildcard certificates with HTTP-01 challenges. To issue wildcard certificates, you must use the DNS-01 challenge.

---

### How HTTP01 validations work

You can read about how the HTTP01 challenge type works on the [Let's Encrypt challenge types page](#).

### Extra options

The HTTP01 Issuer supports a number of additional options. For full details on the range of options available, read the [reference documentation](#).

### servicePort

In rare cases it might be not possible/desired to use NodePort as type for the http01 challenge response service, e.g. because of Kubernetes limit restrictions. To define which Kubernetes service type to use during challenge response specify the following http01 config:

```
http01:
  # Valid values are ClusterIP and NodePort
  serviceType: ClusterIP
```

By default type NodePort will be used when you don't set http01 or when you set serviceType to an empty string. Normally there's no need to change this.

## 3.3.4 Debugging failing Orders

This guide is still in the process of being written.

Please check the [Order resource reference docs](#) to understand how to debug ACME Orders & Challenges when you are having issues.

### Common problems

---

**Todo:** fill in this section with a new header for each issue that we see commonly occurring.

---

## 3.4 Backing up and restoring

If you need to uninstall cert-manager, or transfer your installation to a new cluster, you can backup all of cert-manager's configuration in order to later re-install.

### 3.4.1 Backing up

To backup all of your cert-manager configuration resources, run:

```
kubectl get -o yaml \
  --all-namespaces \
  issuer,clusterissuer,certificates,orders,challenges > cert-manager-backup.yaml
```

If you are transferring data to a new cluster, you may also need to copy across additional Secret resources that are referenced by your configured Issuers, such as:

#### CA Issuers

- The root CA Secret referenced by `issuer.spec.ca.secretName`

#### Vault Issuers

- The token authentication Secret referenced by `issuer.spec.vault.auth.tokenSecretRef`
- The approle configuration Secret referenced by `issuer.spec.vault.auth.appRole.secretRef`

#### ACME Issuers

- The ACME account private key Secret referenced by `issuer.acme.privateKeySecretRef`
- Any Secrets referenced by DNS providers configured under the `issuer.acme.dns01.providers` field

### 3.4.2 Restoring

In order to restore your configuration, you can simply `kubectl apply` the files created above after installing cert-manager.

```
kubectl apply -f cert-manager-backup.yaml
```

If you have migrated from an old cluster, you will need to make sure to run a similar `kubectl apply` command to restore your Secret resources too.

## 3.5 Upgrading cert-manager

This section contains information on upgrading cert-manager. It also contains documents detailing breaking changes between cert-manager versions, and information on things to look out for when upgrading.

**Note:** Before performing upgrades of cert-manager, it is advised to take a backup of all your cert-manager resources just in case an issue occurs whilst upgrading. You can read how to backup and restore cert-manager in the [Backing up and restoring](#) guide.

---

### 3.5.1 Upgrading with Helm

If you installed cert-manager using Helm, you can easily upgrade using the Helm CLI.

---

**Note:** Before upgrading, please read the relevant instructions at the links below for your from and to version.

---

Once you have read the relevant upgrading notes and taken any appropriate actions, you can begin the upgrade process like so - replacing `<release_name>` with the name of your Helm release for cert-manager (usually this is `cert-manager`) and replacing `<version>` with the version number you want to install:

```
# Install the cert-manager CustomResourceDefinition resources before
# upgrading the Helm chart
kubectl apply \
  -f https://raw.githubusercontent.com/jetstack/cert-manager/<version>/deploy/
  ↪manifests/00-crds.yaml

# Ensure the local Helm chart repository cache is up to date
helm repo update

# If you are upgrading from v0.5 or below, you should manually add this
# label to your cert-manager namespace to ensure the `webhook component`_
# can provision correctly.
kubectl label namespace cert-manager certmanager.k8s.io/disable-validation=true

helm upgrade --version <version> <release_name> jetstack/cert-manager
```

This will upgrade you to the latest version of cert-manager, as listed in the [‘Jetstack Helm chart repository’](#).

---

**Note:** You can find out your release name using `helm list | grep cert-manager`.

---

### 3.5.2 Upgrading using static manifests

If you installed cert-manager using the [static deployment manifests](#), you can upgrade them in a similar way to how you first installed them.

---

**Note:** Before upgrading, please read the relevant instructions at the links below for your from and to version.

---

Once you have read the relevant notes and taken any appropriate actions, you can begin the upgrade process like so - replacing `<version>` with the version number you want to install:

```
# If you are upgrading from v0.5 or below, you should manually add this
# label to your cert-manager namespace to ensure the `webhook component`_
# can provision correctly.
kubectl label namespace cert-manager certmanager.k8s.io/disable-validation=true
```

(continues on next page)



(continued from previous page)

```
kubectl apply \  
  -f https://raw.githubusercontent.com/jetstack/cert-manager/<version>/deploy/  
↪manifests/cert-manager.yaml
```

**Note:** If you are running `kubectl v1.12` or below, you will need to add the `--validate=false` flag to your `kubectl apply` command above else you will receive a validation error relating to the `caBundle` field of the `ValidatingWebhookConfiguration` resource. This issue is resolved in Kubernetes 1.13 onwards. More details can be found in [kubernetes/kubernetes#69590](https://github.com/kubernetes/kubernetes/issues/69590).

## Upgrading from v0.2 to v0.3

During the v0.3 release, a number of breaking changes were made that require you to update either deployment configuration and runtime configuration (e.g. Certificate, Issuer and ClusterIssuer resources).

After reading these instructions, you should then proceed to upgrade cert-manager according to your deployment configuration (e.g. using `helm upgrade` if installing via Helm chart, or `kubectl apply` if installing with raw manifests).

A brief summary:

- Supporting resources for ClusterIssuers (e.g. signing CA certificates, or ACME account private keys) will now be stored in the same namespace as cert-manager, instead of kube-system in previous versions (#329, @munnerz)
- Switch to ConfigMaps instead of Endpoints for leader election (#327, @mikebryant)
- Removing support for ACMEv1 in favour of ACMEv2 (#309, @munnerz)
- Removing ingress-shim and compiling it into cert-manager itself (#502, @munnerz)
- Change to the default behaviour of ingress-shim. It now generates Certificates with the `ingressClass` field set instead of the `ingress` field. This will mean users of ingress controllers that assign a single IP to a single Ingress (e.g. the GCE ingress controller) will no longer work without adding a new annotation to your ingress resource.

## Supporting resources for ClusterIssuers moving into the cert-manager namespace

In the past, the cert-manager controller was hard coded to look for supplemental resources, such as Secrets containing DNS provider credentials, in the kube-system namespace.

We now store these resources in the same namespace as the cert-manager pod itself runs within.

When upgrading, you should make sure to move any of these supplemental resources into the cert-manager deployment namespace, or otherwise deploy cert-manager into kube-system itself.

You can also change the 'cluster resource namespace' when deploying cert-manager:

With the helm chart: `--set clusterResourceNamespace=kube-system`.

Or if using the static deployment manifests, by adding the `--cluster-resource-namespace` flag to the `args` field of the cert-manager container.

### Switch to ConfigMaps instead of Endpoints for leader election

cert-manager-controller performs leader election to allow you to run ‘hot standby’ replicas of cert-manager.

In the past, we used Endpoint resources to perform this election. The new best practice is to use ConfigMap resources in order to reduce API overhead in large clusters.

As such, v0.3 switches us to use ConfigMap resources for leader election.

During the upgrade, you should first scale your cert-manager-controller deployment to 0 to ensure no other replicas of cert-manager are running when the new v0.3 deployment starts:

```
kubectl scale --namespace <deployment-namespace> --replicas=0 deployment <cert-  
↪manager-deployment-name>
```

### Removing support for ACMEv1 in favour of ACMEv2

The ACME v2 specification is now in production with Let’s Encrypt. In order to support this new spec, which includes support for wildcard certificates, we have removed support for the v1 protocol altogether.

If you have any ACME Issuer or ClusterIssuer resources, you should update the server fields of these to the new ACMEv2 endpoints.

For example, if you have a Let’s Encrypt production issuer, you should update the server URL:

```
apiVersion: certmanager.k8s.io/v1alpha1  
kind: Issuer  
...  
spec:  
  acme:  
    # server: https://acme-v01.api.letsencrypt.org/directory  
    server: https://acme-v02.api.letsencrypt.org/directory # we switch 'v01' to 'v02'
```

### Removing ingress-shim and compiling it into cert-manager itself

In v0.3 we removed the ingress-shim component and instead now compile in its functionality into the main cert-manager binary.

This change also introduces a change to the way you configure default Issuers and ClusterIssuers at deployment time.

The deployment documentation has been updated accordingly, but instead of setting `ingressShim.extraArgs={--default-issuer-name=letsencrypt-pod}` there are now dedicated Helm chart fields:

```
--set ingressShim.defaultIssuerName=letsencrypt-prod \  
--set ingressShim.defaultIssuerKind=ClusterIssuer
```

### Change to the default behaviour of ingress-shim

In the past, when using ingress-shim, we set the `ingress` field on the Certificate resource to trigger cert-manager to edit the specified Ingress resource to solve the challenge.

The alternate option is to set the `ingressClass` field, which causes cert-manager to create temporary Ingress resources to solve the challenge. This behaviour provides better compatibility with ingress controllers like `nginx-ingress`.

In v0.3 we have changed the default behaviour of ingress-shim to set the `ingressClass` field instead of `ingress`. This will cause validations for ingress controllers like `ingress-gce` to fail without additional configuration in your Ingress resources annotations.

Add the follow annotation to your Ingress resources if you are using the GCE ingress controller, in addition to the usual ingress-shim annotation(s):

```
certmanager.k8s.io/acme-http01-edit-in-place: "true"
```

### Upgrading from v0.3 to v0.4

There are no special notes or considerations when upgrading from v0.3 to v0.4.

### Upgrading from v0.4 to v0.5

Version 0.5 of cert-manager introduces a new ‘webhook’ component, which is used by the Kubernetes apiserver to validate our CRD resource types.

This should help in future to reduce errors caused by misconfigured Certificate and Issuer resources.

When upgrading from a previous release using Helm, it is **essential** that you perform one extra step before upgrading.

### Disabling resource validation on the cert-manager namespace

Before upgrading, you should add the `certmanager.k8s.io/disable-validation: "true"` label to the `cert-manager` namespace.

This will allow the system resources that cert-manager requires to bootstrap TLS to be created in its own namespace.

### Upgrading from v0.5 to v0.6

**Warning:** If you are upgrading from a release older than v0.5, please read the [Upgrading from older versions using Helm](#) note at the bottom of this document!

The upgrade process from v0.5 to v0.6 should be fairly seamless for most users. As part of the new release, we have changed how we ship the CustomResourceDefinition resources that cert-manager needs in order to operate (as well as introducing two **new** CRD types).

Depending on the way you have installed cert-manager in the past, your upgrade process will slightly vary:

### Upgrading with the Helm chart

If you have previously deployed cert-manager v0.5 using the Helm installation method, you will now need to perform one extra step before upgrading.

Due to issues with the way Helm handles CRD resources in Helm charts, we have now moved the installation of these resources into a separate YAML manifest that must be installed with `kubectl apply` before upgrading the chart.

You can follow the [regular upgrade guide](#) as usual in order to upgrade from v0.5 to v0.6.

### Upgrading with static manifests

The static manifests have moved into the `deploy/manifests` directory for this release.

We now also no longer ship different manifests for different configurations, in favour of a single `cert-manager.yaml` file which should work for all Kubernetes clusters from Kubernetes v1.9 onwards.

You can follow the *regular upgrade guide* as usual in order to upgrade from v0.5 to v0.6.

### Upgrading from older versions using Helm

If you are upgrading from a version **older than v0.5** and **have installed with Helm**, you will need to perform a fresh installation of cert-manager due to issues with the Helm upgrade process. This will involve the **removal of all cert-manager custom resources**. This **will not** delete the Secret resources being used by your apps.

Before upgrading you will need to:

1. Read and follow the *backup guide* to create a backup of your configuration.
2. Delete the existing cert-manager Helm release (replacing ‘cert-manager’ with the name of your Helm release):

```
# Uninstall the Helm chart
$ helm delete --purge cert-manager

# Ensure the cert-manager CustomResourceDefinition resources do not exist:
$ kubectl delete crd \
  certificates.certmanager.k8s.io \
  issuers.certmanager.k8s.io \
  clusterissuers.certmanager.k8s.io
```

3. Perform a fresh install (as per the *installation guide*):

```
# Install the cert-manager CRDs
$ kubectl apply \
  -f https://raw.githubusercontent.com/jetstack/cert-manager/release-0.7/deploy/
  ↪ manifests/00-crds.yaml

# Update helm repository cache
$ helm repo update

# Install cert-manager
$ helm install \
  --name cert-manager \
  --namespace cert-manager \
  --version v0.6.6 \
  stable/cert-manager
```

4. Follow the steps in the *restore guide* to restore your configuration.
5. Verify that your Issuers and Certificate resources are ‘Ready’:

```
$ kubectl get clusterissuer,issuer,certificates --all-namespaces
NAMESPACE   NAME                                     READY   SECRET
↪ AGE
cert-manager cert-manager-webhook-ca                 True    cert-manager-webhook-ca
↪ 1m
cert-manager cert-manager-webhook-webhook-tls       True    cert-manager-webhook-
↪webhook-tls 1m
```

(continues on next page)

(continued from previous page)

example-com	example-com-tls	True	example-com-tls	↵
↵	11s			



---

## Reference documentation

---

This section contains detailed reference documentation about cert-manager's types and how it operates. It also includes some simple example configurations in order to help users activate advanced functionality of cert-manager.

Step by step user guides and tutorials can be found in the *tutorials* section.

### 4.1 Certificates

cert-manager has the concept of 'Certificates' that define a desired X.509 certificate. A Certificate is a namespaced resource that references an Issuer or ClusterIssuer for information on how to obtain the certificate.

A simple Certificate could be defined as:

```
1 apiVersion: certmanager.k8s.io/v1alpha1
2 kind: Certificate
3 metadata:
4   name: acme-crt
5 spec:
6   secretName: acme-crt-secret
7   dnsNames:
8     - foo.example.com
9     - bar.example.com
10  acme:
11    config:
12      - http01:
13        ingressClass: nginx
14        domains:
15          - foo.example.com
16          - bar.example.com
17  issuerRef:
18    name: letsencrypt-prod
19    # We can reference ClusterIssuers by changing the kind here.
20    # The default value is Issuer (i.e. a locally namespaced Issuer)
21  kind: Issuer
```

This Certificate will tell cert-manager to attempt to use the Issuer named `letsencrypt-prod` to obtain a certificate key pair for the `foo.example.com` and `bar.example.com` domains. If successful, the resulting key and certificate will be stored in a secret named `acme-crt-secret` with keys of `tls.key` and `tls.crt` respectively. This secret will live in the same namespace as the `Certificate` resource.

The `dnsNames` field specifies a list of [Subject Alternative Names](#) to be associated with the certificate. If the `commonName` field is omitted, the first element in the list will be the common name.

The referenced Issuer must exist in the same namespace as the Certificate. A Certificate can alternatively reference a `ClusterIssuer` which is non-namespaced.

### 4.1.1 Certificate Duration and Renewal Window

cert-manager Certificate resources also support custom validity durations and renewal windows.

**Important:** The backend service implementation can choose to generate a certificate with a different validity period than what is requested in the issuer.

Although the duration and renewal periods are specified on the Certificate resources, the corresponding Issuer or ClusterIssuer must support this.

The table below shows the support state of the different backend services used by issuer types:

Issuer	Description
ACME	Only 'renewBefore' supported
CA	Fully supported
Vault	Fully supported (although the requested duration must be lower than the configured Vault role's TTL)
Self Signed	Fully supported
Venafi	Fully supported

The default duration for all certificates is 90 days and the default renewal windows is 30 days. This means that certificates are considered valid for 3 months and renewal will be attempted within 1 month of expiration.

The `duration` and `renewBefore` parameters must be given in the [golang parseDuration string format](#).

### Example Usage

Here an example of an issuer specifying the duration and renewal window.

The certificate from the previous section is extended with a validity period of 24 hours and to begin trying to renew 12 hours before the certificate expiration.

```
1  apiVersion: certmanager.k8s.io/v1alpha1
2  kind: Certificate
3  metadata:
4    name: example
5  spec:
6    secretName: example-tls
7    duration: 24h
8    renewBefore: 12h
9    dnsNames:
10   - foo.example.com
11   - bar.example.com
12   issuerRef:
```

(continues on next page)



(continued from previous page)

```

13   name: my-internal-ca
14   kind: Issuer

```

## 4.2 Orders

Order resources are used by the ACME issuer to manage the lifecycle of an ACME ‘order’ for a signed TLS certificate.

When a Certificate resource is created that references an ACME issuer, cert-manager will create an Order resource in order to obtain a signed certificate.

As an end-user, you will never need to manually create an Order resource. Once created, an Order cannot be changed. Instead, a new Order resource must be created.

### 4.2.1 Debugging Order resources

In order to debug why a Certificate isn’t being issued, we can first run `kubectl describe` on the Certificate resource we’re having issues with:

```

$ kubectl describe certificate example-com

...
Events:
  Type     Reason          Age   From           Message
  ----     -
  Normal   Generated       1m    cert-manager   Generated new private key
  Normal   OrderCreated    1m    cert-manager   Created Order resource "example-com-
↪1217431265"

```

We can see here that Certificate controller has created an Order resource to request a new certificate from the ACME server.

Orders are a useful source of information when debugging failures issuing ACME certificates. By running `kubectl describe order` on a particular order, information can be gleaned about failures in the process:

```

$ kubectl describe order example-com-1248919344

...
Reason:
State:      pending
URL:        https://acme-v02.api.letsencrypt.org/acme/order/41123272/265506123
Events:
  Type     Reason    Age   From           Message
  ----     -
  Normal   Created   1m    cert-manager   Created Challenge resource "example-com-
↪1217431265-0" for domain "test1.example.com"
  Normal   Created   1m    cert-manager   Created Challenge resource "example-com-
↪1217431265-1" for domain "test2.example.com"

```

Here we can see that cert-manager has created two Challenge resources in order to fulfil the requirements of the ACME order to obtain a signed certificate.

You can then go on to run `kubectl describe challenge example-com-1217431265-0` to further debug the progress of the Order.

Once an Order is successful, you should see an event like the following:

```
$ kubectl describe order example-com-1248919344

...
Reason:
State:          valid
URL:           https://acme-v02.api.letsencrypt.org/acme/order/41123272/265506123
Events:
  Type      Reason      Age   From           Message
  ----      -
  Normal    Created     72s   cert-manager   Created Challenge resource "example-com-
↪1217431265-0" for domain "test1.example.com"
  Normal    Created     72s   cert-manager   Created Challenge resource "example-com-
↪1217431265-1" for domain "test2.example.com"
  Normal    OrderValid  4s    cert-manager   Order completed successfully
```

If the Order is not completing successfully, you can debug the challenges for the Order by running `kubectl describe` on the Challenge resource.

For more information on debugging Challenge resources, read the [challenge reference docs](#).

## 4.3 Challenges

Challenge resources are used by the ACME issuer to manage the lifecycle of an ACME ‘challenge’ that must be completed in order to complete an ‘authorization’ for a single DNS name/identifier.

When an **Order** resource is created, the order controller will create Challenge resources for each DNS name that is being authorized with the ACME server.

As an end-user, you will never need to manually create a Challenge resource. Once created, a Challenge cannot be changed. Instead, a new Challenge resource must be created.

### 4.3.1 Challenge lifecycle

After a Challenge resource has been created, it will be initially queued for processing. Processing will not begin until the challenge has been ‘scheduled’ to start. This scheduling process prevents too many challenges being attempted at once, or multiple challenges for the same DNS name being attempted at once. For more information on how challenges are scheduled, read the [challenge scheduling](#) section.

Once a challenge has been scheduled, it will first be ‘synced’ with the ACME server in order to determine its current state. If the challenge is already valid, its ‘state’ will be updated to ‘valid’, and also set `status.processing = false` to ‘unschedule’ itself.

If the challenge is still ‘pending’, the challenge controller will ‘present’ the challenge using the configured solver, one of HTTP01 or DNS01. Once the challenge has been ‘presented’, it will set `status.presented=true`.

Once ‘presented’, the challenge controller will perform a ‘self check’ to ensure that the challenge has ‘propagated’ (i.e. the authoritative DNS servers have been updated to respond correctly, or the changes to the ingress resources have been observed and in-use by the ingress controller).

If the self check fails, cert-manager will retry the self check with a fixed 10 second retry interval. Challenges that do not ever complete the self check will continue retrying until the user intervenes.

Once the self check is passing, the ACME ‘authorization’ associated with this challenge will be ‘accepted’ (TODO: add link to accepting challenges section of ACME spec).

The final state of the authorization after accepting it will be copied across to the Challenge’s `status.state` field, as well as the ‘error reason’ if an error occurred whilst the ACME server attempted to validate the challenge.

Once a Challenge has entered the `valid`, `invalid`, `expired` or `revoked` state, it will set `status.processing=false` to prevent any further processing of the ACME challenge, and to allow another challenge to be scheduled if there is a backlog of challenges to complete.

### 4.3.2 Challenge scheduling

Instead of attempting to process all challenges at once, challenges are ‘scheduled’ by cert-manager.

This scheduler applies a cap on the maximum number of simultaneous challenges as well as disallows two challenges for the same DNS name and solver type (`http-01` or `dns-01`) to be completed at once.

The maximum number of challenges that can be processed at a time is 60 as of [ddff78](#).

### 4.3.3 Debugging Challenge resources

In order to determine why an ACME Certificate is not being issued, we can debug using the ‘Challenge’ resources that cert-manager has created.

In order to determine which Challenge is failing, you can run `kubectl get challenges`:

```
$ kubectl get challenges
```

NAME	AGE	STATE	DOMAIN	REASON
example-com-1217431265-0	22s	pending	example.com	Waiting <b>for</b> dns-01 challenge propagation

This shows that the challenge has been presented using the DNS01 solver successfully and now cert-manager is waiting for the ‘self check’ to pass.

You can get more information about the challenge by using `kubectl describe`:

```
$ kubectl describe challenge example-com-1217431265-0
```

```
...
Status:
  Presented:  true
  Processing: true
  Reason:    Waiting for dns-01 challenge propagation
  State:     pending
Events:
  Type     Reason      Age   From           Message
  ----     -
  Normal   Started     19s   cert-manager   Challenge scheduled for processing
  Normal   Presented   16s   cert-manager   Presented challenge using dns-01 challenge mechanism
```

Progress about the state of each challenge will be recorded either as Events or on the Challenge’s status block (as shown above).

### 4.3.4 Troubleshooting failing challenges

---

**Todo:** add section describing common issues and resolutions when challenges are failing

---

## 4.4 Issuers

Issuers (and *ClusterIssuers*) represent a certificate authority from which signed x509 certificates can be obtained, such as Let's Encrypt. You will need at least one Issuer or ClusterIssuer in order to begin issuing certificates within your cluster.

An example of an Issuer type is ACME. A simple ACME issuer could be defined as:

```

1 apiVersion: certmanager.k8s.io/v1alpha1
2 kind: Issuer
3 metadata:
4   name: letsencrypt-prod
5   namespace: edge-services
6 spec:
7   acme:
8     # The ACME server URL
9     server: https://acme-v02.api.letsencrypt.org/directory
10    # Email address used for ACME registration
11    email: user@example.com
12    # Name of a secret used to store the ACME account private key
13    privateKeySecretRef:
14      name: letsencrypt-prod
15    # Enable HTTP01 validations
16    http01: {}

```

This is the simplest of ACME issuers - it specifies no DNS-01 challenge providers. HTTP-01 validation can be performed through using Ingress resources by enabling the HTTP-01 challenge mechanism (with the `http01: {}` field). More information on configuring ACME Issuers can be found [here](#).

### 4.4.1 Namespacing

An Issuer is a namespaced resource, and it is not possible to issue certificates from an Issuer in a different namespace. This means you will need to create an Issuer in each namespace you wish to obtain Certificates in.

If you want to create a single issuer than can be consumed in multiple namespaces, you should consider creating a *ClusterIssuer* resource. This is almost identical to the Issuer resource, however is non-namespaced and so it can be used to issue Certificates across all namespaces.

### 4.4.2 Ambient Credentials

Some API clients are able to infer credentials to use from the environment they run within. Notably, this includes cloud instance-metadata stores and environment variables. In cert-manager, the term 'ambient credentials' refers to such credentials. They are always drawn from the environment of the 'cert-manager-controller' deployment.

#### Example Usage

If cert-manager is deployed in an environment with ambient AWS credentials, such as with a [kube2iam](#) role, the following ClusterIssuer would make use of those credentials to perform the ACME DNS01 challenge with route53.

```

1 apiVersion: certmanager.k8s.io/v1alpha1
2 kind: ClusterIssuer
3 metadata:
4   name: letsencrypt-prod

```

(continues on next page)

(continued from previous page)

```

5 spec:
6   acme:
7     server: https://acme-v02.api.letsencrypt.org/directory
8     email: user@example.com
9     privateKeySecretRef:
10      name: letsencrypt-prod
11     dns01:
12       providers:
13         - name: route53
14           route53:
15             region: us-east-1

```

It is important to note that the `route53` section does not specify any `accessKeyID` or `secretAccessKeySecretRef`. If either of these are specified, ambient credentials will not be used.

### When are Ambient Credentials used

Ambient credentials are supported for the 'route53' ACME DNS01 challenge provider.

They will only be used if no credentials are supplied, even if the supplied credentials are invalid.

By default, ambient credentials may be used by ClusterIssuers, but not regular issuers. The `--issuer-ambient-credentials` and `--cluster-issuer-ambient-credentials=false` flags on cert-manager may be used to override this behavior.

Note that ambient credentials are disabled for regular Issuers by default to ensure unprivileged users who may create issuers cannot issue certificates using any credentials cert-manager incidentally has access to.

### 4.4.3 Supported Issuer types

cert-manager has been designed to support pluggable Issuer backends. The currently supported Issuer types are:

Name	Description
<i>ACME</i>	Supports obtaining certificates from an ACME server, validating with HTTP01 or DNS01
<i>CA</i>	Supports issuing certificates using a simple signing keypair, stored in a Secret in the Kubernetes API server
<i>Vault</i>	Supports issuing certificates using HashiCorp Vault.
<i>Self signed</i>	Supports issuing self signed certificates
<i>Venafi</i>	Supports issuing certificates from Venafi Cloud & TPP

Each Issuer resource is of one, and only one type. The type of an Issuer is inferred by which field it specifies in its spec, such as `spec.acme` for the ACME issuer, or `spec.ca` for the CA based issuer.

## 4.5 ClusterIssuers

ClusterIssuers are a resource type similar to *Issuers*. They are specified in exactly the same way, but they do not belong to a single namespace and can be referenced by Certificate resources from multiple different namespaces.

They are particularly useful when you want to provide the ability to obtain certificates from a central authority (e.g. Letsencrypt, or your internal CA) and you run single-tenant clusters.

The docs for Issuer resources apply equally to ClusterIssuers.

You can specify a ClusterIssuer resource by changing the `kind` attribute of an Issuer to `ClusterIssuer`, and removing the `metadata.namespace` attribute:

```
apiVersion: certmanager.k8s.io/v1alpha1
kind: ClusterIssuer
metadata:
  name: letsencrypt-prod
spec:
  ...
```

We can then reference a ClusterIssuer from a Certificate resource by setting the `spec.issuerRef.kind` field to `ClusterIssuer`:

```
apiVersion: certmanager.k8s.io/v1alpha1
kind: Certificate
metadata:
  name: my-certificate
  namespace: my-namespace
spec:
  secretName: my-certificate-secret
  issuerRef:
    name: letsencrypt-prod
    kind: ClusterIssuer
  ...
```

When referencing a `Secret` resource in `ClusterIssuer` resources (eg `apiKeySecretRef`) the `Secret` needs to be in the same namespace as the `cert-manager` controller pod. You can optionally override this by using the `--cluster-resource-namespace` argument to the controller.

For more information on configuring Issuer resources, see the [Issuers](#) reference documentation.

## 4.6 API documentation

## 5.1 Develop with minikube

Minikube is a tool to quickly provision a local Kubernetes cluster on many platforms. It can be used to test and develop cert-manager. This guide will walk you through getting started using Minikube for development.

### 5.1.1 Start minikube

First, run minikube, and configure your local kubectl command to work with minikube; minikube typically does this automatically.

```
# Check your locally installed minikube version
$ minikube version
minikube version: v0.25.0

# Start a local cluster
# If using Minikube v0.25.0 or older:
$ minikube start --extra-config=apiserver.Authorization.Mode=RBAC
# Otherwise:
$ minikube start

# Verify it works. This should output a local apiserver IP
$ kubectl cluster-info

# Create a cluster role binding so Tiller has cluster-admin access rights
$ kubectl create clusterrolebinding default-admin --clusterrole=cluster-admin --
↪serviceaccount=kube-system:default

# Install helm
$ helm init
```

### 5.1.2 Install local development tools

You will need the following tools to build cert-manager:

- Bazel
- Docker (and enable for non-root user)

These instructions have only been tested on Linux; Windows and MacOS may require further changes.

If you need to add dependencies, you will additionally need:

- Git
- Mercurial

You can then run `bazel run //hack:update-deps` to regenerate any dependencies, and `bazel build :images` to build the docker images.

### 5.1.3 Build a dev version of cert-manager

```
# Configure your local docker client to use the minikube docker daemon
$ eval "$(minikube docker-env)"

# Build cert-manager binaries and docker images. Full output omitted for brevity
$ make build
Successfully tagged quay.io/jetstack/cert-manager-controller:canary
```

### 5.1.4 Deploy that version with helm

```
# Install custom resources before running helm
$ kubectl apply -f deploy/manifests/00-crds.yaml

# IMPORTANT: if you are deploying into a namespace that already exists,
# you MUST ensure the namespace has an additional label on it in order for
# the deployment to succeed
$ kubectl label namespace <deployment-namespace> certmanager.k8s.io/disable-
↪validation="true"

# Install our freshly built cert-manager image
$ helm install \
  --set image.tag=canary \
  --set image.pullPolicy=Never \
  --name cert-manager \
  ./deploy/charts/cert-manager
```

From here, you should be able to do whatever manual testing or development you wish to.

### 5.1.5 Deploy a new version

In general, upgrading can be done simply by running `make build`, and then deleting the deployed pod using `kubectl delete pod`.

However, if you make changes to the helm chart or wish to change the controller's arguments, such as to change the logging level, you may also update it with the following:



```
helm upgrade \
  cert-manager \
  --reuse-values \
  --set extraArgs="{-v=5}"
  --set image.tag=build
  ./contrib/charts/cert-manager
```

## 5.2 Running end-to-end tests

cert-manager has an end-to-end test suite that verifies functionality against a real Kubernetes cluster.

This document explains how you can run the end-to-end tests yourself. This is useful when you have added or changed functionality in cert-manager and want to verify the software still works as expected.

### 5.2.1 Requirements

Currently, a number of tools **must** be installed on your machine in order to run the tests:

- `bazel` - As with all other development, Bazel is required to actually build the project as well as end-to-end test framework. Bazel will also retrieve appropriate versions of any other dependencies depending on what ‘target’ you choose to run.
- `docker` - We provision a whole Kubernetes cluster within Docker, and so an up to date version of Docker must be installed. The oldest Docker version we have tested is 17.09.
- `kubectl` - If you are running the tests on Linux, this step is technically not required. For non-Linux hosts (i.e. OSX), you will need to ensure you have a relatively new version of kubectl available on your PATH.
- An internet connection - tests require access to DNS, and optionally Cloudflare APIs (if a Cloudflare API token is provided).

Bazel, Docker and Kubectl should be installed through your preferred means.

### 5.2.2 Run end-to-end tests

You can run the end-to-end tests by executing the following:

```
./hack/ci/run-e2e-kind.sh
```

The full suite may take up to 10 minutes to run. You can monitor output of this command to track progress.

## 5.3 Contributing DNS01 providers

### 5.3.1 WARNING

Because of the overwhelming number of PRs for new DNS providers, We’re changing how we handle the DNS01 contributions. See [this post](#) on the mailing list for more information.

Steps to add a `FOODNS` DNS-01 provider:

1. Create a new package under `pkg/issuer/acme/dns/foodns`. This is where all the code to interact with the DNS providers API will live.

2. Implement functions to match the solver interface (`Present`, `CleanUp` and `Timeout`). Use an existing provider for reference. Most of the cert-manager providers are based off <https://github.com/xenolf/lego>, so if lego supports the DNS provider you want to add, it's fairly easy to copy it over and make modifications to fit with the cert-manager codebase. Examples of the changes required:
  - replace uses of `github.com/xenolf/lego/acme` with `github.com/jetstack/cert-manager/pkg/issuer/acme/dns/util`.
  - replace uses of `github.com/xenolf/lego/log` with `github.com/golang/glog`.
  - remove references to `github.com/xenolf/lego/platform/config/env`. cert-manager does not use environment variables for internal configuration, so calls to this package should not be required.
3. Add unit test coverage for this package.
4. Add your provider configuration types to the API (located in `pkg/apis/certmanager/v1alpha1/types.go`) and regenerate code (run `./hack/update-codegen.sh`). New API types should have an associated short documentation string, which is added to the reference API documentation (run `./hack/update-reference-docs-dockerized.sh` to update the API documentation).
5. Register the provider in `pkg/issuer/acme/dns`:
  - The constructor for the provider needs adding to `dnsProviderConstructors`,
  - `solverForIssuerProvider` must be updated to handle retrieving any information for the new provider (for example, fetching credentials from a secret) and constructing a new instance of the provider.
6. Add coverage for the provider to `pkg/issuer/acme/dns/dns_test.go`.
7. Add example configuration for the new provider to `docs/reference/issuers/acme/dns01/index.rst`. The more information here the better, this example and corresponding documentation should inform users how to use and configure this backend, as well as mentioning any nuances with using this particular provider.
8. Test your provider out against a real account, and make sure you can issue a Certificate.
9. Submit your new provider to cert-manager!

Things to watch out for:

- Assume that at any point the cert-manager process may restart. Make sure values required for operations like `CleanUp` are not solely stored in memory.

## 5.4 DCO Sign off

All authors to the project retain copyright to their work. However, to ensure that they are only submitting work that they have rights to, we are requiring everyone to acknowledge this by signing their work.

Any copyright notices in this repo should specify the authors as “the Jetstack cert-manager contributors”.

To sign your work, just add a line like this at the end of your commit message:

```
Signed-off-by: Joe Bloggs <joe@example.com>
```

This can easily be done with the `--signoff` option to `git commit`. You can also mass sign-off a whole PR with `git rebase --signoff master`, replacing `master` with the branch you are creating a pull request again if not `master`.

By doing this you state that you certify the following (from <https://developercertificate.org/>):

```
Developer Certificate of Origin
Version 1.1
```

```
Copyright (C) 2004, 2006 The Linux Foundation and its contributors.
1 Letterman Drive
Suite D4700
San Francisco, CA, 94129
```

```
Everyone is permitted to copy and distribute verbatim copies of this
license document, but changing it is not allowed.
```

```
Developer's Certificate of Origin 1.1
```

```
By making a contribution to this project, I certify that:
```

- (a) The contribution was created in whole or in part by me and I have the right to submit it under the open source license indicated in the file; or
- (b) The contribution is based upon previous work that, to the best of my knowledge, is covered under an appropriate open source license and I have the right under that license to submit that work with modifications, whether created in whole or in part by me, under the same open source license (unless I am permitted to submit under a different license), as indicated in the file; or
- (c) The contribution was provided directly to me by some other person who certified (a), (b) or (c) and I have not modified it.
- (d) I understand and agree that this project and the contribution are public and that a record of the contribution (including all personal information I submit with it, including my sign-off) is maintained indefinitely and may be redistributed consistent with this project or the open source license(s) involved.

## 5.5 Release process

This document aims to outline the process that should be followed for cutting a new release of cert-manager.

### 5.5.1 Minor releases

A minor release is a backwards-compatible 'feature' release. It can contain new features and bugfixes.

#### Release schedule

We aim to cut a new minor release once per month. The rough goals for each release are outlined as part of a GitHub milestone. We cut a release even if some of these goals are missed, in order to keep up release velocity.

### Process

---

**Note:** This process document is WIP and may be incomplete

---

The process for cutting a minor release is as follows:

1. Ensure upgrading document exists in docs/admin/upgrading
2. Create a new release branch (e.g. `release-0.5`)
3. Push it to the `jetstack/cert-manager` repository
4. Create a pull-request updating the Helm chart version and merge it:
  - Update contrib/charts/cert-manager/README.md
  - Update contrib/charts/cert-manager/Chart.yaml
  - Update contrib/charts/cert-manager/values.yaml
  - Update contrib/charts/cert-manager/requirements.yaml
  - Update contrib/charts/cert-manager/webhook/Chart.yaml
  - Update contrib/charts/cert-manager/webhook/values.yaml
  - Run `helm dep update` in the contrib/charts/cert-manager directory
  - Run `./hack/update-deploy-gen.sh` in the root of the repository
5. Gather release notes since the previous release:
  - Run `relnotes -repo cert-manager -owner jetstack release-0.5`
  - Write up appropriate notes, similar to previous releases
6. Submit the Helm chart changes to the upstream `helm/charts` repo:

```
TARGET_REPO_REMOTE=upstream \  
SOURCE_REPO_REMOTE=upstream \  
SOURCE_REPO_REF=release-0.5 \  
GITHUB_USER=munnerz \  
./hack/create-chart-pr.sh
```
7. Iterate on review feedback (hopefully this will be minimal) and submit changes to `master` of cert-manager, performing a rebase of release-x.y and re-run of the `create-chart-pr.sh` script after each cycle to gather more feedback.
8. Create a new tag taken from the release branch, e.g. `v0.5.0`.

### 5.5.2 Patch releases

A patch release contains critical bugfixes for the project. They are managed on an ad-hoc basis, and should only be required when critical bugs/regressions are found in the release.

We will only perform patch release for the **current** version of cert-manager.

Once a new minor release has been cut, we will stop providing patches for the version before it.

## Release schedule

Patch releases are cut on an ad-hoc basis, depending on recent activity on the release branch.

## Process

---

**Note:** This process document is WIP and may be incomplete

---

Bugs that need to be fixed in a patch release should be cherry picked into the appropriate release branch using the `./hack/cherry-pick-pr.sh` script in this repository.

The process for cutting a patch release is as follows:

1. Create a PR against the **release branch** to bump the chart version:
  - Update contrib/charts/cert-manager/README.md
  - Update contrib/charts/cert-manager/Chart.yaml
  - Update contrib/charts/cert-manager/values.yaml
  - Update contrib/charts/cert-manager/requirements.yaml
  - Update contrib/charts/cert-manager/webhook/Chart.yaml
  - Update contrib/charts/cert-manager/webhook/values.yaml
  - Run `helm dep update` in the contrib/charts/cert-manager directory
  - Run `./hack/update-deploy-gen.sh` in the root of the repository
2. Submit the Helm chart changes to the upstream `helm/charts` repo:

```
TARGET_REPO_REMOTE=upstream \
SOURCE_REPO_REMOTE=upstream \
SOURCE_REPO_REF=release-0.5 \
GITHUB_USER=munnerz \
./hack/create-chart-pr.sh
```

3. Iterate on review feedback (hopefully this will be minimal) and submit changes to `master` of cert-manager, performing a rebase of release-x.y and re-run of the `create-chart-pr.sh` script after each cycle to gather more feedback.
4. Gather release notes since the previous release:
  - Run `relnotes -repo cert-manager -owner jetstack release-0.5`
  - Write up appropriate notes, similar to previous patch releases
5. Create a new tag taken from the release branch, e.g. `v0.5.1`.

## 5.6 Generating Documentation

The documentation is generated from [reStructured Text](#) by [Sphinx](#) (via [Read The Docs](#)). If you're unfamiliar with [reStructured Text](#), the files typically have the extension `.rst`. You can find more details in the [reStructured Text Basics](#).

### 5.6.1 Installation instructions

To install the sphinx tools, you'll need `python` (and `pip`) installed.:

```
.. code-block: shell
```

```
pip install --user -r requirements.txt
```

### 5.6.2 Generating documentation locally

You can generate the documentation locally with the following command:

This will create documentation in the `_build` directory which you can open with your browser.

Note that you do not need to add these files to your git client, as *Read The Docs* will generate the HTML on the fly.