

目录

目录	1
Kubernetes迁移	3
1.1 迁移背景	3
1.2 迁移流程	3
2.1 应用及配置迁移方案概述	3
2.2 使用须知	3
2.3 迁移过程	4
2.3.1 环境准备	4
2.3.1.1 下载安装velero客户端	4
2.3.1.2 创建对象存储bucket	4
2.3.1.3 获取ak/sk	4
2.3.1.4 配置集群的config信息	4
2.3.1.5 安装velero服务	4
2.3.1.6 验证	5
2.3.2 自建kubernetes集群备份	5
2.3.2.1 设置 annotate	5
2.3.2.2 创建backup	5
2.3.3 金山云kubernetes集群恢复	5
2.3.3.1 设置restic helper	5
2.3.3.2 创建 restore	6
2.3.3.3 验证	6
2.3.4 清理删除velero服务（如需删除重新安装时使用）	6
Traefik-ingress部署与使用	6
使用Ingress前置条件	6
创建测试应用	8
Ingress配置策略	9
同一域名下，不同的URL的路径转发到不同服务上	9
不同的域名转发到不同的服务	10
Nginx-ingress部署与使用	10
Nginx-ingress controller部署	10
创建测试应用	18
Ingress配置策略	20
同一域名下，不同的URL的路径转发到不同服务上	20
不同的域名转发到不同的服务	20
HTTPS访问	21
准备证书	21
创建Secret资源	22
创建Ingress规则	22
通过Nginx-ingress实现灰度发布	22
前提条件	22
Ingress配置策略	22
部署示例	23
创建测试应用	23
配置Ingress	24
基于服务权重进行流量切分	24
验证访问情况	25
在KCE集群中对Pod进行带宽限速	25
备注	25
使用方式	25
验证	25

如何选择Containerd和Docker	26
容器运行时简介	26
如何选择Containerd和Docker?	26
Containerd和Docker常用命令对比	26
镜像相关功能	26
容器相关功能	26
POD相关功能	27
调用链对比	27

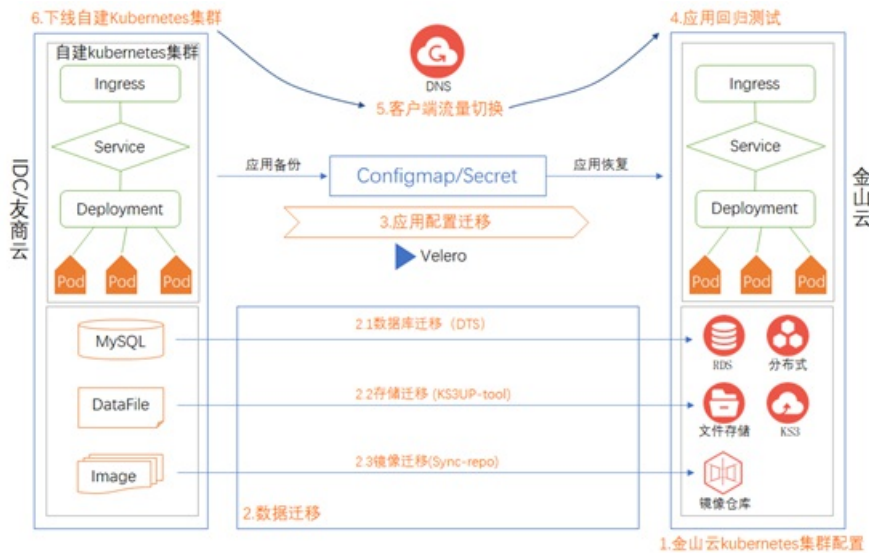
Kubernetes迁移

1.1 迁移背景

进入云架构时代以来，随着公有云成为越来越多客户的选择，云迁移成为必不可少的服务。

在当下，容器具有其跨平台、资源利用率高、敏捷、具有可持续部署和测试等诸多优势，被越来越多的企业接纳。容器结合公有云的弹性和云网络的优势，可谓是完美的结合，越来越多的企业会考虑在云上使用容器集群，本文就以企业自建kubernetes上云迁移至金山云kubernetes集群为例，进行容器迁移上云的方案介绍。

1.2 迁移流程



在迁移前，需确定数据传输方式，如果采用专线或VPN传输，需提前做好专线施工、网络调试，本文主要介绍源端和目标端网络连通后的迁移过程。具体过程介绍如下：

1. 在金山云部署kubernetes集群，并进行集群资源配置。
2. 数据迁移：数据主要包括数据库数据、文件存储、容器镜像等。
 - 对于数据库数据迁移，金山云提供数据传输服务（DTS），通过DTS可实现全量、增量同步。
 - 对于文件存储迁移，金山云提供KS3Up-tool工具，可实现PB级文件数据迁移至金山云对象存储（KS3）上，KS3Up-tool支持将金山云KS3、阿里云OSS、腾讯云COS、百度云、七牛云和 AWS S3 设置为迁移源。
 - 对于容器镜像迁移，如果使用的是Harbor，Harbor提供了迁移能力，可以实现自动迁移到金山云容器镜像仓库；如果镜像规模较小，也可以通过pull/push进行迁移。
3. 应用及配置迁移，该部分迁移主要通过velero工具实现，金山云通过开源velero工具进行二次研发，可实现源端集群配置顺利迁移至金山云kubernetes集群。关于Velero的详细介绍：<https://velero.io/>。
4. 迁移后进行应用回归测试，验证相关服务是否启动运行正常。
5. 灰度切量，完成业务流量逐步切换至金山云容器集群。
6. 业务在新环境运行稳定后，可逐步下线源端集群，完成迁移。

2.1 应用及配置迁移方案概述

具有容器服务集群备份或者迁移需求的场景下，一般考采用Velero集成Restic工具实现，该工具适用以下场景：

- 集群升级前，做集群备份。
- 为保障集群高可靠，做定期备份。
- 开发、测试、生产环境之间，通过备份实现集群的迁移。
- 外部Kubernetes集群迁入至金山云容器服务。

Velero是一个提供 Kubernetes 集群和持久卷的备份、迁移以及灾难恢复等的开源工具。金山云开发了针对Velero的插件以及Velero集成Restic方案，可以实现容器服务的备份、迁移。本文主要介绍Velero集成Restic方案实现Kubernetes集群迁移。Velero支持集成Restic工具来备份和还原存储卷，Restic除了支持块存储之外，还支持更多类型的存储，比如 NFS, Emptydir、Local以及其他不支持快照的任何存储类型。当前Restic不支持HostPath类型的存储，支持Local类型的PV存储卷。

2.2 使用须知

- 在做迁移使用时，建议Kubernetes同版本下迁移(未来Velero新版本会支持跨版本迁移)。
- 推荐排除备份velero和kube-system命名空间下的资源，因为kube-system下都是系统服务不需要备份；velero命名空间下是部署的 velero 服务，也不需要备份。
- Velero的Restic集成需要Kubernetes [MountPropagation功能]，该功能在Kubernetes v1.10.0和更高版本中默认启用。

- 跨厂商迁移使用Velero的Restic集成，由于不同的云厂商，后端的存储基础设施是不一样的，因此在金山云kubernetes集群恢复时，需要先创建一个相同名称的StorageClass来屏蔽对底层存储基础设施差异的感知。

2.3 迁移过程

2.3.1 环境准备

请参考以下步骤在自建kubernetes集群及金山云kubernetes集群中部署velero。

2.3.1.1 下载安装velero客户端

```
wget https://ks3-cn-beijing.ksyun.com/velero/velero-v1.2.0-linux-amd64.tar.gz
```

解压：

```
tar -zxvf velero-v1.2.0-linux-amd64.tar.gz
```

将 velero二进制文件移到PATH环境变量定义的目录下：

```
cp velero /bin
```

验证是否安装成功：

```
velero -h
```

2.3.1.2. 创建对象存储bucket

金山云KS3上创建2个bucket。

2.3.1.3. 获取ak/sk

获取ak/sk，ak/sk的主要目的是实现对金山云对象存储、块存储、块存储快照的操作权限。

2.3.1.4. 配置集群的config信息

连接本地集群无需单独指定，如果在本地执行目标端velero安装，则进行配置。容器服务的config文件可以通过[集群管理 > 集群详情页 > 集群基本信息 > 获取集群config](#)获得。

说明：velero默认会从\$HOME/.kube目录下查找文件名为 config 的文件，如果将在目标端安装，则需要指定 config的路径。命令中添加--kubeconfig，即为config的路径。

2.3.1.5. 安装velero服务

创建credentials-velero文件并设置ak/sk值。

```
vim /data/credentials-velero
```

```
[default]
```

```
ksyun_access_key_id = <your access_key>  
ksyun_access_key_secret = <your secret_key>
```

velero服务安装：

```
velero install --provider ksyun \  
  --image hub.kce.ksyun.com/ksyun/velero:v1.3.0-beta.2 \  
  --plugins hub.kce.ksyun.com/ksyun/velero-plugin-ksyun:v1.2.0 \  
  --kubeconfig <KUBECONFIG> \  
  --bucket <YOUR_BUCKET> \  
  --backup-location-config region=<YOUR_REGION>,resticRepoPrefix=<resticRepoPrefix> \  
  --secret-file ./credentials-velero \  
  --use-volume-snapshots=false \  
  --use-restic
```

运行安装示例，注：示例中的两处bucket不能一样：

```
velero install --provider ksyun \  
  --image hub.kce.ksyun.com/ksyun/velero:v1.3.0-beta.2 \  
  --plugins hub.kce.ksyun.com/ksyun/velero-plugin-ksyun:v1.2.0 \  
  --kubeconfig /data/soukubeconfig \  
  --bucket test-b \  
  --backup-location-config region=cn-beijing-6,resticRepoPrefix=ks3:ks3-cn-beijing.ksyun.com/test-a \  
  --secret-file /data/credentials-velero \  
  --use-volume-snapshots=false \  
  --use-restic
```

```
--use-restic
```

2.3.1.6. 验证

执行：

```
kubectl logs deployment/velero -n velero
```

查看是否有错误。执行 `kubectl get pod -n velero -o wide` 查看 namespace 为 `velero` 的 pod 是否都启动并 `running` 正常。当 `restic-xxx pod` 显示 `CrashLoopBackOff`，需要修改其中的存储卷位置信息：

```
kubectl get daemonset -n velero
```

执行：

```
kubectl edit ds restic -n velero
```

把 `volumes` 下的 `kubelet` 目录配置由 `/var/lib/kubelet/pods` 改成 `/data/kubelet/pods`，再次看 pod 运行是否正常。

2.3.2. 自建kubernetes集群备份

2.3.2.1. 设置 annotate

使用 `restic` 备份时，需要明确指定需要备份的存储卷 pod。velero 通过 pod 的 `annotation` 来过滤需要备份存储卷的 pod。对于需要备份存储卷的 pod，我们需要执行如下命令设置 `annotate`。

```
kubectl -n YOUR_POD_NAMESPACE annotate pod/YOUR_POD_NAME backup.velero.io/backup-volumes=YOUR_VOLUME_NAME_1,YOUR_VOLUME_NAME_2,...
```

示例：

```
kubectl get pod -o yaml
```

查看 `volumes` 的名字。

```
kubectl -n default annotate pod/nginx-7bb7cd8db5-c9fjf backup.velero.io/backup-volumes=default-token-bs7q8
```

注：`-n` 指定 namespace，`backup-volumes=volumes` 中的 name。

2.3.2.2. 创建backup

```
velero backup create NAME OPTIONS...
```

示例：

```
velero backup create nginxbk20200313 --wait
```

2.3.3. 金山云kubernetes集群恢复

2.3.3.1. 设置restic helper

恢复资源时，会启动一个 `restic helper` 容器来帮助数据的恢复，可以通过 `configmap` 来自定义配置该容器的配置，如需要配置则需要在执行恢复操作之前先部署如下 `ConfigMap`。

```
apiVersion: v1
kind: ConfigMap
metadata:
  # any name can be used; Velero uses the labels (below)
  # to identify it rather than the name
  name: restic-restore-action-config
  # must be in the velero namespace
  namespace: velero
  # the below labels should be used verbatim in your
  # ConfigMap.
  labels:
    # this value-less label identifies the ConfigMap as
    # config for a plugin (i.e. the built-in restic restore
    # item action plugin)
    velero.io/plugin-config: ""
    # this label identifies the name and kind of plugin
    # that this ConfigMap is for.
    velero.io/restic: RestoreItemAction
data:
  # The value for "image" can either include a tag or not;
  # if the tag is *not* included, the tag from the main Velero
```

```
# image will automatically be used.
image: hub.kce.ksyun.com/ksyun/velero-restic-restore-helper:v1.2.0

# "cpuRequest" sets the request.cpu value on the restic init containers during restore.
# If not set, it will default to "100m". A value of "0" is treated as unbounded.
cpuRequest: 200m

# "memRequest" sets the request.memory value on the restic init containers during restore.
# If not set, it will default to "128Mi". A value of "0" is treated as unbounded.
memRequest: 128Mi

# "cpuLimit" sets the request.cpu value on the restic init containers during restore.
# If not set, it will default to "100m". A value of "0" is treated as unbounded.
cpuLimit: 200m

# "memLimit" sets the request.memory value on the restic init containers during restore.
# If not set, it will default to "128Mi". A value of "0" is treated as unbounded.
memLimit: 128Mi
```

2.3.3.2. 创建 restore

```
velero restore create --from-backup BACKUP_NAME OPTIONS...
```

示例：

```
velero restore create restore-20200313 --from-backup bk20200313 --wait
velero restore get restore-20200313
```

注： a、不指定create名，系统会自动生成，建议按照自己的规范命名，后面方便查看恢复的结果。 b、此处验证的是完全备份和完全恢复，恢复过程可能会提示部分失败，原因为目标端已经存在，可忽略。

2.3.3.3. 验证

```
kubectl get pod -n default -o wide
```

发现已经恢复并运行，可以通过访问pod应用进行验证。至此kubernetes应用及配置迁移完成。

2.3.4. 清理删除velero服务（如需删除重新安装时使用）

```
kubectl delete namespace/velero clusterrolebinding/velero
kubectl delete crds -l component=velero
```

Traefik-ingress 部署与使用

Ingress是kubernetes集群中授权入站连接到达集群服务的规则集合，您可以通过配置转发规则，实现不同 URL 可以访问到集群内不同的 Service，以实现HTTP层的业务路由机制。

使用Ingress前置条件

为了使Ingress正常工作，集群内必须部署Ingress Controller，以实现为后端所有的service提供一个统一的入口。这里我们使用[Traefik](#)作为集群内的Ingress Controller，traefik部署的YAML如下：

```
kind: ClusterRole
apiVersion: rbac.authorization.k8s.io/v1beta1
metadata:
  name: traefik-ingress-controller
rules:
  - apiGroups:
    - ""
    resources:
    - services
    - endpoints
    - secrets
    verbs:
    - get
    - list
    - watch
  - apiGroups:
    - extensions
    resources:
    - ingresses
    verbs:
    - get
    - list
    - watch
---
```

```
kind: ClusterRoleBinding
apiVersion: rbac.authorization.k8s.io/v1beta1
metadata:
  name: traefik-ingress-controller
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: traefik-ingress-controller
subjects:
- kind: ServiceAccount
  name: traefik-ingress-controller
  namespace: kube-system
---
```

```
apiVersion: v1
kind: ServiceAccount
metadata:
  name: traefik-ingress-controller
  namespace: kube-system
---
```

```
kind: DaemonSet
apiVersion: apps/v1
metadata:
  name: traefik-ingress-controller
  namespace: kube-system
  labels:
    k8s-app: traefik-ingress-lb
spec:
  selector:
    matchLabels:
      k8s-app: traefik-ingress-lb
      name: traefik-ingress-lb
  template:
    metadata:
      labels:
        k8s-app: traefik-ingress-lb
        name: traefik-ingress-lb
    spec:
      nodeSelector:
        kubernetes.io/role: "node"
      tolerations:
      - operator: Exists
      serviceAccountName: traefik-ingress-controller
      terminationGracePeriodSeconds: 60
      containers:
      - image: hub.kce.ksyun.com/ksyun/traefik:v1.6.5-mp
        name: traefik-ingress-lb
        securityContext:
          capabilities:
            drop:
            - ALL
            add:
            - NET_BIND_SERVICE
          args:
            - --api
            - --kubernetes
            - --logLevel=INFO
            - --entryPoints=Name:https Address::443 TLS
            - --entryPoints=Name:http Address::80
            - --defaultentrypoints=https,http
        ---
```

```
kind: Service
apiVersion: v1
metadata:
  name: traefik-ingress-service
  namespace: kube-system
spec:
  selector:
    k8s-app: traefik-ingress-lb
  type: LoadBalancer
  ports:
    - protocol: TCP
      port: 80
      name: web
    - protocol: TCP
      port: 443
      name: tls
    - protocol: TCP
      port: 8080
      name: admin
```

为了让traefik服务在集群外可达，我们这里将traefik-ingress-controller对应的Service的访问类型设置为LoadBalancer。

查看traefik的部署情况：

```
[root@vm10-0-33-13 ~]# kubectl get ds -n kube-system | grep traefik
traefik-ingress-controller      2          2          2          2          2          kubernetes.io/role=node    3m16s
```

查看对应的service：

```
[root@vm10-0-33-13 ~]# kubectl get svc -n kube-system | grep traefik
traefik-ingress-service    LoadBalancer    10.254.67.8      120.92.123.155    80:32676/TCP, 443:31720/TCP, 8080:31840/TCP    105m
```

这里，traefik-ingress-controller服务通过金山云的负载均衡暴露到公网，从这里看到它同时启动了80、8080和443三个端口，80和443对应的服务端口，8080 对应的 UI 端口，用户可以通过LB的IP:8080访问traefik的UI界面。

创建测试应用

以下创建两个应用，用于测试。

hello-world.yaml如下：

```
apiVersion: extensions/v1beta1
kind: Deployment
metadata:
  name: hello-world
spec:
  replicas: 1
  template:
    metadata:
      labels:
        app: hello-world
    spec:
      containers:
        - name: hello-world
          image: hub.kce.ksyun.com/kingsoft/hello-world:latest
---
```

```
apiVersion: v1
kind: Service
metadata:
  labels:
    app: hello-world
  name: hello-world-svc
spec:
  ports:
    - port: 80
      protocol: TCP
      targetPort: 80
  selector:
    app: hello-world
  type: ClusterIP
```

hello-k8s.yaml如下：

```
apiVersion: extensions/v1beta1
kind: Deployment
metadata:
  name: hello-k8s
spec:
  replicas: 1
  template:
    metadata:
      labels:
        app: hello-k8s
    spec:
      containers:
        - name: hello-k8s
          image: hub.kce.ksyun.com/kingsoft/hello-k8s:latest
---
```

```
apiVersion: v1
kind: Service
metadata:
  labels:
    app: hello-k8s
  name: hello-k8s-svc
spec:
  ports:
    - port: 8080
      protocol: TCP
      targetPort: 8080
```



```
selector:
  app: hello-k8s
  type: ClusterIP
```

创建对应的deploy和服务:

```
[root@vm10-0-33-13 hello]# kubectl create -f hello-k8s.yaml
deployment.extensions/hello-k8s created
service/hello-k8s-svc created
```

```
[root@vm10-0-33-13 hello]# kubectl create -f hello-world.yaml
deployment.extensions/hello-world created
service/hello-world-svc created
```

```
[root@vm10-0-33-13 hello]# kubectl get deploy
NAME           DESIRED   CURRENT   UP-TO-DATE   AVAILABLE   AGE
hello-k8s      1         1         1             1           5m2s
hello-world    1         1         1             1           4m50s
```

```
[root@vm10-0-33-13 hello]# kubectl get svc
NAME           TYPE           CLUSTER-IP      EXTERNAL-IP   PORT(S)    AGE
hello-k8s-svc  ClusterIP      10.254.131.29   <none>        8080/TCP   5m31s
hello-world-svc ClusterIP      10.254.244.96   <none>        80/TCP     5m19s
kubernetes     ClusterIP      10.254.0.1       <none>        443/TCP    52d
```

Ingress配置策略

为了支持灵活的分发策略，ingress策略可以按照多种分发方式进行配置，下面对几种常见的ingress转发策略简单介绍。

同一域名下，不同的URL的路径转发到不同服务上

这种配置常用于一个网站通过不同的路径提供不同服务的场景。

通过如下的访问配置:

- 对 <http://my.k8s.traefik/hello-k8s> 的访问将被路由到后端名为“hello-k8s-svc”的Service。
- 对 <http://my.k8s.traefik/hello-world> 的访问将被路由到后端名为“hello-world-svc”的Service。

ingress.yaml如下:

```
apiVersion: extensions/v1beta1
kind: Ingress
metadata:
  name: my-k8s-traefik
  annotations:
    kubernetes.io/ingress.class: traefik
    traefik.frontend.rule.type: PathPrefixStrip
spec:
  rules:
    - host: my.k8s.traefik
      http:
        paths:
          - path: /hello-world
            backend:
              serviceName: hello-world-svc
              servicePort: 80
          - path: /hello-k8s
            backend:
              serviceName: hello-k8s-svc
              servicePort: 8080
```

创建ingress策略:

```
[root@vm10-0-33-13 hello]# kubectl create -f ingres.yaml
ingress.extensions/my-k8s-traefik created
```

```
[root@vm10-0-33-13 hello]# kubectl get ingress
NAME           HOSTS           ADDRESS   PORTS   AGE
my-k8s-traefik my.k8s.traefik      80       73s
```

备注:

- 这里我们将自有域名my.k8s.traefik解析到负载均衡的IP。
- 这里我们根据路径来转发，需要指明 rule 为 PathPrefixStrip，配置为 traefik.frontend.rule.type: PathPrefixStrip。

在浏览器的访问验证如下:



不同的域名转发到不同的服务

这种配置常用于一个网站通过不同的域名或者虚拟主机名提供不同的服务的场景。

通过如下的访问配置：

- 对 <http://traefik.hello.k8s> 的访问将被路由到后端名为“hello-k8s-svc”的Service。
- 对 <http://traefik.hello.world> 的访问将被路由到后端名为“hello-world-svc”的Service。

ingress2.yaml如下：

```
apiVersion: extensions/v1beta1
kind: Ingress
metadata:
  name: my-k8s-traefik-1
  annotations:
    kubernetes.io/ingress.class: traefik
spec:
  rules:
  - host: traefik.hello.k8s
    http:
      paths:
      - path: /
        backend:
          serviceName: hello-k8s-svc
          servicePort: 8080
  - host: traefik.hello.world
    http:
      paths:
      - path: /
        backend:
          serviceName: hello-world-svc
          servicePort: 80
```

```
[root@vm10-0-33-13 hello]# kubectl create -f ingress2.yaml
ingress.extensions/my-k8s-traefik-1 created
```

```
[root@vm10-0-33-13 hello]# kubectl get ingress
NAME                                HOSTS                                ADDRESS  PORTS  AGE
my-k8s-traefik-1                   traefik.hello.k8s,traefik.hello.world  80      21s
```

在浏览器的访问验证如下：



我们可以通过traefik的UI来查看上面配置的ingress规则，如图：



更多traefik的特性，请参考[Kubernetes Ingress Controller](#)。

Nginx-ingress 部署与使用

本文将为您介绍基于金山云容器服务搭建 Nginx-ingress 服务的操作指南。

Nginx-ingress controller 部署

使用Nginx-ingress服务的前提是在集群内部署nginx-ingress controller，controller部署的nginx-ingress-install.yaml如下：

```
apiVersion: v1
kind: ServiceAccount
metadata:
  labels:
    helm.sh/chart: ingress-nginx-3.10.1
    app.kubernetes.io/name: ingress-nginx
    app.kubernetes.io/instance: ingress-nginx
    app.kubernetes.io/version: 0.41.2
    app.kubernetes.io/managed-by: Helm
```

```
  app.kubernetes.io/component: controller
  name: ingress-nginx
  namespace: kube-system
---
apiVersion: v1
kind: ConfigMap
metadata:
  labels:
    helm.sh/chart: ingress-nginx-3.10.1
    app.kubernetes.io/name: ingress-nginx
    app.kubernetes.io/instance: ingress-nginx
    app.kubernetes.io/version: 0.41.2
    app.kubernetes.io/managed-by: Helm
    app.kubernetes.io/component: controller
  name: ingress-nginx-controller
  namespace: kube-system
data:
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
  labels:
    helm.sh/chart: ingress-nginx-3.10.1
    app.kubernetes.io/name: ingress-nginx
    app.kubernetes.io/instance: ingress-nginx
    app.kubernetes.io/version: 0.41.2
    app.kubernetes.io/managed-by: Helm
  name: ingress-nginx
rules:
- apiGroups:
  - ''
  resources:
  - configmaps
  - endpoints
  - nodes
  - pods
  - secrets
  verbs:
  - list
  - watch
- apiGroups:
  - ''
  resources:
  - nodes
  verbs:
  - get
- apiGroups:
  - ''
  resources:
  - services
  verbs:
  - get
  - list
  - update
  - watch
- apiGroups:
  - extensions
  - networking.k8s.io # k8s 1.14+
  resources:
  - ingresses
  verbs:
  - get
  - list
  - watch
- apiGroups:
  - ''
  resources:
  - events
  verbs:
  - create
  - patch
- apiGroups:
  - extensions
  - networking.k8s.io # k8s 1.14+
  resources:
  - ingresses/status
  verbs:
  - update
- apiGroups:
  - networking.k8s.io # k8s 1.14+
  resources:
  - ingressclasses
  verbs:
  - get
```

```
- list
- watch
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  labels:
    helm.sh/chart: ingress-nginx-3.10.1
    app.kubernetes.io/name: ingress-nginx
    app.kubernetes.io/instance: ingress-nginx
    app.kubernetes.io/version: 0.41.2
    app.kubernetes.io/managed-by: Helm
  name: ingress-nginx
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: ingress-nginx
subjects:
- kind: ServiceAccount
  name: ingress-nginx
  namespace: kube-system
---
apiVersion: rbac.authorization.k8s.io/v1
kind: Role
metadata:
  labels:
    helm.sh/chart: ingress-nginx-3.10.1
    app.kubernetes.io/name: ingress-nginx
    app.kubernetes.io/instance: ingress-nginx
    app.kubernetes.io/version: 0.41.2
    app.kubernetes.io/managed-by: Helm
    app.kubernetes.io/component: controller
  name: ingress-nginx
  namespace: kube-system
rules:
- apiGroups:
  - ''
  resources:
  - namespaces
  verbs:
  - get
- apiGroups:
  - ''
  resources:
  - configmaps
  - pods
  - secrets
  - endpoints
  verbs:
  - get
  - list
  - watch
- apiGroups:
  - ''
  resources:
  - services
  verbs:
  - get
  - list
  - update
  - watch
- apiGroups:
  - extensions
  - networking.k8s.io # k8s 1.14+
  resources:
  - ingresses
  verbs:
  - get
  - list
  - watch
- apiGroups:
  - extensions
  - networking.k8s.io # k8s 1.14+
  resources:
  - ingresses/status
  verbs:
  - update
- apiGroups:
  - networking.k8s.io # k8s 1.14+
  resources:
  - ingressclasses
  verbs:
  - get
  - list
```

```

    - watch
- apiGroups:
  - ''
  resources:
    - configmaps
  resourceNames:
    - ingress-controller-leader-nginx
  verbs:
    - get
    - update
- apiGroups:
  - ''
  resources:
    - configmaps
  verbs:
    - create
- apiGroups:
  - ''
  resources:
    - endpoints
  verbs:
    - create
    - get
    - update
- apiGroups:
  - ''
  resources:
    - events
  verbs:
    - create
    - patch
---
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  labels:
    helm.sh/chart: ingress-nginx-3.10.1
    app.kubernetes.io/name: ingress-nginx
    app.kubernetes.io/instance: ingress-nginx
    app.kubernetes.io/version: 0.41.2
    app.kubernetes.io/managed-by: Helm
    app.kubernetes.io/component: controller
  name: ingress-nginx
  namespace: kube-system
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: Role
  name: ingress-nginx
subjects:
- kind: ServiceAccount
  name: ingress-nginx
  namespace: kube-system
---
apiVersion: v1
kind: Service
metadata:
  labels:
    helm.sh/chart: ingress-nginx-3.10.1
    app.kubernetes.io/name: ingress-nginx
    app.kubernetes.io/instance: ingress-nginx
    app.kubernetes.io/version: 0.41.2
    app.kubernetes.io/managed-by: Helm
    app.kubernetes.io/component: controller
  name: ingress-nginx-controller-admission
  namespace: kube-system
spec:
  type: ClusterIP
  ports:
    - name: https-webhook
      port: 443
      targetPort: webhook
  selector:
    app.kubernetes.io/name: ingress-nginx
    app.kubernetes.io/instance: ingress-nginx
    app.kubernetes.io/component: controller
---
apiVersion: v1
kind: Service
metadata:
  labels:
    helm.sh/chart: ingress-nginx-3.10.1
    app.kubernetes.io/name: ingress-nginx
    app.kubernetes.io/instance: ingress-nginx
    app.kubernetes.io/version: 0.41.2

```

```
  app.kubernetes.io/managed-by: Helm
  app.kubernetes.io/component: controller
name: ingress-nginx-controller
namespace: kube-system
spec:
  type: LoadBalancer
  ports:
    - name: http
      port: 80
      protocol: TCP
      targetPort: http
    - name: https
      port: 443
      protocol: TCP
      targetPort: https
  selector:
    app.kubernetes.io/name: ingress-nginx
    app.kubernetes.io/instance: ingress-nginx
    app.kubernetes.io/component: controller
---
apiVersion: apps/v1
kind: Deployment
metadata:
  labels:
    helm.sh/chart: ingress-nginx-3.10.1
    app.kubernetes.io/name: ingress-nginx
    app.kubernetes.io/instance: ingress-nginx
    app.kubernetes.io/version: 0.41.2
    app.kubernetes.io/managed-by: Helm
    app.kubernetes.io/component: controller
  name: ingress-nginx-controller
  namespace: kube-system
spec:
  selector:
    matchLabels:
      app.kubernetes.io/name: ingress-nginx
      app.kubernetes.io/instance: ingress-nginx
      app.kubernetes.io/component: controller
  revisionHistoryLimit: 10
  minReadySeconds: 0
  template:
    metadata:
      labels:
        app.kubernetes.io/name: ingress-nginx
        app.kubernetes.io/instance: ingress-nginx
        app.kubernetes.io/component: controller
    spec:
      dnsPolicy: ClusterFirst
      containers:
        - name: controller
          image: hub.kce.ksyun.com/yangxiaohua/ingress-nginx:v0.41.2
          imagePullPolicy: IfNotPresent
          lifecycle:
            preStop:
              exec:
                command:
                  - /wait-shutdown
          args:
            - /nginx-ingress-controller
            - --election-id=ingress-controller-leader
            - --ingress-class=nginx
            - --configmap=$(POD_NAMESPACE)/ingress-nginx-controller
            - --validating-webhook=:8443
            - --validating-webhook-certificate=/usr/local/certificates/cert
            - --validating-webhook-key=/usr/local/certificates/key
      securityContext:
        capabilities:
          drop:
            - ALL
          add:
            - NET_BIND_SERVICE
        runAsUser: 101
        allowPrivilegeEscalation: true
      env:
        - name: POD_NAME
          valueFrom:
            fieldRef:
              fieldPath: metadata.name
        - name: POD_NAMESPACE
          valueFrom:
            fieldRef:
              fieldPath: metadata.namespace
        - name: LD_PRELOAD
          value: /usr/local/lib/libmimalloc.so
```

```
livenessProbe:
  httpGet:
    path: /healthz
    port: 10254
    scheme: HTTP
  initialDelaySeconds: 10
  periodSeconds: 10
  timeoutSeconds: 1
  successThreshold: 1
  failureThreshold: 5
readinessProbe:
  httpGet:
    path: /healthz
    port: 10254
    scheme: HTTP
  initialDelaySeconds: 10
  periodSeconds: 10
  timeoutSeconds: 1
  successThreshold: 1
  failureThreshold: 3
ports:
- name: http
  containerPort: 80
  protocol: TCP
- name: https
  containerPort: 443
  protocol: TCP
- name: webhook
  containerPort: 8443
  protocol: TCP
volumeMounts:
- name: webhook-cert
  mountPath: /usr/local/certificates/
  readOnly: true
resources:
  requests:
    cpu: 100m
    memory: 90Mi
nodeSelector:
  kubernetes.io/os: linux
serviceAccountName: ingress-nginx
terminationGracePeriodSeconds: 300
volumes:
- name: webhook-cert
  secret:
    secretName: ingress-nginx-admission
---
apiVersion: admissionregistration.k8s.io/v1
kind: ValidatingWebhookConfiguration
metadata:
  labels:
    helm.sh/chart: ingress-nginx-3.10.1
    app.kubernetes.io/name: ingress-nginx
    app.kubernetes.io/instance: ingress-nginx
    app.kubernetes.io/version: 0.41.2
    app.kubernetes.io/managed-by: Helm
    app.kubernetes.io/component: admission-webhook
  name: ingress-nginx-admission
webhooks:
- name: validate.nginx.ingress.kubernetes.io
  matchPolicy: Equivalent
  rules:
    - apiGroups:
      - networking.k8s.io
      apiVersions:
      - v1beta1
      operations:
      - CREATE
      - UPDATE
      resources:
      - ingresses
  failurePolicy: Fail
  sideEffects: None
  admissionReviewVersions:
  - v1
  - v1beta1
  clientConfig:
    service:
      namespace: kube-system
      name: ingress-nginx-controller-admission
      path: /networking/v1beta1/ingresses
---
apiVersion: v1
kind: ServiceAccount
```

```
metadata:
  name: ingress-nginx-admission
  annotations:
    helm.sh/hook: pre-install,pre-upgrade,post-install,post-upgrade
    helm.sh/hook-delete-policy: before-hook-creation,hook-succeeded
  labels:
    helm.sh/chart: ingress-nginx-3.10.1
    app.kubernetes.io/name: ingress-nginx
    app.kubernetes.io/instance: ingress-nginx
    app.kubernetes.io/version: 0.41.2
    app.kubernetes.io/managed-by: Helm
    app.kubernetes.io/component: admission-webhook
  namespace: kube-system
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
  name: ingress-nginx-admission
  annotations:
    helm.sh/hook: pre-install,pre-upgrade,post-install,post-upgrade
    helm.sh/hook-delete-policy: before-hook-creation,hook-succeeded
  labels:
    helm.sh/chart: ingress-nginx-3.10.1
    app.kubernetes.io/name: ingress-nginx
    app.kubernetes.io/instance: ingress-nginx
    app.kubernetes.io/version: 0.41.2
    app.kubernetes.io/managed-by: Helm
    app.kubernetes.io/component: admission-webhook
rules:
- apiGroups:
  - admissionregistration.k8s.io
  resources:
  - validatingwebhookconfigurations
  verbs:
  - get
  - update
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: ingress-nginx-admission
  annotations:
    helm.sh/hook: pre-install,pre-upgrade,post-install,post-upgrade
    helm.sh/hook-delete-policy: before-hook-creation,hook-succeeded
  labels:
    helm.sh/chart: ingress-nginx-3.10.1
    app.kubernetes.io/name: ingress-nginx
    app.kubernetes.io/instance: ingress-nginx
    app.kubernetes.io/version: 0.41.2
    app.kubernetes.io/managed-by: Helm
    app.kubernetes.io/component: admission-webhook
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: ingress-nginx-admission
subjects:
- kind: ServiceAccount
  name: ingress-nginx-admission
  namespace: kube-system
---
apiVersion: rbac.authorization.k8s.io/v1
kind: Role
metadata:
  name: ingress-nginx-admission
  annotations:
    helm.sh/hook: pre-install,pre-upgrade,post-install,post-upgrade
    helm.sh/hook-delete-policy: before-hook-creation,hook-succeeded
  labels:
    helm.sh/chart: ingress-nginx-3.10.1
    app.kubernetes.io/name: ingress-nginx
    app.kubernetes.io/instance: ingress-nginx
    app.kubernetes.io/version: 0.41.2
    app.kubernetes.io/managed-by: Helm
    app.kubernetes.io/component: admission-webhook
  namespace: kube-system
rules:
- apiGroups:
  - ''
  resources:
  - secrets
  verbs:
  - get
  - create
---
```



```
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: ingress-nginx-admission
  annotations:
    helm.sh/hook: pre-install,pre-upgrade,post-install,post-upgrade
    helm.sh/hook-delete-policy: before-hook-creation,hook-succeeded
  labels:
    helm.sh/chart: ingress-nginx-3.10.1
    app.kubernetes.io/name: ingress-nginx
    app.kubernetes.io/instance: ingress-nginx
    app.kubernetes.io/version: 0.41.2
    app.kubernetes.io/managed-by: Helm
    app.kubernetes.io/component: admission-webhook
  namespace: kube-system
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: Role
  name: ingress-nginx-admission
subjects:
- kind: ServiceAccount
  name: ingress-nginx-admission
  namespace: kube-system
---
```

```
apiVersion: batch/v1
kind: Job
metadata:
  name: ingress-nginx-admission-create
  annotations:
    helm.sh/hook: pre-install,pre-upgrade
    helm.sh/hook-delete-policy: before-hook-creation,hook-succeeded
  labels:
    helm.sh/chart: ingress-nginx-3.10.1
    app.kubernetes.io/name: ingress-nginx
    app.kubernetes.io/instance: ingress-nginx
    app.kubernetes.io/version: 0.41.2
    app.kubernetes.io/managed-by: Helm
    app.kubernetes.io/component: admission-webhook
  namespace: kube-system
spec:
  template:
    metadata:
      name: ingress-nginx-admission-create
      labels:
        helm.sh/chart: ingress-nginx-3.10.1
        app.kubernetes.io/name: ingress-nginx
        app.kubernetes.io/instance: ingress-nginx
        app.kubernetes.io/version: 0.41.2
        app.kubernetes.io/managed-by: Helm
        app.kubernetes.io/component: admission-webhook
    spec:
      containers:
      - name: create
        image: hub.kce.ksyun.com/yangxiaohua/kube-webhook-certgen:v1.5.0
        imagePullPolicy: IfNotPresent
        args:
        - create
        - --host=ingress-nginx-controller-admission,ingress-nginx-controller-admission.$(POD_NAMESPACE).svc
        - --namespace=$(POD_NAMESPACE)
        - --secret-name=ingress-nginx-admission
        env:
        - name: POD_NAMESPACE
          valueFrom:
            fieldRef:
              fieldPath: metadata.namespace
        restartPolicy: OnFailure
        serviceAccountName: ingress-nginx-admission
        securityContext:
          runAsNonRoot: true
          runAsUser: 2000
      ---
apiVersion: batch/v1
kind: Job
metadata:
  name: ingress-nginx-admission-patch
  annotations:
    helm.sh/hook: post-install,post-upgrade
    helm.sh/hook-delete-policy: before-hook-creation,hook-succeeded
  labels:
    helm.sh/chart: ingress-nginx-3.10.1
    app.kubernetes.io/name: ingress-nginx
    app.kubernetes.io/instance: ingress-nginx
    app.kubernetes.io/version: 0.41.2
    app.kubernetes.io/managed-by: Helm
```

```

  app.kubernetes.io/component: admission-webhook
  namespace: kube-system
spec:
  template:
    metadata:
      name: ingress-nginx-admission-patch
      labels:
        helm.sh/chart: ingress-nginx-3.10.1
        app.kubernetes.io/name: ingress-nginx
        app.kubernetes.io/instance: ingress-nginx
        app.kubernetes.io/version: 0.41.2
        app.kubernetes.io/managed-by: Helm
        app.kubernetes.io/component: admission-webhook
    spec:
      containers:
        - name: patch
          image: hub.kce.ksyun.com/yangxiaohua/kube-webhook-certgen:v1.5.0
          imagePullPolicy: IfNotPresent
          args:
            - patch
            - --webhook-name=ingress-nginx-admission
            - --namespace=$(POD_NAMESPACE)
            - --patch-mutating=false
            - --secret-name=ingress-nginx-admission
            - --patch-failure-policy=Fail
          env:
            - name: POD_NAMESPACE
              valueFrom:
                fieldRef:
                  fieldPath: metadata.namespace
      restartPolicy: OnFailure
      serviceAccountName: ingress-nginx-admission
      securityContext:
        runAsNonRoot: true
        runAsUser: 2000

```

执行以下命令安装部署ingress-nginx controller 0.41.2版本:

```
[root@vm10-0-11-170 ~]# kubectl apply -f nginx-ingress-install.yaml
```

#将会在集群中创建如下资源

```

serviceaccount/ingress-nginx created
configmap/ingress-nginx-controller created
clusterrole.rbac.authorization.k8s.io/ingress-nginx created
clusterrolebinding.rbac.authorization.k8s.io/ingress-nginx created
role.rbac.authorization.k8s.io/ingress-nginx created
rolebinding.rbac.authorization.k8s.io/ingress-nginx created
service/ingress-nginx-controller-admission created
service/ingress-nginx-controller created
deployment.apps/ingress-nginx-controller created
validatingwebhookconfiguration.admissionregistration.k8s.io/ingress-nginx-admission created
serviceaccount/ingress-nginx-admission created
clusterrole.rbac.authorization.k8s.io/ingress-nginx-admission created
clusterrolebinding.rbac.authorization.k8s.io/ingress-nginx-admission created
role.rbac.authorization.k8s.io/ingress-nginx-admission created
rolebinding.rbac.authorization.k8s.io/ingress-nginx-admission created
job.batch/ingress-nginx-admission-create created
job.batch/ingress-nginx-admission-patch created

```

获取ingress-nginx-controller服务的IP地址。

注: 本示例中为ingress-nginx-controller创建LoadBalancer类型service, 如需定义LoadBalancer有关更多配置, 可参考[注释列表](#)。

```

[root@vm10-0-11-170 ~]# kubectl get svc -n kube-system | grep ingress-nginx-controller
ingress-nginx-controller      LoadBalancer    10.254.39.165    120.92.xx.xx    80:31004/TCP, 443:30182/TCP    26m
ingress-nginx-controller-admission ClusterIP        10.254.219.22   <none>          443/TCP                        26m

```

通过EXTERNAL-IP (120.92.xx.xx), 外部流量将访问到集群中的ingress-nginx-controller, 进而实现Ingress规则中的路由转发。

创建测试应用

以下创建两个应用, 用于测试。

hello-world.yaml如下:

```

apiVersion: apps/v1
kind: Deployment
metadata:
  name: hello-world

```

```
spec:
  replicas: 1
  selector:
    matchLabels:
      app: hello-world
  template:
    metadata:
      labels:
        app: hello-world
    spec:
      containers:
        - name: hello-world
          image: hub.kce.ksyun.com/kingsoft/hello-world:latest
---
apiVersion: v1
kind: Service
metadata:
  labels:
    app: hello-world
  name: hello-world-svc
spec:
  ports:
    - port: 80
      protocol: TCP
      targetPort: 80
  selector:
    app: hello-world
  type: ClusterIP
```

hello-k8s.yaml如下:

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: hello-k8s
spec:
  replicas: 1
  selector:
    matchLabels:
      app: hello-k8s
  template:
    metadata:
      labels:
        app: hello-k8s
    spec:
      containers:
        - name: hello-k8s
          image: hub.kce.ksyun.com/kingsoft/hello-k8s:latest
---
apiVersion: v1
kind: Service
metadata:
  labels:
    app: hello-k8s
  name: hello-k8s-svc
spec:
  ports:
    - port: 8080
      protocol: TCP
      targetPort: 8080
  selector:
    app: hello-k8s
  type: ClusterIP
```

创建对应的deploy和服务:

```
[root@vm10-0-33-13 hello]# kubectl create -f hello-k8s.yaml
deployment.extensions/hello-k8s created
service/hello-k8s-svc created
```

```
[root@vm10-0-33-13 hello]# kubectl create -f hello-world.yaml
deployment.extensions/hello-world created
service/hello-world-svc created
```

```
[root@vm10-0-33-13 hello]# kubectl get deploy
NAME           DESIRED   CURRENT   UP-TO-DATE   AVAILABLE   AGE
hello-k8s      1         1         1             1           5m2s
hello-world    1         1         1             1           4m50s
```

```
[root@vm10-0-33-13 hello]# kubectl get svc
NAME           TYPE        CLUSTER-IP      EXTERNAL-IP   PORT(S)    AGE
hello-k8s-svc  ClusterIP   10.254.131.29   <none>        8080/TCP   5m31s
hello-world-svc ClusterIP   10.254.244.96   <none>        80/TCP     5m19s
```

kubernetes ClusterIP 10.254.0.1 <none> 443/TCP 52d

Ingress配置策略

为了支持灵活的分发策略，ingress策略可以按照多种分发方式进行配置，下面对几种常见的ingress转发策略简单介绍。

同一域名下，不同的URL的路径转发到不同服务上

这种配置常用于一个网站通过不同的路径提供不同服务的场景。

通过如下的访问配置：

- 对 <http://my.nginx.test/hello-k8s> 的访问将被路由到后端名为“hello-k8s-svc”的Service。
- 对 <http://my.nginx.test/hello-world> 的访问将被路由到后端名为“hello-world-svc”的Service。

ingress.yaml如下：

```
apiVersion: extensions/v1beta1
kind: Ingress
metadata:
  name: nginx-test
  annotations:
    kubernetes.io/ingress.class: nginx
spec:
  rules:
  - host: my.nginx.test
    http:
      paths:
      - path: /hello-world
        backend:
          serviceName: hello-world-svc
          servicePort: 80
      - path: /hello-k8s
        backend:
          serviceName: hello-k8s-svc
          servicePort: 8080
```

创建ingress规则：



```
[root@vm10-0-11-170 ~]# kubectl apply -f ingress.yaml
ingress.extensions/nginx-test created
```

```
[root@vm10-0-11-170 ~]# kubectl get ingress
NAME          HOSTS          ADDRESS          PORTS          AGE
nginx-test    my.nginx.test  10.254.0.1       80             15s
```

备注：

- 这里我们将自有域名my.nginx.test解析到负载均衡的IP。
- Ingress规则与nginx-ingress-controller的对应关系通过注解kubernetes.io/ingress.class: nginx来指定，表示此条ingress规则由nginx-ingress-controller处理。

在浏览器的访问验证如下：



不同的域名转发到不同的服务

这种配置常用于一个网站通过不同的域名或者虚拟主机名提供不同的服务的场景。

通过如下的访问配置：

- 对 <http://nginx.hello.k8s> 的访问将被路由到后端名为“hello-k8s-svc”的Service。
- 对 <http://nginx.hello.world> 的访问将被路由到后端名为“hello-world-svc”的Service。

ingress2.yaml如下:


```
apiVersion: extensions/v1beta1
kind: Ingress
metadata:
  name: nginx-test-2
  annotations:
    kubernetes.io/ingress.class: nginx
spec:
  rules:
    - host: nginx.hello.k8s
      http:
        paths:
          - path: /
            backend:
              serviceName: hello-k8s-svc
              servicePort: 8080
    - host: nginx.hello.world
      http:
        paths:
          - path: /
            backend:
              serviceName: hello-world-svc
              servicePort: 80
```

创建Ingress规则:


```
[root@vm10-0-11-170 ~]# kubectl apply -f ingress2.yaml
ingress.extensions/nginx-test-2 created

[root@vm10-0-11-170 ~]# kubectl get ingress
NAME                HOSTS                ADDRESS      PORTS     AGE
nginx-test-2       nginx.hello.k8s,nginx.hello.world      80         19s
```

在浏览器的访问验证如下:



Hello k8s!



Hello world!

HTTPS访问

当Ingress配置TLS时, 服务将以HTTPS协议的方式对外暴露。

准备证书

这里我们作为测试用例使用自签名证书, 使用如下命令快速创建。

```
[root@vm10-0-11-170 ~]# openssl req -newkey rsa:2048 -nodes -keyout tls.key -x509 -days 365 -out tls.crt
Generating a 2048 bit RSA private key
.....+++
.....+++
writing new private key to 'tls.key'
-----
You are about to be asked to enter information that will be incorporated
into your certificate request.
What you are about to enter is what is called a Distinguished Name or a DN.
There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.
-----
Country Name (2 letter code) [XX]:CN
State or Province Name (full name) []:Beijing
```

```
Locality Name (eg, city) [Default City]:Beijing
Organization Name (eg, company) [Default Company Ltd]:Kingsoft
Organizational Unit Name (eg, section) []:Ksyun
Common Name (eg, your name or your server's hostname) []:nginx.hello.k8s #配置需使用证书的域名
Email Address []:ksyun@kingsoft.com
```

将在当前路径下创建证书(tls.crt)和私钥文件(tls.key)。

创建Secret资源

根据现有证书和私钥创建Secret资源，将会创建一个类型为kubernetes.io/tls的Secret。


```
[root@vm10-0-11-170 ~]# kubectl create secret tls secret-https --key tls.key --cert tls.crt
secret/secret-https created
```

创建Ingress规则

对上述示例中的http访问方式进行升级，实现对nginx.hello.k8s的https访问。ingress-https.yaml如下：

```
apiVersion: extensions/v1beta1
kind: Ingress
metadata:
  name: nginx-test-2
  annotations:
    kubernetes.io/ingress.class: nginx
    nginx.ingress.kubernetes.io/ssl-redirect: "true" #开启重定向功能
spec:
  rules:
  - host: nginx.hello.k8s
    http:
      paths:
      - path: /
        backend:
          serviceName: hello-k8s-svc
          servicePort: 8080
  tls:
  - hosts:
    - nginx.hello.k8s
    secretName: secret-https
```

在浏览器的访问验证如下：



Hello k8s!

更多Nginx ingress controller的特性，请参考[NGINX Ingress Controller](#)。

通过Nginx-ingress实现灰度发布

在对服务进行版本发布或版本升级场景中，常会用到灰度发布、蓝绿发布等发布方式。本文将介绍如何在金山云容器服务中通过Nginx-ingress服务实现应用的灰度发布。

前提条件

集群内完成nginx-ingress-controller的部署，并通过金山云的负载均衡服务暴露到集群外。部署方式可参考[Nginx-ingress 部署与使用](#)。

Ingress配置策略

Ingress controller支持通过配置ingress annotations实现不同场景下的灰度发布和测试。Nginx annotations支持以下四种灰度发布（Canary）规则：

- `nginx.ingress.kubernetes.io/canary-by-header`：基于 Request Header 的流量切分，适用于灰度发布以及 A/B 测试。当 Request Header 设置为 `always`时，请求将会被一直发送到 Canary 版本；当 Request Header 设置为 `never`时，请求不会被发送到 Canary 入口；对于任何其他 Header 值，将忽略 Header，并通过优先级将请求与其他 Canary 规则进行优先级的比较。

- `nginx.ingress.kubernetes.io/canary-by-header-value`: 要匹配的 Request Header 的值, 用于通知 Ingress 将请求路由到 Canary Ingress 中指定的服务。当 Request Header 设置为此值时, 它将被路由到 Canary 入口。该规则允许用户自定义 Request Header 的值, 必须与上一个 annotation (即: `canary-by-header`) 一起使用。
- `nginx.ingress.kubernetes.io/canary-weight`: 基于服务权重的流量切分, 适用于蓝绿部署, 权重范围 0 - 100 按百分比将请求路由到 Canary Ingress 中指定的服务。权重为 0 意味着该金丝雀规则不会向 Canary 入口的服务发送任何请求。权重为 100 意味着所有请求都将被发送到 Canary 入口。
- `nginx.ingress.kubernetes.io/canary-by-cookie`: 基于 Cookie 的流量切分, 适用于灰度发布与 A/B 测试。用于通知 Ingress 将请求路由到 Canary Ingress 中指定的服务的 cookie。当 cookie 值设置为 `always` 时, 它将被路由到 Canary 入口; 当 cookie 值设置为 `never` 时, 请求不会被发送到 Canary 入口; 对于任何其他值, 将忽略 cookie 并将请求与其他金丝雀规则进行优先级的比较。

注意: 当同时使用多个 Canary 规则时, 按如下优先顺序进行排序: `canary-by-header` - > `canary-by-cookie` - > `canary-weight`

部署示例

以下示例中将会部署 `helloworld` 服务的 v1 和 v2 版本, 将 v2 版本作为 canary 版本, 在 ingress 中设置 canary 规则, 实现基于服务权重的流量切分。

创建测试应用

`helloworld-v1.yaml` 如下:

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: hello-world
spec:
  replicas: 1
  selector:
    matchLabels:
      app: hello-world
      version: v1
  template:
    metadata:
      labels:
        app: hello-world
        version: v1
    spec:
      containers:
        - name: hello-world
          image: hub.kce.ksyun.com/kingsoft/hello-world:v1
---
apiVersion: v1
kind: Service
metadata:
  labels:
    app: hello-world
  name: hello-world-svc
spec:
  ports:
    - port: 80
      protocol: TCP
      targetPort: 80
  selector:
    app: hello-world
    version: v1
  type: ClusterIP
```

`helloworld-v2.yaml` 如下:

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: hello-world-v2
spec:
  replicas: 1
  selector:
    matchLabels:
      app: hello-world
      version: v2
  template:
    metadata:
      labels:
        app: hello-world
        version: v2
    spec:
      containers:
```

```

- name: hello-world
  image: hub.kce.ksyun.com/kingsoft/hello-world:v2
---
apiVersion: v1
kind: Service
metadata:
  labels:
    app: hello-world
    name: hello-world-svc-v2
spec:
  ports:
    - port: 80
      protocol: TCP
      targetPort: 80
  selector:
    app: hello-world
    version: v2
  type: ClusterIP

```

创建并验证v1, v2版本服务的部署情况:

```

[root@vm10-0-11-201 ~]# kubectl apply -f helloworld-v2.yaml
deployment.apps/hello-world-v2 created
service/hello-world-svc-v2 created
[root@vm10-0-11-201 ~]# kubectl get deploy
NAME                READY   UP-TO-DATE   AVAILABLE   AGE
hello-world         1/1     1             1           15s
hello-world-v2      1/1     1             1           9s
[root@vm10-0-11-201 ~]# kubectl get svc
NAME                TYPE        CLUSTER-IP      EXTERNAL-IP   PORT(S)    AGE
hello-world-svc     ClusterIP   10.254.200.211  <none>        80/TCP     19s
hello-world-svc-v2 ClusterIP   10.254.14.130   <none>        80/TCP     13s
kubernetes          ClusterIP   10.254.0.1      <none>        443/TCP    7d
[root@vm10-0-11-201 ~]# curl 10.254.200.211
Hello World v1!
[root@vm10-0-11-201 ~]# curl 10.254.14.130
Hello World v2!

```

配置 Ingress

基于服务权重进行流量切分

对v1版本服务进行Ingress配置, 创建helloworld-ingress.yaml:

```

apiVersion: extensions/v1beta1
kind: Ingress
metadata:
  name: hello-world
  annotations:
    kubernetes.io/ingress.class: nginx
spec:
  rules:
    - host: hello.world.test
      http:
        paths:
          - backend:
              serviceName: hello-world-svc
              servicePort: 80

```

对v2版本的canary规则进行配置, 创建weight-ingress.yaml:

```

apiVersion: extensions/v1beta1
kind: Ingress
metadata:
  annotations:
    kubernetes.io/ingress.class: nginx
    nginx.ingress.kubernetes.io/canary: "true"
    nginx.ingress.kubernetes.io/canary-weight: "50"
  name: helloworld-weight
spec:
  rules:
    - host: hello.world.test
      http:
        paths:
          - backend:
              serviceName: hello-world-svc-v2
              servicePort: 80

```

创建Ingress规则:

```

[root@vm10-0-11-201 ~]# kubectl apply -f helloworld-ingress.yaml

```



```

ingress.extensions/hello-world created
[root@vm10-0-11-201 ~]# kubectl apply -f weight-ingress.yaml
ingress.extensions/helloworld-weight created
[root@vm10-0-11-201 ~]# kubectl get ingress
NAME                CLASS    HOSTS                ADDRESS    PORTS    AGE
hello-world         <none>   hello.world.test    80         41s
helloworld-weight   <none>   hello.world.test    80         27s

```

验证访问情况

通过以下命令获取EXTERNAL-IP及访问服务:

```

[root@vm10-0-11-201 ~]# kubectl get svc -n ingress-nginx
NAME                TYPE                CLUSTER-IP    EXTERNAL-IP    PORT(S)                AGE
nginx-ingress       LoadBalancer        10.254.28.54   120.92.xx.xx   80:31741/TCP,443:32754/TCP 3h19m
[root@vm10-0-11-201 ~]# for i in $(seq 1 10); do curl -H "Host: hello.world.test" http://120.92.xx.xx; done;
Hello World v2!
Hello World v2!
Hello World v1!
Hello World v2!
Hello World v2!
Hello World v1!
Hello World v1!
Hello World v1!
Hello World v1!
Hello World v2!
Hello World v2!

```

多次访问能发现约50%的流量会被分发到v2版本服务中。

在KCE集群中对Pod进行带宽限速

金山云容器服务原生支持对Pod进行带宽限速。本文档介绍如何在KCE集群设置Pod带宽限速。

备注

- 对于创建时间在2021-3-16之后的集群，默认支持对Pod带宽进行限速；对于创建时间在2021-3-16之前的集群，如需要使用Pod带宽限速的能力，请联系您的商务申请
- 集群网络插件Flannel和Canal均支持Pod带宽限速的能力

使用方式

新建Pod，通过annotations的方式设置Pod出入带宽

- `kubernetes.io/egress-bandwidth`: 定义Pod的出向带宽，如240M，这里单位是Mbit，
- `kubernetes.io/ingress-bandwidth`: 定义Pod的入向带宽，如400M，这里单位是Mbit

Yaml示例如下:

```

apiVersion: v1
kind: Pod
metadata:
  annotations:
    kubernetes.io/egress-bandwidth: 240M
    kubernetes.io/ingress-bandwidth: 400M
  name: iperf3-4
  namespace: default
spec:
  containers:
  - args:
    - iperf3
    - -s
    image: networkstatic/iperf3
    imagePullPolicy: Always
    name: iperf3-4
    resources: {}
  dnsPolicy: ClusterFirst

```

验证

你可以通过以下两种方式验证:

- 登陆pod所在的节点，执行:

```
tc qdisc show
```

返回如下示例代表设置成功

```
qdisc tbf 1: dev vethbab31be6 root refcnt 2 rate 400Mbit burst 256Mb lat 25.0ms
qdisc ingress ffff: dev vethbab31be6 parent ffff:ffff1 -----
qdisc tbf 1: dev bwp91fff0f8f685 root refcnt 2 rate 240Mbit burst 256Mb lat 25.0ms
```

- 使用iperf3工具验证

```
iperf3 -c <服务 IP> -i 1
```

结果如下:

```
Server listening on 5201
-----
Accepted connection from 10.0.11.85, port 45690
[ 5] local 10.0.11.207 port 5201 connected to 10.0.11.85 port 45692
[ ID] Interval      Transfer    Bandwidth
[ 5]  0.00-1.00    sec   243 MBytes  2.03 Gbits/sec
[ 5]  1.00-2.00    sec  49.2 MBytes  413 Mbits/sec
[ 5]  2.00-3.00    sec  27.3 MBytes  229 Mbits/sec
[ 5]  3.00-4.00    sec  27.3 MBytes  229 Mbits/sec
[ 5]  4.00-5.00    sec  27.4 MBytes  230 Mbits/sec
[ 5]  5.00-6.00    sec  27.3 MBytes  229 Mbits/sec
[ 5]  6.00-7.00    sec  27.3 MBytes  229 Mbits/sec
[ 5]  7.00-8.00    sec  27.3 MBytes  229 Mbits/sec
[ 5]  8.00-9.00    sec  27.3 MBytes  229 Mbits/sec
[ 5]  9.00-10.00   sec  27.3 MBytes  229 Mbits/sec
[ 5] 10.00-10.06   sec   1.49 MBytes  225 Mbits/sec
```

- [ID] Interval Transfer Bandwidth [5] 0.00-10.06 sec 0.00 Bytes 0.00 bits/sec sender [5] 0.00-10.06 sec 512 MBytes 427 Mbits/sec receiver

如何选择Containerd和Docker

容器运行时简介

容器运行时 (Container Runtime) 是kubernetes最重要的组件之一, 负责管理镜像和容器的生命周期。Kubelet通过Container Runtime Interface (CRI) 与容器运行时交互, 来管理镜像和容器。

如何选择Containerd和Docker?

容器服务支持Containerd和Docker两种运行时, 您可根据实际需求选择不同的运行时:

- Containerd运行时: 调用链更短、组件更少、更稳定、占用节点资源更少, 建议您选择Containerd运行时。
- 若您有以下使用场景, 建议您选择Docker运行时:
 - 使用 docker in docker。
 - 在节点使用docker build/push/save/load等命令。
 - 调用docker API。
 - 需要docker compose 或 docker swarm。

Containerd和Docker常用命令对比

说明: Containerd 不支持 docker API 和 docker CLI, 若您有该需求, 可以通过 cri-tool 命令实现类似的功能。

镜像相关功能

命令	Docker	Containerd
查看镜像列表	docker images	crictl images
下载镜像	docker pull	crictl pull
上传镜像	docker push	-
删除镜像	docker rmi	crictl rmi
查看镜像详情	docker inspect	crictl inspect

容器相关功能

命令	Docker	Containerd
查看容器列表	docker ps	crictl ps

创建容器	docker create	crictl create
启动容器	docker start	crictl start
停止容器	docker stop	crictl stop
删除容器	docker rm	crictl rm
查看容器详情	docker inspect	crictl inspect
挂载容器	docker attach	crictl attach
容器内执行命令	docker exec	crictl exec
查看容器日志	docker logs	crictl logs
显示容器资源使用情况	docker stats	crictl stats

POD相关功能

命令	Docker	Containerd
查看pod列表 -		crictl pods
查看pod详情 -		crictl inspectp
运行pod -		crictl runp
停止pod -		crictl stopp

调用链对比

容器运行时	调用链
Docker	kubelet -> docker shim -> dockerd -> containerd -> containerd-shim -> runC容器
Containerd	kubelet -> containerd -> containerd-shim -> runC容器