ATMIC NETWORKS
YOUR GUIDE TO GLOBAL CERTIFIATION

"Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns."

www.atmicnetworks .com

Warning: Keep connected with our support team for latest updates

Question: 1

Refer to Exhibit.

```
Set configuration context;

    f '.uCen'i';n."^ 1 . < kubectl
    config use-context k8s
```

Set Configuration Context:
[student@node-1] $ | kubectl
Config use-context k8s
Context

A web application requires a specific version of redis to be used as a cache.

Task

Create a pod with the following characteristics, and leave it running when complete:

- The pod must run in the web namespace.

The namespace has already been created

- The name of the pod should be cache
- Use the lfccncf/redis image with the 3.2 tag
- Expose port 6379

Answer: See the solution below.

Explanation:

Solution:

SB Readme >_ Web Terminal

□ THE **LI NUX** FOUNDATION

```
studentPnode-1:~$ kubectl run cache --image—Ifccncf/redis:3.2 —port—6379 -n web pod/cache created atudentenode-
1;*$ kubectl get pods -n web NAME     READY STATUS    RESTARTS AGE
cache   0/1    Containercreating     0          6a
student@node~l:^S kubectl get pods -n web NAME       READY    STATUS    RESTARTS  AGE
cache   1/1    Running     0        9a
atudent@node-1:-$ Q
```

Question: 2

Refer to Exhibit.

Set configuration context:

i-tul ' .cr >'-i $ kubectl onng use-context k8s

Context
You are tasked to create a secret and consume the secret in a pod using environment variables as follow:

Task

- Create a secret named another-secret with a key/value pair; key1/value4
- Start an nginx pod named nginx-secret using container image nginx, and add an environment variable exposing the value of the secret key key 1, using COOL_VARIABLE as the name for the environment variable inside the pod

Answer: See the solution below.

Explanation:

Solution:

```
student^      create secret generic some-secret --fron-LiteraL-keyl-walue'
secret/some-secret created
student&node-l:-S kubectl get secret
NAME                  TYPE                                  DATA      AGE
dctault-tokcn-Ikvro   kubcrnctes.io/sexvice-account-token   3         2dllh
some-secret           Opaque                                1         5s
3tudent0node-l:^$ kubecti run nginx-secret –image=nginx –dry-ri
```

```
student9node-l>$ vim nginx secret.yml
```

Readme >_ Web Terminal

□ THE **LINUX** FOUNDATION

```
apiVersion: v1
kind: Pod
metadata:
  creationTimestamp: null
  labels:
    run: nginx-secret
  name: nginx-secret
spec:
  containers:
  - image: nginx
    name: nginx-secret
    resources: {}
  dnsPolicy: Always
  restartPolicy: Always
status: {}
~
~
~
~
~
~
~
~
~
~
~
~
```

```
'nginx secret.yml" 15*L,
253C
```

All |

SB Readme >_ Web Terminal

□ THE **LINUX** FOUNDATION

```
          vl
- Pod
```

```
          nginx-secret
   - - nginx-secret
spec:
```

```
  nginx
nginx-secret
```

```
i COOL VARIABLE
```

```
some-secret
keyl|
```

```
INSERT --
```

16,20

```
student@node-l:*$ kubectl get pods -n web NAME       READY     STATUS    RESTARTS AGE
c8 Readme     >_ Web Terminal          9s
studentgnode-1:*$ kubectl create secret generic some-secret —from-literal=key1-value4 secret/some - secret
created 3tudent0node-l :*»$ kubectl get secret NAME          TYPE    DATA       AGE
default-token   4kvr5  kubernetes.io/service    account-token   3      2dllh
some-secret           Opaque                                 1      5s
atudcnt0node-l:-§ kubectl run nginx-secret --image nginx --dry-run-client -o yaml > nginx_secrct .yml
student@ncde-l:*$ vim nginx secret.yml studcnt(*node-l:-S kubectl create -f nginx_secret.yml pod/nginx-secret
created studentgnode-l!*$ kubectl get pods
NAME           READY    STATUS              RESTARTS   AGE
liveness-http  1/1      Running             0          6h3Sm
nginx-101      1/1      Running             0          6h39m
nginx-secret   0/1      ContainerCreating   0          4s
poller         1/1      Running             0          Sh39m
student0node-1: i *S kubectl get pods
NAME           READY    STATUS     RESTARTS     AGE
livens39-http  1/1      Running 0               6h3Bm
nginx-101      1/1      Running 0               6h39m
nginx-secret   1/1      Running 0               &3
poller         1/1      Running 0               6h39m
studentpnode■1: -S 1
```


Question: 3

Refer to Exhibit.

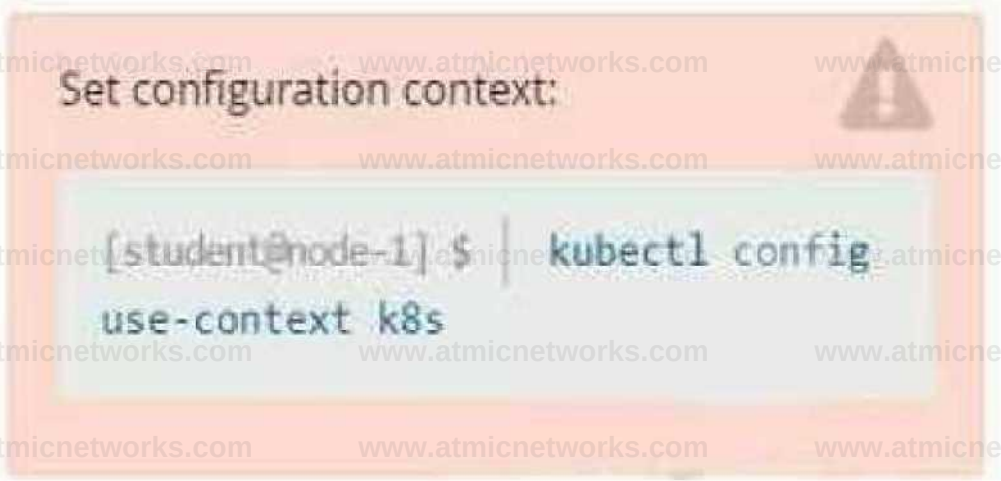

Task

You are required to create a pod that requests a certain amount of CPU and memory, so it gets scheduled to-a node that has those resources available.

- Create a pod named nginx-resources in the pod-resources namespace that requests a minimum of 200m CPU and 1Gi memory for its container
- The pod should use the nginx image
- The pod-resources namespace has already been created

Answer: See the

solution below.

Explanation:

Solution:

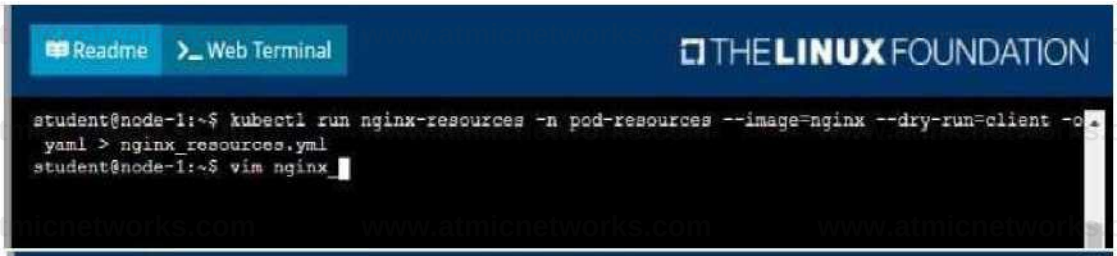

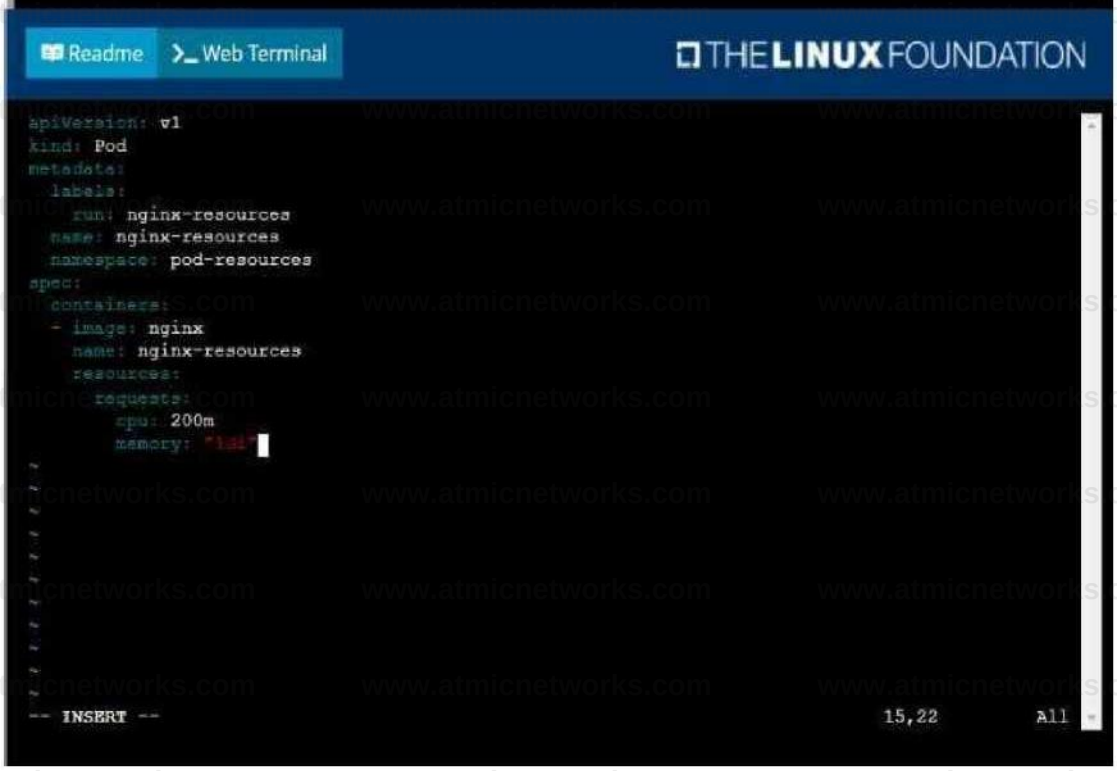

Question: 4

Refer to Exhibit.

Ser configuration context:

[student^vxle-11 : kubectl contig use-context k8s

Context

You are tasked to create a ConfigMap and consume the ConfigMap in a pod using a volume mount. Task Please complete the following:

- Create a ConfigMap named another-config containing the key/value pair: key4/value3
- start a pod named nginx-configmap containing a single container using the nginx image, and mount the key you just created into the pod under directory /also/a/path

Answer: See the solution below.

Explanation:

Solution:

```
3tudenr0node-l:^$ kubectl create configmap another-config —from-literal-key4-value3
configmap/another config created
atudentOnode-l: •*$ kubectl get configmap
NAME             DATA AGE
another-config   1      5s
student0node-l:*$ kubectl run nginx-configmap —image—nginx —dry-run—client -o yaml > ngin conf xgmap■yml
3tudent0node-l:*$ vim ngin_configmap.yrn] *C
student0node-1:*$ mv ngin_configmap.ynil nginx_configraap.yml
3tudcnt0node 1:'$ vim nginx_co|
```

90 Readme >_ Web Terminal

□ **THELINUX** FOUNDATION

```
I ^1
      Pod metadata;
creationTimestamp
labels:
  nginx-configmap nginx-configmap
spec:
  containers:
  - image: nginx
       nginx-configmap resource?■
    .       CluaterFirst
                Always
 statua! |}
```

```
"nginx_configmap. yml" 15L, 2G2C                    1,1          All
```

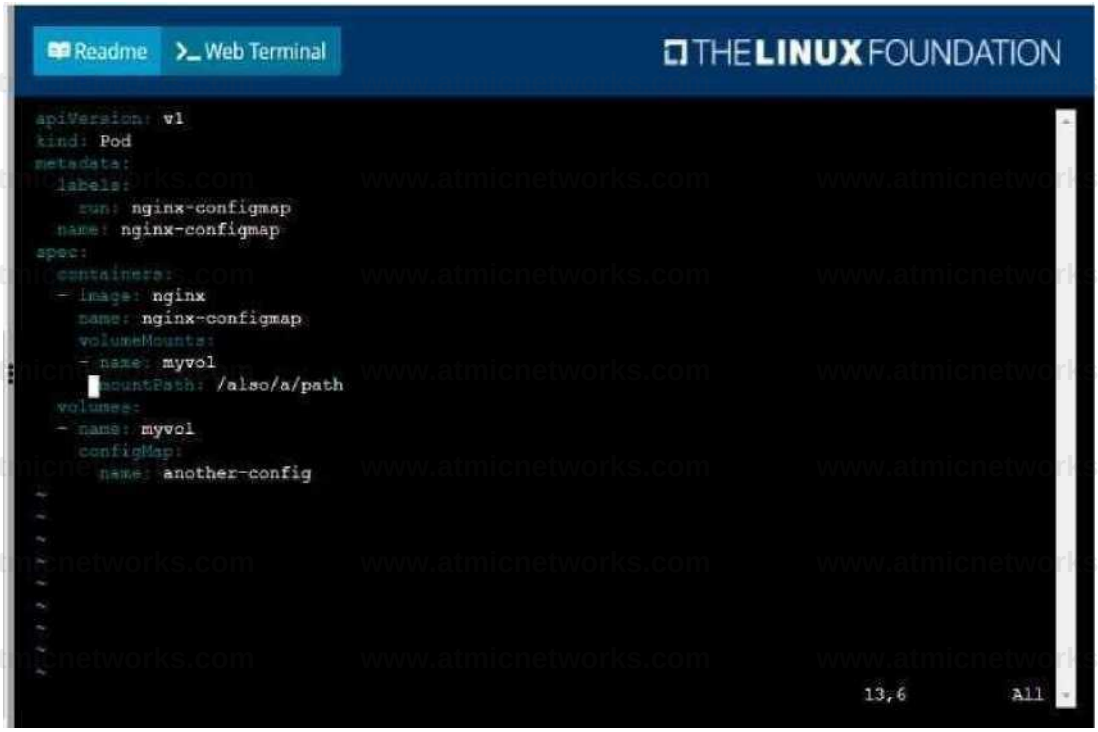

```
student0node-1:*$ kubectl create confignap another-confiq —from-iiteral=key4=value3 configmap/another-config
created                                                                                    |
student0node-l:*$ kubectl get configmap NAME         DATA AGE
another-config    1      5a


student8node-l:^$ kubectl run nginx-configmap —image—nginx —dry-run-client -o yaml > agin confH igmap. yml
student0node-lX'-S vim ngin_configmap.yml *C


student9node-l:*$ mv nginconfigmap.yml nginxconfigmap.yml student^node-l:*$ vim nginx_conf igmap.yml
9tudent9node-l:»$ |
```

```
student0node-1: >>$ kubectl run nginx-configmap —image—nginx —dry-run—client -o yaml > ngin conf igmap . yml
studentSnode-1: ••$ vim ngin_conf igmap. yml ^C
studentQnode-1:-$ mv ngin confignap. yml nginx configmap.yml
3tudcnt(?nodc - 1:-$ vim ngrnx_configmap. yml
5tudent0node-1:*$ kubectl create f nginx_configmap.yml
Error: must specify one of -f and -k
```

```
error: unknown command "f nginx configmap. yml" See 'kubectl create —h' for help and examples studentSnode-1:*$
kubectl create f nginx_configmap.yml errors error validating "nginx configmap.yml": error validating data:
ValidationError(Pod.spec.c ontainers[1]): unknown field "mountPath" in io.kQs.api.core.vl.Container; if you choose
to ignor e these errata, turn validation off with -vulidat^hhe studentSnode-1:*$ vim nginx_configmap.yml I
```

OB Readme >_ Web Terminal

□ THE **LINUX** FOUNDATION

```
student0node-l:-$ kubectl create f nginx configmap.yml Error: must specify one of f and -k

error: unknown command "f nginx_configmap.yml"
See 'kubectl create -h* for help and examples
studentgnode-1 :■*$ kubectl create -f nginx configmap.yml
  error: error validating "nginx_configmap.yml": error validating data: ValidationError(Pod.spec.c ontainers[1]):
unknown field "mountPath" in io.k8s.api.core.vl.Container; if you choose to ignor e these errors, turn validation
off with —validate—false
student8node-l:~$ vim ngrnx_conf igmap. yml
student&node-l:*$ kubectl create -f nginxconfigmap.yml
pod/nginx-configmap created
student0nodc-l: ••$ kubectl get pods
NAME            READY STATUS              RESTARTS AGE
liveneas-http      1/1     Running             0          6h44m
nginx-101          1/1     Running             0          6h45m
nginx-configmap    0/1     containersresting 0            5s
nginx-secret       1/1     Running             0          5m39s
poller             1/1     Running             0          6h44m
student0ncde-l:—$ kubectl get pods
NAME              READY STATUS       RESTARTS AGE
Liveness-http     1/1     Running 0            6h44m
nginx-101         1/1     Running 0            6h45m
nginx-conf igmap  1/1     Running 0            8s
nginx-secret      1/1     Running 0            5m42s
poller            1/1     Running 0            6h45m
student&node-l:*S 1|
```

Question: 5

Refer to Exhibit.

Set configuration context:

i: W-nt^ri^ . kubect) config use-context k8s

Context

Your application's namespace requires a specific service account to be used.

Task

Update the app-a deployment in the production namespace to run as the restrictedservice service account. The service account has already been created.

Answer: See the solution below.

Explanation:

Solution:

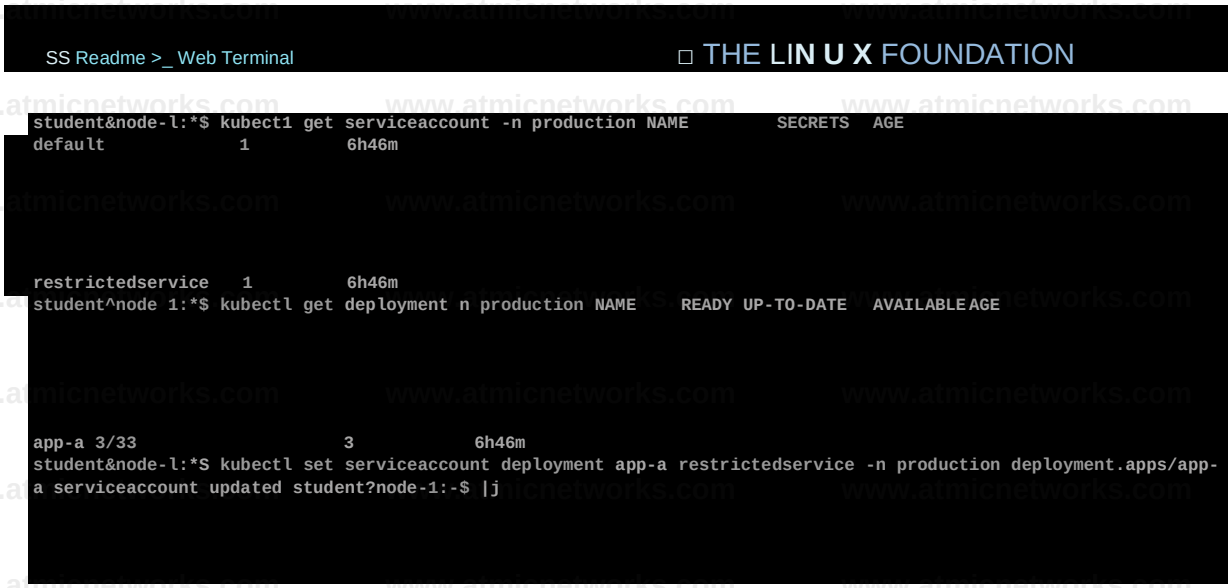

Question: 6

Refer to Exhibit.

Set Configuration Context:
[student@node-1] $ | kubectl
Config use-context k8s
Context
A pod is running on the cluster but it is not responding.

Task
The desired behavior is to have Kubemetes restart the pod when an endpoint returns an HTTP 500 on the /healthz endpoint. The service, probe-pod, should never send traffic to the pod while it is failing. Please complete the following:

- The application has an endpoint, /started, that will indicate if it can accept traffic by returning an HTTP 200. If the endpoint returns an HTTP 500, the application has not yet finished initialization.
- The application has another endpoint /healthz that will indicate if the application is still working as expected by returning an HTTP 200. If the endpoint returns an HTTP 500 the application is no longer responsive.
- Configure the probe-pod pod provided to use these endpoints
- The probes should use port 8080

Answer: See the solution below.

Explanation:

Solution:
To have Kubernetes automatically restart a pod when an endpoint returns an HTTP 500 on the /healthz endpoint, you will need to configure liveness and readiness probes on the pod.
First, you will need to create a livenessProbe and a readinessProbe in the pod's definition yaml file. The livenessProbe will check the /healthz endpoint, and if it returns an HTTP 500, the pod will be restarted. The readinessProbe will check the /started endpoint, and if it returns an HTTP 500, the pod will not receive traffic.
Here's an example of how you can configure the liveness and readiness probes in the pod definition yaml file:
apiVersion: v1
kind: Pod

```
metadata:
name: probe-pod
spec:
containers:
- name: probe-pod
  image: <image-name>
  ports:
  - containerPort: 8080
  livenessProbe:
    httpGet:
      path: /healthz
      port: 8080
    initialDelaySeconds: 15
    periodSeconds: 10
    failureThreshold: 3
  readinessProbe:
    httpGet:
      path: /started
      port: 8080
    initialDelaySeconds: 15
    periodSeconds: 10
    failureThreshold: 3
```

The httpGet specifies the endpoint to check and the port to use. The initialDelaySeconds is the amount of time the pod will wait before starting the probe. periodSeconds is the amount of time between each probe check, and the failureThreshold is the number of failed probes before the pod is considered unresponsive.

You can use kubectl to create the pod by running the following command:

```
kubectl apply -f <filename>.yaml
```

Once the pod is created, Kubernetes will start monitoring it using the configured liveness and readiness probes. If the /healthz endpoint returns an HTTP 500, the pod will be restarted. If the /started endpoint returns an HTTP 500, the pod will not receive traffic.

Please note that if the failure threshold is set to 3, it means that if the probe fails 3 times consecutively it will be considered as a failure.

The above configuration assumes that the application is running on port 8080 and the endpoints are available on the same port.

Question: 7
Refer to Exhibit.

Set configuration context:

[hnl*<rm.<k 1] • kubectl config use-context k8s

Set Configuration Context:
[student@node-1] $ | kubectl
Config use-context k8s
Context
You sometimes need to observe a pod's logs, and write those logs to a file for further analysis.

Task
Please complete the following;

- Deploy the counter pod to the cluster using the provided YAMLspec file at /opt/KDOB00201/counter.yaml
- Retrieve all currently available application logs from the running pod and store them in the file /opt/KDOB0020I/log_Output.txt, which has already been created

Answer: See the solution below.

Explanation:

Solution:
To deploy the counter pod to the cluster using the provided YAML spec file, you can use the kubectl apply command. The apply command creates and updates resources in a cluster.
kubectl apply -f /opt/KDOB00201/counter.yaml
This command will create the pod in the cluster. You can use the kubectl get pods command to check the status of the pod and ensure that it is running.
kubectl get pods

To retrieve all currently available application logs from the running pod and store them in the file /opt/KDOB0020l/log_Output.txt, you can use the kubectl logs command. The logs command retrieves logs from a container in a pod.

kubectl logs -f <pod-name> > /opt/KDOB0020l/log_Output.txt

Replace <pod-name> with the name of the pod.

You can also use -f option to stream the logs.

kubectl logs -f <pod-name> > /opt/KDOB0020l/log_Output.txt &

This command will retrieve the logs from the pod and write them to the /opt/KDOB0020l/log_Output.txt file.

Please note that the above command will retrieve all logs from the pod, including previous logs. If

you want to retrieve only the new logs that are generated after running the command, you can add the --since flag to the kubectl logs command and specify a duration, for example --since=24h for logs generated in the last 24 hours.

Also, please note that, if the pod has multiple containers, you need to specify the container name using -c option.

kubectl logs -f <pod-name> -c <container-name> > /opt/KDOB0020l/log_Output.txt

The above command will redirect the logs of the specified container to the file.

```
studentSnode-l:*S kubectl create -f /opt/KDOB00201/counter.yaml pod/counter created
student?node-l:-$  kubectl  get pods
NAME              READY    STATUS    RESTARTS   AGE
counter           1/1      Sunning   0          10s
livenoss-http     1/1      Running   0          6h45m
nginx-101         1/1      Running   0          6h46m
nginx-configmap   1/1      Running   0          107B
nginx-aecret      1/1      Running   0          7m2 la
poller            1/1      Running   0          6h46m
student@node-l:-$  kubectl  logs counter
1: 2b305101817ae25ca60ac46510fb6dll
 2:  3648cf2eae95ab680dba8fl95f891af4
 3:  G5cBbbd4dbf70bf81f2a0984a3a44ede
 4:  40d3a9c8e4€f5533bb4828fbe5c8d038
 5:  390442d2530a90c3602901e3fe999ac8
 6:  b71d95187417el39effb33a£77681040
 7:  66a8e55a6491e756d2d0549ad6ab90a7
 8:  ff2b3d583b64125d2f9129c443bb37ff
 9:  bGcGal2bGe77944edttbaaal6c242dae4
 10:  bfcc9a894a0604fc4b814b37d0a200a4
studentUnode-1:*$ kubectl logs counter > /opt/KI>0B00201/log_output. txt studentOnode 1: •-$ Q
```

```
studentgnode-1:*$ kubectl logs counter > /opt/KDOB00201Zlog_output.txt
student 0node-1 :-•$ kubectl logs counter > /opt/KD0B00201/log_output. txt
atudcntBncdc 1 ;-S ca|opt/KD0B00201/log_output. txt
```

• Readme >_ Web Terminal □ **THELINUX** FOUNDATION

```
student@node-l!^$ kubectl logs counter > /opt/KDOB00201/log output.txt student8nodo-l:*$ cat
/opt/KDOB00201/log_output.txt 1: 2b305101817ae25ca60ae46510fb6dll 2: 364Bcf2eae95ab6B0dba8fl95f891af4 3:
65c8bbd4dbf70b£81f2a0984a3a44ede 4: 40d3a9cBe46f5533bb482Bfbe5c8d03B 5: 390442d2530a90c3602901e3fe999acB 6;
b71d95187417el39c££b33af77681040 7: €6aBe55a6491e756d2d0549ad6ab90a7 8: ff2b3d583b64125d2f9129c443bb37ff 9:
b6c6al2b€e77944ed8baaaf6c242dae4 10: bfcc9a894a0604fc4b814b37d0a200a4 11: 5493cdl6al790a5£b9512b0c9d4c5ddl 12:
03fl69e93e€143438e6dfe4ecb3cc9ed 13: 764b37fe611373c42d0b47154041f6eb 14: la56fbel896b0ee6394136166281839e 15:
ecc492ebl7715de090c47345a98d98d3 16: 7974a6bac0fb44b6b8bbfc71aa3fb®74 17; 9ae01bcf0174Bbl2cc9f97a5f9f72cd6 18:
23fb22ee34d4272e4c9e005fl774515f 19: ec7ela5d314da9a0ad45d53be5a7acae 20: 0bccdd8ee02cd42029e8162cdlcl197c 21:
d6851ea43546216b95bcb81ced997102 22: 7ed9a3Bea8bf0d86206569481442af44 23: 29b8416ddc63dbfcb987ab3c8198e9fe 24:
lf2062001df51al08ab25010f506716f atudcntBnodc 1;*$ |
```

Question: 8

Refer to Exhibit.

Set configuration context:

i <ahte'ii-T<'-'-I j < kubectl config use-context k8s

Context

It is always useful to look at the resources your applications are consuming in a cluster.

Task

- From the pods running in namespace cpu-stress , write the name only of the pod that is consuming the most CPU to file /opt/KDOBG030I/pod.txt, which has already been created.

Answer: See the solution below.

Explanation:

Solution:

Question: 9

Context

Anytime a team needs to run a container on Kubernetes they will need to define a pod within which to run the container.

Task

Please complete the following:

- Create a YAML formatted pod manifest

/opt/KDPD00101/podl.yml to create a pod named app1 that runs a container named app1cont using image lfccncf/arg-output

with these command line arguments: -lines 56 -F

* Create the pod with the kubect1 command using the YAML file created in the previous step

• When the pod is running display summary data about the pod in JSON format using the kubect1 command and redirect the output to a file named /opt/KDPD00101/out1.json

- All of the files you need to work with have been created, empty, for your convenience

When qieaQng your pod, you do not need to specify a container command, only args,

Answer: See the solution below.

Explanation:

Solution:

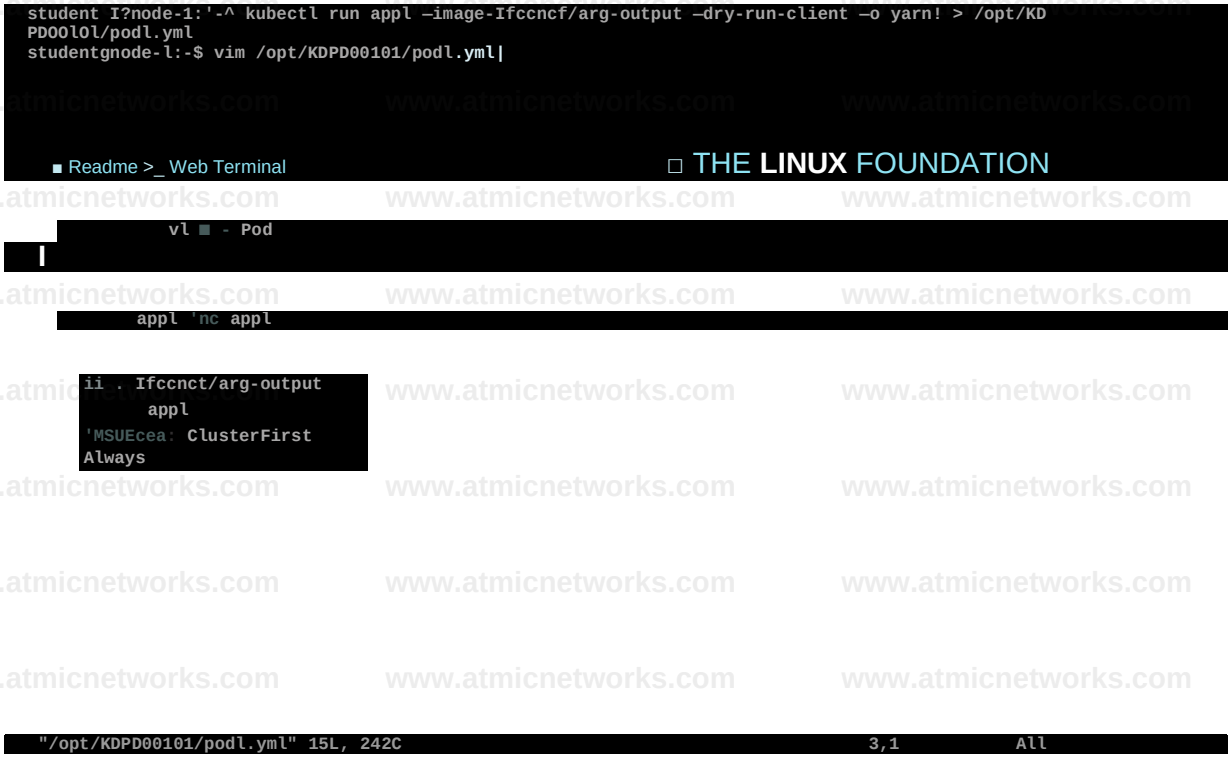

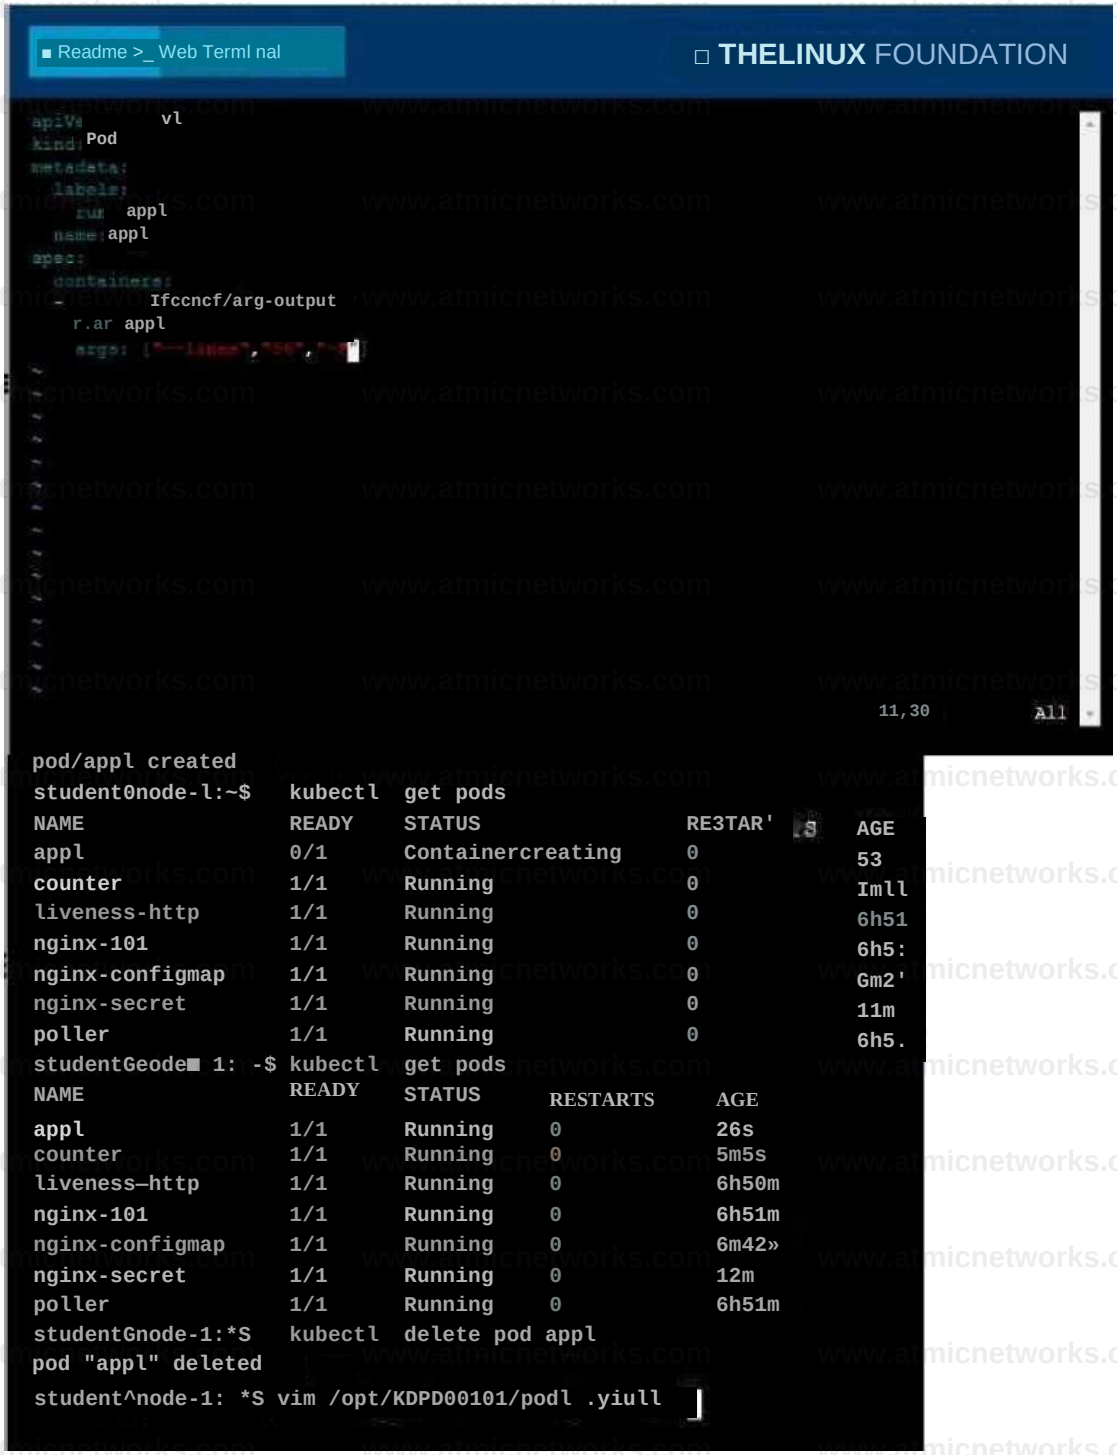
■ Readme >_ Web Terml nal
□ THELINUX FOUNDATION
vl
Pod
appl
appl
Ifccncf/arg-output
r.ar appl
11,30
All
pod/appl created
student0node-l:~$ kubectl get pods
NAME READY STATUS RE3TAR' 3 AGE
appl 0/1 Containercreating 0 53
counter 1/1 Running 0 Imll
liveness-http 1/1 Running 0 6h51
nginx-101 1/1 Running 0 6h5:
nginx-configmap 1/1 Running 0 Gm2'
nginx-secret 1/1 Running 0 11m
poller 1/1 Running 0 6h5.
studentGeode■ 1: -$ kubectl get pods
NAME READY STATUS RESTARTS AGE
appl 1/1 Running 0 26s
counter 1/1 Running 0 5m5s
liveness—http 1/1 Running 0 6h50m
nginx-101 1/1 Running 0 6h51m
nginx-configmap 1/1 Running 0 6m42»
nginx-secret 1/1 Running 0 12m
poller 1/1 Running 0 6h51m
studentGnode-1:*S kubectl delete pod appl
pod "appl" deleted
student^node-1: *S vim /opt/KDPD00101/podl .yiull

■ Readme >_ Web Terminal

```
nginx-configmap     1/1     Running             0                2
nginx-secret        1/1     Running             0                li
poller              1/1     Running             0                6h
student ("node -1:  kubectl get pods
NAME                READY   STATUS     RESTARTS     AGE
appl                1/1     Running 0               26a
counter             1/1     Running 0               5m5s
liveness-http       1/1     Running 0               6h50m
nginx-101           1/1     Running 0               6h51m
nginx-configmap     1/1     Running 0               6m42s
nginx-secret        1/1     Running 0               12m
poller              1/1     Running 0               6h51m
| 3tudentl?node-l:->$ kubectl delete pod appl
pod "appl" deleted
student?node-l:-$   vim. /opt/RDPD00101/podl.yml
student&nc-de-1:  -$ kubectl create f /opt/KDPD00101/podl.yml
pod/appl created
student0node-l:~$
NAME
ªPPi                kubectl get pods
counter liveness-   READY   STATUS     RESTARTS AGE
http nginx-101      1/1     Running    0        20B
nginx-configmap     1/1     Running    0        6m57s
nginx-secret poller 1/1     Running    0        6h52m
student^node-1:^S   1/1     Running    0        6h53m
                    1/1     Running    0        8m34s
                    1/1     Running    0        14m
                    1/1     Running    0        6h53m
                    kubectl get pod appl -o json > |
```

• Readme >_ Web Terminal □ THELINUX FOUNDATION

```
poller              1/1     Running             0          6h51m
student0node-l:-S   kubectl get poda
NAME                READY   STATUS     RESTARTS    AGE
appl                1/1     Running 0              26s
counter             1/1     Running 0              5m5a
livenesa-http       1/1     Running 0              6h50m
nginx-101           1/1     Running 0              6h51m
nginx-configmap     1/1     Running 0              6m42a
nginx-secret        1/1     Running 0              12m
poller              1/1     Running 0              6h51m
atudentSnode-1:»S   kubectl delete pod appl
pod "appl" deleted
```

```
NAME                READY   STATUS     RESTARTS   AGE
ªPPl                1/1     Running    0          20a
counter             1/1     Running    0          Gm 5 7 3
livenesa-http       1/1     Running    0          6h52m
nginx-101           1/1     Running    0          6h53m
nginx-configmap     1/1     Running    0          8m34s
nginx-secret        1/1     Running    0          14m
poller              1/1     Running    0          Gh53m
studentAnode-1:-S   kubectl get pod  appl -o json  > /opt/KDPD00101/outl.json
```

```
studentSnode-1:-S vim /opt/KDPD00101/podl.yml
ntudent0node-l:*S kubectl create -f /opt/KDPD00101/podl.yml pod/appl created
student&node !:■•$ kubectl get pods
```

Question: 10

Refer to Exhibit.

```
student0node-1:•»$
studentSnode -1: ••$
Q
```

Set configuration context:

j • tuHdit^* de*l l . kubectl config use-context k8s

Task
Create a new deployment for running.nginx with the following parameters;

- Run the deployment in the kdpd00201 namespace. The namespace has already been created
- Name the deployment frontend and configure with 4 replicas
- Configure the pod with a container image of lfccncf/nginx:1.13.7
- Set an environment variable of NGINX PORT=8080 and also expose that port for the container above

Answer: See the solution below.

Explanation:

Solution:

```
BP Readme >_ Web Terminal                                   □ THE LIN UX FOUNDATION

studentOnoda-1:*$ kubectl create deployment api —image=lfccncf/nginx:1.13.7-alpine —replicas-4


 -n kdpd00201 -dry-run-client o yaml > nginx_deployment.yml
studentOnode-1:*$ vim nginxdeployment.yml
atudent0node-l: •*$ kubectl create nginx deployment, yml
Error: must specify one of -f and -k

error: unknown command "nginx_depioyment.yml"
See 'kubectl create -h* for help and examples


student@node-l: •*$ kubectl create -f nginx deployment. yml
error: error validating "ngrnx_deployment.yml": error validating data: ValidationError(Deploymen
t.spec.template.spec): unknown field "env" in io.k8s.api.core.vl.PodSpec; if you choose to ignor e these errors,
turn validation off with —validate—false student^node-1:^$ vim nginx_deployraent.yml
student9node-l:*$ kubectl create -f nginx_deployment.yml
deployment.apps/api created


studentdnode-1;*$ kubectl get pods n kdpd00201


KAME                   READY   STATUS    RESTARTS   AGE
api-745677f7dc-7hnvm   1/1     Running   0          13s
api■745677f7dc-9q5vp   1/1     Running   0          13s
api-745677f7dc-fd4gk   1/1     Running   0          13s
api — 745677f7dc—      1/1     Running   0          13s
studentSnode-l:'$ []
```

Question: 11

Refer to Exhibit.

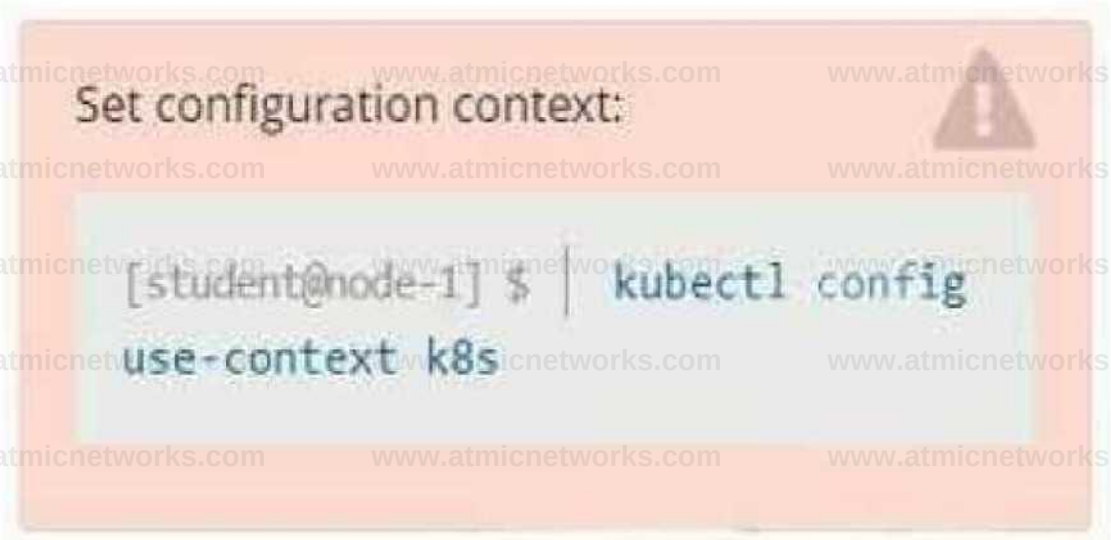

Context

As a Kubernetes application developer you will often find yourself needing to update a running application.

Task

Please complete the following:

- Update the app deployment in the kdpd00202 namespace with a maxSurge of 5% and a maxUnavailable of 2%
- Perform a rolling update of the web1 deployment, changing the lfccncf/ngmx image version to 1.13
- Roll back the app deployment to the previous version

Answer: See the solution below.

Explanation:

Solution:

W Readme >_ Web Terminal □ THE LI N UX FOUNDATION

```
student.Qnode-1:^$ kubectl edit deployment app —n kdpd00202|
```

■ Readme >_ Web Terminal □ THELINUX FOUNDATION

```
          Idfa2527-5c61-46a9-8dd3-e24643d3cel4

      replicas ■
      revisisnKistoryLimit;

      mar -zhLaheis: nginx
      strategy:

                5%

          RollingUpdate

        metadata: oreationTunestanip-: null

            nginx

          containers:
            :      Ifccncf/nginx:1.13
                •    . TfNotPresent
            • "ii IT nginx


          pretocDl TCP
    :^l|
```

BO Readme >_ Web Terminal □ THELINUX FOUNDATION

```
student^node—1: •*$ kubectl edit deployment app -n kdpd00202 deployment.apps/app edited
studentBnode-l:«S kubectl rollout status deployment app -n kdpd00202
Waiting for deployment "app" rollout to finish: € out of 10 new replicas have been updated..
Waiting for deployment "app" rollout to finish: 7 out of 10 new replicas have been updated..
Waiting for deployment "app" rollout to finish: 7 out of 10 new replicas have been updated..
Waiting for deployment "app" rollout to finish: 7 out of 10 new replicas have been updated..
Waiting for deployment "app" rollout to finish: 8 out of 10 new replicas have been updated..
Waiting for deployment "app" rollout to finish: 8 out of 10 new replicas have been updated..
Waiting for deployment "app" rollout to finish: B out of 10 new replicas have been updated..
Waiting for deployment "app" rollout to finish: 8 out of 10 new replicas have been updated..
Waiting for deployment "app" rollout to finish: 9 out of 10 new replicas have been updated..
Waiting for deployment "app" rollout to finish: 9 out of 10 new replicas have been updated..
Waiting for deployment "app" rollout to finish: 9 out of 10 new replicas have been updated..
Waiting for deployment "app" rollout to finish: 1 old replicas are pending termination...
 Waiting for deployment "app" rollout to finish: 8 of 10 updated replicas are available... Waiting for deployment
 "app" rollout to finish: 9 of 10 updated replicas are available... deployment "app" successfully rolled out
 studentBnode~l:*S kubectl rollout undo deployment app ~n kdpd00202 deployment.apps/app rolled back student6nodo-
 l:'$ kubectl rollout status deployment app -n kdpd00202

 5tudent?node-l:*$ kubectl rollout status deployment app ~n kdpd00202

  Waiting for   deployment "app"  rollout to   finish:€   outof 10   newreplicas have been updated...
  Waiting for   deployment "app"  rollout to   finish;6   outof 10   newreplicas have been updated...
  Waiting for   deployment "app"  rollout to   finish:6   outof 10   newreplicas have been updated...
  Waiting for   deployment "app"  rollout to   finish:€   outof 10   newreplicas have been updated...
  Waiting for   deployment "app"  rollout to   finish:7   outof 10   newreplicas have been updated...
  Waiting for   deployment "app"  rollout to   finish:7   outof 10   newreplicas have been updated...
  Waiting for   deployment "app"  rollout to   finish:9   outof 10   newreplicas have been updated...
  Waiting for   deployment "app"  rollout to   finish:9   outof 10   newreplicas have been updated...
  Waiting for   deployment "app"  rollout to   finish:9   outof 10   newreplicas have been updated...
  Waiting for   deployment "app"  rollout to   finish:1   old replicas are pending termination...
  Waiting for   deployment "app"  rollout to   finish:1   old replicas are pending termination—
  Waiting for   deployment "app"  rollout to   finish:1   old replicas are pending termination...
```

```
Waiting for   deployment "app"  rollout to   finish:8   of 10 updated replicas are available...
Waiting for   deployment "app"  rollout to   finish:9   of 10 updated replicas are available...

deployment "app" successfully rolled out student^node 1:*S |
```

Question: 12

Refer to Exhibit.

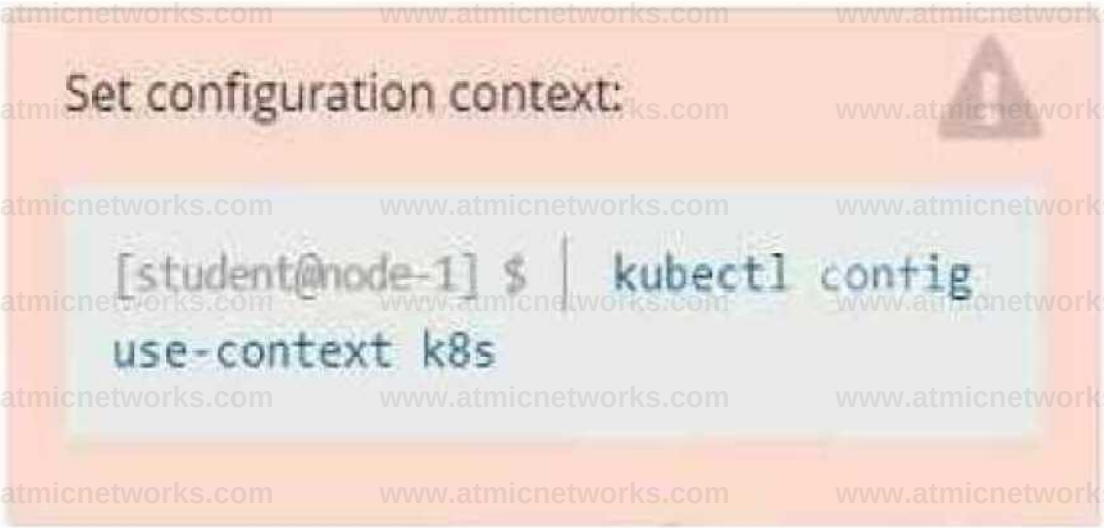

Context

You have been tasked with scaling an existing deployment for availability, and creating a service to expose the deployment within your infrastructure.

Task

Start with the deployment named kdsn00101-deployment which has already been deployed to the namespace kdsn00101 . Edit it to:

- Add the func=webFrontEnd key/value label to the pod template metadata to identify the pod for the service definition
- Have 4 replicas

Next, create ana deploy in namespace kdsn00l01 a service that accomplishes the following:

- Exposes the service on TCP port 8080
- is mapped to me pods defined by the specification of kdsn00l01-deployment
- Is of type NodePort
- Has a name of cherry

Answer: See the solution below.

Explanation:

Solution:

```
  student8node-1;*$ kubectl edit deployment kdsnOOl01-deployment -n kdsnOOlOl deployment.appa/EdanO0101-deployment
edited student0node*l;*$ kubectl get deployment kdsnOOlOl-deployment n kdenOOlOl NAME       READY    UP-TO-DATE
                      AVAILABLE             AGE
kdsn.00101-deployment   4/4     4            4           7hl7m
  student0node~l: ••$ kubectl expose deployment kdsnOOl01-deployment n kdsnOOlOl --type NodePort -port 8080 —name
cherry service/cherry exposed
```

Question: 13

Refer to Exhibit.

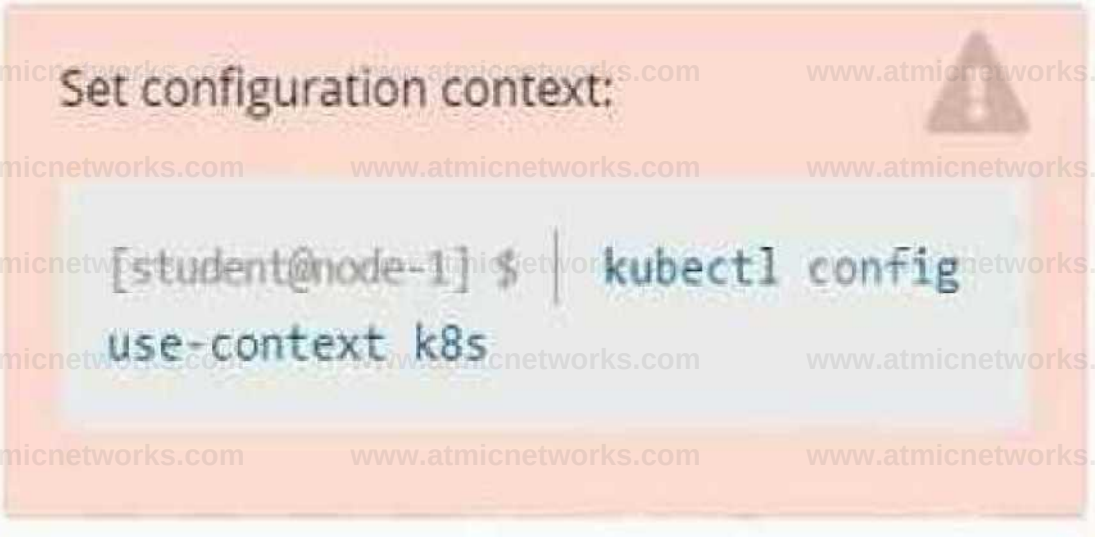

Set Configuration Context:

[student@node-1] $ | kubectl

Config use-context k8s

Context

A container within the poller pod is hard-coded to connect the nginxsvc service on port 90 . As this port changes to 5050 an additional container needs to be added to the poller pod which adapts the container to connect to this new port. This should be realized as an ambassador container within the pod.

Task

- Update the nginxsvc service to serve on port 5050.
- Add an HAproxy container named haproxy bound to port 90 to the poller pod and deploy the enhanced pod. Use the image haproxy and inject the configuration located at /opt/KDMC00101/haproxy.cfg, with a ConfigMap named haproxy-config, mounted into the container so that haproxy.cfg is available at /usr/local/etc/haproxy/haproxy.cfg. Ensure that you update the args of the poller container to connect to localhost instead of nginxsvc so that the connection is correctly proxied to the new service endpoint. You must not modify the port of the endpoint in poller's args . The spec file used to create the initial poller pod is available in /opt/KDMC00101/poller.yaml

Answer: See the solution below.

Explanation:

Solution:

To update the nginxsvc service to serve on port 5050, you will need to edit the service's definition yaml file. You can use the kubectl edit command to edit the service in place.

kubectl edit svc nginxsvc

This will open the service definition yaml file in your default editor. Change the targetPort of the service to 5050 and save the file.

To add an HAproxy container named haproxy bound to port 90 to the poller pod, you will need to edit the pod's definition yaml file located at /opt/KDMC00101/poller.yaml.

You can add a new container to the pod's definition yaml file, with the following configuration: containers:

- name: haproxy
image: haproxy
ports:
- containerPort: 90
volumeMounts:
- name: haproxy-config
mountPath: /usr/local/etc/haproxy/haproxy.cfg
subPath: haproxy.cfg
args: ["haproxy", "-f", "/usr/local/etc/haproxy/haproxy.cfg"]

This will add the HAproxy container to the pod and configure it to listen on port 90. It will also mount the ConfigMap haproxy-config to the container, so that haproxy.cfg is available at /usr/local/etc/haproxy/haproxy.cfg.

To inject the configuration located at /opt/KDMC00101/haproxy.cfg to the container, you will need to create a ConfigMap using the following command:

kubectl create configmap haproxy-config --from-file=/opt/KDMC00101/haproxy.cfg

You will also need to update the args of the poller container so that it connects to localhost instead of nginxsvc. You can do this by editing the pod's definition yaml file and changing the args field to args: ["poller","--host=localhost"].

Once you have made these changes, you can deploy the updated pod to the cluster by running the following command:

kubectl apply -f /opt/KDMC00101/poller.yaml

This will deploy the enhanced pod with the HAproxy container to the cluster. The HAproxy container will listen on port 90 and proxy connections to the nginxsvc service on port 5050. The poller container will connect to localhost instead of nginxsvc, so that the connection is correctly proxied to the new service endpoint.

Please note that, this is a basic example and you may need to tweak the haproxy.cfg file and the args based on your use case.

Question: 14

Refer to Exhibit.

Set configuration context:

[student@node-1] $ kubectl contig use-context k8s

Context
Developers occasionally need to submit pods that run periodically.

Task
Follow the steps below to create a pod that will start at a predetermined time and]which runs to completion only once each time it is started:

- Create a YAML formatted Kubernetes manifest /opt/KDPD00301/periodic.yaml that runs the following shell command: date in a single busybox container. The command should run every minute and must complete within 22 seconds or be terminated oy Kubernetes. The Cronjob namp and container name should both be hello
- Create the resource in the above manifest and verify that the job executes successfully at least once

Answer: See the solution below.

Explanation:

Solution:

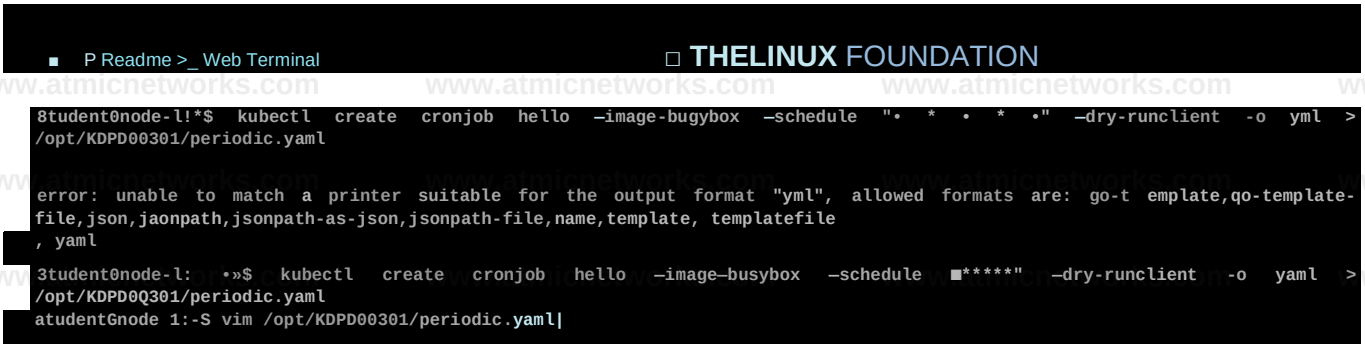

Question: 15
Refer to Exhibit.

Set configuration context*

^Ld^iiUni Je-l] S kubectl config u se-context nk8s

Task
A deployment is falling on the cluster due to an incorrect image being specified. Locate the deployment, and fix the problem.

Answer: See the solution below

Explanation:

create deploy hello-deploy --image=nginx --dry-run=client -o yaml > hello-deploy.yaml

Update deployment image to nginx:1.17.4: kubectl set image deploy/hello-deploy nginx=nginx:1.17.4

Question: 16

Refer to Exhibit.

Set configuration context:

stuj (’mu! 1] < kubectl config use-context nk8s

Set Configuration Context:

[student@node-1] $ | kubectl

Config use-context k8s

Task

You have rolled out a new pod to your infrastructure and now you need to allow it to communicate with the web and storage pods but nothing else. Given the running pod kdsn00201 -newpod edit it to use a network policy that will allow it to send and receive traffic only to and from the web and storage pods.

All work on this item should be conducted in the kdsn00201 namespace.

All required NetworkPolicy resources are already created and ready for use as appropriate. You should not create, modify or delete any network policies whilst completing this item.

Answer: See the solution below.

Explanation:

To allow a pod to send and receive traffic only to and from specific pods, you can use network policies in Kubernetes.

First, you will need to create a network policy that defines the allowed traffic. You can create a network policy yaml file with the following rules:

```
apiVersion: networking.k8s.io/v1
kind: NetworkPolicy metadata:
name: newpod-network-policy namespace: default
spec:
podSelector:
matchLabels:
    app: kdsn00201-newpod ingress:
- from:
  - podSelector: matchLabels: app: web
  - podSelector:
    matchLabels:
      app: storage
```

This policy will only allow incoming traffic to the pod with the label app=kdsn00201-newpod from pods with the label app=web or app=storage. If you have different labels on your web and storage pods please update the matchLabels accordingly.

Once you have created the network policy, you can apply it to the cluster by running the following command:

```
kubectl apply -f <network-policy-file>.yaml
```

This will apply the network policy to the cluster, and the newpod will only be able to send and receive traffic to and from the web and storage pods.

Please note that, NetworkPolicy resource is not available by default, you need to enable the NetworkPolicy feature on your Kubernetes cluster. This feature is enabled by default on some clusters and must be explicitly enabled on others. You can check if NetworkPolicy is available by running the command kubectl api-versions | grep networking

Also, you need to ensure that the pods that you want to allow traffic to and from are running on the same namespace.

Question: 17

Refer to Exhibit.

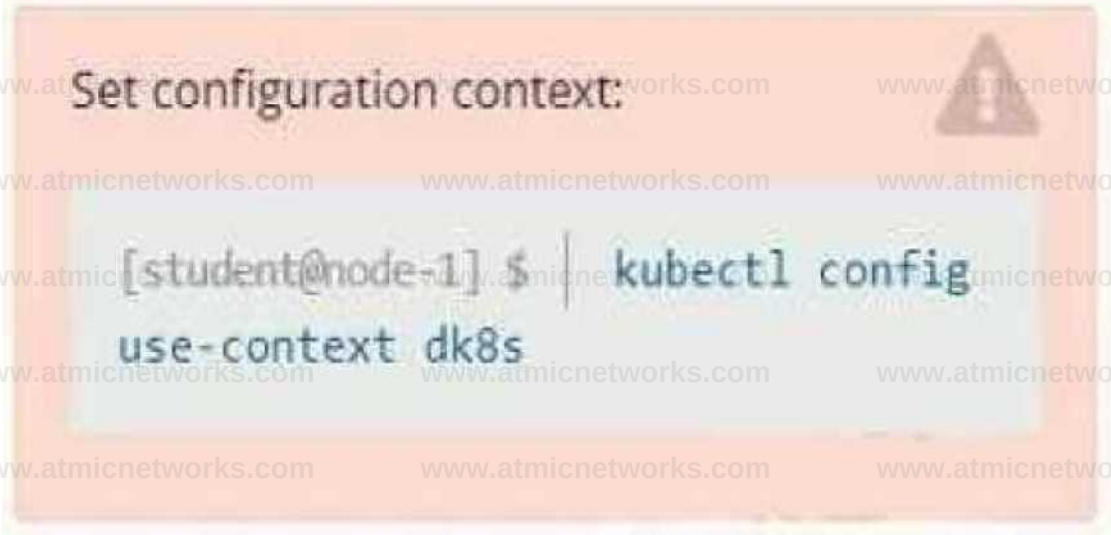

Set Configuration Context:
[student@node-1] $ | kubectl
Config use-context k8s
Context
A user has reported an aopticauon is unteachable due to a failing livenessProbe .

Task
Perform the following tasks:

- Find the broken pod and store its name and namespace to /opt/KDOB00401/broken.txt in the format: <namespace>/<pod>

<namespace>/<pod>

The output file has already been created

- Store the associated error events to a file /opt/KDOB00401/error.txt, The output file has already been created. You will need to use the -o wide output specifier with your command
- Fix the issue.

The associated deployment could oe running in an/of the following namespaces.

- qa
- test
- production
- alan

Answer: See the solution below.

Explanation:

To find the broken pod and store its name and namespace to /opt/KDOB00401/broken.txt, you can use the kubectl get pods command and filter the output by the status of the pod.

```
kubectl get pods --field-selector=status.phase=Failed -o jsonpath='{.items[*].metadata.namespace}/{.items[*].metadata.name}' > /opt/KDOB00401/broken.txt
```

This command will list all pods with a status of Failed and output their names and namespaces in the format <namespace>/<pod>. The output is then written to the /opt/KDOB00401/broken.txt file.

To store the associated error events to a file /opt/KDOB00401/error.txt, you can use the kubectl describe command to retrieve detailed information about the pod, and the grep command to filter the output for error events.

```
kubectl describe pods <pod-name> --namespace <pod-namespace> | grep -i error -B5 -A5 > /opt/KDOB00401/error.txt
```

Replace <pod-name> and <pod-namespace> with the name and namespace of the broken pod you found in the previous step.

This command will output detailed information about the pod, including error events. The grep command filters the output for lines containing "error" and also prints 5 lines before and after the match.

To fix the issue, you need to analyze the error events and find the root cause of the issue.

It could be that the application inside the pod is not running, the container image is not available, the pod has not enough resources, or the liveness probe configuration is incorrect.

Once you have identified the cause, you can take appropriate action, such as restarting the application, updating the container image, increasing the resources, or modifying the liveness probe configuration.

After fixing the issue, you can use the kubectl get pods command to check the status of the pod and ensure

Question: 18

Refer to Exhibit.

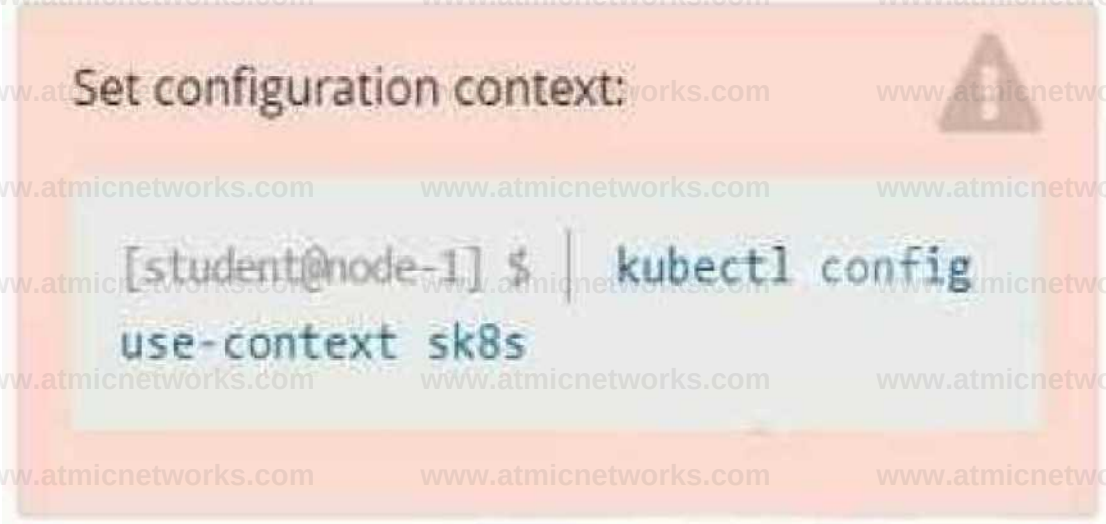

Context

A project that you are working on has a requirement for persistent data to be available.

Task

To facilitate this, perform the following tasks:

- Create a file on node sk8s-node-0 at /opt/KDSP00101/data/index.html with the content Acct=Finance
- Create a PersistentVolume named task-pv-volume using hostPath and allocate 1Gi to it, specifying that the volume is at /opt/KDSP00101/data on the cluster's node. The configuration should specify the access mode of ReadWriteOnce . It should define the StorageClass name exam for the PersistentVolume , which will be used to bind PersistentVolumeClaim requests to this PersistenetVolume.
- Create a PefsissentVolumeClaim named task-pv-claim that requests a volume of at least 100Mi

and specifies an access mode of ReadWriteOnce

- Create a pod that uses the PersistentVolmeClaim as a volume with a label app: my-storage-app

mounting the resulting volume to a mountPath /usr/share/nginx/html inside the pod

You can access sk8snode-O by issuing the following command:

/■ fu.Pnti>i^ 1] : ssh sk8 s-node 6

Ensure that you return to the base node (with hostname node-1) once you have completed your work on sk8s-node-0

©copy

Answer: See the solution below.

Explanation:

Solution:

BP Readme >_ Web Terminal □ THE **LINUX** FOUNDATION

```
student 0node-l:-v$ kubectl config use-context sk8s Switched to context "sk8a".
student9node-l:~S |
```

• Readme >_ Web Terminal

□ THE **LINUX** FOUNDATION

```
•  Documentation: https://help.ubuntu.coin
•  Management:     https://landscape.canonical.com
•  Support:       https://ubuntu.com/advantage

  System information as of Fxi Oct 9 08:52:09 UTC 2020

  System load:   2.02               Users logged in:         0
  Usage of /:   10.3% of 242.29GB IP      address for ethO:   10.250.3.115
  Memory usage:   2%                IP    address for dockerO: 172.17.0.1
  Swap usage:   0%                  IP    address for cniO:   10.244.1.1
  Processes:    38

•  Kubemetes 1.19 is out! Get it in one command with:

     sudo snap install microk8s — channel-1.19 --classic

   https://microk8a.io/ has docs and details.

7 packages can be updated.
1 update is a security update.

New release '20.04.1 LTS' available.
Run 'do-release-upgrade' to upgrade to it.


studentPskSs-node 0:-$ |
```

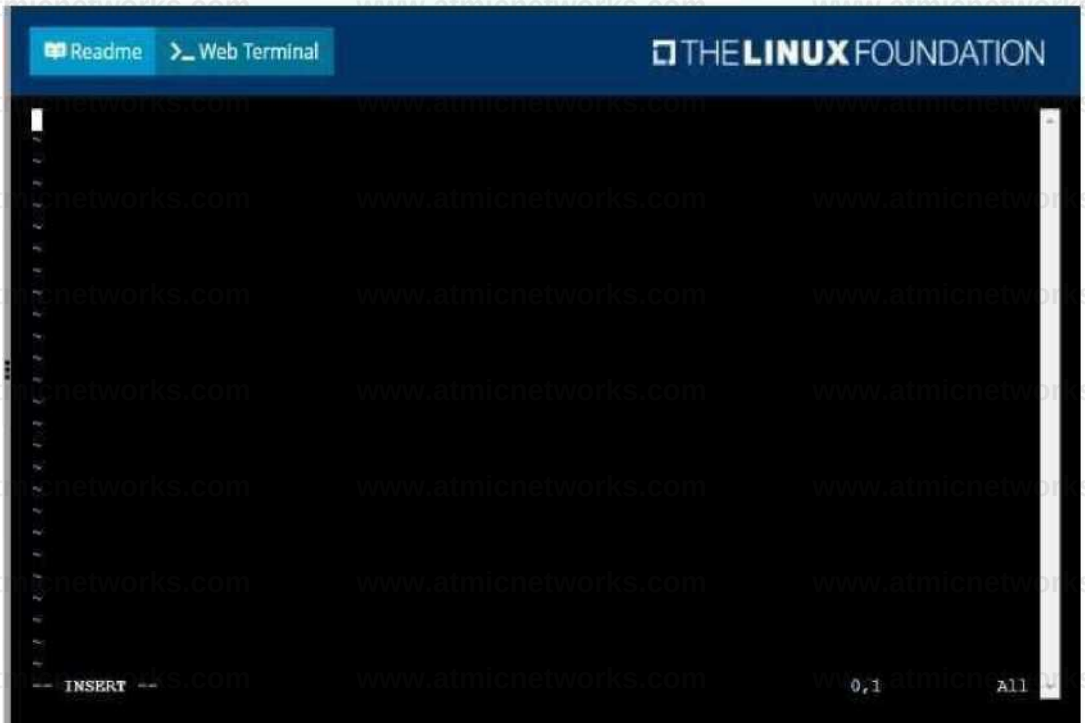

• Readme >_ Web Terminal □ THELINUX FOUNDATION

```
. vl PersiatentVolume metadata: nr: task-pv-volume
apse;
```

```
:*iii 1G1 acccssModes: - ReadWriteOnce storage hwtFath:
/opt/KDSP00101/data Director^
```

SB Readme >_ Web Terminal □ THELINUX FOUNDATION

```
vl PersxstentVolumoClaim met a da t a: i task-pv-claim
aseeaaModea: ReadWrxteOnce resources. requests: rt:raz- 100M1 - - storage
```

```
student0skSa-node-O:*$ kubectl create -f pv.yml persistentvolume/task-pv volume created
student0sk8a-node-Q:*$ kubectl create -f pvc.ynl
persistentvolumeclaxm/task-pv-claim created
studentPsk8s - node-0:-$ kubectl get pv NAME CAPACITY ACCESS MODES RECLAIM
POLICY STATUS CLAIM STO
RAGECLASS REASON AGE
task-pv-volume iGi RWO Retain Bound defaul t/task-pv-claim sto
rage Ils
student^skSs-node-0:-$ kubectl get pvc NAME STATUS VOLUME CAPACITY ACCESS MODES STORAGECLASS AGE
taek-pv-claim Bound task-pv-volume IGi RWO storage 9s
studentOskSs-node-O:-$ vim pod.ymlB
```


```
studentSskBs-node-O:*? kubectl create -f pod.yml pod/mypod created

atudcntfiakSa'nodc 0:--$ kubectl get|
```

■ Readme >_ Web Terminal □ THELINUX FOUNDATION

```
student0ak8s-node-O!*$ kubectl get pods NAME READY STATUS RESTARTS AGE
mypod 0/1 Containercreating 0 4s

student(?ak8s-node-0: ■*$ kubectl get pods NAME READY STATUS RESTARTS AGE
```

Question: 19

Refer to Exhibit.

Ser configuration context:

[snincntlfrwh- |j ; kubectl config use-context k8s

Given a container that writes a log file in format A and a container that converts log files from format A to format B, create a deployment that runs both containers such that the log files from the first container are converted by the second container, emitting logs in format B.

Task:

- Create a deployment named deployment-xyz in the default namespace, that:
- Includes a primary lfccncf/busybox:1 container, named logger-dev
- includes a sidecar lfccncf/fluentd:v0.12 container, named adapter-zen
- Mounts a shared volume /tmp/log on both containers, which does not persist when the pod is deleted
- Instructs the logger-dev container to run the command

```
while true; do
echo "i luv cncf" » /
tmp/log/input.log;
sleep 10;
done
```

which should output logs to /tmp/log/input.log in plain text format, with example values:

```
i luv cncf
i luv cncf
i luv cncf
```

- The adapter-zen sidecar container should read /tmp/log/input.log and output the data to /tmp/log/output.* in Fluentd JSON format. Note that no knowledge of Fluentd is required to complete this task: all you will need to achieve this is to create the ConfigMap from the spec

file provided at /opt/KDMC00102/fluentd-configma p.yaml , and mount that ConfigMap to /fluentd/etc in the adapter-zen sidecar container

Answer: See the solution below.

Explanation:

Solution:

SB Readme >_ Web Termi nal □THE **LINUX** FOUNDATION

```
        apps/vl Deployment
B-*a data• rr<j- _TL
  label 3:
        deployment-xyz
       deployment-xyz
replicas:
selector:
  match7,mbel 3;
        deployment-xyz
strategy; I
template:
  metadata:
    crentionT ma stamp ; *
        deployment-xyz
  3f.ee;
    ennramexa ■
          Ifccncf/busybox: 1
    busybox resources: |)
3t 3V15 . I I
"deployment_xyz.yml* 241, 4340                    3,1           All |
```

CP Readme >_ Web Terminal □ THE **LI N UX** FOUNDATION

```
  1 Deployment
metsoars: labels:
        deployment-xyz
deployment-xyz spec:
replicas
selector:
  ma t enZ-aba 13:
deployment-xyz
template;
  metadata:
    labels;
          deployment-xyz
  apes:
    volumes:
  - t     myvo11
    emprylix: containera: -
          Ifccncf/busybox:1
          logger-dev
    voluweNounts:
  - r. my veil /tmp/log
- ■        Ifccncf/fluentd:v0.12
    nar.* adapterflzen
3 lines yanked
                                          27,22         Bot |
```

BB Readme >_ Web Terminal □ **THELINUX** FOUNDATION

```
replicas: 1 selector;
  mtchLabela: deployment—xyz template:
  metadata.: LakeLc: deployment·xyz spec : volumes: myvoll enpty^iz: i • containers 1 -1fccncf/busybox:1
      .i   logger—dev
                    r •
  vc.uxeHouc.za myvoll i   /tmp/log
-       Ifccncf/fluentd:vO.12
        adapter-zen : .nnians -it vc .luxe Mount e -   -.- myvoll
    r -. • /tmp/log
                                                  29,83         Bot
```

• Readme >_ Web Terminal □ **THELINUX** FOUNDATION

```
meta datai labels;
        deployment-xyz spec: volumes -        - myvoll
```

```
eapryDxr: myvol2 confi qHap: . logconf coptainers;
-  .     Ifccncf/buaybox:1
         logger-dev

voluraeMounto:

-       myvoll
        /tmp/log
       1fccncf/fluentd:v0.12 adapter—zen rcroona     ,        ,
-  ɔ 4 mse-Mounta.

-       myvol1
        /tmp/log
-       myvol2
        Zfluentd/etJ
                                                    37,33         Bot
```

```
3tudentSnode-l:»$ kubectl create -f deploymentxyz.yml deployment.appa/depioymcnt-xyz created
student@node~lt*$ kubectl get deployment NAME          READY    UP-TO-DATE     AVAILABLE
                  AGE
deployment-xyz    0/1      1             0            5a


studentSnode-1:*$ kubectl get deployment NAME          READY    UP-TO-DATE     AVAILABLE
                  AGE
deployment-xyz    0/1      1             0             9a


student?node-l:~$ kubectl get deployment NAME          READY    UP-TO-DATE     AVAILABLE
                  AGE
deployment-xyz    1/1      1             1            123
Btudent0node-1:~$ H
```

Question: 20

Refer to Exhibit.

You must switch to the correct

cluster/conf iguration context Failure to do so may result in a zero score.

ii [1] : i kubectl config use c

ontext skAs

Task

A Deployment named backend-deployment in namespace staging runs a web application on port 8081.

< *The Deployment's manifest files con be found ot*

^/spicy-pikathu/backend-Oeplayment.yciml

Modify :he Deployment specifying a readiness prooe us ng path /heahh/.

Set initialDelaySecond*. to B and peoodSeconds to 5

Answer: See the solution below.

Explanation:

Solution:

```
File  Edit  View  Terminal  Tabs  Help
Warning: Permanently added '172.31.17.21' (ECDSA) to the list of known hosts.

The programs included with the Ubuntu system are free software;
the exact distribution terms for each program are described in the
individual files in /usr/share/doc/*/copyright.

Ubuntu comes with ABSOLUTELY NO WARRANTY, to the extent permitted by
applicable law.

candidate@node-1:~$ vi ~/spicy-pikachu/backend-deployment.yaml
candidate@node-1:~$ kubectl config use-context sk8s
Switched to context "sk8s".
candidate@node-1:~$ vim .vimrc
candidate@node-1:~$ vim ~/spicy-pikachu/backend-deployment.yaml
```

```
File  Edit  View  Terminal  Tabs  Help
apiVersion: apps/v1
kind: Deployment
metadata:
  name: backend-deployment
  namespace: staging
spec:
  selector:
    matchLabels:
      app: nginx
  replicas: 3
  template:
    metadata:
      labels:
        app: nginx
    spec:
      containers:
        - name: nginx
          image: nginx:1.14.2
          ports:
            - containerPort: 8081
          readinessProbe:
              initialDelaySeconds: 8
              periodSeconds: 5
              httpGet:
                 path: /healthz
                 port: 8081
          volumeMounts:
            - mountPath: /etc/nginx/conf.d/
              name: config
            - mountPath: /usr/share/nginx/html/
              name: www
-- INSERT --                                                    26,28         Top
```

```
File  Edit  View  Terminal  Tabs  Help
Warning: Permanently added '172.31.17.21' (ECDSA) to the list of known hosts.

The programs included with the Ubuntu system are free software;
the exact distribution terms for each program are described in the
individual files in /usr/share/doc/*/copyright.

Ubuntu comes with ABSOLUTELY NO WARRANTY, to the extent permitted by
applicable law.

candidate@node-1:~$ vi ~/spicy-pikachu/backend-deployment.yaml
candidate@node-1:~$ kubectl config use-context sk8s
Switched to context "sk8s".
candidate@node-1:~$ vim .vimrc
candidate@node-1:~$ vim ~/spicy-pikachu/backend-deployment.yaml
candidate@node-1:~$ kubectl  apply  -f ~/spicy-pikachu/backend-deployment.yaml
deployment.apps/backend-deployment configured
candidate@node-1:~$ kubectl get  pods -n staging
NAME                                  READY   STATUS    RESTARTS   AGE
backend-deployment-59d449b99d-cxct6   1/1     Running   0          20s
backend-deployment-59d449b99d-h2zjq   0/1     Running   0          9s
backend-deployment-78976f74f5-b8c85   1/1     Running   0          6h40m
backend-deployment-78976f74f5-flfsj   1/1     Running   0          6h40m
candidate@node-1:~$ kubectl get deploy -n staging
NAME                 READY   UP-TO-DATE   AVAILABLE   AGE
backend-deployment   3/3     3            3           6h40m
candidate@node-1:~$ kubectl get deploy -n staging
NAME                 READY   UP-TO-DATE   AVAILABLE   AGE
backend-deployment   3/3     3            3           6h41m
candidate@node-1:~$ vim ~/spicy-pikachu/backend-deployment.yaml
```

Question: 21

Refer to Exhibit.

You must switch to the correct duster/configuration context Failure to do so may result in a zero score.

ii [1] : i' kubectl config use c ontext skHs

Task:

Update the Deployment app-1 in the frontend namespace to use the existing ServiceAccount app.

Answer: See the solution below.

Explanation:

Solution:

Question: 22

Refer to Exhibit.

You must switch to the correct cluster/configu ration context. Failure to do so may result tn a zero score.

ii ■ > i fr kubuctl config use ont^xt skHs

Task:

A pod within the Deployment named buffale-deployment and in namespace gorilla is logging errors.

1) Look at the logs identify errors messages.

Find errors, including User “system:serviceaccount:gorilla:default” cannot list resource “deployment” [...] in the namespace “gorilla”

2) Update the Deployment buffalo-deployment to resolve the errors in the logs of the Pod.

The buffalo-deployment ‘S manifest can be found at -/prompt/escargot/buffalo-deployment.yaml

Answer: See the solution below.

Explanation:

Solution:

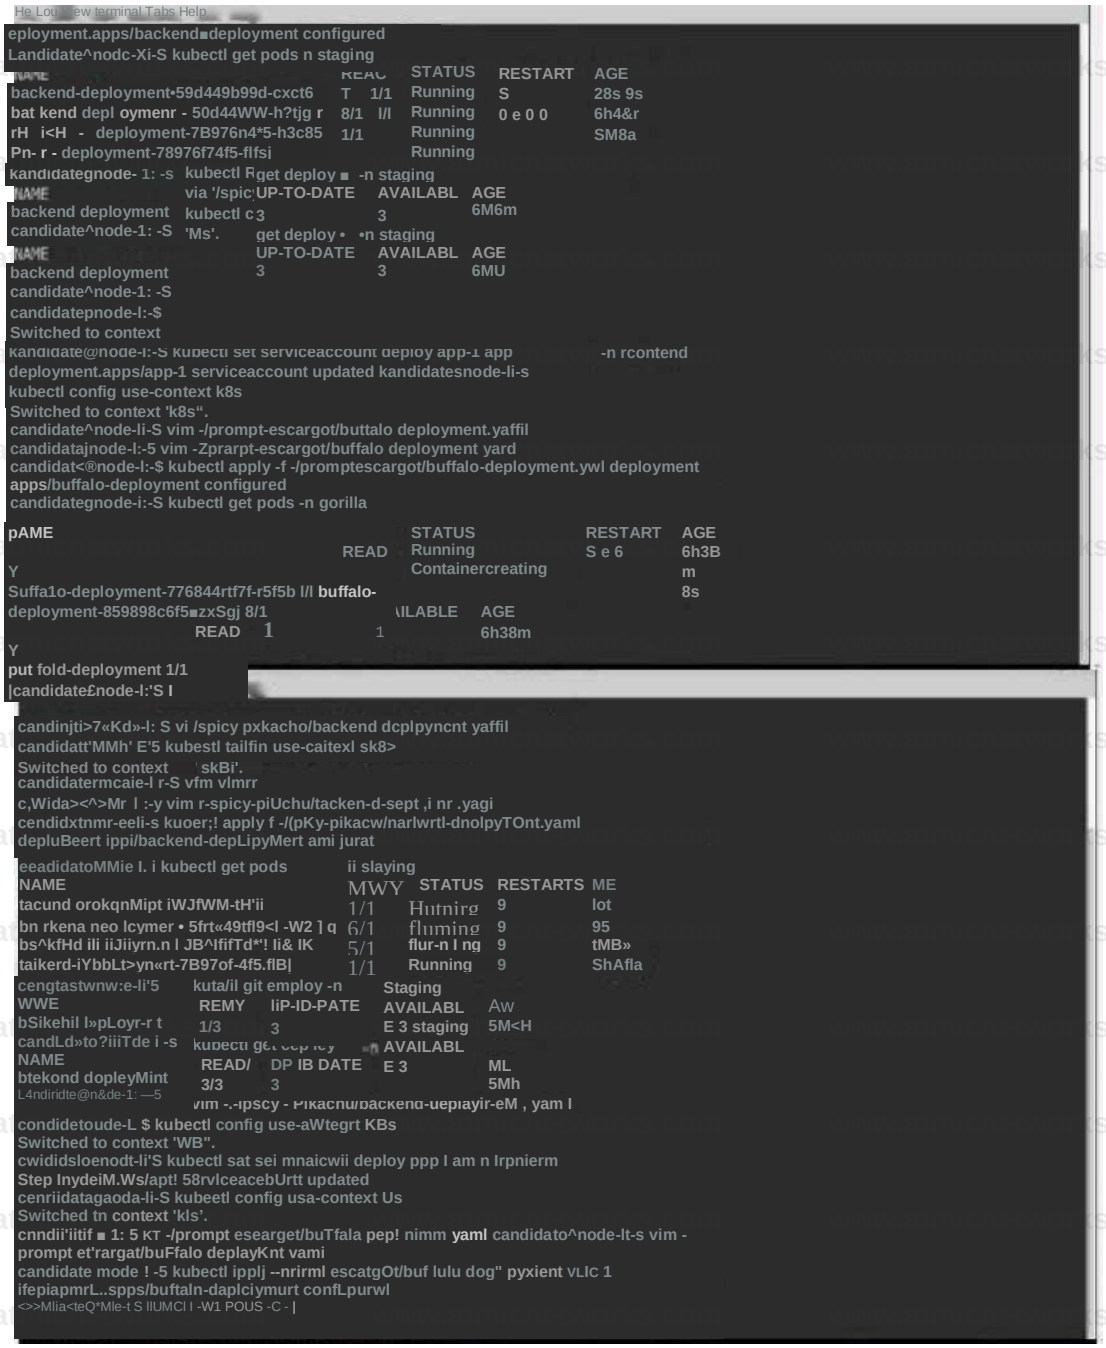

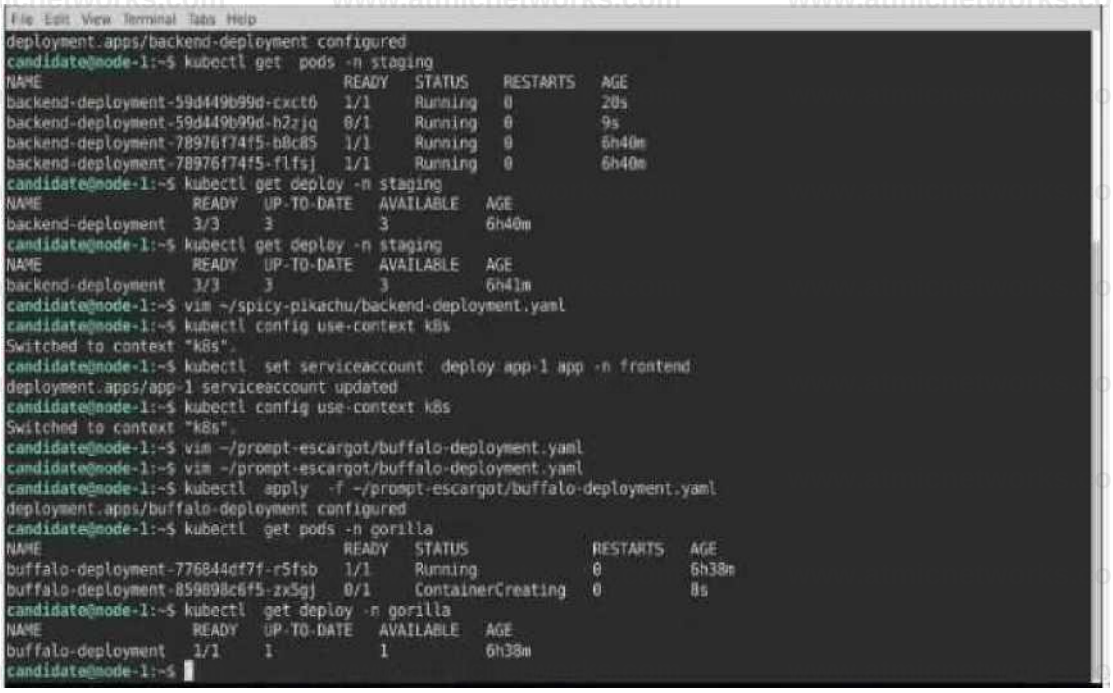

```
File  Edit  View  Terminal  Tabs  Help
deployment.apps/backend-deployment configured
candidate@node-1:~$ kubectl get  pods -n staging
NAME                                  READY   STATUS    RESTARTS   AGE
backend-deployment-59d449b99d-cxct6   1/1     Running   0          20s
backend-deployment-59d449b99d-h2zjq   0/1     Running   0          9s
backend-deployment-78976f74f5-b8c85   1/1     Running   0          6h40m
backend-deployment-78976f74f5-flfsj   1/1     Running   0          6h40m
candidate@node-1:~$ kubectl get deploy -n staging
NAME                 READY   UP-TO-DATE   AVAILABLE   AGE
backend-deployment   3/3     3            3           6h40m
candidate@node-1:~$ kubectl get deploy -n staging
NAME                 READY   UP-TO-DATE   AVAILABLE   AGE
backend-deployment   3/3     3            3           6h41m
candidate@node-1:~$ vim ~/spicy-pikachu/backend-deployment.yaml
candidate@node-1:~$ kubectl config use-context k8s
Switched to context "k8s".
candidate@node-1:~$ kubectl  set serviceaccount  deploy app-1 app -n frontend
deployment.apps/app-1 serviceaccount updated
candidate@node-1:~$ kubectl config use-context k8s
Switched to context "k8s".
candidate@node-1:~$ vim ~/prompt-escargot/buffalo-deployment.yaml
candidate@node-1:~$ vim ~/prompt-escargot/buffalo-deployment.yaml
candidate@node-1:~$ kubectl  apply  -f ~/prompt-escargot/buffalo-deployment.yaml
deployment.apps/buffalo-deployment configured
candidate@node-1:~$ kubectl  get pods -n gorilla
NAME                                  READY   STATUS              RESTARTS   AGE
buffalo-deployment-776844df7f-r5fsb   1/1     Running             0          6h38m
buffalo-deployment-859898c6f5-zx5gj   0/1     ContainerCreating   0          8s
candidate@node-1:~$ kubectl  get deploy -n gorilla
NAME                 READY   UP-TO-DATE   AVAILABLE   AGE
buffalo-deployment   1/1     1            1           6h38m
candidate@node-1:~$
```

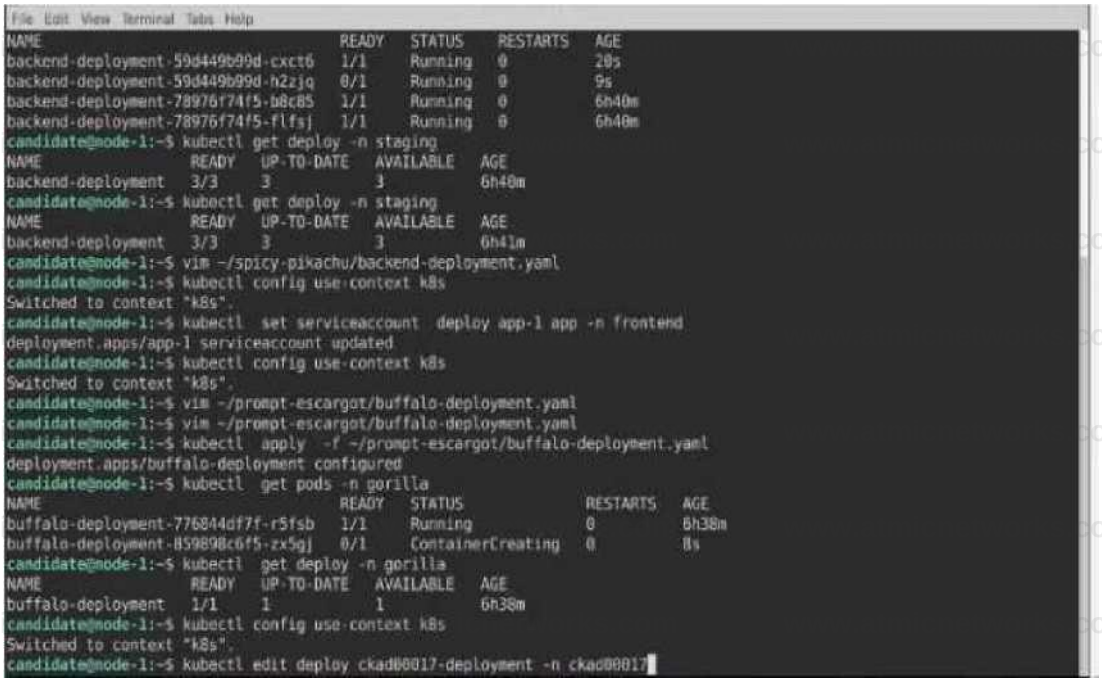

```
File  Edit  View  Terminal  Tabs  Help
NAME                                  READY   STATUS    RESTARTS   AGE
backend-deployment-59d449b99d-cxct6   1/1     Running   0          20s
backend-deployment-59d449b99d-h2zjq   0/1     Running   0          9s
backend-deployment-78976f74f5-b8c85   1/1     Running   0          6h40m
backend-deployment-78976f74f5-flfsj   1/1     Running   0          6h40m
candidate@node-1:~$ kubectl get deploy -n staging
NAME                 READY   UP-TO-DATE   AVAILABLE   AGE
backend-deployment   3/3     3            3           6h40m
candidate@node-1:~$ kubectl get deploy -n staging
NAME                 READY   UP-TO-DATE   AVAILABLE   AGE
backend-deployment   3/3     3            3           6h41m
candidate@node-1:~$ vim ~/spicy-pikachu/backend-deployment.yaml
candidate@node-1:~$ kubectl config use-context k8s
Switched to context "k8s".
candidate@node-1:~$ kubectl  set serviceaccount  deploy app-1 app -n frontend
deployment.apps/app-1 serviceaccount updated
candidate@node-1:~$ kubectl config use-context k8s
Switched to context "k8s".
candidate@node-1:~$ vim ~/prompt-escargot/buffalo-deployment.yaml
candidate@node-1:~$ vim ~/prompt-escargot/buffalo-deployment.yaml
candidate@node-1:~$ kubectl  apply  -f ~/prompt-escargot/buffalo-deployment.yaml
deployment.apps/buffalo-deployment configured
candidate@node-1:~$ kubectl  get pods -n gorilla
NAME                                  READY   STATUS              RESTARTS   AGE
buffalo-deployment-776844df7f-r5fsb   1/1     Running             0          6h38m
buffalo-deployment-859898c6f5-zx5gj   0/1     ContainerCreating   0          8s
candidate@node-1:~$ kubectl  get deploy -n gorilla
NAME                 READY   UP-TO-DATE   AVAILABLE   AGE
buffalo-deployment   1/1     1            1           6h38m
candidate@node-1:~$ kubectl config use-context k8s
Switched to context "k8s".
candidate@node-1:~$ kubectl edit deploy ckad00017-deployment -n ckac00017
```

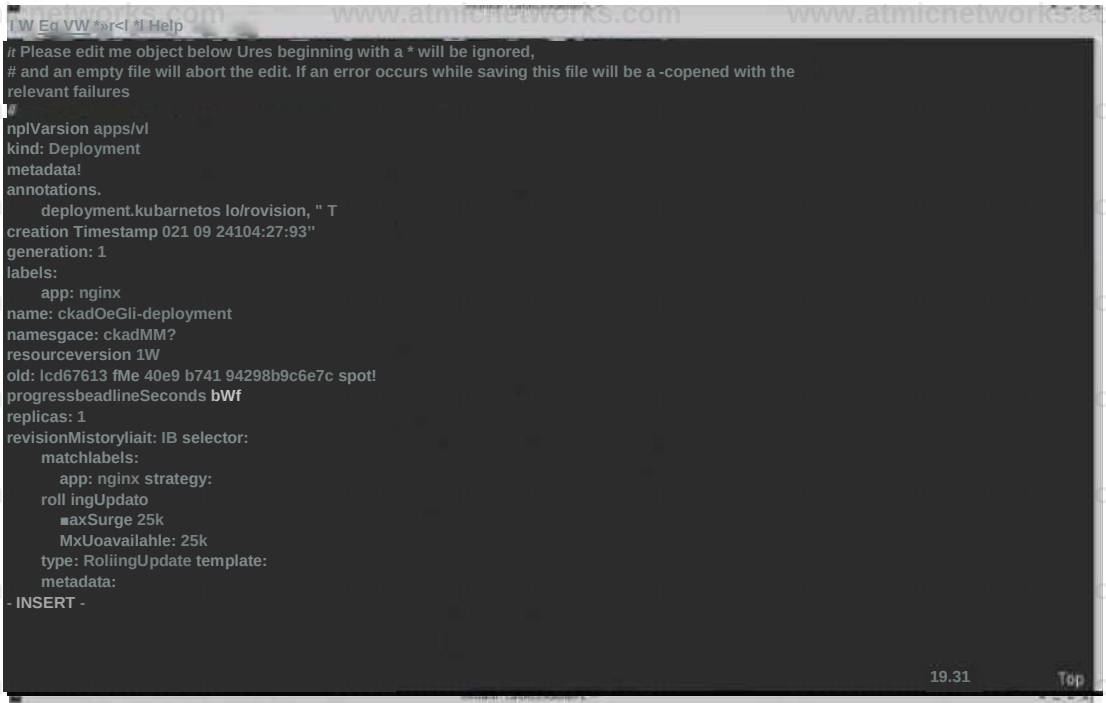
I W Eg VW *»r<I *I Help
it Please edit me object below Ures beginning with a * will be ignored,
and an empty file will abort the edit. If an error occurs while saving this file will be a -copened with the
relevant failures
#
nplVarsion apps/vl
kind: Deployment
metadata!
annotations.
deployment.kubarnetos lo/rovision, " T
creation Timestamp 021 09 24104:27:93"
generation: 1
labels:
app: nginx
name: ckadOeGli-deployment
namesgace: ckadMM?
resourceversion 1W
old: lcd67613 fMe 40e9 b741 94298b9c6e7c spot!
progressbeadlineSeconds bWf
replicas: 1
revisionMistoryliait: IB selector:
matchlabels:
app: nginx strategy:
roll ingUpdato
■axSurge 25k
MxUoavailahle: 25k
type: RoliingUpdate template:
metadata:
- INSERT -
19.31
Top

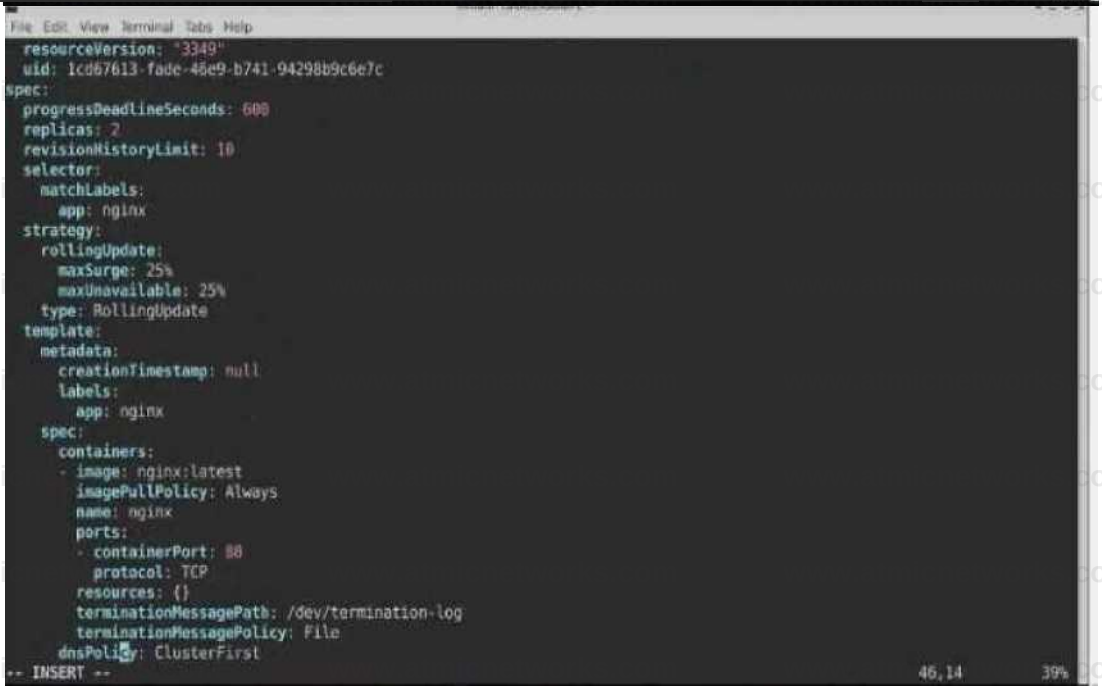
File Edit View Terminal Tabs Help
resourceVersion: "3349"
uid: 1cd67613-fade-40e9-b741-94298b9c6e7c
spec:
progressDeadlineSeconds: 600
replicas: 2
revisionHistoryLimit: 10
selector:
matchLabels:
app: nginx
strategy:
rollingUpdate:
maxSurge: 25%
maxUnavailable: 25%
type: RollingUpdate
template:
metadata:
creationTimestamp: null
labels:
app: nginx
spec:
containers:
- image: nginx:latest
imagePullPolicy: Always
name: nginx
ports:
- containerPort: 80
protocol: TCP
resources: {}
terminationMessagePath: /dev/termination-log
terminationMessagePolicy: File
dnsPolicy: ClusterFirst
-- INSERT --
46,14
39%

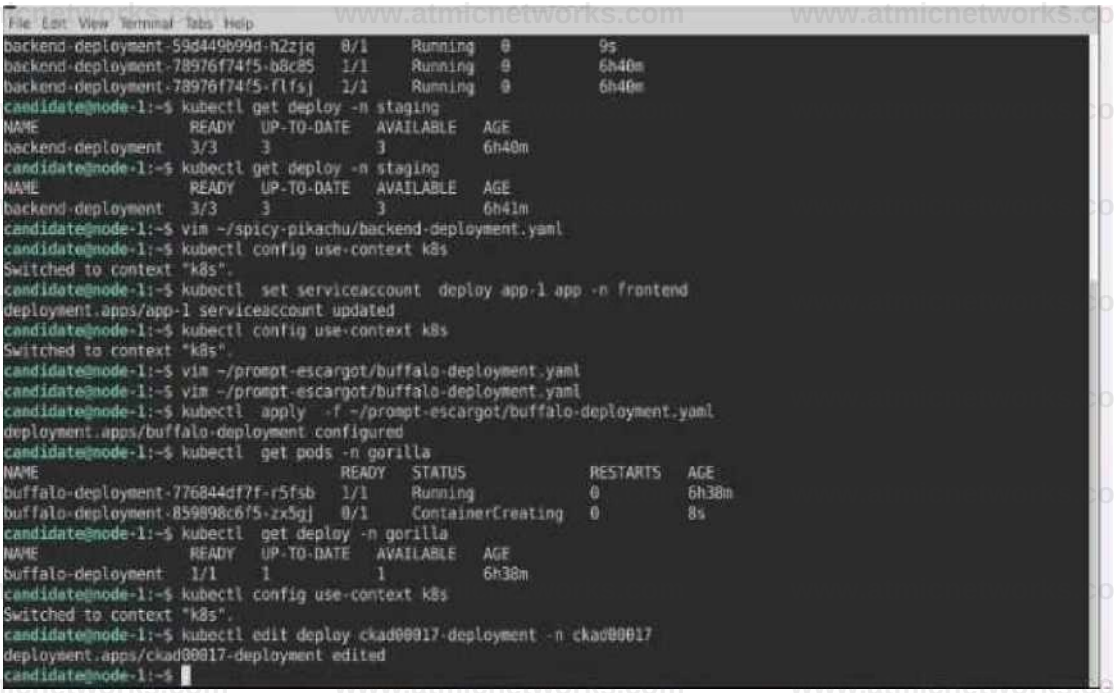

```
File  Edit  View  Terminal  Tabs  Help
backend-deployment-59d449b99d-h2zjq   0/1     Running   0          9s
backend-deployment-78976f74f5-b8c85   1/1     Running   0          6h40m
backend-deployment-78976f74f5-flfsj   1/1     Running   0          6h40m
candidate@node-1:~$ kubectl get deploy -n staging
NAME                 READY   UP-TO-DATE   AVAILABLE   AGE
backend-deployment   3/3     3            3           6h40m
candidate@node-1:~$ kubectl get deploy -n staging
NAME                 READY   UP-TO-DATE   AVAILABLE   AGE
backend-deployment   3/3     3            3           6h41m
candidate@node-1:~$ vim ~/spicy-pikachu/backend-deployment.yaml
candidate@node-1:~$ kubectl config use-context k8s
Switched to context "k8s".
candidate@node-1:~$ kubectl  set serviceaccount  deploy app-1 app -n frontend
deployment.apps/app-1 serviceaccount updated
candidate@node-1:~$ kubectl config use-context k8s
Switched to context "k8s".
candidate@node-1:~$ vim ~/prompt-escargot/buffalo-deployment.yaml
candidate@node-1:~$ vim ~/prompt-escargot/buffalo-deployment.yaml
candidate@node-1:~$ kubectl  apply  -f ~/prompt-escargot/buffalo-deployment.yaml
deployment.apps/buffalo-deployment configured
candidate@node-1:~$ kubectl  get pods -n gorilla
NAME                                 READY   STATUS              RESTARTS   AGE
buffalo-deployment-776844df7f-r5fsb  1/1     Running             0          6h38m
buffalo-deployment-859898c6f5-zx5gj  0/1     ContainerCreating   0          8s
candidate@node-1:~$ kubectl  get deploy -n gorilla
NAME                 READY   UP-TO-DATE   AVAILABLE   AGE
buffalo-deployment   1/1     1            1           6h38m
candidate@node-1:~$ kubectl config use-context k8s
Switched to context "k8s".
candidate@node-1:~$ kubectl edit deploy ckad00017-deployment -n ckad00017
deployment.apps/ckad00017-deployment edited
candidate@node-1:~$
```

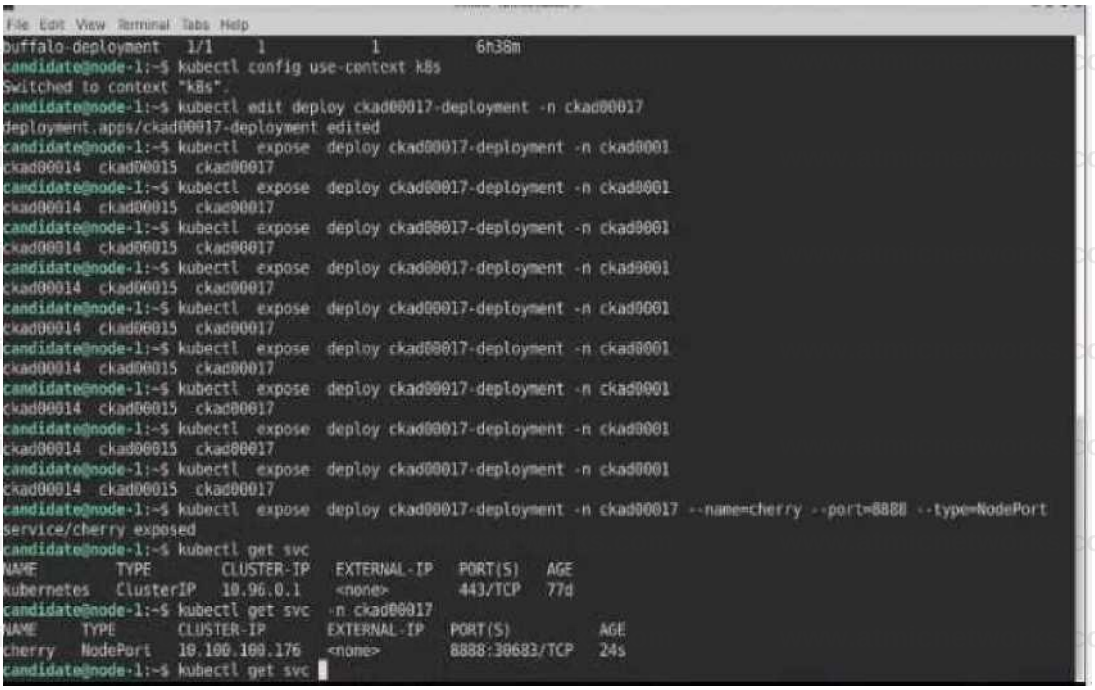

```
File  Edit  View  Terminal  Tabs  Help
buffalo-deployment   1/1     1            1           6h38m
candidate@node-1:~$ kubectl config use-context k8s
Switched to context "k8s".
candidate@node-1:~$ kubectl edit deploy ckad00017-deployment -n ckad00017
deployment.apps/ckad00017-deployment edited
candidate@node-1:~$ kubectl  expose  deploy ckad00017-deployment -n ckad0001
ckad00014  ckad00015  ckad00017
candidate@node-1:~$ kubectl  expose  deploy ckad00017-deployment -n ckad0001
ckad00014  ckad00015  ckad00017
candidate@node-1:~$ kubectl  expose  deploy ckad00017-deployment -n ckad0001
ckad00014  ckad00015  ckad00017
candidate@node-1:~$ kubectl  expose  deploy ckad00017-deployment -n ckad0001
ckad00014  ckad00015  ckad00017
candidate@node-1:~$ kubectl  expose  deploy ckad00017-deployment -n ckad0001
ckad00014  ckad00015  ckad00017
candidate@node-1:~$ kubectl  expose  deploy ckad00017-deployment -n ckad0001
ckad00014  ckad00015  ckad00017
candidate@node-1:~$ kubectl  expose  deploy ckad00017-deployment -n ckad0001
ckad00014  ckad00015  ckad00017
candidate@node-1:~$ kubectl  expose  deploy ckad00017-deployment -n ckad0001
ckad00014  ckad00015  ckad00017
candidate@node-1:~$ kubectl  expose  deploy ckad00017-deployment -n ckad00017 --name=cherry --port=8888 --type=NodePort
service/cherry exposed
candidate@node-1:~$ kubectl get svc
NAME         TYPE        CLUSTER-IP    EXTERNAL-IP   PORT(S)    AGE
kubernetes   ClusterIP   10.96.0.1     <none>        443/TCP    77d
candidate@node-1:~$ kubectl get svc -n ckad00017
NAME     TYPE       CLUSTER-IP       EXTERNAL-IP   PORT(S)          AGE
cherry   NodePort   10.100.100.176   <none>        8888:30683/TCP   24s
candidate@node-1:~$ kubectl get svc
```

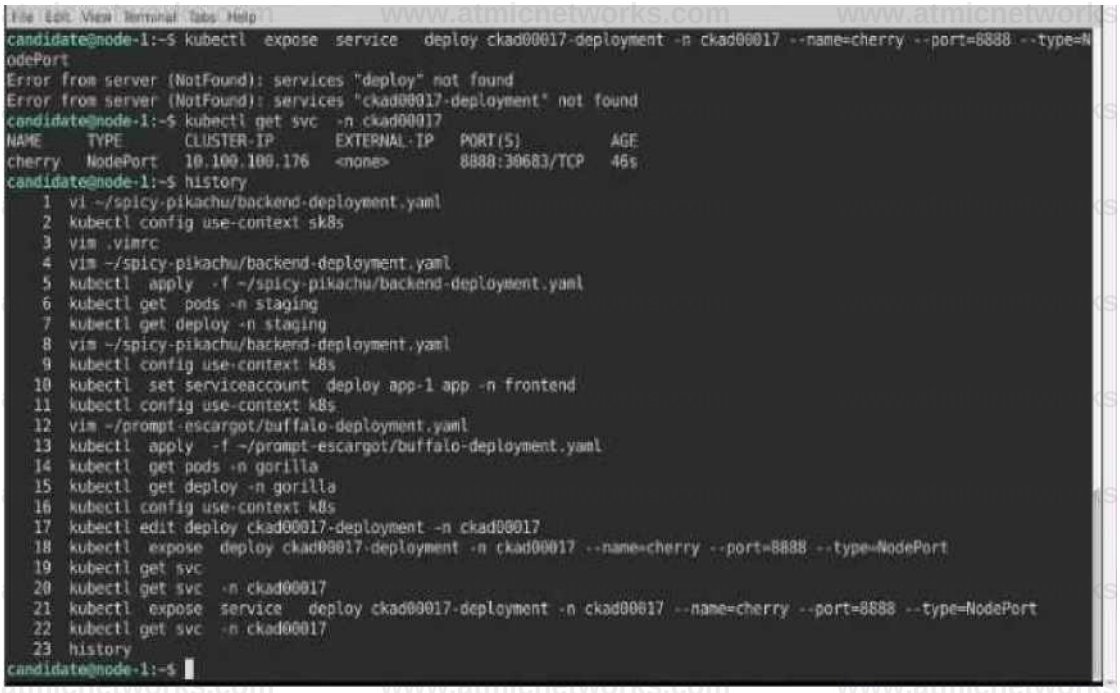

Question: 23

Refer to Exhibit.

You must switch to the correct cluster/configuration context. Failure to do may result tn a zero score.

ii ■ > i fi kubectl config [1]

ont^xt skHs

Task:

1) First update the Deployment cka00017-deployment in the ckad00017 namespace:

*To run 2 replicas of the pod

*Add the following label on the pod:

Role userUI

2) Next, Create a NodePort Service named cherry in the ckad00017 nmespace exposing the ckad00017-deployment Deployment on TCP port 8888

Answer: See the solution below.

Explanation:

Solution:

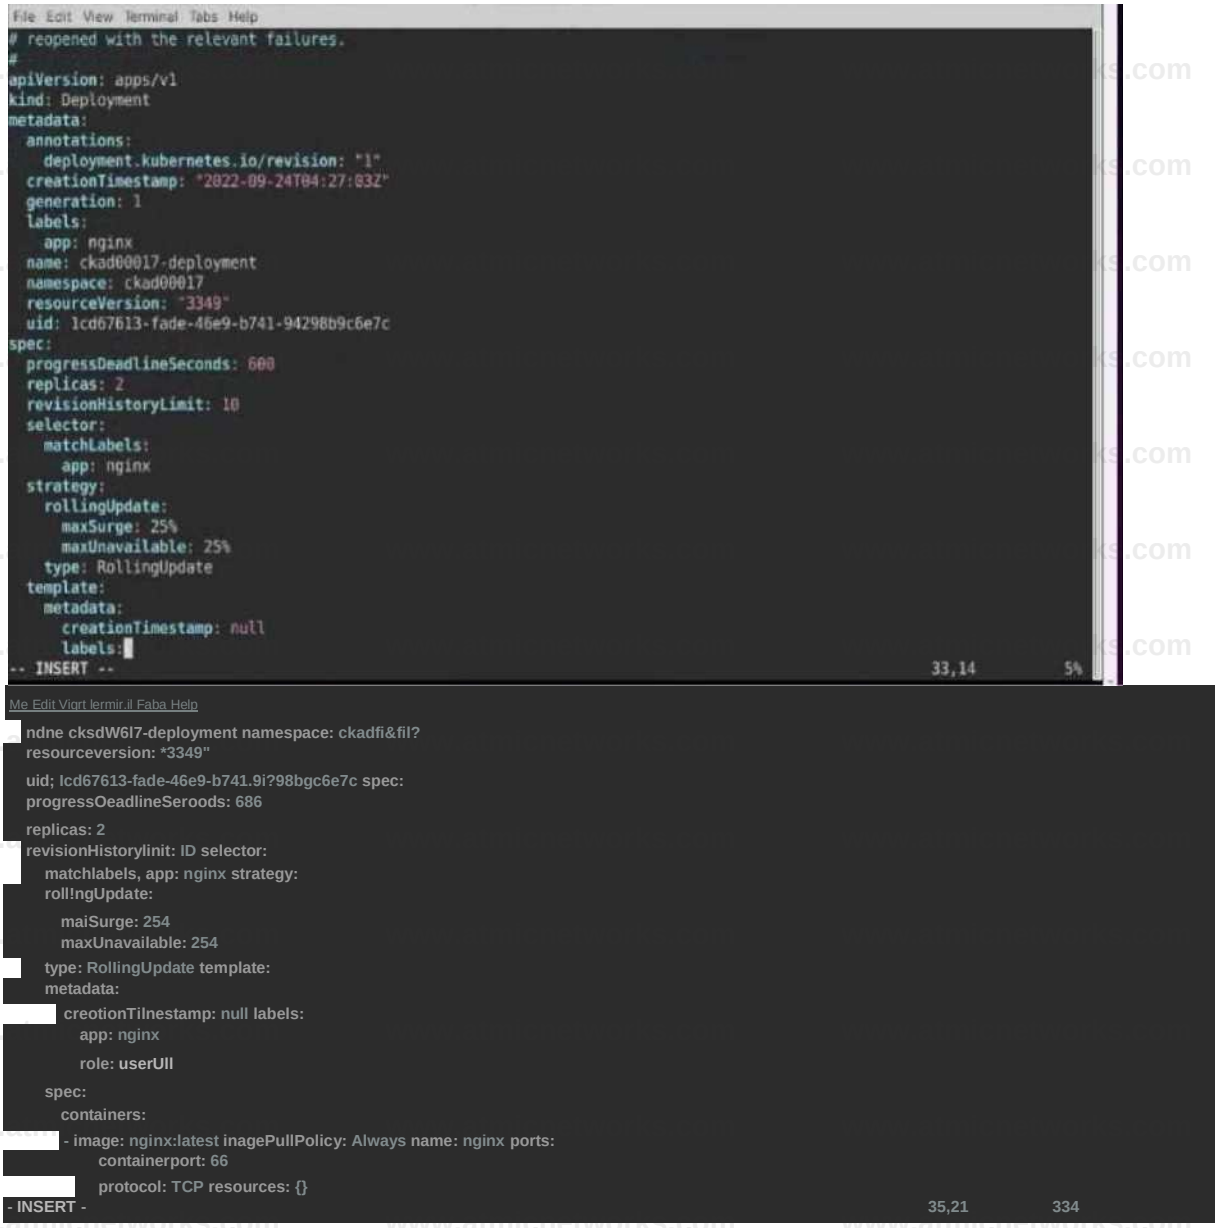

```
File  Edit  View  Terminal  Tabs  Help
# reopened with the relevant failures.
#
apiVersion: apps/v1
kind: Deployment
metadata:
  annotations:
    deployment.kubernetes.io/revision: "1"
  creationTimestamp: "2022-09-24T04:27:03Z"
  generation: 1
  labels:
    app: nginx
  name: ckad00017-deployment
  namespace: ckad00017
  resourceVersion: "3349"
  uid: 1cd67613-fade-46e9-b741-94298b9c6e7c
spec:
  progressDeadlineSeconds: 600
  replicas: 2
  revisionHistoryLimit: 10
  selector:
    matchLabels:
      app: nginx
  strategy:
    rollingUpdate:
      maxSurge: 25%
      maxUnavailable: 25%
    type: RollingUpdate
  template:
    metadata:
      creationTimestamp: null
      labels:
-- INSERT --                                                    33,14         5%
```



```
Me Edit Vigrt lermir.il Faba Help
  ndne cksdW6l7-deployment namespace: ckadfi&fil?
  resourceversion: *3349"
  uid; lcd67613-fade-46e9-b741.9i?98bgc6e7c spec:
  progressOeadlineSeroods: 686
  replicas: 2
  revisionHistorylinit: ID selector:
    matchlabels, app: nginx strategy:
    roll!ngUpdate:
      maiSurge: 254
      maxUnavailable: 254
    type: RollingUpdate template:
    metadata:
      creotionTilnestamp: null labels:
        app: nginx
        role: userUll
    spec:
      containers:
      - image: nginx:latest inagePullPolicy: Always name: nginx ports:
          containerport: 66
          protocol: TCP resources: {}
- INSERT -                                                    35,21         334
```

```
I * Ie Edit View terminal abs Help
backend Deployment-59d449D99d-n2ijq      9/1      Running     9          9s
backend-Deployment 78916174'5-b8c85      1/1      Running     0          SMto
backend deployment 78976174'5-fits]      1/1      Running     0          6M0n
candidate^nade-1;-S kubectl get deploy -n staging KAME      READY      UP-TO-DATE          AVAILABLE AGE
backend -deployment    3/3      3                3           6h46m
candidate^node-I:-S kubectl get deploy -n staging NAME      READY      UP-TO-DATE          AVAILABLE AGE
backend-deployment     3/3      3                3           6h41m
candidateanode-1:-S vim -/spicy-pikachu/backend-deployment.yanl candidate®node-I:~5 kubectl config use-context k8s Switched to context 'kas".
candidateanode-1:-S kubectl set serviceaccount deploy app 1 apo -n frontend deployment.apps/app-1 serviceaccount updated tandidate^node-1:-$ kubectl
contig use-context k8s Switched to context "k8s".
candidateanode-I:-$ vim -/prompt-escargot/buffalp-deployment .yawl candidategnoda-1:s vim -/prompt■ escargot/buf falo deployment.yMl candidate0node-I:-5
kubectl apply f -/pronpt-escargot/buffalo-deployment.ywl deployment .apps/bu!lain-deplbyWnL configured candidateinode-I:-s kubectl get pods -n gorilka
KAME                                      READY   STATUS              RESTARTS  AGE
buffalo deployment■776844df?f-rSfsb       1/1     Running             0         6038m
buffalodeployment-859898c6fS-zxSgJ        fl/1    Containercreating   0         8s
»ndidate@ncde-I:-S kubectl get deploy -n gorilla KAME      READY      UP-TO-DATE          AVAILABLE AGE
buffalo-deployment     1/1      1                1           6h38m
candidate^node-1:-S kubectl config use context k8s Switched to context "kas" candidate0node-I:-S kubectl edit deploy ckadSMl?-deployment -n ckadiMl?
deployment .apps/ckadSM17-deployment edited candidatejnode-li-S |
```

```
He Ice   View lermiA&I aba Help
candidate$node-I:~S kubectl get pods -n gorilla
NAME                                      READY  STATUS
                                          RESTARTS
buffalo-deployment 776844uf7f rSIsb       1/1      Running                0
but falo-oeployment-859898c6t5-tx5gj      8/1      Containertreating      e
candidategnode-I:-$ kubectl get deploy -n gorilla NAMEREADY UP-TO-DATE
                          AVAILABLE              AGE
buffalo-deployment     1/11                      1           6938m
candidate@node-I:~S kubectl config use-context kgs Switched to context "kSs".
candidategnode-1:-I kubectl edit deploy ckadMS 17-deployment -n ckadW17
deployment.apps/ckad80817-   deployment edited
candidate®node-I:-S    kubectl    expose deploy ckadOOOL? deployment -n cliad6901
ckadeeoi4 ckadmis ckauOBOl?
candidate$node-I:-S    kubectl    expose deploy ckad0D917-deployment -n ckad0001
ckad00D14      cknd00915      ckat!B0017
candidategnode-I:~3    kubectl    expose deploy ckad99017-deployment -n ckad0001
ckadS0O14      ckadQGDlS      ckao0O017
candidategncde-lf-S    kubectl    expose deploy ckad0G017-deployment -n ckadBSSI
ckad00O14      ckad00B15      ckadBB617
candidategnode-L:-$    kubectl    expose deploy ckad9901?-i!eploywnt -n ckadOBBl
ckadB0Bl4      ckad099l5      ckan00B17
candidate@node-I:-$    kubectl    expose c deploy ckaHKIT-deployeent -n ckadBBel
<0000314       ckadG0915      ckad90917
candidategnode-li-S    kubectl    expose deploy ckadB0B17-deployment -n ckad0001
ckadoesu       ckadG0915      ckao9G017
candidataanode-1:-s    kubectl    expose deploy ckad90017-deployment -n ckadSGSI
ckadWeiA       ckad60915      ckanGG017
candidate@node-I:-$    kubectl    expose deploy ckad00017 deployment n ckadBOGI
c<0000314      ckadmis        ckab00017
candidate^node-li-S    kubectl    expose deploy ck«|90917-deployment -n ckad00E17 name=cherry -port-8888 type=NodePort
service/cherry exposed candidate^node-
1:~S |
```

```
candidatepncde-I:-S kubectl get sue
NAME          TYPE          CLUSTER IP     EXTERNAL IP POfIT(S)     AGE
kubernetes    ClusterIP     19 96.8.1      <none> 443/TCP           77d
candidate^ncde-li-S kubectl get sue -n ckWM017
NAME     TYPE          CLUSTER-IP        EXTERNAL   IP PORTIS)          AGE
cherry   NodePort      10,100,180.176    <none> 8888:30683/TCP          24s
randidate?node-I:-s kubectl expose Service deploy (kadODBl7-deployment n Ckad00G17 -name-Cherry --port-8888 -type-N
odePort
Error from server <Not Found): services "deploy' not found
Error from server (NotFcund): services "-cka<fM917 deployment' not found
candidate@node-I:-S kubectl get sve -n ckadM017
NAME     TYPE          CLUSTER-IP        EXTERNAL   IP PORTIS)          AGE
cherry   NodePort      18,100.100.176    <nones 8888:36683/TCP          46s
candidate3node-I:-S |
```

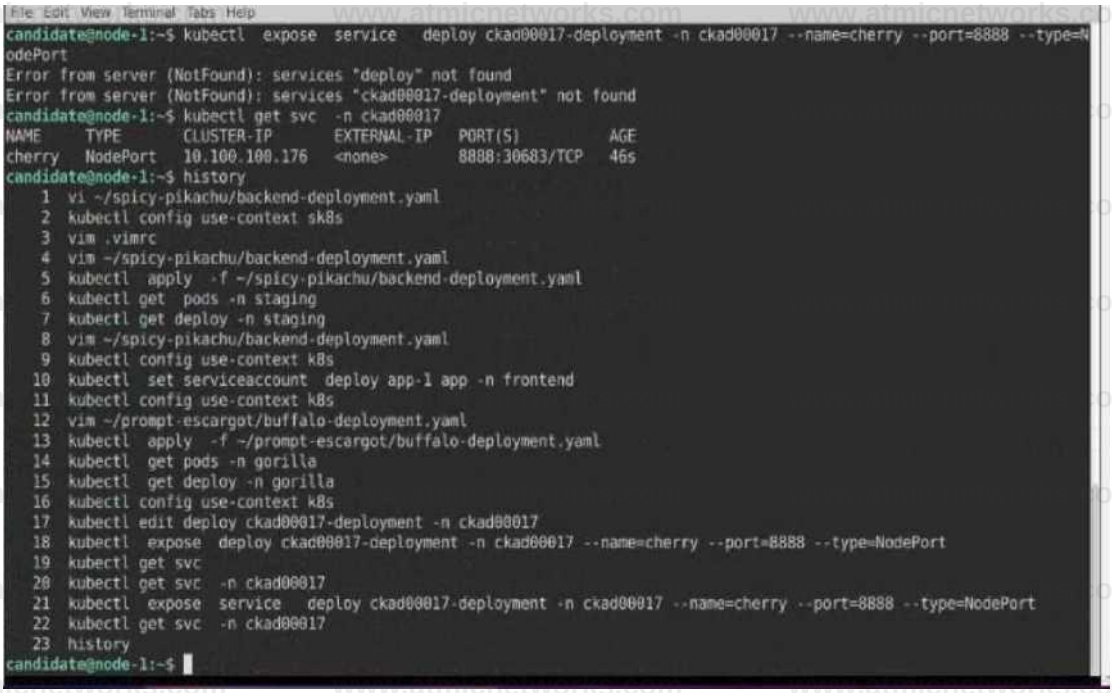

Question: 24
Refer to Exhibit.

You must switch to the correct duster/configuration context. Failure to do so may result tn a zero score.

>!; ; ' 1 fi kubectl config use ont^xt skHs

Task:
Create a Pod named nginx resources in the existing pod resources namespace.
Specify a single container using nginx:stable image.
Specify a resource request of 300m cpus and 1G1 of memory for the Pod's container.

Answer: See the solution below.

Explanation:

Solution:

```
candidat^node-r-S kubectl config use context k8s
Switched to context *kBsH
candidate@node*l>$ kubectl run nginx-resources -n pod-resources i«age*nginx:stable -dry-runaclient -o yaml > hw.yaml candidate@node-l:-S vim hw.yatl |
```

```
F it tail View terminal labs Help
ipiVersion: vl
tind: Pod metadata:
creation! lines tamp: null labels:
    run: nginx-resources name: nginx-resources namespace: pod - res ou rices spec:
containers;
* Image: nginx:stable name: nginx*resources resources:
```

```
requests;
    cpu: 300*
    memory: '1G1“
```

```
«i
```

```
candidate@node-l:-$ kubectl config use context k8s
Switched to context 'k8s"
candida te@node-l:~$ kubectl run nginx-resources -n pod-resources --iMge»nginx: stable --dry-runnel lent -□ yaml > hw.yawl candidate@node-l!'$ vim
hw.yaml candidate^node-l:-$ kubectl create f hw.yanl
pod/nginx resources created
candidate^nade-k-S kubectl get pods n pod resources
MAME                READY  STATUS    RESTARTS   AGE
nginx-resources     1/1    Running   9          13s
candidate@node-l:*$ kubectl describe pods *n pod-resources |
```

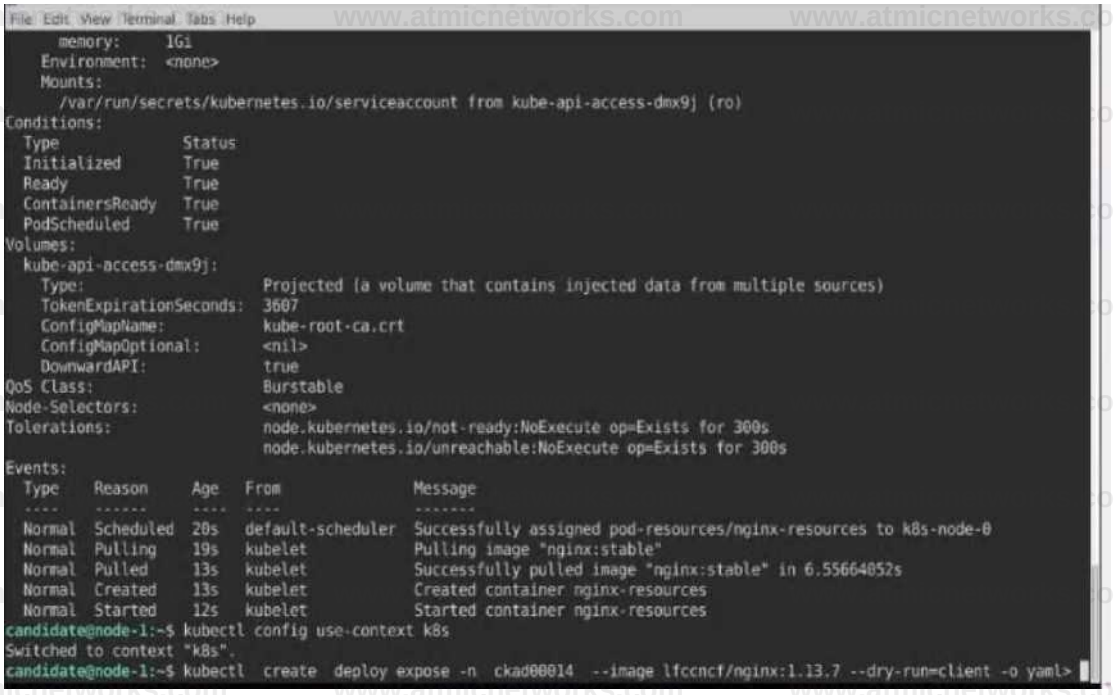

Question: 25

Refer to Exhibit.

You must switch to the correct duster/configuration context Failure to do so may result in a zero score.

► Il ' I i, kubectl cwfig usi ontext skHs

Task:

Create a Deployment named expose in the existing ckad00014 namespace running 6 replicas of a
Pod. Specify a single container using the ifccncf/nginx: 1.13.7 image
Add an environment variable named NGINX_PORT with the value 8001 to the container then expose port 8001

Answer: See the solution below.

Explanation:

Solution:

```
candidate inode-1: -s kutw ' l t orii it) tise-contexf k8s
Switched to context 'kes".
cand Ida tetgn ode-1 :~$ kubectl create deploy expose n ckadOMld -image lfccncf/nginx: 1.13.7 ■-dry-run=cllent -o yail> d ep.yaml candidat^node-l;-$
candidatetlnode-l:~$
candidategnode-1:-$ candidate^node-l:-$

candidategnode-1: *$
candidate^node-1:-$
candidate^node-1:*$
cand idate@node-1: -$
candidat&3node-l:-$
candidatconode-1: -$
candidategnode-l:~$
candida toy node-1: -$
cand idatevin ode-1;    l
```

```
Me toil View (enrunai tabs Help
■piVersion: apps/vl
kind: Deployment
neladata
creationTimestamp: null
labels:
    app: expose
name: expose
namespace: ckadO0614
spec:
replicas: 6
selector:
    matchlabels:
      app: expose
strategy: {}
template:
    metadata:
      crealionTimestamp: null
      labels:
        app: expose
    spec:
      containers:


    ■ image: lfccncf/nginx:1.13,7 name: nginx
      ports:


      - containerFort: 8661 eW!
        name: NGINX PORT
        value: *8001“
```

```
File  Edit  View  Terminal  Tabs  Help
candidate@node-1:~$ kubectl config use-context k8s
Switched to context "k8s".
candidate@node-1:~$ kubectl  create  deploy expose -n  ckad00014  --image lfccncf/nginx:1.13.7 --dry-run=client -o yaml> d
ep.yaml
candidate@node-1:~$
candidate@node-1:~$
candidate@node-1:~$
candidate@node-1:~$

candidate@node-1:~$
candidate@node-1:~$
candidate@node-1:~$
candidate@node-1:~$
candidate@node-1:~$
candidate@node-1:~$
candidate@node-1:~$
candidate@node-1:~$
candidate@node-1:~$ vim dep.yaml
candidate@node-1:~$ kubectl  create  -f dep.yaml
deployment.apps/expose created
candidate@node-1:~$ kubectl get pods -n ckad00014
NAME                      READY   STATUS              RESTARTS   AGE
expose-85dd99d4d9-25675   0/1     ContainerCreating   0          6s
expose-85dd99d4d9-4fhcc   0/1     ContainerCreating   0          6s
expose-85dd99d4d9-fld7j   0/1     ContainerCreating   0          6s
expose-85dd99d4d9-tt6rm   0/1     ContainerCreating   0          6s
expose-85dd99d4d9-vjd8b   0/1     ContainerCreating   0          6s
expose-85dd99d4d9-vtzpq   0/1     ContainerCreating   0          6s
candidate@node-1:~$ kubectl  get deploy -n ckad00014
NAME     READY   UP-TO-DATE   AVAILABLE   AGE
expose   6/6     6            6           15s
candidate@node-1:~$
```

Question: 26

Refer to Exhibit.

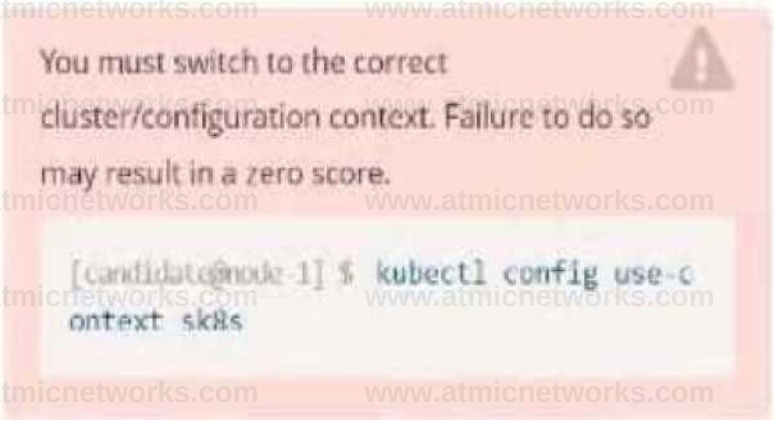

Task:

1) Fix any API depreciation issues in the manifest file -/credible-mite/www.yaml so that this application can be deployed on cluster K8s.

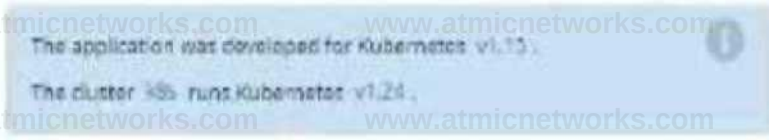

2) Deploy the application specified in the updated manifest file -/credible-mite/www.yaml in namespace cobra

Answer: See the solution below.

Explanation:

Solution:

```
candidate@node-1:~$ kubectl config use-context k8s
Switched to context "k8s".
candidate@node-1:~$ vim -/credible-mite/www.yaml
```

```
Me Lait View terminal Tabs Heip
dpiVersion; agps/vl
kind: Deployment
metadata:
  name: www-deployment
  namespace: cobra
spec
  replicas: 3
  selector:
        matchLabels:
              app: nginx
  template:
    metadata:
      labels:
        app: nginx
    spec:
      containers:
        - name: nginx
          image: "nginx:stable"


          ports:
              containerport: 86 volumeMounts:
              mountPath: /var/log/nginx name: logs


          env:
~
~
:wq
```

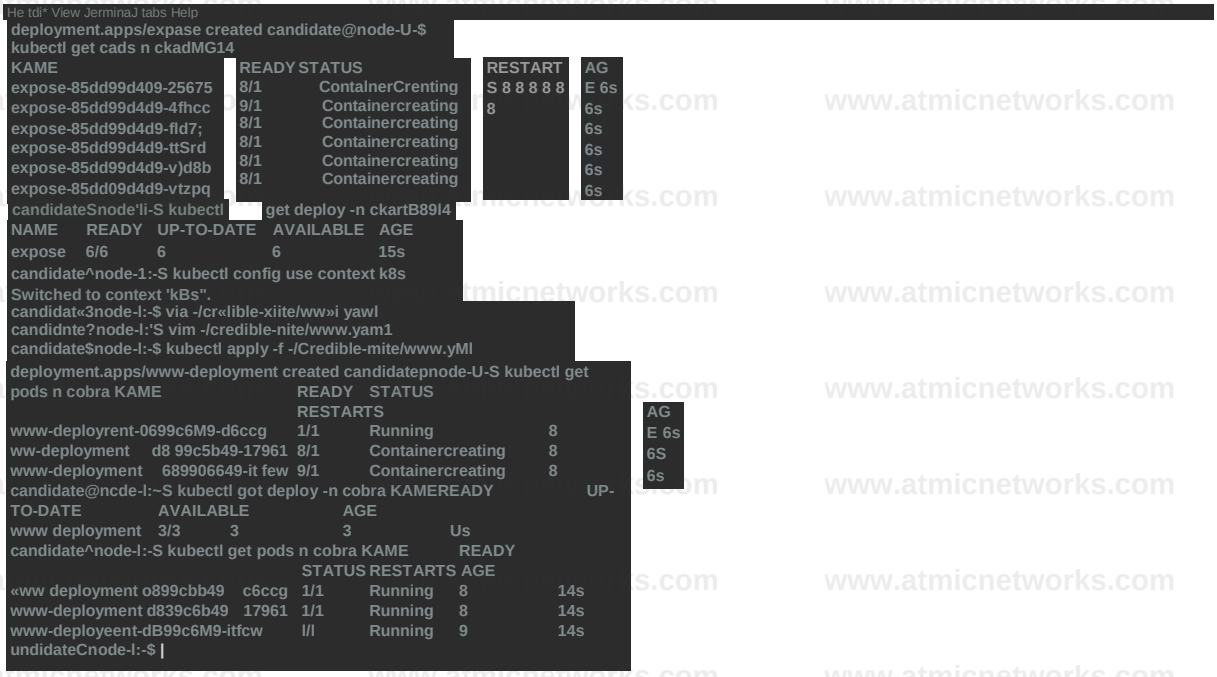

Question: 27

Refer to Exhibit.

You must switch to the correct duster/configuration context. Failure to do so may result tn a zero score.

ii ■ ■ i i- kubucll config usi
ont^xt skHs

Task:

1) Create a secret named app-secret in the default namespace containing the following single keyvalue pair:

Key3: value1

2) Create a Pod named ngnix secret in the default namespace.Specify a single container using the nginx:stable image.

Add an environment variable named BEST_VARIABLE consuming the value of the secret key3.

Answer: See the solution below.

Explanation:

Solution:

```
cdndiddt^node-k'S kubectl config use context k8s
Switched to context 'kBs".
candidateynode-1:-s kubectl create secret generic app-secret -n default --from-literal keyS-valuel
```

```
sec ret/app- secret created
candidate@node-l:-$ kubectl get secrets
NAME            TYPE      DATA   AGE
app secret      Opaque    1      4s
candidate^nodc-1:-S kubectl run nginx secret -n default i@age-nglnx:stable dry-run-cllent -o yaml> sec.yaffil candidate@node-l:-S vim sec.yawl |
```

```
le Lon Vie* lermind fabs Help
apiVersion: vl kind: Pod metadata:
creationTimestamp: null labels:
      run: nginx secret name: nginx-secret namespace: default spec;

containers;
      image: nginx:stable
      name: nginx-sec ret wv;
            - naw BEST VARIABLE vat MF ram:
                  secretKeyRef:
                        name: app-sec ret key: key!
```

```
-'I
```

```
candidate^node-ii-s kubectl contig use context KBs
Switched to context 'k8s".
candidate$node-l;^$ kubectl create secret generic app-secret -n default - f ram-literal-*.cy3*value 1
secret/app-secret created
candidate@node-l:-$ kubectl get secrets
NAME            TYPE      DATA   AGE
app sec ret     Opaque    1      4s
candidate$node-l:-S kubectl run nginx-secret -n default --image-nginx:stable --dry-run-client -□ yaml> sec.yaml candidatebnode-1:-$ vim sec.yaml
candidateanode-l:-$ kubectl create f sec.yaml
pod/nginx secret created
candidate^node-l:-$ kubectl get pods
NAME            READY    STATUS    RESTARTS  AGE
nginx-secret    1/1      Running   0         7s
candidate@node-l:-$ |
```

Question: 28

Refer to Exhibit.

You must switch to (he correct cluster/configuration context. Failure to do may result in a zero score.

HI.. -' I f, kubectl config ontext skHs

Task:

1- Update the Propertunel scaling configuration of the Deployment web1 in the ckad00015 namespace setting maxSurge to 2 and maxUnavailable to 59

2- Update the web1 Deployment to use version tag 1.13.7 for the lfconf/nginx container image.

3- Perform a rollback of the web1 Deployment to its previous version

Answer: See the solution below.

Explanation:

Solution:

```
candida te^node -1 >$ kubectl config use-context k8s Switched to context "Ws".
candidatacnode-l:-$ kubectl edit deploy webl -n ckad666is|
```

```
[f'le Edit View fermir-al 'abs Hetp
        app: nginx
    strategy:
      rollingUpdate:
        maxSurge: 2%
        maxUnavailable: 5%
      type: RolUngUpdate
    template:
      metadata:
        creationTimestaop null
        labels:
          app: nginx
      spec:
        containers:
        • image. Ifccncf/nginx:1. LI.[
          imagePullPolicy: IlNotPresent
          name: nginx
          ports:
            containerport: 80

            protocol: TCP
          resources: {}
          terminatiunflessagePath: /dev/terminal ion- log terminationHessagePoUcy: File
        dnsPolicy: ClusterFirst restartPolicy: Always schedulerName: default-scheduler securitycontext: {} termnationGracePeriodSeconds: 38 status:
    availableReplicas: 2
    conditions:
      lastIransitionTime: "2G22 89 24104:26:412
```

```
Lie icit Vie* fermir-al 'aba Help

Pitched to context *k8s'.
:andidateSnode-l:-S kubectl create secret generic ano secret n default from literal=key3*valuel iecret/app-secret created :andidate^no4e-l>S kubectl get
secrets WIE   TYPE        DATA  AGE
ipp-secret    Opaque      I     As
tandidategnode-lt-S kubectl run nginx secret -n default ■■iaage=nginx:stable dry run=cllent -o yaml> sec.yard :andidate§node-lt~S vim sec.yaml
:ondidate^node-l:'S kubectl create -f sec.yarn! jod/nglnx-secret created candidategn«fe-l;-5 kubectl get pods WIE READY       STATUS      RESTARTS  AGE
iginx secret 1/1             Running     9                 7s
:armidate@node-l:-"S kubectl contig use-tortext k8s Switched to context 'kss".
andidatagoMSe-li'S kubectl edit deploy wbl -n ckadBeOlS
ieployment.apps/webl edited
:andidate@nade-l:-s kubectl rollout status deploy webl -n ckadGUOlS ieploygient 'webl" successfully rolled out landidate^node-li'S kubectl rollout undo
deploy webl -n ckadCODlS leploynent.apps/webl rolled back :andidate@node-l:-S kubectl rollout history deploy webl -n ckadQDOlS leployaent.apps/webl
IEVISI0N CHANGE-CAUSE !              <none>
I             <none>

Mndidateinade-1:-$ kubectl get rs -n ckadMSIS
WIE                  DESIRED   CURRENT READY   AGE
vttbl 56f9Sbci>79    0         0       fi      63s
rebl-8S775b6t79      2         2       2       6hS3m
:andidate3node-l:--S |
```

Question: 29

Refer to Exhibit.

Yau must switch ta the correct dustt" -'configuration context. Failure to do so may result in a zero score.

•-I i.. ' I f kubectl config UM ontext skHs

Task:

Modify the existing Deployment named broker-deployment running in namespace quetzal so that its containers.

1) Run with user ID 30000 and
2) Privilege escalation is forbidden

The broker-deployment is manifest file can be found at:

~/daring-moccasin/broker-deployment.yaml

Answer: See the solution below.

Explanation:

Solution:

```
candidateenode-li^S kubectl config use-context kas Switched to context *k8s".
candidate@node-l>S vim |
```

```
le toil view leiminal abs Help
        containers: nMM: broker image: redis:alpine ports:
          - containerport; 6379 securityContext:
```

```
runAsUser: 390^0 privileged! false
```

```
candiciate^node*l:~$ kubectl config use context k8s
Switched to context 'k8s\
candidateanode*!:-$ vim -/danng*ffioccasin/broker-deployment yawl candidate@node-l:-5 kubectl apply -f -/daring-moccasin/broker-deployment.yaml
deployment.apps/broker-deployment configured candidategnode-lt-S kubectl get pods -n quetzal NAME              READY       STATUS      RESTARTS AGE
broker-deployment -6544606094-868p6      1/1         Running    6               30s
broker-deployment-65446d6d94-8dn7l       l/l         Running    6               32s
broker deployment•65446d6c94p4h4l        1/1         Running    0               31s
candidate@node-l:-$ kubectl get deploy n quetzal MAME   READY            UP-TO DATE AVAILABLE              AGE
broker-deployment      3/3         3                3               7h3m
candidate^node-lt-S |
```

Question: 30

Refer to Exhibit.

You must switch to the correct cluster^conf igu ration context. Failure to do so may result in a zero score.

[fl L > I f, kubectl config use ontext skRs

Task:

Update the Pod ckad00018-newpod in the ckad00018 namespace to use a NetworkPolicy allowing the Pod to send and receive traffic only to and from the pods web and db

/< i 'c-ru rcr 'iffi^e'kPolVMt Mм# sitai^ b"i CMW

• ■ t .'■ wr <•••.:». r-rir,o ...-i-*,, '>?.vo k^*il^y*f •»*> HIJM .

(heui*. т[1]#? coy vse jt"i Sewir**UKJw

Answer: See the solution below.

Explanation:

Solution:

```
candidate@node-1:~$ kubectl config use-context nk8s
Switched to context "nk8s".
candidate@node-1:~$ kubectl  describe  netpol -n ckad00018
```

```
File  Edit  View  Terminal  Tabs  Help
Name:         all-access
Namespace:    ckad00018
Created on:   2022-09-24 04:27:37 +0000 UTC
Labels:       <none>
Annotations:  <none>
Spec:
  PodSelector:     all-access=true
  Allowing ingress traffic:
    To Port: <any> (traffic allowed to all ports)
    From: <any> (traffic not restricted by source)
  Allowing egress traffic:
    To Port: <any> (traffic allowed to all ports)
    To: <any> (traffic not restricted by destination)
  Policy Types: Ingress, Egress


Name:         default-deny
Namespace:    ckad00018
Created on:   2022-09-24 04:27:37 +0000 UTC
Labels:       <none>
Annotations:  <none>
Spec:
  PodSelector:     <none> (Allowing the specific traffic to all pods in this namespace)
  Allowing ingress traffic:
    <none> (Selected pods are isolated for ingress connectivity)
  Not affecting egress traffic
  Policy Types: Ingress
candidate@node-1:~$ kubectl  label  pod ckad00018-newpod -n ckad00018   web-access=true
pod/ckad00018-newpod labeled
candidate@node-1:~$ kubectl  label  pod ckad00018-newpod -n ckad00018   db-access=true
pod/ckad00018-newpod labeled
candidate@node-1:~$
```

Question: 31

Refer to Exhibit.

You must switch to the correct duster/confjguration context Failure to do so may result in a zero score.

(fi ■ ' • I f, kubectl cwfig use ontext skRs

Context

You are asked to prepare a Canary deployment for testing a new application release.

Task:

A Service named krill-Service in the goshark namespace points to 5 pod created by the Deployment named current-krill-deployment

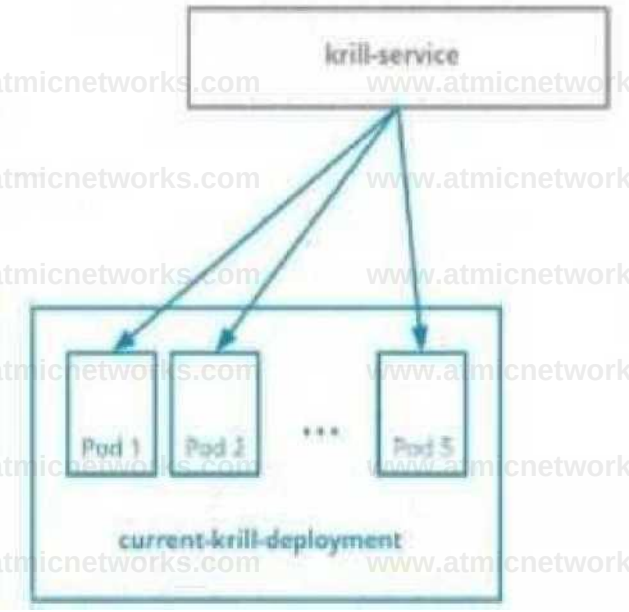

1) Create an identical Deployment named canary-kill-deployment, in the same namespace.
2) Modify the Deployment so that:

-A maximum number of 10 pods run in the goshawk namespace.

-40% of the krill-service 's traffic goes to the canary-krill-deployment pod(s)

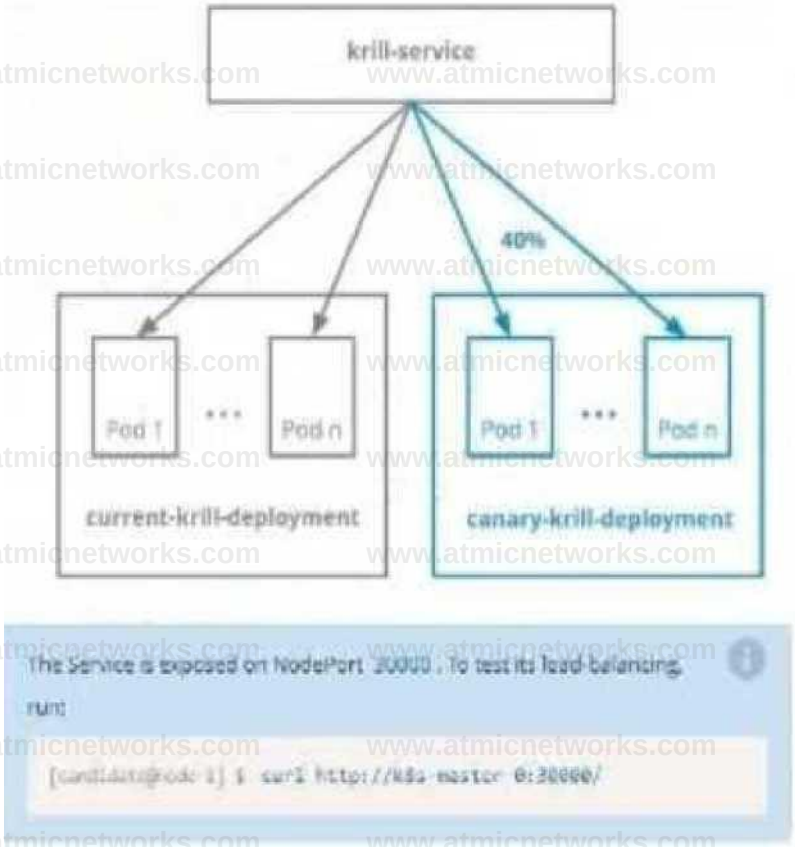

Answer: See the solution below.

Explanation:

Solution:

```
undiilate3imle-l;-/iiuiune-storliS kubect I. stale deploy tanary-krill-deploymenl --replicas <1 -n goshawk deployment apps/canary-krilt-deployment scaled
candidat«sn6de-l:-/numane-stork5 kubectl get deploy n goshawk
NAIE                        READY   UP-TO-DATE   AVAILABLE A6E
canary krill deployment     4/4     4            4           46s
current-krill-deployment    5/5     S            5           7h22»
candidate$node-l:~/hunane-storkS wget hltps7/k8s.b/Bcaiples/|
```

He tai¹. View terminal >abs Help

```
candidate@node-]:-/huBan«-storkS wget https://k8s.io/exaniples/admln/re5ource/qijota pod.yaffil
■2022 69 24 11:43:51 https://k85.io/dxamples/adtiin/resource/quota pod.yanil
Resolving kSs.io lkBs.io)... 34.10?.204.266. 2600:1901:9:2613::
Connecting to kOs.10 (kOs.io)134.107.204,2061:443... connected.
HTTP request sent, awaiting response... 301 Moved Permanently
Location: httpsJ/kubemetes.io/exatiples/admin/resource/quota-pod.yaml [following]
   2022-OS 2a 11:43:52- https://kuhernetes.iO/examptes/admln/resource/quota-pod.yaml
Resolving kubernetes.io ikubernetes.io)... 14?.75.49.148
Connecting to kubernetes.io (kubernetes.io)|147.75.40.148| 443 .. connected.
HTTP request sent, awaiting response... 200 Ok
Length: 99 [application/x-yail]
Basing to: 'quota pod.yaml'

quota-pod.yaml                  Ma|m»..:»1-1.m.jm..n<M....u-j.i....m..u.uJ|————————————————————————99 —.-KB/s in 0S

2922 09-24 11:43:52 115.0 MB/si - 'quota-pod.yml' saved (98/MJ

condidate@node-l:-/hunane-stork$ vim quota■pod.yaml |
r ie tan view ruminal labs H«p


2022 09-24 11:43:52 (15 0 Mb/s] ■ quota pod.yaml' saved [99/90]

candidateSnotle-l;~/hWMne-storkS Rib quota-pod.ya»l
candidate3ncde-l:-/hunane-storks kubectl create -f quota-pod.yml resourcequota/pod-demo created
candidate®node-l:~/liu»ane-stom kubectl get quota n go No resources found in go namespace.
candidatejnode-li'/huMfie-storks kubectl get quota -n goshawk
NAME       AGE    REQUEST LIMIT
pod-demo   19s    pods: 9/19
candidat«gnodd-l:-/hUMM-it«rk$ curl http://k8s-naster 0:39999/
current-krill-deployaent fb7c7995c kvtjr
app.kuberne t as.io/nabe*"current"
```

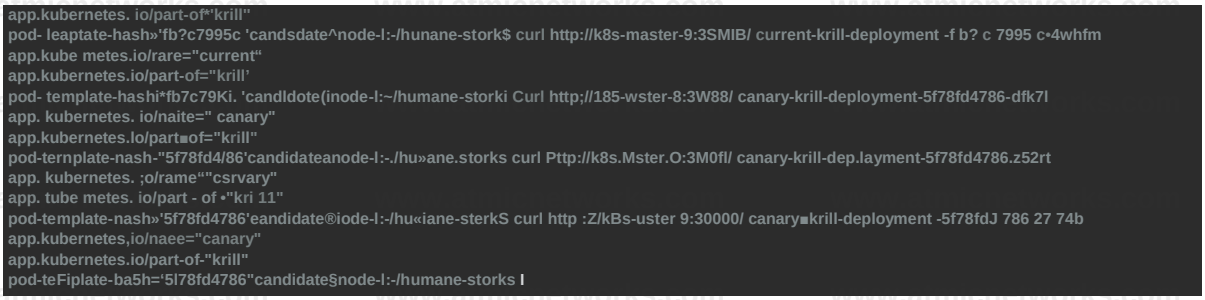

Question: 32

Refer to Exhibit.

Yau must switch to the correct duster/configuration context. Failure to d< may result in a zero score.

>ii i •' 1 I, kubectl config ont^xt skHs

Task:
The pod for the Deployment named nosql in the craytisn namespace fails to start because its container runs out of resources.
Update the nosol Deployment so that the Pod:

1) Request 160M of memory for its Container
2) Limits the memory to half the maximum memory constraint set for the crayfah name space.

- .htrC'tjf'^^ji' •|F»'-i<)<H"«

Answer: See the solution below.

Explanation:

Solution:

```
candidate@node-1:~$ kubectl config use-context k8s
Switched to context "k8s".
candidate@node-1:~$ vim ~/chief-cardinal/nosql.yaml
```

File Edit View Terminal Tabs Help

```
---
apiVersion: apps/v1
kind: Deployment
metadata:
  name: nosql
  namespace: crayfish
  labels:
    app.kubernetes.io/name: nosql
    app.kubernetes.io/component: backend
spec:
  selector:
    matchLabels:
      app.kubernetes.io/name: nosql
      app.kubernetes.io/component: backend
  replicas: 1
  template:
    metadata:
      labels:
        app.kubernetes.io/name: nosql
        app.kubernetes.io/component: backend
    spec:
      containers:
        - name: mongo
          image: mongo:4.2
          args:
            - --bind_ip
            - 0.0.0.0
          ports:
            - containerPort: 27017
~
~
-- INSERT --                                                    12,1          All
```

He EflilView lermiral Tabs Help

```
naw: mongo image: monger4.2 args:
  * -'bindip
  • 0.0.B.G
ports: ■ containerPort: 27017
resources:
    requests:
        memory: "IfjMi" limits:
        memory: "320R '
```

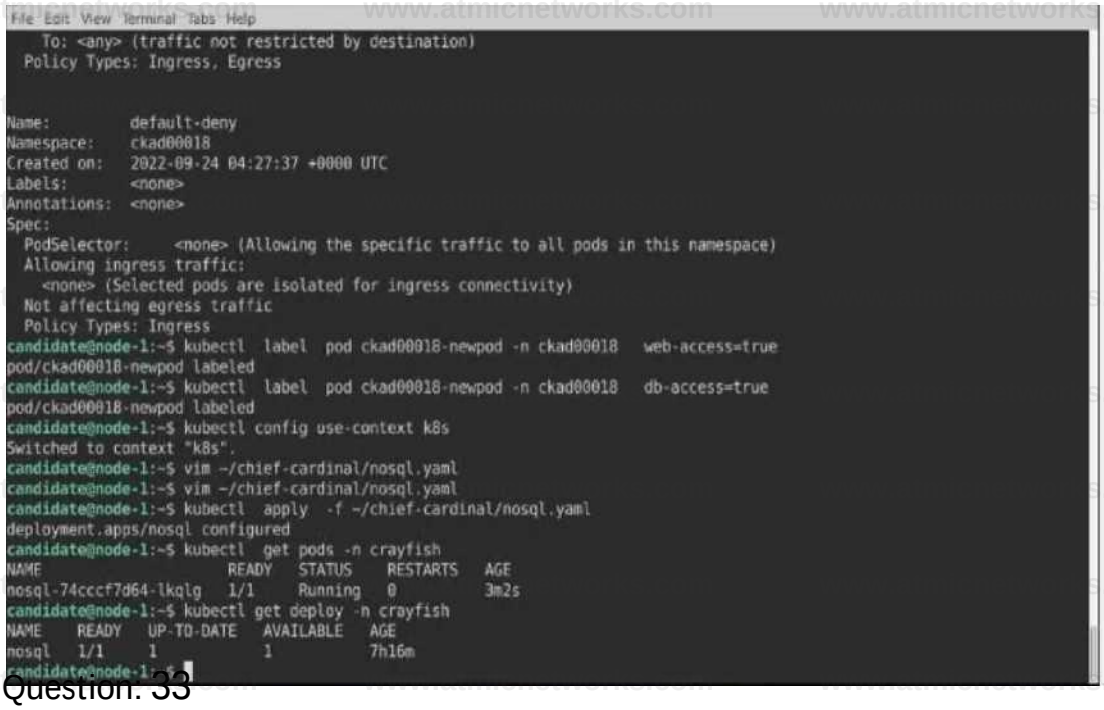

Question: 33

Refer to Exhibit.

Task:

A Dockerfile has been prepared at -/human-stork/build/Dockerfile

1) Using the prepared Dockerfile, build a container image with the name macque and lag 3.0. You may install and use the tool of your choice.

2) Using the tool of your choice export the built container image in OC-format and store it at - /human stork/macque 3.0 tar

Answer: See the solution below.

Explanation:

Solution:

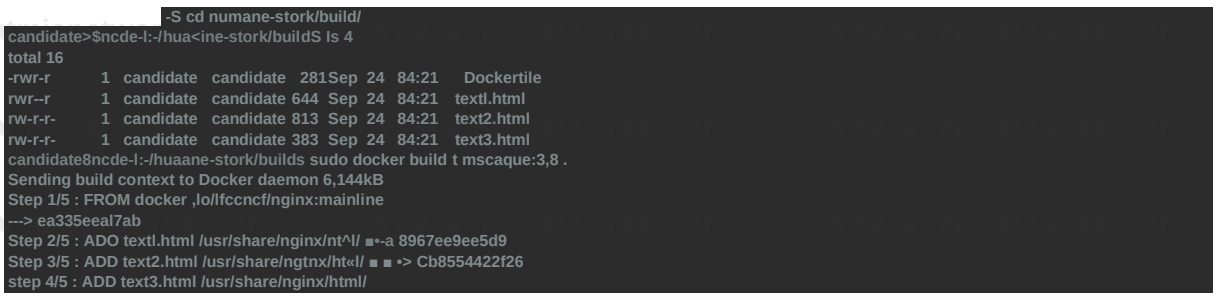

```
—* 62eB78ab821e
Step 5/5 i COPY text2.html /usr/share/nginx/html/index.html
      -> 331C8a94372c
SuccesSfully built 331c8a94372c
Successfully tagged macaque:3.0
candidate^node 1:-?humane-stork/builds stido docker save macaque:3,0 > -/humane storx/macague-3.0. tar candidate^node-l:-/humane-stork/buildS cd ,.
candidat^no<le-l;-/humane-storks Is -l
total 142532
drwxr xr x 2 candidate candidate 4896 Sep 24 84:21 build
iw-rw-r - 1 candidate candidate 145948672 Sep 24 11:39 macaque-3,0.tar
candidate^ncxie-li-Zhumane-storks |
```

```
F •« tort View lirmwil CMn Help
pod/ckgdDMlfl-newpod labeled
caodldatapnode-lS kubectl label pod ckadOMlB newpod n ckadOMU db occessutrue pod/cUdDBOlSnewpod labeled candidatepnode-l:-) kubectl contig us»-
conte»t kBs Switched to context -kBs". caadidatasnode-l:-! via -/chief -cardinal/dMql .yaal candidate^nodeli-S via -/chief cardioal/nosol yaal candidate^node-
li-s kubectl apply •' -/chief-caroinal/nosql.yaal doployacnt apps/nosql configured candidatepnode-l:-s kubectl get pods -n crayfish NAME     READY STATUS
                                RESTARTS AGE
MHqV74cccfMM-lkql,           1/1          Running    8                3a2s
cnndidatapnadeli.S kubectl get deploy n crayfish NAME           READY       UP-TO-DATE                AVA1LABL1 AGE
nosql    l/l            1                  1                 tn lha
candidatednode-l:-s cd hmanpxtort/tunld/ candidate0noda-l:-/huaana.itork/balld5 Is -l total 16
r« '       r  -  1   candidate    candidate281  Sep  24    04 21  Docker tile
rw r       r     1   candidate    candidate644  Sep  24    04:21     trxtl.html
•rw-r-     r--   1 candidate       candidate813 Sep   24  84:21    text2.html
rw r       r     1   candidate    candidate383  Sep  24    04:31    text3.html
candidatepnode-l:-/huaane'Stork/build5 sudo docker build t aacaque:3.0 Sending build context to Docker daemon 6 144U Step 1/5 : FROM docker.
io/ltccnc1/nqmx:aalnlme

■ •■> «335aM17ab

Step 2/5 : ADO textl.html Zusr/share/ngin,/hta1/ ■ > 8867ee9ee5d0

Step 3/5 : AW tcrt2.html /uor/fhare/ngiM/htal/ * 166554422126
Step 4/5 ADD text3.html ZusrZshareZnglnxZhlnlZ
```

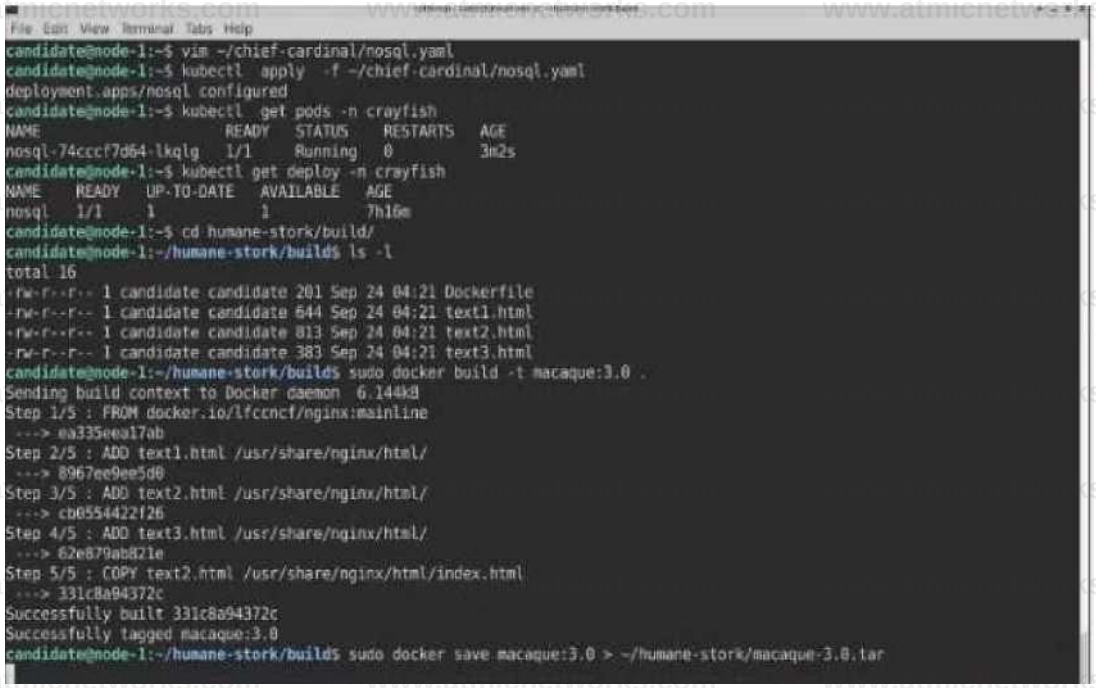
File Edit View Terminal Tabs Help
candidate@node-1:~$ vim ~/chief-cardinal/nosql.yaml
candidate@node-1:~$ kubectl apply -f ~/chief-cardinal/nosql.yaml
deployment.apps/nosql configured
candidate@node-1:~$ kubectl get pods -n crayfish
NAME READY STATUS RESTARTS AGE
nosql-74cccf7d64-lkqlg 1/1 Running 0 3m2s
candidate@node-1:~$ kubectl get deploy -n crayfish
NAME READY UP-TO-DATE AVAILABLE AGE
nosql 1/1 1 1 7h16m
candidate@node-1:~$ cd humane-stork/build/
candidate@node-1:~/humane-stork/build$ ls -l
total 16
-rw-r--r-- 1 candidate candidate 201 Sep 24 04:21 Dockerfile
-rw-r--r-- 1 candidate candidate 644 Sep 24 04:21 text1.html
-rw-r--r-- 1 candidate candidate 813 Sep 24 04:21 text2.html
-rw-r--r-- 1 candidate candidate 383 Sep 24 04:21 text3.html
candidate@node-1:~/humane-stork/build$ sudo docker build -t macaque:3.0 .
Sending build context to Docker daemon 6.144kB
Step 1/5 : FROM docker.io/lfccncf/nginx:mainline
---> ea335eea17ab
Step 2/5 : ADD text1.html /usr/share/nginx/html/
---> 8967ee9ee5d0
Step 3/5 : ADD text2.html /usr/share/nginx/html/
---> cb0554422f26
Step 4/5 : ADD text3.html /usr/share/nginx/html/
---> 62e879ab821e
Step 5/5 : COPY text2.html /usr/share/nginx/html/index.html
---> 331c8a94372c
Successfully built 331c8a94372c
Successfully tagged macaque:3.0
candidate@node-1:~/humane-stork/build$ sudo docker save macaque:3.0 > ~/humane-stork/macaque-3.0.tar

Question: 34

which one is the correct option

A. Get the pods with label env=dev and also output the labels
B. kubectl get pods -lenv=dev--show-labels
C. kubectl get pods -lenv=dev--show-commend
D. kubectl get pods -lenv=dev--show-caption

Answer: B

Question: 35

Delete the service and create the service again with type NodePort

A. delete the service
kubectl delete svc nginx// create service with expose command
kubectl expose po nginx --port=80 --type=NodePort
B. get the clusterIP from this command
kubectl get svc nginx -o wide// create temporary busybox to check the nodeport
kubectl run busybox --image=busybox --restart=Never -it --rm -- wget -o- <Cluster IP>:80
C. None of the above
D. All of the above

Answer: A

Question: 36

Find out the label of the pod and verify the service has the same label

A. get the pod with labelsbr>kubectl get po nginx --show-labels// get the service and check the selector columnbr>kubectl get svc my-service -o wide
B. kubectl get svc my-service of cloud
C. kubectl get svc my-windows systems

Answer: A

Question: 37

which one is the correct option

A. List all the config maps in the cluster
B. kubectl get cm or
kubectl get configmap
C. kubect2 get config cloud
D. kubectl get config system

Answer: B

Question: 38

Check what all are the options that we can configure with readiness and liveness probes

A. kubectl explain Pod.spec.containers.livenessProbe kubectl explain

Pod.spec.containers.readinessProbe
B. kubect2 explain Pod.spec.containers.livenessProbe kubect2 explain Pod.spec.containers.readinessProbe C. None of the above
D. All of the above
Answer: A

Question: 39

which one is the correct option?

A. Verify the configmap we just created has this data
B. you will see under datakubectl get cm -o yaml or kubectl describe cm
C. kubect2 describe command line
D. kubect1 describe cloud computing
Answer: B

Question: 40

which one is the correct option?

A. Create a configmap called myconfigmap with literal value appname=myapp
B. kubectl create cm myconfig windows --from-literal=appname=mywebsite
C. kubect2 create cm command
D. kubectl create cm myconfigmap --from-literal=appname=myapp
Answer: D

Question: 41

which one is the correct option

A. List all the configmaps in the cluster
B. kubectl get cm or kubectl get configmap
C. kubect2 get config cloud
D. kubectl get config system
Answer: B

Question: 42

which one is the correct option?

A. Create a hostPath PersistentVolume named task-pv-volume with storage 10Gi, access modes ReadWriteOnce, storageClassName manual, and volume at /mnt/data and verify
B. Kubect1 create -f task -pv-volume.yaml kubect1 get pv

apiVersion: v1
kind: PersistentVolume

metadata: name: task-pv-volume
labels: type: local
spec:
storageClassName: manual
capacity: storage: 10Gi
accessModes: -
ReadWriteOnce
hostPath: path: "/mnt/data"

C. Kubect2 create -f task -pv-volume.yaml
kubect2 get pv
D. None of the above

Answer: B

Question: 43

which one is the correct option?

A. List Persistent Volumes in the cluster
B. kubectl get pv
C. kubectl getpost
D. kubectl getpublish

Answer: B

Question: 44

which one is the correct option?

A. list all the namespace in the cluster
B. kubectl get namespaces kubectl get ns
C. kubect2 get namespaces kubect2 get ns
D. kubernetes1 get namespaces **Answer: B**

Question: 45

which one is the correct option? Get the pods with label env=dev

A. kubectl get pods -l env=dev
B. kubeclt get pods --show-labels
C. All of the above
D. None of the above

Answer: A

Question: 46

which one is the correct option? Verify all the pods are created with correct labels

A. kubeclt get pods --show-labels
B. kubeclt get pods --show-title

C. kubeclt get pods -- show-comment

Answer: A

Question: 47

which one is the correct option?

A. Get the pods with label information
B. kubectl get pods --show-labels
C. Kubect1 got post
D. kubect2 -show-caption

Answer: B

Question: 48

which one is the correct option?

A. Run command ls in the third container busybox3 of the above pod
B. kubect2 exec busybox -c busybox3 -- that
C. kubectl exec busybox -c busybox3 -- ls
D. None of the above

Answer: C

Question: 49

which one is the correct option?

A. Check the previous logs of the second container busybox2 if any
B. kubect2 logs busybox -c busybox4 --next
C. kubectl logs busybox -c busybox2 --previous
D. all of the above

Answer: C

Question: 50

which one is the correct option?

A. Check the logs of each container that you just created
B. kubect1 logs busybox -c busybox1
kubect1 logs busybox -c busybox2
kubect1 logs busybox -c busybox
C. kubect2 logs busybox -c busybox
D. kubect200 logs busybox -c busybox2
kubect100 logs busybox -c busybox1

Answer: B

Question: 51

which one is the correct option?

A. List all the pods in the particular namespace

B. kubect1 get po -n <namespace name>
C. kubect2 get
D. All of the above

Answer: B

Question: 52

which one is the correct option?

A. List all the pods in all namespaces
B. Kubect1 get po -- all-namespace
C. Kubect2 get po -- complete-namespace
D. None of the above

Answer: B

Question: 53

Create the temporary busybox pod and hit the service. Verify the service that it should return the nginx page index.html.

A. get the clusterIP from this command
kubectl get svc nginx -o wide// create temporary busybox to check the nodeport
kubectl run busybox --image=busybox --restart=Never -it --rm -- wget -o- <Cluster IP>:80
B. delete the service
kubectl delete svc nginx// create service with expose command
kubectl expose po nginx --port=80 --type=NodePort
C. All of the above

Answer: A

Question: 54

Create a NetworkPolicy which denies all ingress traffic
A. apiVersion: networking.k8s.io/v1
kind: NetworkPolicy
metadata:
name: default-deny
spec:
podSelector: {}
policyTypes:
- Ingress
B. delete the service
kubectl delete svc my-service// create the service again
kubectl expose po nginx --port=80 --target-port=9376// verify the label
kubectl get svc -l app=my-nginx
C. None of the above

Answer: A

Question: 55

Delete the service and create the service with kubectl expose command and verify the label

A. delete the service
kubectl delete svc my-service// create the service again
kubectl expose po nginx --port=80 --target-port=9376// verify the label
kubectl get svc -l app=my-nginx
B. delete the service
kubectl delete svc nginx// create service with expose command
kubectl expose po nginx --port=80 --type=NodePort
C. None of the above

Answer: A

Question: 56

Create an nginx pod with containerPort 80 and it should check the pod running at endpoint /healthz on port 80 and it should only receive traffic only it checks the endpoint / on port 80. verify the pod.

A. kubectl run nginx --image=nginx --restart=Never --port=80 --dry-run -o yaml > nginx-pod.yaml// add the livenessProbe and readiness section and create kubectl create -f nginx-pod.yaml// verify kubectl describe pod nginx | grep -i readiness kubectl describe pod nginx | grep -i liveness

apiVersion: v1
kind: Pod
metadata: creationTimestamp: null
labels: run: nginx
name: nginx
spec: containers: -
image: nginx
name: nginx
ports:
- containerPort: 80
livenessProbe:
httpGet: path: /healthz
port: 80
readinessProbe:
httpGet: path:
/port: 80
resources: {}
dnsPolicy: ClusterFirst
restartPolicy: Neve
rstatus: {}
B. kubect2 describe pod nginx | grep -i liveness
C. None of the above
D. All of the above

Answer: A

Question: 57

Output the YAML file for the service account we just created

A. kubectl get sa admin -o yaml
B. Kubect1 got sa
C. kubect got command line

Answer: A

Question: 58

Create a service account called admin

A. kubectl create sa admin
B. create admin admin
C. kubect5 admin sa

Answer: A

Question: 59

List all the service accounts in all namespaces

A. kubectl get sa --all-namespaces
B. kubectl get sa --complete-namespaces
C. None of the above
D. All of the above

Answer: A

Question: 60

Create a file called config.txt with two values key1=value1 and key2=value2 and verify the file

A. cat >> config.txt << EOF
key1=value1
key2=value2
EOFcat config.txt
B. kubectl create cm keyvalcfgmap --from-file=config.txtkubectl get cm keyvalcfgmap -o yaml
C. None of the above
D. All of the above

Answer: A

Question: 61

Create PersistentVolume named task-pv-volume with storage 10Gi, access modes ReadWriteOnce, storageClassName manual, and volume at /mnt/data and Create a PersistentVolumeClaim of at least 3Gi storage and access mode ReadWriteOnce and verify status is Bound

A. kubectl create -f task-pv-volume.yaml kubectl create -f task-pv-claim.yamlkubectl get pv kubectl get pvc
B. kubectl create -f task-pv-pod.yaml
C. None of the above
D. All of the above

Answer: A

Question: 62

Delete persistent volume and PersistentVolumeClaim we just created

A. kubectl delete pvc task-pv-claim kubectl delete pv task-pv-volume

B. kubectl create -f task-pv-claim.yamlkubectl get pvc
C. None of the above
D. All of the above

Answer: A

Question: 63

Output the YAML file of the above cronjob

A. kubectl get cj date-job -o yaml
B. kubectl get pokubectl logs date-job-<jobid>-<pod>
C. All of the above
D. None of the above

Answer: A

Question: 64

Delete the job we just created

A. kubectl delete job hello-job
B. kubectl create -f hello-job.yaml
C. None of the above
D. All of the above

Answer: A

Question: 65

Clean the cluster by deleting deployment and hpa you just created

A. kubectl delete deploy webappkubectl delete hpa webapp
B. kubectl create job nodeversion --image=node -- node -vkubectl get job -w
C. kubectl get pod
D. None of the above
E. All of the above

Answer: A

Question: 66

Remove all the pods that we created so far

A. kubectl delete po --all
B. kubectl create -f webapp.yaml
C. None of the above
D. All of the above

Answer: A

Question: 67

Show metrics of the above pod containers and puts them into the file.log and verify

A. kubectl top pod busybox --containers// putting them into file kubectl top pod busybox --containers > file.log cat file.log
B. kubectl exec busybox -c busybox3 -- ls
C. None of the above
D. All of the above

Answer: A

Question: 68

List all the pods sorted by name ?

A. kubectl get pods --sort-by=.metadata.name
B. kubectl get pods--sort-by=.metadata.creationTimestamp
C. None of the above
D. All of the above

Answer: A

Question: 69

Create a busybox pod and echo message ‘How are you’ and have it deleted immediately

A. notice the --rm flag
kubectl run busybox --image=nginx --restart=Never -it --rm -- echo "How are you"
B. create a pod
kubectl run nginx --image=nginx --restart=Never --port=80// List the pod with different verbosity
kubectl get po nginx --v=7
kubectl get po nginx --v=8
kubectl get po nginx --v=9
C. None of the above
D. All of the above

Answer: A

Question: 70

Delete the pod you just created ?

A. kubectl delete po nginxkubectl delete -f nginx-pod.yaml
B. kubectl describe pod nginx
C. None of the above
D. All of the above

Answer: A

Question: 71

Output the yaml file of the pod you just created without the cluster-specific information

A. kubectl get po nginx -o yaml --export
B. kubectl describe pod nginx
C. None of the above
D. All of the above

Answer: A

Question: 72

List all the pods showing name and namespace with a json path expression

A. kubectl space ['metadata.name', 'metadata.namespace']}"
B. kubectl get pods -o=jsonpath="{.items[*]['metadata.name', 'metadata.namespace']}"
C. None of the above

Answer: B

Question: 73

What is the List all the services in the particular namespace

A. kubectl get svc -n <namespace name>
B. kubectl get got-n <space name>
C. None of the above
D. All of the above

Answer: A

Question: 74

Create a Pod with an image Redis and configure a volume that lasts for the lifetime of the Pod

A. emptyDir is the volume that lasts for the life of the podkubectl create -f redis-storage.yaml B. kubectl delete pvc task-pv-claim kubectl delete pv task-pv-volume
C. None of the above
D. All of the above

Answer: A

Question: 75

Delete the CronJob and verify all the associated jobs and pods are also deleted.

A. kubectl delete cj date-job// verify pods and jobs
kubectl get po
kubectl get job
B. kubectl get job
kubectl get pokubectl logs date-job-<jobid>-<pod>
C. None of the above

Answer: A

Question: 76

Resume the rollout of the deployment

A. kubectl rollout history deploy webappkubectl rollout history deploy webapp --revision=9
B. kubectl rollout resume deploy webapp
C. None of the above

Answer: B

Question: 77

Get the deployment you just created with labels

A. kubectl get deploy webapp --show-labels
B. kubectl delete po --all
C. None of the above

Answer: A

Question: 78

select the correct option

A. Label the node (minikube if you are using) nodeName=nginxnode
B. kubectl label node minikube nodeName=nginxnode
C. kubectl get nodes --show-labels
D. None of the above

Answer: B

Question: 79

Let’s add the label app=nginx for all the pods and verify

A. kubectl label pod nginx-dev{1..3} app=nginx
kubectl label pod nginx-prod{1..2} app=nginxkubectl get po --show-labels
B. kubectl label pod/nginx-dev3 env=uat --overwrite
kubectl get pods --show-labels
C. None of the above

Answer: A

Question: 80

Get the pods with label env

A. kubectl get pods -L env
B. kubectl get pods -l 'env in (dev,prod)'
C. All of the above

Answer: A

Question: 81

If pod crashed check the previous logs of the pod

A. kubectl logs busybox -p
B. kubectl push busybox -d
C. None of the above

Answer: A

Question: 82

Check the Image version without the describe command ?

A. kubectl get po nginx -o jsonpath='{.spec.containers[].image}{"\ n"}'
B. kubect0 get po nginx -1 jsonpath='{.spec.containers[].gallery}{"\ n"}'

C. None of the above
D. All of the above

Answer: A

Question: 83

Create the nginx pod with version 1.17.4 and expose it on port 80

A. kubectl run nginx --image=nginx:1.17.4 --restart=Never --port=80 B. kubectl run nginx --image=nginx:1.17.4 --open=Never --port=180 C. None of the above

Answer: A

Question: 84

Which Kubernetes resource is used to ensure that a specified number of pod replicas are running at all times, even in the case of node failures or application crashes?

A. Services
B. Deployments
C. ConfigMaps
D. Pods

Answer: B

Question: 85

In Kubernetes, what is the primary purpose of a Pod?

A. To define the desired state of a single container
B. To define the desired state of one or more containers
C. To manage worker nodes in a cluster
D. To expose services externally to the cluster **Answer: B**

Question: 86

What command can be used to view the logs of a specific container within a multi-container pod?

A. kubectl logs pod-name
B. kubectl describe pod pod-name
C. kubectl logs pod-name -c container-name
D. kubectl get logs pod-name container-name

Answer: C

Question: 87

In a multi-container pod, what happens if one container fails while the other containers are running?

A. All containers are automatically restarted
B. Only the failed container is restarted
C. All containers are terminated

Answer: B

Question: 88

How can containers within a multi-container pod communicate with each other?

A. By using inter-process communication (IPC)
B. By using separate ports for each container
C. By using shared volumes and the localhost network interface
D. By using the pod's IP address and port mapping

Answer: C

Question: 89

Which of the following is a common use case for multi-container pods in Kubernetes?

A. Running completely unrelated applications in the same pod
B. Running multiple instances of the same application for load balancing
C. Running two containers with different Linux distributions in the same pod
D. Running containers from different namespaces in the same pod

Answer: B

Question: 90

In a multi-container pod, how do containers share the same network namespace?

A. By default, all containers in the same pod share the network namespace
B. They don't share the network namespace; each container has its own
C. It's not possible to share the network namespace in multi-container pods
D. By specifying the same "networkMode" in each container's configuration **Answer: A**

Question: 91

Which field in a Kubernetes Pod definition file is used to specify the list of containers within a pod?

A. containers
B. pods
C. spec
D. multi-containers

Answer: A

Question: 92

What is the primary benefit of using multi-container pods in Kubernetes?

A. Improved pod scalability
B. Enhanced pod security
C. Simplified pod management
D. Better communication and resource sharing between containers

Answer: D

Question: 93

What is a multi-container pod in Kubernetes?

A. Apodwith multiple worker nodes
B. Apodwith multiple containers sharingthe samenetworknamespace
C. Apodwith multiple instances of the same container
D. Apodwith multiple containers sharingthe samevolumeandnetwork namespace

Answer: D

Question: 94

Get the logs of the job just created

A. kubectl logs <pod name> // created from the job
B. kubectl logs <push name> // created from the task C. All of the above
D. None of the above

Answer: A

Question: 95

Watch the job that runs 10 times one by one and verify 10 pods are created and delete those after it's completed ?

A. kubectl get job -w
kubectl get pokubectl delete job hello-job
B. kubectl create job hello-job --image=busybox --dry-run -o yaml -- echo "Hello I am from job" > hello-job.yaml// edit the yaml file to add parallelism: 10
C. kubectl create -f hello-job.yaml
D. None of the above
E. All of the above

Answer: B

Question: 96

Create an nginx pod with a yaml file with label my-nginx and expose the port 80

A. kubectl run nginx --image=nginx --restart=Never --port=80 --dry-run -o yaml > nginx.yaml// edit the

label app: my-nginx and create the pod kubectl create -f nginx.yaml

apiVersion: v1 kind: Pod
metadata: creationTimestamp: null labels:
app: my-nginx
name: nginx
spec:
containers: - image: nginx name: nginx ports: -
containerPort: 80
resources: {}
dnsPolicy: ClusterFirst restartPolicy: Never status: {}

B. kubectl create -f nginx-svc.yaml
C. None of the above
D. All of the above

Answer: A

Question: 97

Create a pod with an image alpine which executes this command ”while true; do echo ‘Hi I am from alpine’; sleep 5; done” and verify and follow the logs of the pod ?

A. create the pod
kubectl run hello --image=alpine --restart=Never -- /bin/sh -c "while true; do echo 'Hi I am from Alpine'; sleep 5;done"// verify and follow the logs
kubectl logs --follow hello
B. kubectl get events --sort-by=.metadata.creationTimestamp// putting them into file.log
kubectl get events --sort-by=.metadata.creationTimestamp > file.logcat file.log
C. None of the above
D. All of the above

Answer: A

Question: 98

Select the correct the option ?

A. copy the logs of the above pod to the busybox-logs.txt and verify ?
B. kubectl logs busybox > busybox-logs.txtcat busybox-logs.txt
C. kubectl run busybox --image=busybox --restart=Never -- /bin/sh -c "echo I am from busybox pod; sleep 3600;"kubectl logs busybox
D. None of the above

E. All of the above

Answer: B

Question: 99

Create a secret mysecret with values user=myuser and password=mypassword

A. kubectl create secret generic my-secret --from-literal=username=user --from-literal=password=mypassword
B. kubectl get secret --all-namespaces
C. None of the above
D. All of the above

Answer: A

Question: 100

Create a configmap named keyvalcfgmap and read data from the file config.txt and verify that configmap is created correctly ?

A. kubectl create cm keyvalcfgmap --from-file=config.txtkubectl get cm keyvalcfgmap -o yaml
B. cat >> config.txt << EOF
key1=value1
key2=value2
EOFcat config.txt
C. None of the above
D. All of the above

Answer: A

Question: 101

Get the complete details of the pod you just created

A. kubect2 define got ngi
B. kubectl describe pod nginx
C. None of the above
D. All of the above

Answer: B

Question: 102

What Kubernetes resource is used to expose a set of pods as a network service within the cluster?

A. Config Map
B. Service
C. Network Policy
D. Deployment

Answer: B

Question: 103

What is the role of "seccomp" in Kubernetes security?

A. It's a networking policy enforcement tool
B. It's a pod communication protocol
C. It's a security feature that restricts the actions a container can perform
D. It's a role-based access control mechanism

Answer: C

Question: 104

Which Kubernetes API resource is used to manage access to the Kubernetes API server itself?

A. Service Account
B. Role Binding
C. Cluster Role
D. Pod Security Policy

Answer: C

Question: 105

What is the purpose of a "Pod Security Policy" in Kubernetes?

A. To restrict pod communication within the same namespace
B. To enforce security-related policies on pods
C. To manage the deployment of pods
D. To define resource limits for pods

Answer: B

Question: 106

What is the primary function of Kubernetes ServiceAccounts?

A. To give pods an identity and set of permissions within the cluster
B. To define how a pod communicates with other services
C. To grant access to external networks
D. To provide container runtime environment variables

Answer: A

Question: 107

Kubernetes resource is responsible for managing and distributing sensitive credentials to pods?

A. Role
B. Secret
C. Persistent Volume
D. Config Map

Answer: B

Question: 108

Select the correct option

A. Update the deployment with the image version 1.16.1 and verify the image and also check the rollout history ?
B. kubectl set image deploy/webapp nginx=nginx:1.16.1kubectl describe deploy webapp | grep Imagekubectl rollout history deploy webapp

C. kubectl rollout undo deploy webapp --to-revision=3kubectl describe deploy webapp | grep Imagekubectl rollout status deploy webapp
D. None of the above
E. All of the above

Answer: B

Question: 109

Choose the correct option

A. Verify the job and the associated pod is created and check the logs as well ?
B. kubectl create job hello-job --image=busybox --dry-run -o yaml -- echo "Hello I am from job
C. kubectl get job
kubectl get pokubectl logs hello-job-*
D. None of the above
E. All of the above

Answer: C

Question: 110

Select the correct option

A. Delete the above pod and create again from the same yaml file and verifies there is no file.txt in the path /data/redis ?
B. kubectl delete pod rediskubectl create -f redis-storage.yaml kubectl exec -it redis-storage /bin/sh cat /data/redis/file.txt // file doesn't exist
C. kubectl create -f task-pv-volume.yaml kubectl create -f task-pv-claim.yamlkubectl get pv kubectl get pvc
D. None of the above
E. All of the above

Answer: B

Question: 111

Choose the correct option ?

A. Get the memory and CPU usage of all the pods and find out top 3 pods which have the highest usage and put them into the cpu-usage.txt file
B. get the top 3 hungry pods
kubectl top pod --all-namespaces | sort --reverse --key 3 --numeric | head -3// putting into file
kubectl top pod --all-namespaces | sort --reverse --key 3 --numeric | head -3 > cpu-usage.txt// verify cat cpu-usage.txt
C. kubectl get po -n namespace2// update the image
kubectl set image pod/pod2 pod2=nginx -n namespace2// verify again
kubectl get po -n namespace2
D. None of the above
E. All of the above

Answer: B

Question: 112

Select the correct option ?

A. List all the events sorted by timestamp and put them into file.log and verify ?
B. kubectl get events --sort-by=.metadata.creationTimestamp// putting them into file.log
kubectl get events --sort-by=.metadata.creationTimestamp > file.logcat file.log
C. kubectl run busybox --image=busybox --restart=Never -- /bin/sh -c "echo I am from busybox pod; sleep 3600;"kubectl logs busybox
D. None of the above
E. All of the above

Answer: B

Question: 113

Choose the correct option ?

A. List all the service accounts in the default namespace
B. kubectl get sa
C. kubectl get sa --all-namespaces
D. ---all-namespaces

Answer: B

Question: 114

Select the correct option ? Exec into the above pod and create a file named file.txt with the text 'This is called the file' in the path /data/redis and open another tab and exec again with the same pod and verifies file exist in the same path ?

A. first terminal
kubectl exec -it redis-storage /bin/sh
cd /data/redis
echo 'This is called the file' > file.txt//open another tab

kubectl exec -it redis-storage /bin/sh
cat /data/redis/file.txt
B. emptyDir is the volume that lasts for the life of the podkubectl create -f redis-storage.yaml
C. None of the above

Answer: A

Question: 115

Create an nginx pod in a default namespace and verify the pod running

A. creating a pod
kubectl run nginx --image=nginx --restart=Never// List the pod
kubectl get po
B. kubectl get pods -o=jsonpath="{.items[*]['metadata.name', 'metadata.namespace']}"
C. None of the above

D. All of the above

Answer: A

Question: 116

how Get the IP Address of the pod you just created

A. kubectl get po nginx -o wide
B. kubect2 post nginx -1 world wide web
C. get ip address

Answer: A

Question: 117

Check the connection of the nginx pod from the busybox pod ?

A. kubectl get po nginx -o wide// check the connection
kubectl exec -it busybox -- wget -o- <IP Address>
B. kubectl run busybox --image=nginx --restart=Never -it -- echo "How are you"kubectl delete po busybox
C. None of the above
D. All of the above

Answer: A

Question: 118

Select correct option ?

A. Exec into both containers and verify that main.txt exist and query the main.txt from sidecar container with curl localhost ?
B. exec into main container
kubectl exec -it multi-cont-pod -c main-container -- sh
cat /var/log/main.txt// exec into sidecar container
kubectl exec -it multi-cont-pod -c sidecar-container -- sh
cat /usr/share/nginx/html/index.html// install curl and get default page
kubectl exec -it multi-cont-pod -c sidecar-container -- sh
apt-get update && apt-get install -y curl
curl localhost
C. create an initial yaml file with this
kubectl run multi-cont-pod --image=busbox --restart=Never --dry-run -o yaml > multi-container.yaml// edit the yml as below and create it
kubectl create -f multi-container.yamlkubectl get po multi-cont-pod
D. None of the above
E. All of the above

Answer: B

Question: 119

Create 5 nginx pods in which two of them is labeled env=prod and three of them is labeled env=dev

A. kubectl run nginx-dev1 --image=nginx --restart=Never --labels=env=dev
kubectl run nginx-dev2 --image=nginx --restart=Never --labels=env=dev
kubectl run nginx-dev3 --image=nginx --restart=Never --labels=env=devkubectl run nginx-prod1 --image=nginx --restart=Never --labels=env=prod
kubectl run nginx-prod2 --image=nginx --restart=Never --labels=env=prod
B. kubeclt get pods --show-labels
C. All of the above
D. None of the above

Answer: A

Question: 120

Get the pods with labels env=dev and env=prod

A. kubectl get pods -L env
B. kubectl get pods -l 'env in (dev,prod)'
C. None of the above
D. All of the above

Answer: B

Question: 121

Select the correct option

A. Verify the pod that it is scheduled with the node selector
B. kubectl describe po nginx | grep Node-Selectors
C. kubectl describe po nginx | grep Labels
D. None of the above
E. All of the above

Answer: B

Question: 122

Choose the correct option

A. Delete the deployment you just created and watch all the pods are also being deleted
B. kubectl delete deploy webappkubectl get po -l app=webapp -w
C. kubectl rollout status deploy webapp
D. kubectl rollout status deploy System
E. None of the above

Answer: B

Question: 123

choose the correct option

A. Check the history of the specific revision of that deployment
B. kubectl rollout history deploy webapp --revision=7
C. kubectl rollout undo deploy webapp
kubectl rollout status deploy webappkubectl get pods
D. None of the above
E. All of the above

Answer: B

Question: 124

Select correct option

A. Copy the above YAML file to hello-job.yaml file and create the job ?
B. kubectl create job hello-job --image=busybox --dry-run -o yaml -- echo "Hello I am from job" > hello-job.yamlkubectl create -f hello-job.yaml
C. kubectl get job
kubectl get pokubectl logs hello-job-*
D. None of the above
E. All of the above

Answer: B

Question: 125

Select the correct option ? List the secrets in all namespaces

A. kubectl get secret --all-namespaces
B. kubectl create secret generic my-secret --from-literal=username=user --from-literal=password=mypassword
C. None of the above
D. All of the above

Answer: A

Question: 126

Which type of Kubernetes Service exposes a pod to the external world?

A. Cluster IP
B. NodePort Load Balancer
C. External IP
D. **Answer: C**

Question: 127

In Kubernetes, what is an Ingress resource used for?

A. To define networking policies for pods
B. To configure container runtime settings
C. To expose a service to external traffic and manage routing
D. To manage resource allocation for pods

Answer: C

Question: 128

What is the primary function of Kubernetes DNS (Domain Name System) within the cluster?

A. To resolve external domain names
B. To provide IP addresses for nodes
C. To manage network policies
D. To resolve service names to IP addresses within the cluster

Answer: D

Question: 129

Which of the following is true about the "ClusterIP" service type in Kubernetes?

A. It exposes the service externally using a load balancer
B. It creates a service without any IP address
C. It allocates a unique IP address per pod
D. It exposes the service on a cluster-internal IP only

Answer: D

Question: 130

What is the purpose of a "NodePort" service type in Kubernetes?

A. To provide a fixed IP address for pods
B. To expose a service externally on a static port across all nodes
C. To define networking policies for pods
D. To manage routing between different services **Answer: B**

Question: 131

In Kubernetes networking, what does a "Pod-to-Pod" communication refer to?

A. Communication between different clusters
B. Communication between different nodes within the same cluster
C. Communication between containers within the same pod
D. Communication between different pods within the same cluster

Answer: D

Question: 132

which one of the following can be considered as the primary data store of kubernetes

A. Pod
B. Node
C. Etcd
D. None of above

Answer: C

Question: 133

What is the purpose of a Kubernetes ConfigMap?

A. To define security policies for a pod
B. To manage container runtime environment
C. To store sensitive data like passwords and API keys
D. To store configuration data separate from application code

Answer: D

Question: 134

are the types of kubernetes pods

A. Single container Pod
B. Dubulecontainer Pods
C. Multicontainer Pods D. Both A and C

Answer: D

Question: 135

A set is used to keep replica pods stable

A. Replication
B. Replica
C. Both A and B
D. None of above

Answer: B

Question: 136

Are the tools that are used for container monitoring

A. Advisor
B. Influx DB
C. Heapster
D. All of the above

Answer: D

Question: 137

GKE stand for

A. Google Cluster Engine
B. GoogleKubernetes Engine
C. Google Container Engine
D. None of the above

Answer: B

Question: 138

what are the main benefits that deployment offer that Replication Controllers do not

A. Strong update and rollback model
B. Client side update androllbacks
C. A declarative interface
D. Desired state

Answer: A

Question: 139

which of the following is kuberntes controllers

A. Replicaset
B. Deployment
C. Namespace
D. Both Replicaset and Deployment

Answer: D

Question: 140

which of the following forms the core kubernetes object

A. Pod
B. Service
C. Volume
D. All of the Above

Answer: D

Question: 141

Cronjobs in kubernetes run in

A. GMT only
B. UTC only
C. Local time Zone
D. Server NTP setting

Answer: B

Question: 142

is responsible for health check of the pods running on individual nodes

A. KubectI
B. Kubetet
C. Kube Scheduler
D. Kube Controller manager

Answer: B

Question: 143

Which of the following process runs on kubernetes no master node

A. Kube- proxy
B. Kube –apiserver
C. Both kube-proxy & Kube –apiserver
D. None of above

Answer: A

Question: 144

Kub –apiserver on kubenetes master is designed to scale

A. Vertically
B. Horizontally
C. Both Vertically and Horizontally
D. None of above

Answer: B

Question: 145

To create a new deployment in kubernetes use the command

A. Kubectl run
B. Kubernetes get deployment
C. Kubernetes set deployment
D. None of above

Answer: A

Question: 146

what is the basic operational unit of kubernetes

A. Task
B. Pod
C. Nodes
D. Container

Answer: B

Question: 147

which of the following runs on each node and ensures containers are running in a pod

A. Pod
B. Etcd
C. Kubelet
D. Scheduler

Answer: C

Question: 148

Replication Controllers and Deployment Controllers are part of

A. Kubeadm
B. Etcd manager
C. API Controller Manager
D. Master Controller Manager

Answer: D

Question: 149

Kubernetes was developed by

A. Google
B. IBM
C. Microsoft

D. None of above

Answer: A

Question: 150

what is kubernetes

A. Portable platform
B. Extensible platform
C. Open Source platform
D. All of the above

Answer: D

Question: 151

what is the latest version of Kubernetes

A. 1.1
B. 2.0
C. 1.19
D. 1.20

Answer: D

Question: 152

In Kubernetes Node is

A. A worker machine
B. A tool for starting a kubernetes cluster on a local machine
C. Developing tool
D. None of above

Answer: A

Question: 153

what can you deploy on kubernetes

A. Containers
B. Virtual Machine System processes None of above
C. **Answer: A**
D.

Question: 154

At its core , kubernetes is a platform for

A. Packaging software in containers
B. Provisioning machines
C. Running and scheduling container application on a cluster
D. None of above

Answer: C

Question: 155

manages the assigning nodes to pods depending on resource availability

A. Etcd
B. Kubectl
C. Scheduler
D. None of above

Answer: C

Question: 156

kubernetes is written in which languages

A. Go
B. C++
C. Python
D. JavaScript

Answer: A

Question: 157

Which of the following command is used to create kuberntes services

A. Kubectl run
B. Kubectl deploy
C. Kubectl expose
D. None of the above

Answer: C

Question: 158

In Kubernetes, what is an example of a Pod-level security mechanism?

A. Network Policies
B. Config Maps
C. Service Accounts
D. Ingress Controllers

Answer: C

Question: 159

What Kubernetes resource is used to enforce network rules that allow or deny communication between pods?

A. ConfigMap
B. Service
C. Network Policy
D. Deployment

Answer: C

Question: 160

Which of the following authentication mechanisms is commonly used for securing access to a Kubernetes cluster?

A. OAuth
B. JWT
C. Basic Auth
D. HMAC

Answer: A

Question: 161

What is the purpose of Role-Based Access Control (RBAC) in Kubernetes?

A. To control access to network resources
B. To define pod configuration parameters
C. To manage load balancer settings
D. To control access to cluster resources based on roles and permissions **Answer: D**

Question: 162

What is the role of the "Service Discovery" feature in Kubernetes?

A. To provide load balancing for services
B. To allocate IP addresses to pods
C. To automatically detect and register new services in the cluster
D. To manage network policies for pods

Answer: C

Question: 163

Which programming language is commonly used for writing Kubernetes Operators?

A. Java
B. Python
C. Ruby
D. Go

Answer: D

Question: 164

In Kubernetes, which type of object is used to store sensitive information like passwords and API tokens?

A. Secret
B. Key Vault
C. Credential Store
D. Private Data

Answer: A

Question: 165

Which command-line tool is commonly used to interact with a Kubernetes cluster?

A. docker
B. kubectl
C. kubecontrol
D. helm

Answer: B

Question: 166

What is the primary purpose of an Ingress resource in Kubernetes?

A. To manage container networking within a pod
B. To define pod resource limits and requests
C. To expose services to the external world with routing rules
D. To store sensitive configuration data securely

Answer: C

Question: 167

Which Kubernetes resource is responsible for scaling the number of pod replicas based on CPU or memory usage?

A. Pod
B. Replica Set
C. Deployment
D. Horizontal Pod Autoscaler

Answer: D

Question: 168

What is a Kubernetes Persistent Volume (PV) used for?

A. To store ephemeral data within a pod
B. To define network policies for pods
C. To store configuration data for applications
D. To provide persistent storage for pods

Answer: D

Question: 169

Which Kubernetes object is used to manage and persist configuration data separately from the application code?

A. Pod
B. Deployment
C. Config Map

D. Secret

Answer: C

Question: 170

What is the purpose of a Kubernetes Service?

A. To define the deployment strategy for pods B. To store configuration data for applications C. To expose a set of pods as a network service D. To manage the container runtime environment **Answer: C**

Question: 171

Which Kubernetes resource is used to define how a container should run within a pod?

A. Deployment
B. Service
C. Config Map
D. Persistent Volume Claim

Answer: A

Question: 172

Which of the following is a common container runtime used in Kubernetes?

A. Docker
B. Virtual Box
C. VMware
D. Hyper-V

Answer: A

Question: 173

What does CKAD stand for in the context of Kubernetes certification?

A. CertifiedKubernetes Administrator
B. CertifiedKubernetes Application Developer
C. CertifiedKubernetes Architect
D. CertifiedKubernetes Developer

Answer: B

Question: 174

Which Kubernetes object is used to run a one-time job to completion and then terminate?

A. Replication Controller
B. Job
C. Cron Job
D. Deployment

Answer: B

Question: 175

What is the role of a Kubernetes Operator?

A. To perform automated security audits on Kubernetes clusters.
B. To manage the scheduling of pods across nodes in a cluster.
C. To automate the deployment and management of complex applications.
D. To provide an interactive command-line interface for Kubernetes administrators.

Answer: C

Question: 176

What is the primary purpose of a Kubernetes Service?

A. To store and manage application logs.
B. To automatically provision and manage virtual machines.
C. To provide network connectivity and load balancing to a set of pods.
D. To automate the process of building container images

Answer: C

Question: 177

Which Kubernetes resource is used to automatically scale the number of pods based on observed CPU or memory usage?

A. Deployment
B. Stateful Set
C. Daemon Set
D. Horizontal Pod Auto scaler

Answer: D

Question: 178

What is a Kubernetes Secret?

A. A secure channel for communication between nodes in a cluster.
B. A resource used to manage user authentication and access control.
C. An encrypted storage volume for sensitive data in a pod.
D. A type of Kubernetes service used for exposing applications externally.

Answer: C

Question: 179

In Kubernetes, what is the purpose of a StatefulSet?

A. To manage stateless applications.
B. To provide a consistent network identity to pods.
C. To automate the scaling of pods based on CPU usage.
D. To manage stateful applications with unique identities and stable network addresses.

Answer: D

Question: 180

What is a Kubernetes Pod?

A. A unit of deployment that contains one or more containers.
B. A virtual machine in a Kubernetes cluster.
C. A load balancer for distributing traffic to services.
D. An external storage system used by Kubernetes.

Answer: A

Question: 181

What is the primary purpose of a Kubernetes ConfigMap?

A. Storing binary data in a Kubernetes cluster.
B. Managing access control for Kubernetes resources.
C. Storing configuration data separate from application code.
D. Defining security policies for Kubernetes pods.

Answer: C

Question: 182

What does the term "Ingress" refer to in Kubernetes?

A. A Kubernetes object used to define rules for routing external traffic to services.
B. A Kubernetes resource used to manage containerized applications.
C. A tool for monitoring resource utilization in Kubernetes clusters.
D. An API server that manages all Kubernetes control plane components.

Answer: A

Question: 183

What is the primary purpose of a Kubernetes Deployment?

A. To create and manage multiple instances of a pod.
B. To manage access control for Kubernetes resources.
C. To define configuration settings for a Kubernetes cluster.
D. To provide load balancing for external traffic.

Answer: A

Question: 184

Which of the following programming languages is commonly used to create Kubernetes applications?

A. Java
B. C++
C. Python
D. Ruby

Answer: C

Question: 185

Which of the following programming languages is commonly used to create Kubernetes applications?

A. Java C++ Python Ruby
B. **Answer: C**
C.
D.

Question: 186

What is the primary purpose of a Kubernetes Deployment object?

A. To create persistent storage
B. To define the layout of the Kubernetes cluster
C. To manage the deployment of containerized applications
D. To manage network routing within the cluster

Answer: C

Question: 187

Which Kubernetes resource is used to create a load balancer service?

A. Pod
B. Ingress

C. Deployment
D. Service

Answer: D

Question: 188

In Kubernetes, what is a "Namespace" used for?

A. To define resource limits for a container
B. To group and isolate resources within a cluster
C. To define network policies for a service
D. To specify the number of replicas for a deployment

Answer: B

Question: 189

What Kubernetes object is responsible for automatically scaling the number of pods based on CPU or memory usage?

A. Service
B. Ingress
C. Horizontal Pod Autoscaler
E. Config Map

Answer: C

Question: 190

Which Kubernetes feature is used for storing sensitive information like passwords or API keys securely?

A. ConfigMap
B. Secret
C. Persistent Volume
D. Resource Quota

Answer: B

Question: 191

What is the purpose of a "Readiness Probe" in a Kubernetes pod?

A. To determine when a pod is ready to be terminated
B. To periodically check if a container is alive and ready to accept traffic
C. To allocate CPU and memory resources to a container
D. To define environment variables for a pod

Answer: B

Question: 192

Which Kubernetes resource is used to manage and store configuration data separately from the application code? a

A. Pod
B. ConfigMap
C. Service
D. Persistent Volume Claim

Answer: B

Question: 193

What is the primary purpose of a "Job" resource in Kubernetes?

A. To define long-running services
B. To manage storage volumes
C. To run a task to completion in a pod
D. To create load balancers

Answer: C

Question: 194

You need to ensure that your Kubernetes application can handle increased load without causing performance degradation. Which of the following approaches would best help you achieve horizontal scaling of your application?

A. Increase the CPU and memory requests for the pods.
B. Use a Horizontal Pod Autoscaler (HPA).
C. Deploy the application using StatefulSets.
D. Use a NodePort service to expose your application.

Answer: B

Question: 195

You are deploying an application that runs a container requiring elevated privileges to perform certain operations. How can you configure the Pod to run securely while providing the necessary privileges? Select two correct answers.

A. Create a ServiceAccount with elevated permissions and assign it to the Pod
B. Set the securityContext to runAsNonRoot and add necessary capabilities to the container.
C. Use a PodSecurityPolicy to allow privileged access and set the securityContext of the Pod to enable privileged mode.
D. Run the container as the root user without any additional security configurations.
E. Set the securityContext of the Pod to enable privileged mode and add only the necessary capabilities required by the application.

Answer: C,E

Question: 196

You have deployed a service called my-app-service in your Kubernetes cluster, which is not accessible from outside the cluster. You want to expose this service on port 8080 to the external world. Which configuration should you use to achieve this?

A. Use a ClusterIP Service type and configure a NetworkPolicy to allow external traffic.
B. Use a headless Service and expose it via a NodePort.
C. Update the Service type to LoadBalancer and set the external port to 8080.
D. Change the Service type to NodePort and expose port 8080.

Answer: C

Question: 197

You are deploying a logging solution for your application running in Kubernetes. The application itself does not handle log aggregation. You decide to use a sidecar container to collect and send the logs to a centralized logging service. Which design pattern is being used when you deploy a logging agent as a separate container within the same Pod as your application?

A. Init Container Pattern
B. Sidecar Pattern
C. Ambassador Pattern
D. Builder Pattern

Answer: B

Question: 198

You have a Kubernetes Deployment for a web application, and you need to update the application to a new version without downtime. Which Kubernetes primitives and techniques should you use to manage this rolling update? Select two correct answers.

A. Use kubectl set image to update the image of the Deployment
B. Manually scale down the old Deployment and scale up the new one
C. Use a ConfigMap to update the image version
D. Delete the old Deployment and create a new one with the updated image
E. Modify the Deployment manifest to update the image and apply it using kubectl apply

Answer: A,E

Question: 199

You need to restrict traffic in your Kubernetes cluster so that only pods in the app namespace can communicate with pods in the db namespace on port 5432. Which NetworkPolicy should you apply in the db namespace?

A. ANetworkPolicy thatallows all ingress traffic to the db namespace.
B. ANetworkPolicy thatallows ingress from theapp namespace to the db namespace onport 5432.
C. ANetworkPolicy thatallows ingress from theapp namespace to the db namespace onall ports.
D. ANetworkPolicy thatallows egress from thedb namespace to the app namespace onport 5432.

Answer: B

Question: 200

You have a Kubernetes cluster where the application components need to be deployed with specific resource limits and affinity rules to ensure optimal performance. The application consists of a web server, a database, and a cache. The web server should only run on nodes labeled with "type=frontend," and the database should only run on nodes with "type=backend." The cache can run on any node but must

always have at least 1 CPU and 2Gi of memory. How would you define the configuration for the web server in this scenario?

A. Use a DaemonSet with nodeSelector for "type=frontend" and specify resource limits in the Pod spec.
B. Use a Job with nodeSelector for "type=frontend" and specify resource requests and limits.
C. Use a Deployment with nodeSelector for "type=frontend" and specify resource limits in the Pod spec.
D. Use a StatefulSet with nodeSelector for "type=frontend" and specify resource requests and limits.

Answer: C

Question: 201

How can a ConfigMap be used to set environment variables in a Pod?

A. By mounting it as a volume
B. By referencing it in the Pod's environment field
C. By creating a Secret from it
D. By applying it to the kube-apiserver configuration

Answer: B

Question: 202

You have a Kubernetes application that comprises a frontend, a backend, and a Redis database. The frontend and backend should be exposed to external traffic, while Redis should only be accessible by the backend. Additionally, you need to ensure that the backend can handle rolling updates without downtime. Which Kubernetes resource configurations should you use to meet these requirements?

A. Use LoadBalancer for all services and set the backend Deployment strategy to RollingUpdate
B. Use Ingress for frontend and backend, ClusterIP for Redis, and set RollingUpdate strategy for the backend Deployment
C. Use NodePort for all services and set the backend Deployment strategy to Recreate
D. Use DaemonSet for the backend and frontend, StatefulSet for Redis

Answer: B

Question: 203

You need to store database credentials securely in a Kubernetes cluster and ensure they are consumed by the application pods at runtime. Which method should you use to achieve this?

A. Define a StatefulSet to manage the application pods and ensure secure access to the credentials.
B. Store the database credentials in a ConfigMap and mount it as a volume in the application pods.
C. Create a Secret to store the database credentials and mount it as environment variables in the application pods.
D. Include the database credentials in the Deployment manifest and pass them as environment variables.

Answer: C

Question: 204

In your Kubernetes cluster, multiple teams are deploying applications, and you need to ensure that no single team consumes all the cluster resources. Which combination of Kubernetes features can help you achieve this goal?

A. Deploy DaemonSets in each namespace to control resource usage and enforce limits.
B. Use ResourceQuotas and LimitRanges to control the total resource usage and set default resource limits per namespace.
C. Configure HorizontalPodAutoscalers (HPA) to automatically scale pods based on resource usage.
D. Use ConfigMaps to store resource limits and apply them to each pod.

Answer: B

Question: 205

You are tasked with creating a Kubernetes deployment for an application that requires high availability. The application should have a minimum of 3 replicas running at all times, and it should be distributed across different nodes to ensure resilience. Additionally, you need to ensure that rolling updates are performed with zero downtime. Which of the following configuration options best meets these requirements?

A. Set replicas to 5 and use a Recreate strategy.
B. Set replicas to 3 and use a RollingUpdate strategy with maxUnavailable set to 0 and maxSurge set to 1.
C. Set replicas to 4 and use a RollingUpdate strategy with maxUnavailable set to 1 and maxSurge set to 2.
D. Set replicas to 2 and use a Recreate strategy.

Answer: B

Question: 206

Which multi-container pod design pattern is used to manage logging, monitoring, or communication services for the main application container?

A. Sidecar container
B. Init container
C. Ambassador container
D. Replication controller

Answer: A

Question: 207

You have a Kubernetes Deployment running a web application. During a rolling update, you notice that the application is becoming unresponsive. Which action can you take to mitigate this issue while ensuring minimal downtime?

A. Increase the maxUnavailable field in the Deployment strategy.
B. Increase the maxSurge field in the Deployment strategy.

C. Delete the current Deployment and create a new one with the updated image.
D. Scale down the Deployment to zero and then scale it back up.

Answer: B

Question: 208

Which of the following Kubernetes resources is primarily used to deploy a scalable web application with multiple replicas?

A. Service
B. Deployment
C. PersistentVolumeClaim
D. ConfigMap

Answer: B

Question: 209

You have deployed a web application in a Kubernetes cluster, and you need to ensure that it doesn't consume more resources than allowed. You want to set the resource limits for CPU and memory in the deployment manifest. What should you do?

A. Add limits and requests fields under spec in the deployment manifest.
B. Set resource limits in the service manifest.
C. Use kubectl set resources to define resource limits directly.
D. Add limits and requests fields under resources in the container spec.

Answer: D

Question: 210

Which type of probe would you configure in a Kubernetes pod to ensure that a container is ready to start accepting traffic?

A. Health Probe
B. Readiness Probe
C. Liveness Probe
D. Startup Probe

Answer: B

Question: 211

You need to prevent any single Pod in a namespace from consuming excessive resources and ensure balanced resource distribution. What Kubernetes object should you configure?

A. Configure a NetworkPolicy to limit resource usage per Pod.
B. Create a LimitRange object in the namespace to set default CPU and memory requests and limits for Pods.
C. Use a PodDisruptionBudget to control resource usage per Pod.
D. Create a ResourceQuota object in the namespace with specified CPU and memory limits.

Answer: B

Question: 212

You performed a rolling update on your Deployment, but the new version has a critical bug. What Kubernetes primitives and techniques should you use to rollback to the previous version? Select two correct answers.

A. Use kubectl rollout restart deployment/<deployment-name>
B. Edit the Deployment manifest to use the previous image version and apply it
C. Use kubectl delete deployment/<deployment-name> and recreate it with the previous version
D. Scale the current Deployment down to zero and manually create Pods with the previous version
E. Use kubectl rollout undo deployment/<deployment-name>

Answer: B,E

Question: 213

Which Kubernetes resource can be used to implement a canary deployment by incrementally increasing the traffic to the new version of an application?

A. StatefulSet
B. Ingress
C. ConfigMap
D. Service

Answer: B

Question: 214

You are managing a Kubernetes cluster where a microservices application consists of multiple services. One of the services, payment-service, needs to handle increased load during peak times. You decide to use Horizontal Pod Autoscaling (HPA) to scale the service based on CPU utilization. The paymentservice should scale between 2 and 10 replicas, maintaining an average CPU utilization of 60%. How would you configure the HPA to meet the above requirements?

A. Create an HPA with minReplicas: 2, maxReplicas: 10, and targetCPUUtilizationPercentage: 60. B. Create an HPA with minReplicas: 1, maxReplicas: 10, and targetCPUUtilizationPercentage: 50. C. Create an HPA with minReplicas: 1, maxReplicas: 12, and targetCPUUtilizationPercentage: 70. D. Create an HPA with minReplicas: 2, maxReplicas: 8, and targetCPUUtilizationPercentage: 60.

Answer: A

Question: 215

You are tasked with deploying a highly available web application on Kubernetes. The application consists of a frontend and a backend, and you want to ensure the frontend service is always available even if some pods fail. Which of the following options best achieves this?

A. Create a StatefulSet for the frontend with a PersistentVolume.
B. Create aDeployment for the frontend with a ReplicaSet and use a Serviceof typeClusterIP.
C. Create aDeployment for the frontend with a ReplicaSet and use a Serviceof typeNodePort.

D. Create aDaemonSet for the frontend with a Service of type ClusterIP.

Answer: B

Question: 216

You need to store database credentials securely and access them from a Kubernetes pod in a way that they are available as environment variables to your application. How should you configure this?

A. Create a ConfigMap and reference it in the pod's environment variables.
B. Create a Secret and reference it in the pod's environment variables.
C. Store the credentials in a plain text file in the container image.
D. Store the credentials directly in the pod's environment variables.

Answer: B

Question: 217

You are deploying a database application in Kubernetes that requires persistent storage for database files and ephemeral storage for temporary cache data. Which volume types should you use for these requirements? Select two correct answers.

A. PersistentVolumeClaim (PVC) for database files
B. HostPath for temporary cache data
C. ConfigMap for database files
D. EmptyDir for database files
E. EmptyDir for temporary cache data

Answer: A,E

Question: 218

You need to perform a canary deployment for a new version of your application to test it with a small percentage of users before fully rolling it out. Which approach using Kubernetes primitives would allow you to achieve this?

A. Use a StatefulSet to deploy the new version and configure a Service to split traffic between the old and new versions.
B. Use a ConfigMap to store the new version's configuration and a Service to route traffic based on ConfigMap values.
C. Use a Deployment to create a new version of the application and a Service to route 90% of traffic to the old version and 10% to the new version.
D. Use a Job to deploy the new version and a Service to route traffic based on the Job completion status.

Answer: C

Question: 219

You are managing a Kubernetes cluster with a microservices-based application. One of the services, inventory-service, is experiencing intermittent issues where pods are terminating unexpectedly. You need to investigate the issue and determine why the pods are failing. Which built-in Kubernetes CLI tool

command would you use to gather detailed information about the terminated pods in the inventory-service?

A. kubectl get pod inventory-service -o yaml
B. kubectl describe pod -l app=inventory-service
C. kubectl top pod -l app=inventory-service
D. kubectl logs -l app=inventory-service

Answer: B

Question: 220

Which component in Kubernetes is responsible for making decisions about whether a request to the API server is allowed to proceed, based on policies defined by the cluster administrators?

A. Admission Controller
B. Kubelet
C. API Server
D. Scheduler

Answer: A

Question: 221

Which of the following best describes a Kubernetes Operator?

A. A tool for managing cluster nodes
B. A custom controller that manages complex applications on behalf of the user
C. A method for configuring network policies
D. A configuration file for setting resource quotas

Answer: B

Question: 222

You need to perform a zero-downtime deployment for a critical application in your Kubernetes cluster. Which deployment strategy should you use to ensure that the application remains available to users during the update process?

A. Blue-Green Deployment Strategy
B. StatefulSet Deployment
C. Recreate Deployment Strategy
D. RollingUpdate Deployment Strategy

Answer: D

Question: 223

You have a Kubernetes application deployed in a namespace called prod. The application consists of multiple microservices, and one of the microservices, microservice-A, needs to communicate securely with another microservice, microservice-B. How would you configure the communication to ensure it is secure?

A. Use a ConfigMap to store the credentials needed for microservice-A to access microservice-B.
B. Use a ServiceAccount for microservice-A and configure RBAC to allow access to microservice-B.

C. Use mutual TLS (mTLS) to encrypt the traffic between microservice-A and microservice-B.
D. Use a NetworkPolicy to restrict traffic to only allow microservice-A to communicate with microservice-B.

Answer: C

Question: 224

You are deploying a web application that requires a primary container to serve HTTP requests and a secondary container to handle log aggregation from the primary container. The secondary container should run alongside the primary container within the same Pod. Which multi-container Pod design pattern should you use?

A. Init Container
B. Ambassador
C. Adapter
D. Sidecar

Answer: D

Question: 225

You are developing a logging system where logs need to be stored persistently and shared between multiple Pods in a Kubernetes cluster. Which type of volume configuration should you use to achieve this?

A. ConfigMap
B. PersistentVolumeClaim (PVC) with ReadWriteMany (RWX) access mode

C. HostPath
D. EmptyDir

Answer: B

Question: 226

You are deploying an application in a Kubernetes cluster that requires different configurations for development, staging, and production environments. What is the best practice for managing these environment-specific configurations?

A. Use a single ConfigMap for all environments and switch configurations based on the pod's label.
B. Use environment variables defined in the Deployment manifest.
C. Create separate namespaces for each environment and hard-code configurations in each namespace.
D. Store configurations in secrets for each environment.

Answer: B

Question: 227

You are managing a Kubernetes cluster where several applications are running. One of your applications, app-A, is experiencing performance issues due to resource contention with other applications. To resolve this, you need to properly set resource requests and limits for app-A. Which of

the following steps should you take? Select two correct answers.

A. Set a CPU request and limit for app-A to ensure it gets enough CPU resources.
B. Increase the replica count of app-A to distribute the load.
C. Disable resource limits to allow app-A to use as much CPU and memory as needed.
D. Set a memory request and limit for app-A to ensure it gets enough memory resources.
E. Use kubectl top to monitor resource usage of all pods in the cluster.

Answer: A,D

Question: 228

You need to deploy a pod that requires elevated privileges to access certain system resources. Which configuration should you add to the pod spec to allow a specific container to run with these elevated privileges?

A. Set allowPrivilegeEscalation: true in the pod's metadata.
B. Add runAsUser: 0 in the pod's spec.
C. Set privileged: true in the container's securityContext.
D. Set readOnlyRootFilesystem: true in the container's securityContext.

Answer: C

Question: 229

You are managing a Kubernetes cluster with a multi-tier application composed of frontend, backend, and database services. The frontend service is experiencing high latency and timeout issues. To

troubleshoot and resolve this issue, which of the following steps should you take? Select two correct answers.

A. Scale up the frontend deployment to increase the number of replicas.
B. Analyze the logs of the frontend pods using kubectl logs <frontend-pod>.
C. Inspect the resource usage of the frontend pods using kubectl top pods.
D. Check the network policies between the frontend and backend services.
E. Update the frontend service to use a NodePort instead of a ClusterIP.

Answer: B,C

Question: 230

You have a multi-tier application with frontend, backend, and database services running in separate pods. You need to ensure that the frontend pods can only communicate with the backend pods, and the backend pods can only communicate with the database pods. Which configuration should you apply?

A. Create a ServiceAccount for each tier and use RoleBindings to restrict communication.
B. Use a ConfigMap to define the allowed IP ranges for each tier and apply it to the pods.
C. Define a PodSecurityPolicy to control the communication between the pods.
D. Create a NetworkPolicy for frontend pods that allows egress to backend pods and another for backend pods that allows egress to database pods.
E.

Answer: D

Question: 231

You are deploying a critical application on Kubernetes that consists of multiple services. The application is currently deployed using an older version of a container image for the web-service that contains security vulnerabilities. You need to update this web-service to use a newer, secure image version. Which command should you use to update the web-service deployment to use the secure/web-service:v2.1.3 image in the production namespace?

A. kubectl set image deployment/web-service web-service=secure/web-service:v2.1.3 --namespace=production
B. kubectl set image deployments/web-service web-service=secure/web-service:v2.1.3 --namespace=production
C. kubectl update deployment web-service --image=secure/web-service:v2.1.3 --namespace=production
D. kubectl set image deployment web-service secure/web-service:v2.1.3 --namespace=production

Answer: A

Question: 232

You are tasked with deploying an application that requires a specific configuration file to be present before it can start. This configuration file needs to be fetched from a remote server and placed in a specific directory within the container file system. Which type of container should you use to fetch and prepare the configuration file before starting the main application container?

A. Job
B. Init Container
C. Sidecar Container
D. DaemonSet

Answer: B

Question: 233

In Kubernetes, which component is primarily responsible for authenticating API requests?

A. etcd
B. kubelet
C. kube-apiserver
D. kube-scheduler

Answer: C

Question: 234

You need to automate the management of a database in your Kubernetes cluster, including tasks like backups, updates, and failover. Which resource should you use to achieve this?

A. PersistentVolume
B. Operator
C. PodDisruptionBudget
D. CustomResourceDefinition (CRD)

Answer: B

Question: 235

You need to securely manage sensitive data such as database passwords and API keys in your Kubernetes cluster. Which of the following methods should you use? Select two correct answers.

A. Use plain text files mounted as volumes to store sensitive data.
B. Use a third-party secret management tool like HashiCorp Vault integrated with Kubernetes.
C. Use Kubernetes Secrets to store sensitive data and reference them in your pods.
D. Store sensitive data in a ConfigMap and reference them in your pods.
E. Store sensitive data in environment variables directly in the deployment YAML files.

Answer: B,C

Question: 236

Which of the following commands is used to build a Docker image from a Dockerfile for deploying an application on Kubernetes?

A. docker run myapp
B. kubectl apply -f Dockerfile
C. docker build -t myapp .
D. kubectl build myapp

Answer: C

Question: 237

You are deploying a multi-tier application with a frontend, backend, and database. The frontend should only be accessible externally, the backend should only be accessible by the frontend, and the database should only be accessible by the backend. Which network policy configurations would you use to enforce these rules? Select two correct answers.

A. Configure a network policy to allow frontend access from all sources and deny backend access from all sources except the frontend, and allow database access from all sources except the backend.
B. Use a network policy to allow all traffic from the frontend to the backend and from the backend to the database without any restrictions.
C. Implement a network policy to deny all traffic by default and then create specific rules to allow traffic from the frontend to the backend and from the backend to the database.
D. Allow all traffic within the cluster by default and use firewalls to restrict access from outside sources.
E. Create a network policy to allow frontend access from all sources, another to allow backend access from the frontend, and another to allow database access from the backend.

Answer: C,E

Question: 238

You need to ensure fair resource distribution among multiple teams using a shared Kubernetes cluster. You decide to implement resource quotas. Which of the following actions should you take to achieve this? Select two correct answers.

A. Set resource requests and limits on individual pods to control their resource usage.
B. Define ResourceQuotas for each namespace to limit the total amount of CPU and memory.
C. Manually monitor resource usage and inform teams when they exceed their allocated resources.
D. Use LimitRange to set default resource requests and limits for all pods in a namespace.

Answer: B,D

Question: 239

You want to restrict the capabilities of a container in your Kubernetes pod to enhance security. How can you remove all Linux capabilities except for NET_BIND_SERVICE from the container?

A. Define runAsUser: 1000 in the pod's spec.
B. Use the capabilities field in the container's securityContext to add NET_BIND_SERVICE and drop all others.
C. Configure a PodSecurityPolicy to restrict capabilities.
D. Set allowPrivilegeEscalation: false in the container's securityContext.

Answer: B

Question: 240

You need to configure your Kubernetes cluster to require all API requests to be authenticated using a specific authentication method. Which configuration should you modify?

A. Admission Controller
B. Pod Security Policy
C. Kubeconfig file
D. API Server Authentication Configuration

Answer: D

Question: 241

You have a Kubernetes service that needs to be accessible within the cluster on a specific port for a specific application. You also want to ensure it is not exposed externally. Which type of service and configuration should you use?

A. Use a ClusterIP service with the specific port defined.
B. Use a NodePort service with the specific port defined.
C. Use a Headless service with a specific port defined.
D. Use a LoadBalancer service with internal annotations.

Answer: A

Question: 242

You need to restrict a container's access to certain system resources for security purposes. Which securityContext settings should you configure in the Pod specification? Select two correct answers.

A. Set runAsUser to 0 and allowPrivilegeEscalation to false.
B. Set privileged to true and runAsUser to a non-root UID.
C. Set runAsUser to a non-root UID and allowPrivilegeEscalation to false.
D. Set readOnlyRootFilesystem to true and runAsNonRoot to true.

E. Add the SYS_ADMIN capability and set runAsNonRoot to true.

Answer: C,D

Question: 243

You are managing a Kubernetes cluster and have deployed a microservices application consisting of multiple services. You notice that one of the services, "auth-service", is experiencing intermittent failures due to high CPU usage. You need to automatically scale the "auth-service" based on CPU usage to handle the load effectively. Which of the following configurations should you apply to achieve this?

A. Create a CustomResourceDefinition (CRD) to define a new resource type for scaling "auth-service".
B. Add more replicas to the "auth-service" Deployment manually.
C. Create a HorizontalPodAutoscaler (HPA) resource for "auth-service" and set the target CPU utilization.
D. Create a PodDisruptionBudget (PDB) for "auth-service" to prevent too many pods from being evicted at once.

Answer: C

Question: 244

You need to monitor the resource usage of the pods in a namespace called production to ensure they are not overusing CPU and memory resources. You want to get a quick overview of the current resource usage for all pods in this namespace. Which built-in Kubernetes CLI tool command would you use to achieve this?

A. kubectl describe pods -n production
B. kubectl get events -n production
C. kubectl top pods -n production
D. kubectl get pods -n production -o wide

Answer: C

Question: 245

You are deploying a microservices application on Kubernetes where each service needs access to different configurations. One of your services, "payment-service", requires database connection strings and API keys that frequently change. How can you manage and inject these configurations into the "payment-service" containers efficiently?

A. Define a PersistentVolume and store the configuration data in it.
B. Configure a StatefulSet for "payment-service" to manage the configuration changes.
C. Use a ConfigMap to store the configuration data and mount it as environment variables in the "payment-service" pods.
D. Include the configuration data directly in the container image of "payment-service".

Answer: C

Question: 246

Your Kubernetes cluster has a deployment with multiple replicas of a web application. To monitor the application’s performance and quickly identify any issues, you decide to implement logging and metrics collection. Which of the following solutions best addresses these requirements?

A. Use Fluentd for metrics collection and Prometheus for log aggregation.
B. Use Heapster for log aggregation and Kibana for metrics collection.
C. Use Prometheus for metrics collection and Fluentd to collect and forward logs to Elasticsearch.
D. Use Grafana for metrics collection and Kibana for log aggregation.

Answer: C

Question: 247

You have a Kubernetes cluster running several applications using deprecated API versions that will be removed in the upcoming Kubernetes release. You need to ensure your applications continue to work after the cluster upgrade. What is the best approach to handle deprecated APIs in your applications?

A. Ignore the deprecation warnings and rely on Kubernetes backward compatibility.
B. Upgrade the cluster first and then handle the migration to supported API versions.
C. Migrate your resources to use the supported API versions before upgrading the cluster.
D. Disable the deprecated API versions in the cluster to see which applications fail and then migrate those.

Answer: C

Question: 248

You have deployed an application using a Deployment with multiple replicas. Each replica needs to persist data to a shared NFS storage. However, the application pods are failing to start with errors indicating they cannot mount the storage. What could be the possible reason?

A. The NFS storage is configured with a PersistentVolume and PersistentVolumeClaim but lacks the ReadWriteMany access mode.
B. The application pods are not configured with the correct resource limits.
C. The Deployment does not specify a ServiceAccount.
D. The NFS server is not reachable from the Kubernetes nodes.

Answer: D

Question: 249

You need to configure a Kubernetes ServiceAccount to grant specific permissions to an application running in your cluster. Which of the following steps is NOT necessary when creating and using a ServiceAccount for this purpose?

A. Specify the ServiceAccount in the Pod definition under the spec.serviceAccountName field.
B. Bind the ServiceAccount to a Role or ClusterRole using a RoleBinding or ClusterRoleBinding.
C. Create a Secret and manually mount it into the Pod to provide the ServiceAccount credentials.
D. Create a ServiceAccount using kubectl create serviceaccount.

Answer: C

Question: 250

You have a Kubernetes cluster with a web-app Deployment running in the production namespace. You need to expose this application to the internet using an Ingress resource with TLS termination and the domain example.com. What steps should you take to correctly configure this Ingress resource?

A. Annotate the web-app Deployment with ingress.example.com/tls=enabled to automatically generate the Ingress resource.
B. Create an Ingress resource in the production namespace with a host rule for example.com and specify the TLS secret.
C. Deploy a ClusterIP Service for web-app and configure an Ingress resource with the domain example.com and a TLS secret.
D. Create a LoadBalancer Service type for web-app and attach an Ingress resource with a host rule for example.com.

Answer: C

Question: 251

You need to monitor the health and performance of your Kubernetes cluster and applications. Which of the following tools and practices should you implement for effective observability and maintenance? Select two correct answers.

A. Deploy Grafana to visualize metrics collected by Prometheus.
B. Use Prometheus to collect and query metrics from the cluster.
C. Implement liveness and readiness probes for your application pods.
D. Enable Kubernetes auditing to log all API requests and responses.
E. Use kubectl delete pod to restart malfunctioning pods.

Answer: A,B

Question: 252

You have a Kubernetes deployment with multiple replicas of a microservice running in your cluster. The microservice depends on a Redis cache. Recently, you've noticed increased response times and occasional errors when the microservice tries to access Redis. What steps should you take to troubleshoot and resolve the issue? Select two correct answers.

A. Deploy a sidecar container for logging in the Redis pod.
B. Check the resource usage of the Redis pod.
C. Update the Redis container to the latest version.
D. Examine the network policies affecting Redis.
E. Restart the microservice pods.

Answer: B,D

Question: 253

You need to securely store and manage sensitive information such as database credentials in your Kubernetes cluster. Which Kubernetes resource should you use?

A. Secret
B. Ingress
C. ConfigMap
D. Volume

Answer: A

Question: 254

You need to deploy an application using Kustomize with different configurations for development and production environments. The production environment requires a higher number of replicas and different resource limits compared to the development environment. How can you achieve this using Kustomize?

A. Use a single kustomization.yaml file and modify it manually for each environment before deploying.
B. Use environment variables in the kustomization.yaml file to switch between configurations.
C. Use a base directory with common configurations and overlay directories for development and production with respective kustomization.yaml files.
D. Use different base directories for development and production with separate kustomization.yaml files.

Answer: C

Question: 255

You are running a build job in Kubernetes that requires temporary space for compiling source code and intermediate artifacts. The data does not need to persist after the job completes. Which volume type is most appropriate for this scenario?

A. ConfigMap
B. EmptyDir
C. HostPath
D. PersistentVolumeClaim (PVC)

Answer: B

Question: 256

You are designing a Kubernetes application that requires persistent storage for user-generated content and ephemeral storage for logs that can be lost when the Pod is deleted. Which storage solutions should you implement? Select two correct answers.

A. ConfigMap for user-generated content
B. EmptyDir for logs
C. PersistentVolumeClaim (PVC) for user-generated content
D. Secret for logs
E. EmptyDir for user-generated content

Answer: B,C

Question: 257

You are implementing a blue/green deployment strategy to update your application with zero downtime. Which steps and Kubernetes primitives should you use to achieve this?

A. Create a new Deployment for the new version, then update the Service to point to the new Deployment once the new version is verified.
B. Use a StatefulSet to deploy the new version and update the Service selector to point to the new StatefulSet.

C. Use a DaemonSet to deploy the new version across all nodes, then delete the old DaemonSet once the new version is verified.
D. Scale the old Deployment down to zero replicas and scale the new Deployment up to the desired number of replicas.

Answer: A

Question: 258

What is the primary purpose of an init container in a Kubernetes pod?

A. To initialize some prerequisite tasks before the main application container starts.
B. To restart the main application container in case of failure.
C. To handle logging and monitoring for the main application container.
D. To run continuously alongside the main application container for managing communication.

Answer: A

Question: 259

You need to deploy a complex stateful application that requires custom logic for its lifecycle management, including backups and updates. Which approach should you take to simplify this deployment and management?

A. Use a Kubernetes Deployment with liveness and readiness probes.
B. Utilize a StatefulSet and manually handle backups and updates.
C. Create a Custom Resource Definition (CRD) and implement a controller to manage the application.
D. Deploy the application using a DaemonSet to ensure one instance per node.

Answer: C

Question: 260

You have a multi-container pod, and both containers need to access a shared set of sensitive configuration data stored in a Kubernetes Secret. What is the best way to make this data available to both containers?

A. Include the sensitive data directly in the application code.
B. Create a ConfigMap and mount it as a volume in each container.
C. Mount the Secret as a volume and share it between the containers.
D. Store the sensitive data in environment variables directly in the deployment spec.

Answer: C

Question: 261

You have a Kubernetes cluster where a backend service named payment-service depends on a database service named db-service. You need to ensure that payment-service only starts after db-service is fully operational. Which Kubernetes feature will you use to achieve this?

A. Horizontal Pod Autoscaler
B. Readiness Probes
C. Init Containers

D. Liveness Probes

Answer: C

Question: 262

You need to deploy an application where a database initialization script must run before the main application container starts. The initialization script should run to completion before the main container starts. Which multi-container Pod design pattern should you use?

A. Ambassador
B. Init Container
C. Sidecar
D. Adapter

Answer: B

Question: 263

Your application deployed on Kubernetes is experiencing intermittent failures. You need to inspect the logs of a specific container within a pod to diagnose the issue. Which command will help you accomplish this?

A. kubectl describe pod [pod-name]
B. kubectl get logs [pod-name]
C. kubectl logs [pod-name] -c [container-name]
D. kubectl top pod [pod-name]

Answer: C

Question: 264

You want to perform a rolling update on a Deployment to minimize user disruption. How should you configure the Deployment's update strategy to ensure a balance between availability and update speed?

A. Set maxUnavailable to 1 and maxSurge to 1.
B. Set maxUnavailable to 0 and maxSurge to 1.
C. Set maxUnavailable to 50% and maxSurge to 50%.
D. Set maxUnavailable to 100% and maxSurge to 0.

Answer: A

Question: 265

You have deployed an application in Kubernetes, which requires different configurations for development and production environments. You decide to use ConfigMaps to manage these configurations. How would you update the running pods to use the new ConfigMap if you change the configuration in the ConfigMap?

A. Use the kubectl rollout restart deployment <deployment-name> command.
B. Patching the deployment using kubectl patch deployment <deployment-name>.
C. Manually restart each pod using kubectl delete pod <pod-name> command.
D. Delete the ConfigMap and recreate it with the new configurations.

Answer: A

Question: 266

You are tasked with ensuring that a pod runs with a specific user ID and has a read-only root filesystem for security purposes. How should you configure this in the pod manifest?

A. Configure securityContext in the namespace default security policy.
B. Define securityContext with runAsUser and readOnlyRootFilesystem under the spec of the pod.
C. Set runAsUser and readOnlyRootFilesystem under containers.

D. Define securityContext with runAsUser and readOnlyRootFilesystem under the metadata of the pod.

Answer: B

Question: 267

You are developing a Kubernetes application that needs to handle high traffic and should automatically scale based on the CPU usage. Which of the following configurations is the most appropriate to achieve this?

A. Use a Job to run pods that process tasks and then terminate.
B. Use a Deployment with a Horizontal Pod Autoscaler (HPA) configured to scale based on CPU usage.
C. Use a DaemonSet to ensure that a pod runs on every node.
D. Use a StatefulSet with a fixed number of replicas.

Answer: B

Question: 268

Which deployment strategy involves running two separate environments (one for the current version and one for the new version) and switching traffic from the old to the new version all at once?

A. Canary
B. Rolling Update
C. Blue/Green
D. Recreate
E.

Answer: C

Question: 269

You need to expose a set of pods running a backend service to the internet, but only through HTTPS, and want to ensure the backend pods receive traffic from clients on the original client IP. Which of the following configurations should you use?

A. Use a NodePort Service and manually configure the nodes' firewall rules to allow HTTPS traffic.
B. Create a Service of type LoadBalancer and use a NetworkPolicy to restrict traffic to HTTPS.
C. Deploy an Ingress resource with TLS termination and configure the backend Service with externalTrafficPolicy: Local.
D. Use a ClusterIP Service with a NetworkPolicy that only allows HTTPS traffic from the internet.

Answer: C

Question: 270

You have a Kubernetes cluster and want to deploy a database that requires complex configuration and management, such as automated backups and failover. You have decided to use a Kubernetes Operator to manage this database. Which of the following steps should you take to deploy and use the Operator effectively?

A. Define a CustomResourceDefinition (CRD) for the database and manage it manually.

B. Create a StatefulSet for the database and manually configure backup scripts.
C. Install the Operator from the OperatorHub and create a CustomResource (CR) instance for the database.
D. Use a DaemonSet to ensure the database is running on all nodes for high availability.

Answer: C

Question: 271

You have a Kubernetes cluster running a web application with multiple microservices. Each microservice is deployed in a different namespace for better isolation. One of the microservices, payment-service, needs to communicate with another microservice, user-service, deployed in a different namespace. Both microservices use internal DNS names for communication. What is the correct way to allow payment-service to resolve and communicate with user-service across namespaces?

A. Use a headless Service to expose user-service and allow direct pod-to-pod communication.
B. Use the fully qualified domain name (FQDN) of user-service in the payment-service configuration.
C. Configure a NetworkPolicy that allows traffic between the two namespaces.
D. Create a ServiceEntry resource in Istio for cross-namespace communication.

Answer: B

Question: 272

You need to ensure that certain security policies are enforced for all incoming requests to create or modify pods in your Kubernetes cluster. Which type of admission controller should you configure?

A. ServiceAccount
B. ValidatingAdmissionWebhook
C. MutatingAdmissionWebhook
D. ResourceQuota

Answer: B

Question: 273

You have an application running in a Kubernetes cluster and want to implement a canary deployment strategy to gradually roll out a new version. Which sequence of steps using Kubernetes primitives would achieve this?

A. Update the existing Deployment with the new version and immediately scale it up.
B. Use a Job to deploy the new version and monitor the Job's status to decide on the rollout.
C. Deploy the new version as a separate Deployment, create a new Service to balance traffic between the old and new Deployments, and adjust traffic gradually.
D. Create a StatefulSet for the new version and configure a Service to switch between the StatefulSet and the existing Deployment.

Answer: C

Question: 274

You are deploying an application that processes video files. This application has a primary container that performs the processing and a secondary container that uploads the processed files to a cloud storage service. Which multi-container Pod design pattern best fits this scenario?

A. PreStop Hook
B. Adapter
C. Ambassador
D. Sidecar

Answer: D

Question: 275

You are tasked with deploying an application using Helm, but you need to override some default values provided in the Helm chart to meet your specific requirements. Which command would you use to deploy the Helm chart with your custom values?

A. helm install my-release my-chart --set key1=value1,key2=value2
B. helm template my-release my-chart --set key1=value1,key2=value2
C. helm create my-release --set key1=value1,key2=value2
D. helm upgrade my-release my-chart --set key1=value1,key2=value2

Answer: A

Question: 276

You need to ensure that a set of Pods handling sensitive financial transactions are not co-located on the same node to ensure high availability and fault tolerance. Which Kubernetes feature should you use to achieve this?

A. StatefulSet
B. Pod Anti-Affinity
C. Taints and Tolerations
D. NodeSelector

Answer: B

Question: 277

Which multi-container pod design pattern would you use if you need to provide a logging service that collects and processes logs from a primary application container in the same pod?

A. Adapter Pattern
B. Sidecar Pattern
C. Ambassador Pattern
D. Replication Controller Pattern

Answer: B

Question: 278

You are managing a Kubernetes cluster where a deployment of a critical microservice has been experiencing intermittent crashes. These crashes are impacting the overall application performance and availability. What steps should you take to identify and resolve the issue? Select two correct answers.

A. Scale down the number of replicas to reduce load.
B. Increase the CPU and memory resources allocated to the microservice.
C. Check the logs of the microservice pods.
D. Update the microservice image to the latest version.
E. Enable debug mode in the microservice configuration.

Answer: C,E

Question: 279

What is the main benefit of using a rolling update strategy with Kubernetes Deployments?

A. It provides zero downtime by incrementally updating the application.
B. It allows the new version of the application to be tested on all nodes simultaneously.
C. It completely deletes the old version before starting the new one.
D. It instantly replaces all old pods with new ones.

Answer: A

Question: 280

Which command is used to build a Docker image from a Dockerfile?

A. docker create -f Dockerfile
B. docker run -t myimage .
C. docker-compose build myimage
D. docker build -t myimage .

Answer: D

Question: 281

Your Kubernetes cluster is running version 1.20, and you have deployed several applications that utilize the Ingress API. You plan to upgrade your cluster to version 1.22. What should you be aware of regarding API deprecations in this upgrade process? Select two correct answers.

A. Pod Security Policies are deprecated in 1.22.
B. All your existing Ingress resources will continue to work without any changes.
C. The Ingress API version extensions/v1beta1 is removed in 1.22.
D. The extensions/v1beta1 API group is fully removed in Kubernetes 1.20.
E. You should update your Ingress resources to use networking.k8s.io/v1

Answer: C,E

Question: 282

You need to configure a Kubernetes application to use a ConfigMap for environment-specific settings. Which steps should you follow to correctly create and apply a ConfigMap to your application? Select two

correct answers.

A. Create a ConfigMap using kubectl create configmap and reference it in the Pod definition under envFrom.
B. Create a ConfigMap using kubectl create configmap and reference it in the Pod definition under volumeMounts.
C. Write the configuration settings directly into the Pod definition under env.
D. Use a ConfigMap to store configuration settings but access it directly from within the application code without referencing it in the Pod definition.
E. Create a ConfigMap using kubectl create configmap and reference it in the Pod definition under volumes and volumeMounts.

Answer: A,E

Question: 283

Which of the following Kubernetes components is used to enforce policies for what users and service accounts are allowed to do?

A. Role-Based Access Control (RBAC)
B. kube-dns
C. Network Policies
D. Pod Security Policies

Answer: A

Question: 284

Your Kubernetes cluster is experiencing high latency issues, and upon investigation, you find that a particular pod is consuming excessive CPU and memory resources. You want to limit the resources for this pod to prevent it from affecting other applications. Which Kubernetes object will you use to enforce these limits?

A. ConfigMap
B. LimitRange
C. HorizontalPodAutoscaler
D. ResourceQuota

Answer: B

Question: 285

You are managing a web application running in a Kubernetes cluster. You need to deploy a new version of the application without any downtime. Which Kubernetes feature will help you achieve a seamless update?

A. Rolling Update
B. Recreate
C. Canary Deployment
D. Blue-Green Deployment

Answer: A

Question: 286

You need to ensure that a container in your Kubernetes pod runs as a non-root user, but you want to allow privilege escalation for specific operations. How can you achieve this?

A. Add privileged: true in the pod's securityContext.
B. Set runAsNonRoot: true and allowPrivilegeEscalation: true in the container's securityContext.
C. Set runAsUser: 0 in the container's securityContext.
D. Set readOnlyRootFilesystem: true in the container's securityContext.

Answer: B

Question: 287

You need to expose two applications, app1 and app2, running in the same namespace staging using the same Ingress controller. The applications should be accessible via app1.example.com and app2.example.com respectively. How should you configure the Ingress resource?

A. Annotate both applications' Deployments with ingress.host=app1.example.com and ingress.host=app2.example.com.
B. Use a LoadBalancer Service type for each application and configure DNS entries for app1.example.com and app2.example.com.
C. Create a single Ingress resource with two host rules, one for app1.example.com and another for app2.example.com.
D. Create two separate Ingress resources, each with its own host rule for app1.example.com and app2.example.com.

Answer: C

Question: 288

Your Kubernetes cluster runs a web application with multiple microservices, each in its own namespace. One of the microservices, auth-service, needs to communicate securely with user-service using HTTPS, and you must ensure that it retries the connection up to 5 times if it fails. Additionally, you need to configure a circuit breaker pattern to handle potential service failures gracefully. Which combination of configuration steps best achieves these requirements?

A. Use a Kubernetes Service for HTTP and configure a ConfigMap with retry logic.
B. Configure auth-service with a PersistentVolume and use Kubernetes ConfigMap for retries.
C. Use a Kubernetes Service for HTTPS and an Ingress controller for retries and circuit breaking.
D. Use Istio for mutual TLS, configuring destination rules for retries and circuit breaking.

Answer: D

Question: 289

You are managing a Kubernetes cluster that hosts a microservices-based application. The application consists of multiple services that need to communicate with each other. Additionally, you need to ensure that the application can scale automatically based on CPU usage, and that it uses secure communication between the services. Which Kubernetes configurations should you use to meet these requirements?

A. Deploy each service using StatefulSets, use HorizontalPodAutoscaler for auto-scaling, and NetworkPolicies for secure communication
B. Deploy each service using StatefulSets, use VerticalPodAutoscaler for auto-scaling, and ServiceAccounts for secure communication
C. Deploy each service using Deployments, use HorizontalPodAutoscaler for auto-scaling, and NetworkPolicies for secure communication
D. Deploy each service using Deployments, use HorizontalPodAutoscaler for auto-scaling, and Ingress for secure communication
Answer: C

Question: 290

You need to ensure that your Kubernetes pod can start with guaranteed resources and does not exceed a certain amount of CPU and memory usage. How should you define this in the pod manifest?

A. Set requests for CPU and memory under resources in the container spec.
B. Set limits for CPU and memory under resources in the container spec.
C. Set both requests and limits for CPU and memory under resources in the container spec.
D. Use kubectl set resources to configure in the live pod.
Answer: C

Question: 291

Which Kubernetes resource is used to dynamically provision persistent storage for a pod?

A. EmptyDir
B. ConfigMap
C. PersistentVolumeClaim
D. VolumeMount
Answer: C

Question: 292

You have exposed your application using a Service of type ClusterIP. However, the application is not accessible from within the cluster. What could be the issue?

A. The Service is of type ClusterIP and not NodePort.
B. The DNS server in the cluster is not configured properly.
C. The kube-proxy component is not running on the nodes.
D. The pods backing the Service are not labeled correctly to match the Service selector. **Answer: D**

Question: 293

You are managing a Kubernetes cluster where a crucial service, payment-service, is deployed as a Deployment with multiple replicas. Recently, users have reported intermittent failures when trying to process payments. You suspect that one or more pods of the payment-service are encountering issues. You need to identify which pods are failing and why. Which built-in Kubernetes CLI tool command would you use to get detailed information about the payment-service pods, including their status, recent events, and any errors?

A. kubectl logs deployment/payment-service
B. kubectl describe pod -l app=payment-service
C. kubectl get events --namespace=default
D. kubectl get pods -l app=payment-service

Answer: B

Question: 294

You are tasked with deploying an application that needs to maintain its state across restarts and updates. Which Kubernetes resource should you use to ensure that each instance of the application gets a unique but stable network identity and persistent storage?

A. Job
B. Deployment
C. StatefulSet
D. DaemonSet

Answer: C

Question: 295

You are deploying a stateful application on Kubernetes which consists of a StatefulSet with three replicas. The application requires that each pod must successfully complete an initialization script before it can serve traffic. Additionally, the application should not receive any traffic until it has passed a readiness check, which involves making an HTTP GET request to /health. How would you configure the StatefulSet to ensure the pods meet these requirements?

A. Use a lifecycle preStop hook for the initialization script and configure a livenessProbe with httpGet to /health.
B. Use an initContainer for the initialization script and configure a livenessProbe with httpGet to /health.
C. Use an initContainer for the initialization script and configure a readinessProbe with exec to run a custom script.
D. Use an initContainer for the initialization script and configure a readinessProbe with httpGet to /health.

Answer: D

Question: 296

You are managing a Kubernetes cluster with a set of microservices that require frequent updates to their configurations. To handle this efficiently, you decide to use Custom Resource Definitions (CRDs). Which of the following actions should you take to implement and manage these configurations using CRDs? Select two correct answers.

A. Directly modify the Kubernetes source code to add your custom configurations.
B. Use kubectl create -f <crd.yaml> to create the CRD in the cluster.
C. Use Kubernetes Operators to automate the management of your custom resources.
D. Define a new Custom Resource Definition (CRD) for your configurations using a YAML file.
E. Manually manage each microservice configuration without using CRDs.

Answer: B,D

Question: 297

You are managing a Kubernetes cluster for a high-traffic web application. The application consists of multiple microservices, each deployed in separate namespaces. Due to increased traffic, one particular microservice named payment-service is experiencing performance degradation. The payment-service is stateless and horizontally scalable. You need to improve its performance by scaling it. Which of the following commands would you use to scale the payment-service deployment to 5 replicas in its namespace payments?

A. kubectl scale deployment payment-service --replicas=5 -n payments
B. kubectl set replicas deployment payment-service --replicas=5
C. kubectl scale deployment payment-service --replicas=5
D. kubectl scale deploy payment-service --replicas=5 -n payments

Answer: A

Question: 298

You are tasked with automating the lifecycle management of a stateful application, including tasks such as automated backups, seamless updates, and health checks. Which Kubernetes resource best suits this requirement?

A. Secret
B. ConfigMap
C. PodSecurityPolicy
D. Operator

Answer: D

Question: 299

You want to implement a canary deployment strategy for your application in Kubernetes to test new features with a small subset of users before rolling out to the entire user base. Which Kubernetes object and configuration will you use to achieve this?

A. Job
B. Deployment with multiple revisions and Service with weighted traffic splitting
C. Deployment with multiple replicasets
D. StatefulSet

Answer: B

Question: 300

You are responsible for deploying a logging solution that should automatically scale based on the volume of logs and manage log retention policies. Which Kubernetes resource should you use to extend the functionality of your cluster to meet these requirements?

A. Set up a ReplicaSet with autoscaling enabled.
B. Create a Custom Resource Definition (CRD) for log management and implement a custom controller.
C. Deploy a StatefulSet with log volume mounts.

D. Use a DaemonSet to ensure log collectors run on every node.

Answer: B

Question: 301

You are deploying a multi-tier web application in Kubernetes consisting of a frontend, backend, and database. The frontend should be accessible externally via a domain name, while the backend and database should only be accessible internally. How should you configure the Services to achieve this? Select two correct answers.

A. Expose the frontend using a NodePort Service, the backend using an Ingress resource, and the database using a ClusterIP Service.
B. Expose the frontend using a ClusterIP Service, the backend using a NodePort Service, and the database using a ClusterIP Service.
C. Expose the frontend using a LoadBalancer Service, the backend using a ClusterIP Service, and the database using a ClusterIP Service.
D. Expose the frontend using an Ingress resource, the backend using a ClusterIP Service, and the database using a ClusterIP Service.

Answer: C,D

Question: 302

You need to configure multiple applications in your Kubernetes cluster with specific configuration files that should be updated dynamically without restarting the pods. What is the most suitable approach to achieve this?

A. Store the configuration files in a PersistentVolume and manually update them when needed.
B. Include the configuration files in the container image and redeploy the pods for every update.
C. Create a ConfigMap with the configuration data and mount it as a volume in the pods.
D. Deploy a new ReplicaSet with the updated configuration files whenever changes are needed.

Answer: C

Question: 303

Your team needs to deploy an application that must scale based on custom metrics, such as the length of a message queue. Which Kubernetes resource should you use to automatically scale your application pods based on these custom metrics?

A. Vertical Pod Autoscaler (VPA)
B. Horizontal Pod Autoscaler (HPA)
C. CustomResourceDefinition (CRD)
D. DaemonSet

Answer: B

Question: 304

You need to deploy a new version of your application with zero downtime using a Blue/Green deployment strategy. What Kubernetes primitives and techniques should you use to achieve this? Select two correct answers.

A. Use Jobs to manage the blue and green deployments
B. Use a RollingUpdate strategy in a single Deployment
C. Use a ConfigMap to manage deployment versions
D. Create two Deployments (blue and green) and use a Service to switch between them
E. Use an Ingress resource to switch traffic between the blue and green Deployments **Answer: D,E**

Question: 305

You need to ensure that a pod in your Kubernetes cluster runs with specific permissions that are different from the default ServiceAccount. How can you configure this?

A. Use a ConfigMap to store the ServiceAccount details.
B. Specify the desired ServiceAccount in the pod spec.
C. Define the permissions in the container spec.
D. Define the ServiceAccount in the namespace spec.
E.

Answer: B

Question: 306

You have a microservice that experiences occasional memory leaks, causing the application to become unresponsive. To mitigate this, you want Kubernetes to restart the pod if the application exceeds a memory usage threshold of 512Mi. How would you configure the Kubernetes pod to handle this requirement?

A. Configure a startupProbe with exec to check memory usage and delay the pod startup until memory usage is below 512Mi.
B. Configure a readinessProbe with exec to check memory usage and remove the pod from the service endpoints if it exceeds 512Mi.
C. Configure a livenessProbe with exec to check memory usage and restart the pod if it exceeds 512Mi.
D. Set resources.limits.memory to 512Mi and configure a livenessProbe with httpGet to /health.

Answer: C

Question: 307

You are managing a Kubernetes cluster for an e-commerce website. The application has three main components: a front-end service, a back-end API, and a database. The front-end service needs to scale based on user traffic, the back-end API requires multiple replicas for high availability, and the database should run on a single node with persistent storage. Which combination of Kubernetes resources should you use to deploy this application?

A. Deployment for front-end, Deployment for back-end, StatefulSet for database
B. Deployment for front-end, StatefulSet for back-end, PersistentVolumeClaim for database
C. DaemonSet for front-end, Deployment for back-end, PersistentVolume for database
D. Deployment for front-end, ReplicaSet for back-end, StatefulSet for database

Answer: A

Question: 308

You are deploying a distributed application with multiple microservices. One of the services, 'Service-A', requires a specific version of a library that conflicts with the version required by 'Service-B'. How can you ensure both services run correctly in your Kubernetes cluster?

A. Create separate Pods for 'Service-A' and 'Service-B', each with its own container image that includes the required library versions.
B. Deploy 'Service-A' and 'Service-B' in the same Pod, using a single container image that includes both versions of the library.
C. Deploy 'Service-A' and 'Service-B' in different namespaces to avoid conflicts.
D. Deploy 'Service-A' and 'Service-B' in the same Pod, using sidecar containers for each service.

Answer: A

Question: 309

Which type of volume is deleted along with the pod and is useful for temporary storage needs during the pod's lifetime?

A. PersistentVolumeClaim
B. PersistentVolume
C. HostPath
D. EmptyDir

Answer: D

Question: 310

You are running multiple applications in your Kubernetes cluster, and you need to monitor the resource usage of all pods in the staging namespace to ensure they are not exceeding their resource limits. You want to quickly get an overview of the CPU and memory usage for these pods.

A. Question: Which built-in Kubernetes CLI tool command would you use to achieve this?
B. kubectl top pods -n staging
C. kubectl get events -n staging
D. kubectl describe pods -n staging
E. kubectl get pods -n staging -o wide

Answer: B

Question: 311

You are designing a Kubernetes application that consists of a frontend web server and a backend database. The frontend and backend need to communicate securely, and you also want to ensure that the backend database is only accessible by the frontend web server. Which of the following configurations should you use in your Kubernetes manifests?

A. Use a Service of type ClusterIP for the backend database and a NetworkPolicy to allow traffic only from the frontend.
B. Use aService of type LoadBalancer for the backend database to allow external access.

C. Use aService of type NodePort for the backend database to ensure it is accessible.
D. Use aService of type ExternalName for the backend database.

Answer: A

Question: 312

You need to deploy an application using a Helm chart from a remote repository. Which sequence of Helm commands should you use to add the repository, update the local repository cache, and install the chart?

A. helm fetch my-repo/my-chart && helm install my-release my-repo/my-chart
B. helm update repo && helm install my-release my-repo/my-chart
C. helm install my-release https://example.com/charts/my-chart
D. helm repo add my-repo https://example.com/charts && helm repo update && helm install my-release my-repo/my-chart

Answer: D

Question: 313

You are tasked with deploying a web application using a Helm chart named myapp from a Helm repository named myrepo. Which of the following commands will correctly deploy the Helm chart?

A. helm install myapp myrepo/myapp
B. helm install --name myapp myrepo/myapp
C. helm deploy myapp --repo myrepo
D. helm install myapp myrepo

Answer: A

Question: 314

You need to deploy an application with different configurations for staging and production environments. The staging environment requires a debug flag to be set and fewer replicas compared to the production environment. How can you achieve this using Kustomize?

A. Use environment variables within the kustomization.yaml file to switch configurations.
B. Use a different kustomization.yaml file for each environment with a base configuration for common settings.
C. Use a base configuration and overlay directories for staging and production with respective kustomization.yaml files.
D. Use Helm charts with values files for each environment.

Answer: C

Question: 315

You have a stateful application running in Kubernetes that requires persistent storage to retain data across pod restarts. Which of the following configurations is the most appropriate to achieve this?

A. Use a Job with a temporary directory volume.
B. Use a ReplicaSet with a hostPath volume.

C. Use a StatefulSet with a PersistentVolumeClaim (PVC).
D. Use a Deployment with an emptyDir volume.

Answer: C

Question: 316

You want to configure a Kubernetes pod to run with a specific user and limit the resources to ensure security and stability. What should you include in the pod spec?

A. Define runAsUser and set limits under resources in the deployment spec.
B. Define securityContext with runAsUser and set requests under resources in the pod spec.
C. Define runAsUser and set limits under spec in the service manifest.
D. Define securityContext with runAsUser and set limits under resources in the container spec.

Answer: D

Question: 317

You need to grant a ServiceAccount in your namespace the permissions to list and get pods. What steps should you follow to achieve this?

A. Create a Role with the necessary permissions and bind it to the ServiceAccount with a RoleBinding.
B. Define the permissions directly in the ServiceAccount spec.
C. Modify the default ServiceAccount to include the necessary permissions.
D. Use a ConfigMap to store the permissions and apply it to the ServiceAccount.

Answer: A

Question: 318

You have deployed an application using a StatefulSet in a Kubernetes cluster. The application requires each pod to have a unique DNS entry to communicate with each other. However, you notice that the pods are not able to resolve each other's DNS names. What could be the issue?

A. The StatefulSet is missing the readiness probe configuration.
B. The StatefulSet pods are not running.
C. The StatefulSet service is created as a ClusterIP service.
D. The StatefulSet service is not created as a headless service.

Answer: D

Question: 319

You are deploying an application consisting of multiple microservices in a Kubernetes cluster. Service A needs to communicate with Service B, which is a StatefulSet with a headless service. However, you notice that Service A cannot resolve the DNS name of Service B. What could be the possible reason?

A. Service B is not exposing the correct ports.
B. Service A does not have the correct service account permissions.
C. Service A is not in the same namespace as Service B.
D. The headless service for Service B is not created properly.

Answer: D

Question: 320

You are designing a Kubernetes application that requires an initial setup script to run before the main application starts. This setup script needs to initialize a database schema. Which multi-container Pod design pattern should you use?

A. Adapter
B. Ambassador
C. Init Container
D. Sidecar

Answer: C

Question: 321

You need to update an already deployed Helm chart named myapp with new configuration values stored in a file named values.yaml. Which command should you use?

A. helm upgrade myapp myrepo/myapp -f values.yaml
B. helm update myapp myrepo/myapp --values values.yaml
C. helm upgrade myapp myrepo --values values.yaml
D. helm update myapp myrepo -f values.yaml

Answer: A

Question: 322

You are deploying an application that requires a stable and unique network identity, as well as persistent storage for each instance. Which Kubernetes resource is designed to handle these requirements?

A. Deployment
B. StatefulSet
C. Job
D. DaemonSet

Answer: B

Question: 323

You want to update your application with zero downtime using a blue/green deployment strategy. How would you implement this using Kubernetes primitives?

A. Scale down the blue Deployment to zero replicas and then create a green Deployment.
B. Create a new Deployment for the green version, update the Service to point to the green Deployment once it's verified, and keep the blue Deployment running for rollback.
C. Create a DaemonSet for the green version and delete the blue DaemonSet once the green version is verified.
D. Use a StatefulSet for the green version and update the Service selector to the new StatefulSet.

Answer: B

Question: 324

You are managing a Kubernetes cluster with a web application that experiences fluctuating traffic. During peak times, the application requires more resources to handle the load, but you want to ensure that during off-peak hours, it doesn't consume unnecessary resources. How can you configure the web application's deployment to handle this scenario effectively?

A. Set resource requests and limits for CPU and memory on the web application's containers.
B. Set up a PersistentVolume for the web application to handle the fluctuating resource needs.
C. Deploy the application with a fixed number of replicas to handle peak traffic.
D. Use a CronJob to scale the deployment up during peak hours and down during off-peak hours.

Answer: A

Question: 325

You are managing a Kubernetes cluster and have a critical application that requires different configurations for staging and production environments. You need to ensure that sensitive information such as database credentials is securely managed and injected into the application without hardcoding them into the container images. Which approach should you use to achieve this?

A. Use environment variables defined in the Deployment manifest to pass sensitive information to the pods.
B. Store sensitive information in ConfigMaps and inject them as environment variables in the pods.

C. Use Kubernetes Secrets to store sensitive information and mount them as environment variables or volumes in the pods.
D. Include the sensitive information directly in the container image and use a version control system to manage changes.

Answer: C

Question: 326

You want to implement a Canary deployment strategy to gradually roll out a new version of your application and monitor its performance. Which Kubernetes primitives and techniques should you use? Select two correct answers.

A. Use a Deployment with multiple ReplicaSets and update the Service selector B. Use multiple Deployments with different labels and a Service to route traffic C. Use a ConfigMap to switch between versions
E. Use an Ingress with traffic splitting based on weights
F. Use a StatefulSet to manage the different versions

Answer: B,D

Question: 327

You have a web application running in a Kubernetes cluster, and you want to ensure it remains healthy and available. You decide to implement liveness and readiness probes. What considerations should you make when configuring these probes? Select two correct answers.

A. Set the liveness probe to check the main application endpoint.
B. Set a short timeout for the readiness probe to quickly detect unready states.
C. Use the same endpoint for both liveness and readiness probes.
D. Set the readiness probe to check a deeper application health endpoint.
E. Configure a long initial delay for the liveness probe to avoid premature restarts.

Answer: D,E

Question: 328

What is the purpose of the COPY instruction in a Dockerfile?

A. To execute a command in the container.
B. To copy files from the host system into the container image.
C. To expose ports in the container.
D. To copy files from the container to the host system.

Answer: B

Question: 329

You are responsible for maintaining a Kubernetes application that handles high transaction volumes. To ensure reliable observability and maintenance, you decide to implement a comprehensive monitoring and logging solution. Which of the following setups will provide the best combination of real-time metrics,

log aggregation, and alerting?

A. Use Grafana for metrics, Fluentd for log collection, and Kibana for alerting.
B. Use Prometheus for metrics, Fluentd for log collection, Grafana for visualization, and Alertmanager for alerting.
C. Use cAdvisor for metrics, Fluentd for log aggregation, and Grafana for alerting.
D. Use Kibana for metrics, Logstash for log collection, and Prometheus for alerting.

Answer: B

Question: 330

Your application needs to handle configuration changes without restarting the Pods. Which approach will allow you to update the configuration of a running application dynamically?

A. PersistentVolumeClaim (PVC) B. ConfigMap with volume mount C. Environment variables
D. Secrets

Answer: B

Question: 331

Your application requires access to sensitive database credentials which need to be stored securely in your Kubernetes cluster. Which of the following methods is the most appropriate for storing and managing these credentials?

A. Store the credentials in a ConfigMap.
B. Store the credentials in the application code repository.
C. Store the credentials in a Secret.

D. Store the credentials as environment variables directly in the Pod specification.

Answer: C

Question: 332

You are tasked with deploying an application that needs environment-specific configurations stored in a ConfigMap. The application should securely access database credentials stored in a Kubernetes Secret. Which approach should you use to inject both ConfigMap and Secret into the application pods?

A. Inject ConfigMap data as environment variables and mount Secrets as volumes.
B. Use ConfigMap to store both environment-specific configurations and database credentials.
C. Mount both ConfigMap and Secret as volumes in the application pods.
D. Use Secrets to store both configurations and credentials, and inject them as environment variables.

Answer: A

Question: 333

You need to deploy a stateful application in your Kubernetes cluster and want to use an operator to manage its lifecycle. Which of the following steps should you take to deploy and use an operator? Select two correct answers.

A. Install the operator from the OperatorHub or another trusted source.
B. Configure the operator to watch for changes in the custom resources it manages.
C. Create a Custom Resource Definition (CRD) for the stateful application.
D. Manually scale the application replicas using kubectl scale commands.
E. Use the Kubernetes Dashboard to monitor and manage the operator.

Answer: A,B

Question: 334

You are tasked with deploying a complex application consisting of multiple microservices. Each microservice is deployed as a separate Kubernetes Deployment. You need to ensure that each Deployment can communicate with others via DNS and also want to minimize downtime during updates. Which two strategies should you employ?

A. Use RollingUpdate strategy for Deployments.
B. Use Kubernetes Services to expose each microservice.
C. Use StatefulSets instead of Deployments.
D. Deploy all microservices in a single Deployment.
E. Set up Pod Disruption Budgets for each microservice.

Answer: B,E

Question: 335

You are tasked with securing access to your Kubernetes cluster. You need to ensure that only specific users can create, update, or delete deployments, and that all resources comply with a set of predefined policies before they are admitted into the cluster. Which combination of strategies should you implement?

A. Enable PodSecurityPolicy for authorization and use ServiceAccounts for authentication.
B. Apply NodeSelector for authentication and use PodDisruptionBudgets for authorization.
C. Implement a custom admission controller webhook for policy enforcement and use Kubernetes RBAC for authorization.
D. Use Kubernetes Role-Based Access Control (RBAC) for authorization and configure Network Policies for admission control.

Answer: C

Question: 336

You have a Kubernetes Deployment running a critical application that occasionally consumes a high amount of memory, causing some Pods to be killed by the Out Of Memory (OOM) killer. How can you prevent this issue by configuring memory requests and limits?

A. Set a memory request equal to the memory limit to avoid any memory usage beyond a fixed value. B. Only set a memory limit and leave the memory request unset.
C. Set a memory request to the minimum amount of memory the application needs to start and run

efficiently, and a limit slightly above the peak memory usage.
D. Set a memory request and limit both to the maximum amount of memory the node can provide.

Answer: C

Question: 337

You need to ensure high availability for a set of Pods in your Kubernetes cluster, ensuring that traffic is distributed evenly across all instances of the Pods. Which type of Service should you use and how should you configure it? Select two correct answers.

A. Use a ClusterIP Service and configure it to point to the set of Pods.
B. Use an Ingress resource and configure it to point to the set of Pods.
C. Use a LoadBalancer Service and configure it to point to the set of Pods.
D. Use a Headless Service without selectors and configure it to point to the set of Pods.
E. Use a NodePort Service and configure it to point to the set of Pods.

Answer: A,C

Question: 338

You are tasked with deploying a new version of your application while minimizing downtime and ensuring a smooth transition for users. Your application is highly sensitive to abrupt changes. Which Kubernetes deployment strategy should you use?

A. Rolling Update Deployment
B. Blue/Green Deployment
C. Recreate Deployment
D. StatefulSet Deployment

Answer: B

Question: 339

You have a Kubernetes cluster with a frontend service in the frontend-namespace and a backend service in the backend-namespace. The frontend service needs to communicate with the backend service. Both namespaces have NetworkPolicies that deny all traffic by default. What is the correct way to allow the frontend service to communicate with the backend service?

A. Create a NetworkPolicy in the backend-namespace that allows ingress traffic from the frontend-namespace.
B. Use the serviceName.namespace.svc.cluster.local address for the backend service in the frontend service configuration.
C. Modify the NetworkPolicy in the frontend-namespace to allow traffic to the backend-namespace.
D. Add an annotation to the frontend service to bypass NetworkPolicies.

Answer: A

Question: 340

You need to create a Kubernetes Secret to store an API key and ensure it is available as an environment variable in a Pod. Which steps should you follow to correctly create and consume this Secret? Select two correct answers.

A. Create a Secret using kubectl create secret generic, then reference it in the Pod definition under env.
B. Write the API key directly into the Pod definition under env.
C. Create a Secret using kubectl create secret generic, then reference it in the Poddefinitionunder envFrom.
D. Store the API key in a Secret and manually pass it to the application at runtime.
E. Create a Secret using kubectl create configmap, then reference it in the Pod definition under envFrom.

Answer: A,C

Question: 341

You need to debug a Kubernetes pod that is not starting correctly. Which command should you use to view the details and events related to this pod?

A. kubectl top pod [pod-name]
B. kubectl logs [pod-name]
C. kubectl describe pod [pod-name]
D. kubectl get events

Answer: C

Question: 342

You have a Kubernetes cluster with multiple nodes, and you need to deploy an application that requires high availability and load balancing. The application should be able to handle traffic spikes and ensure zero downtime during deployments. Which Kubernetes resource and configuration will best achieve this

goal?

A. StatefulSet with Pod Anti-Affinity and HPA
B. CronJob with Parallelism and HPA
C. DaemonSet with HPA
D. Deployment with RollingUpdate strategy and HPA

Answer: D

Question: 343

You need to run a container in your Kubernetes pod with minimal privileges, only allowing it to bind to network ports below 1024. How should you configure this?

A. Set readOnlyRootFilesystem: true in the container's securityContext.
B. Use the capabilities field in the container's securityContext to add NET_BIND_SERVICE and drop all others.
C. Set runAsUser: 1000 in the pod's spec
D. Use allowPrivilegeEscalation: false in the container's securityContext. E.

Answer: B

Question: 344

You have an application deployed in Kubernetes that requires different configurations based on the geographic region of the user. You decide to use ConfigMaps to manage these region-specific configurations. If you need to dynamically update the configuration for all regions without causing downtime, what is the best approach?

A. Create a new ConfigMap with the updated configurations and update the deployment to use it. B. Patch the ConfigMap in the deployment using kubectl patch deployment <deployment-name>. C. Directly edit the ConfigMap using kubectl edit configmap <configmap-name>.
D. Use a Volume and mount the ConfigMap into the pod.

Answer: A

Question: 345

You need to schedule a task that runs a batch job every day at midnight to process data in your Kubernetes cluster. Which Kubernetes resource is most appropriate for this task?

A. StatefulSet
B. DaemonSet
C. Deployment
D. CronJob

Answer: D

Question: 346

Which of the following tools can be used to monitor the performance and health of a Kubernetes cluster?

A. Prometheus

B. Ansible
C. Terraform
D. Kafka

Answer: A

Question: 347

You have a Kubernetes deployment with multiple microservices. One of your microservices, service-A, is experiencing intermittent connectivity issues with service-B. You need to debug and resolve this issue. Which of the following steps should you take? Select two correct answers.

A. Check the network policies applied to service-A and service-B.
B. Enable automatic scaling for service-A based on CPU usage.
C. Increase the replica count of service-B to handle more requests.
D. Inspect the service endpoints and verify if the DNS resolution for service-B is correct.
E. Modify the deployment strategy of service-A to Recreate instead of RollingUpdate.

Answer: A,D

Question: 348

Your organization requires that every pod deployed in the cluster must have specific labels and annotations for monitoring and compliance. Additionally, certain users should only have read access to the pods in the production namespace. How should you configure the cluster to meet these requirements?

A. Use an admission controller webhook to enforce labels and annotations, and configure RBAC to manage read access.
B. Create a PodSecurityPolicy to enforce labels and annotations, and use RoleBindings for read access.
C. Use a ResourceQuota to enforce labels and annotations, and configure NodeSelectors for read access.
D. Implement a NetworkPolicy to enforce labels and annotations, and set up ClusterRoles for read access.

Answer: A

Question: 349

You have a multi-container Pod running a web server and a sidecar container for logging. The logging container needs to access the web server's log files, which are stored in a shared volume. How would you define the Pod specification to ensure the logging container has access to these log files?

A. Use a ConfigMap to store the log files and mount it to both containers.
B. Use a PersistentVolumeClaim to store the log files and mount it to both containers.
C. Create a shared emptyDir volume and mount it to both containers at the same path.
D. Use hostPath to mount the log directory from the Node's filesystem to both containers.

Answer: C

Question: 350

You have deployed a web application using a Deployment and exposed it using a Service of type NodePort. However, the application is still not accessible from outside the cluster. What could be the problem?

A. The kube-dns add-on is not installed in the cluster.
B. The nodes do not have a public IP address.
C. The Service is not annotated with the correct externalTrafficPolicy.
D. The Deployment is missing resource requests and limits.

Answer: B

Question: 351

You are using a Kubernetes deployment strategy that involves initially releasing the new version of an application to a small subset of users to test its stability before rolling it out to the entire user base. Which

deployment strategy are you using?

A. Rolling Update Deployment
B. Recreate Deployment
C. DaemonSet Deployment
D. Canary Deployment

Answer: D

Question: 352

You are designing a stateful application that requires data persistence even if the Pod is restarted or rescheduled. Which type of volume would you use to ensure the data remains available across Pod restarts?

A. EmptyDir
B. PersistentVolumeClaim (PVC)
C. ConfigMap
D. HostPath

Answer: B

Question: 353

You are managing a Kubernetes cluster and need to deploy a multi-tier application that consists of a frontend, backend, and a database. The frontend and backend should be exposed to the internet, but the database should only be accessible by the backend service. Additionally, you need to ensure that any configuration changes are applied correctly without downtime. Which Kubernetes resource configuration should you use to meet the above requirements?

A. LoadBalancer for all services, including the database
B. NodePort services for frontend and backend, and ClusterIP for database
C. Ingress for frontend and backend, and NetworkPolicy for database access
D. StatefulSet for the frontend, backend, and database

Answer: C

Question: 354

You need to restrict traffic in your Kubernetes cluster so that only pods with a specific label can communicate with a backend service. What steps should you take to configure this?

A. Define a PodSecurityPolicy that restricts ingress traffic to the backend service.
B. Set up an Ingress controller to manage traffic to the backend service.
C. Create a NetworkPolicy that allows ingress traffic from pods with the specific label to the backend service.
D. Create a ServiceAccount with the necessary permissions and assign it to the pods.

Answer: C

Question: 355

You are managing a Kubernetes cluster and have implemented Role-Based Access Control (RBAC) for securing your resources. A new developer needs access to view and edit deployments within a specific namespace. Which of the following steps should you take to grant the necessary permissions? Select two correct answers.

A. Bind the Role to the developer using a RoleBinding.
B. Create a ClusterRole with the necessary permissions to view and edit deployments.
C. Create a Role with the necessary permissions to view and edit deployments in the specific namespace.
D. Modify the kube-apiserver configuration to grant the developer access.

Answer: A,C

Question: 356

You are deploying an application in Kubernetes that consists of a Deployment with three replicas. The application has an endpoint /status that returns a status code of 200 when the application is running correctly. Additionally, the application has an initialization process that takes about 60 seconds to complete. You need to ensure that the application does not receive traffic until it is fully initialized and remains available during its lifecycle. How would you configure the probes for this Deployment?

A. Configure both a readinessProbe and a livenessProbe with httpGet to /status, both with an initial delay of 60 seconds.
B. Configure a readinessProbe with httpGet to /status and an initial delay of 60 seconds.
C. Configure a readinessProbe with exec to check process ID and a livenessProbe with httpGet to /status.
D. Configure a livenessProbe with httpGet to /status and an initial delay of 60 seconds. **Answer: A**

Question: 357

You have a namespace production with two services: web-service and db-service. You need to restrict all traffic to db-service except from web-service. How do you achieve this using NetworkPolicies?

A. Use a ClusterIP service for db-service to automatically restrict access.

B. Create a NetworkPolicy that allows all ingress traffic to db-service.
C. Create a NetworkPolicy in the production namespace that allows ingress traffic to db-service only from pods with a specific label.
D. Create a NetworkPolicy in the production namespace that denies all egress traffic from web-service except to db-service.

Answer: C

Question: 358

You are developing an application that requires a Kubernetes job to periodically backup data from a persistent volume. The job should run every night at 2 AM and must complete successfully before the next instance is scheduled. However, if the job fails, it should retry a maximum of three times. How would you configure the Kubernetes job to meet the above requirements?

A. Use a CronJob withschedule: "0 2 * * *",restartPolicy: OnFailure, and backoffLimit: 3.
B. Use a CronJob withschedule: "0 2 * * *", restartPolicy: Never, and backoffLimit: 3.
C. Use a CronJob withschedule: "0 2 * * *",restartPolicy: OnFailure, and backoffLimit: 0.
D. Use a regular Job with restartPolicy: Always and manually configure cron on the host system to run the job at 2 AM.

Answer: A

Question: 359

Your application has been experiencing intermittent performance issues due to resource contention in your Kubernetes cluster. To address this, you need to ensure that each pod in the production namespace has appropriate CPU and memory requests and limits set. Additionally, you want to enforce a maximum resource usage policy for the namespace. What steps should you take?

A. Apply a PodDisruptionBudget to limit the number of pods that can be disrupted at once.
B. Configure Horizontal Pod Autoscaler (HPA) to manage resource usage dynamically.
C. Set resource requests and limits in each pod's specification and apply a ResourceQuota for the production namespace.
D. Use a LimitRange to set default resource requests and limits, and configure a NetworkPolicy for the namespace.

Answer: C

Question: 360

You need to secure communication in your Kubernetes cluster using NetworkPolicies. You have two namespaces: frontend and backend. You want to allow incoming traffic to all pods in the backend namespace only from pods in the frontend namespace. Which NetworkPolicy would you apply?

A. NetworkPolicy that allows egress from backend namespace to frontend namespace.
B. A NetworkPolicy that allows ingress from frontend namespace to backendnamespace.
C. A NetworkPolicy that allows egress from frontend namespace to backendnamespace.
D. A NetworkPolicy that blocks all ingress traffic in the backend namespace.

Answer: B

Question: 361

You are deploying a machine learning application on Kubernetes, and it is crucial to ensure that each pod gets a guaranteed minimum amount of CPU and memory while also capping the maximum usage. Which of the following configurations should you use to achieve this?

A. Set resource requests and limits for CPU and memory in the pod specification.
B. Use HorizontalPodAutoscaler (HPA) to manage CPU and memory allocation.
C. Deploy the application with a StatefulSet to ensure resource allocation.
D. Configure a PodDisruptionBudget (PDB) to guarantee minimum resources for each pod.

Answer: A

Question: 362

You need to define resource requests and limits for a memory-intensive application to ensure it runs efficiently without affecting other applications. Which configuration should you use?

A. Set a memory request but no limit in the Pod definition.
B. Set both memory requests and limits to appropriate values in the Pod definition.
C. Set a high memory limit and no request in the Pod definition.
D. Set memory requests and limits in the Deployment definition and not in the Pod definition.

Answer: B

Question: 363

Your team needs to deploy a microservice-based application where each microservice should be independently updatable without downtime. The application components include a user interface, several backend services, and a database. Which combination of Kubernetes resources should you use to achieve this?

A. Deployment for user interface, Deployment for backend services, StatefulSet for database B. StatefulSet for user interface, Deployment for backend services, Deployment for database
C. Deployment for user interface, ReplicaSet for backend services, Deployment for database
D. DaemonSet for user interface, Deployment for backend services, PersistentVolumeClaim for database

Answer: A

Question: 364

You are managing a Kubernetes application that consists of multiple microservices. One of the microservices, payment-service, experiences intermittent high latency during peak hours. You need to investigate the issue and determine the cause. Which of the following actions should you take first?

A. Check the resource usage (CPU and memory) of the payment-service Pods.
B. Apply a network policy to isolate the payment-service from other microservices.
C. Implement Horizontal Pod Autoscaling (HPA) for the payment-service.
D. Increase the number of replicas for the payment-service.

Answer: A

Question: 365

You need to store sensitive information such as database credentials in a Kubernetes cluster and make it available to a pod. What steps should you follow to securely store and access these credentials in the pod?

A. Store the credentials in the pod's environment variables directly in the deployment manifest.
B. Create a ConfigMap, mount it as a volume in the pod, and use environment variables to access it.
C. Create a Secret, mount it as a volume in the pod, and use environment variables to access it.
D. Create a Secret, mount it as a volume in the pod, and access it directly from the volume.

Answer: D

Question: 366

Which Kubernetes resource should you use to ensure that a specific pod runs on every node in the cluster?

A. Deployment B. StatefulSet C. DaemonSet D. CronJob

Answer: C

Question: 367

You are deploying an application in Kubernetes that needs to interact with the Kubernetes API to list and watch resources. How can you ensure that the application has the necessary permissions to perform these actions securely? Select two correct answers.

A. Create a new ServiceAccount, bind it to a Role with the required permissions, and assign it to the application Pod.
B. Assign the default ServiceAccount to the application Pod.
C. Create a new ServiceAccount and assign it to the application Pod without additional Role or RoleBinding configurations.
D. Create a new ServiceAccount, bind it to a ClusterRole with the required permissions, and assign it to the application Pod.
E. Grant cluster-admin permissions to the default ServiceAccount.

Answer: A,D

Question: 368

You are managing a Kubernetes cluster and need to perform a rolling update on a Deployment to update your application from version 1.0 to version 2.0. Which steps would ensure the update is performed correctly without downtime?

A. Update the image of the Deployment and Kubernetes will automatically perform a rolling update. B. Update the ConfigMap used by the Deployment to point to the new version and restart the Pods. C. Delete all existing Pods manually and create new ones with the updated image.
D. Scale down the current Deployment to zero replicas and then apply the update.

Answer: A

Question: 369

You are deploying a microservices application in a Kubernetes cluster where services need to communicate with each other securely and reliably. Which of the following features would you use to ensure automatic service discovery and load balancing within the cluster?

A. Ingress Controller
B. Service

C. ConfigMap
D. NetworkPolicy

Answer: B

Question: 370

You have a Kubernetes cluster with multiple teams working on different projects. You need to ensure that each team has restricted access to only their own namespaces and that all deployments comply with company policies. Which combination of Kubernetes features should you use to achieve this?

A. Configure HorizontalPodAutoscalers (HPA) to manage resource usage and ensure policy compliance.
B. Use Role-Based Access Control (RBAC) to define roles and role bindings, and Admission Controllers to enforce policies.
C. Use Network Policies to restrict access to namespaces and ensure compliance with company policies.
D. Enable Kubernetes Secrets and ConfigMaps to manage access control and enforce policies. **Answer: B**

Question: 371

You are deploying a multi-tier application in Kubernetes consisting of a frontend, backend, and database. You need to ensure that only the frontend can communicate with the backend, and only the backend can communicate with the database. Which NetworkPolicy configurations would you use to enforce these rules? Select two correct answers.

A. Create a NetworkPolicy for the frontend to allow all outbound traffic and a NetworkPolicy for the backend to allow all inbound traffic.
B. Create a NetworkPolicy for the frontend to allow ingress from all sources and a NetworkPolicy for the backend to allow egress to the database only.
C. Create a NetworkPolicy for the backend to allow ingress from the frontend only, and another NetworkPolicy for the database to allow ingress from the backend only.
D. Create a NetworkPolicy for the frontend to deny all outbound traffic except to the backend, and another NetworkPolicy for the backend to deny all outbound traffic except to the database.

Answer: C,D

Question: 372

Which Kubernetes object is used to set resource requests and limits for pods?

A. ConfigMap
B. ResourceQuota
C. Secret
D. PodSpec

Answer: D

Question: 373

You are designing a container image for a web application that needs to handle configuration changes dynamically without rebuilding the image. Which two practices should you follow to achieve this?

A. Include the configuration file as a part of the build process.
B. Use a ConfigMap to manage configuration data.
C. Hard-code the configuration within the Dockerfile.
D. Embed the configuration file in the container image.
E. Use environment variables for configuration.

Answer: B,E

Question: 374

You need to deploy a stateless web application on a Kubernetes cluster that should be accessible via a domain name. The application also requires some environment variables to be set during the deployment. Which combination of Kubernetes resources should you use?

A. DaemonSet, Ingress, ConfigMap
B. CronJob, Ingress, ConfigMap
C. StatefulSet, Service, Secret
D. Deployment, Service, ConfigMap

Answer: D

Question: 375

You are responsible for deploying a new version of an application to a Kubernetes cluster. The new version introduces changes to the service mesh configuration that need to be rolled out gradually. Which Kubernetes feature allows you to implement a gradual rollout of the new version while ensuring minimal disruption to the existing traffic?

A. Blue-Green Deployment
B. RollingUpdate Deployment Strategy
C. Helm
D. Canary Deployment

Answer: D

Question: 376

Which command is used to create a Deployment in Kubernetes from a YAML file?

A. kubectl create deployment -f deployment.yaml
B. kubectl start deployment.yaml
C. kubectl run -f deployment.yaml
D. kubectl apply -f deployment.yaml

Answer: D

Question: 377

You need to configure a Kubernetes pod to use a specific ServiceAccount that has permissions to access secrets in the namespace. What steps should you follow to ensure this setup?

A. Define the necessary permissions in the pod's metadata.
B. Create a ServiceAccount, bind it to a Role with necessary permissions, and specify it in the pod spec.
C. Create a ConfigMap with the ServiceAccount name and mount it in the pod.
D. Use environment variables to pass the ServiceAccount name to the pod.

Answer: B

Question: 378

You have a Kubernetes application that includes multiple microservices. One of the microservices is experiencing intermittent connectivity issues with its dependent service. You need to diagnose the root cause. Which two of the following commands can help you investigate the issue?

A. kubectl exec -it <pod-name> -- /bin/bash
B. kubectl delete pod <pod-name>
C. kubectl create namespace debug
D. kubectl logs <pod-name>
E. kubectl describe pod <pod-name>

Answer: D,E

Question: 379

You have a Kubernetes cluster with a StatefulSet running a database application that requires persistent storage. Recently, the application has been experiencing performance issues due to high disk I/O. As the CKAD, you need to improve the performance of the database application. Which of the following actions is the best approach to mitigate this issue?

A. Migrate the StatefulSet's PersistentVolume to a faster storage class.
B. Increase the CPU and memory limits for the StatefulSet.
C. Use a sidecar container to offload disk-intensive operations.
D. Scale the StatefulSet horizontally by increasing the number of replicas.

Answer: A

Question: 380

You need to define a new custom resource for your Kubernetes cluster to manage a specific application lifecycle. Which resource should you use to create this custom resource?

A. ResourceQuota
B. ConfigMap
C. CustomResourceDefinition (CRD)
D. Horizontal Pod Autoscaler (HPA)

Answer: C

Question: 381

You need to ensure that your Kubernetes cluster enforces a memory quota for all namespaces to prevent any single namespace from using more than its allocated memory. How would you configure this?

A. Use a NetworkPolicy to control memory usage across namespaces.
B. Create a ResourceQuota object in each namespace with the specified memory limits.
C. Create a LimitRange object with memory limits and apply it to each namespace.
D. Set memory requests and limits at the Pod level without configuring any quota.
E.

Answer: B

Question: 382

You have a Kubernetes cluster that runs a web application which needs to scale based on incoming traffic. During peak hours, the application should handle thousands of requests per second, but during off-peak hours, the traffic drops significantly. You need to ensure high availability and efficient resource usage. Which Kubernetes resource should you use to manage the deployment of this web application?

A. StatefulSet
B. DaemonSet
C. CronJob
D. Deployment

Answer: D

Question: 383

How can you enforce a memory usage limit for all pods in a specific namespace?

A. Define a Secret with memory limits
B. Update the kube-apiserver configuration
C. Create a ResourceQuota in the namespace
D. Set a ConfigMap in the namespace

Answer: C

Question: 384

You need to assign a specific ServiceAccount to a Pod and ensure that the Pod can access a specific Kubernetes API endpoint securely. Which steps should you follow to correctly achieve this? Select two correct answers.

A. Create a ServiceAccount, bind it to a ClusterRole that allows access to the required API endpoint, and specify the ServiceAccount in the Pod specification.
B. Create a ServiceAccount, bind it to a Role that allows access to the required API endpoint, and specify the ServiceAccount in the Pod specification.
C. Create a Role that allows access to the required API endpoint and specify it in the Pod specification.
D. Create a ServiceAccount and specify it in the Pod specification without additional Role or RoleBinding configurations.
E. Use the default ServiceAccount and bind it to a Role with the required permissions.

Answer: A,B

Question: 385

You are responsible for deploying a logging agent across all nodes in a Kubernetes cluster to ensure that logs are collected from every node. This agent should run on every node, including any new nodes added in the future. Which Kubernetes resource should you use to accomplish this?

A. CronJob
B. DaemonSet
C. StatefulSet
D. Deployment

Answer: B

Question: 386

You are tasked with ensuring a multi-container pod in your Kubernetes cluster communicates securely over HTTPS. The pod consists of a frontend and backend container. You need to configure mutual TLS authentication between these containers. Which of the following steps is necessary to achieve this?

A. Configure a NetworkPolicy to allow HTTPS traffic between the containers.
B. Mount a ConfigMap containing the certificates to both containers.
C. Use a PersistentVolume to store the certificates.
D. Set up a LoadBalancer service to expose the backend over HTTPS.

Answer: B

Question: 387

You need to ensure that your application remains available during updates and that only a few Pods are updated at a time. Which Deployment strategy in Kubernetes allows you to perform such an update?

A. Recreate strategy.
B. RollingUpdate strategy.
C. Canary strategy.
D. OnDelete strategy.

Answer: B

Question: 388

You need to pass sensitive API keys to your application running in a Kubernetes pod using environment variables. What is the correct way to achieve this?

A. Add the API keys directly in the pod's environment variables in the deployment spec.
B. Create a ConfigMap, and reference it in the pod's environment variables in the deployment spec.
C. Store the API keys in a plain text file and reference it in the pod's environment variables.
D. Create a Secret, and reference it in the pod's environment variables in the deployment spec.

Answer: D

Question: 389

You have a Kubernetes Deployment running multiple Pods that occasionally fail due to CPU resource contention. To ensure your application runs smoothly, how should you configure the CPU requests and limits for the Pods?

A. Only set a CPU limit and leave the CPU request unset.
B. Do not set any CPU requests or limits, allowing the application to use as much CPU as needed.
C. Set a CPU request to the minimum amount of CPU required for normal operation and a limit to the maximum amount the application might need during peak loads.
D. Set both CPU request and limit to the maximum amount of CPU available on the node

Answer: C

Question: 390

Which Kubernetes resource is most appropriate for running a task periodically, such as a nightly backup?

A. CronJob
B. ReplicaSet
C. Deployment
D. DaemonSet

Answer: A

Question: 391

Your Kubernetes application needs to consume different configurations based on deployment environments (e.g., staging and production). Additionally, the application must securely handle sensitive data such as API keys. How should you configure the pods to meet these requirements?

A. Use a single ConfigMap to store both environment-specific configurations and API keys.
B. Use ConfigMaps for environment-specific configurations and Secrets for API keys.
C. Create multiple ConfigMaps for different environments and use one Secret for all sensitive data.
D. Inject both ConfigMaps and Secrets as environment variables in the pod specification.

Answer: B

Question: 392

You are experiencing an issue where a pod in your Kubernetes cluster is repeatedly crashing. Which command will provide you with the logs of the last terminated container in the pod to help you debug the issue?

A. kubectl logs --previous <pod_name>
B. kubectl exec -it <pod_name> -- /bin/sh
C. kubectl describe pod <pod_name>
D. kubectl get events

Answer: A

Question: 393

You have an application running in a Kubernetes cluster that sometimes becomes unresponsive. The application has an endpoint /healthz that should return a status code of 200 when it is healthy. To ensure the application remains responsive, you want Kubernetes to restart the pod if the application becomes unresponsive. How would you configure the health checks for this application?

A. Configure a readinessProbe with exec to check the application process and an initial delay of 10 seconds.
B. Configure a readinessProbe with httpGet to /healthz and an initial delay of 10 seconds.
C. Configure both a readinessProbe and a livenessProbe with httpGet to /healthz, both with an initial delay of 10 seconds.
D. Configure a livenessProbe with httpGet to /healthz and an initial delay of 10 seconds.

Answer: D

Question: 394

You need to monitor the health and performance of your Kubernetes application, which is running in a production environment. The application consists of several services that need to be monitored for resource usage, latency, and error rates. Which tool or combination of tools would you use to achieve this?

A. Jenkins and Spinnaker
B. Prometheus and Grafana
C. Kubernetes Dashboard and Minikube
D. Helm and Kustomize

Answer: B

Question: 395

Your development team frequently deploys resource-intensive applications in the dev namespace, causing other applications to suffer from resource starvation. To prevent this, you need to enforce policies that control resource allocation in the dev namespace. What should you configure?

A. Define a ResourceQuota for the dev namespace and apply a LimitRange for default resource requests and limits.
B. Use a PersistentVolumeClaim to manage storage resources.
C. Create a ServiceAccount with limited resource permissions.
D. Configure a Horizontal Pod Autoscaler to scale down pods when resources are limited.

Answer: A

Question: 396

You need to secure your Kubernetes cluster by implementing proper authentication, authorization, and admission control mechanisms. Which of the following steps should you take to enhance security? Select two correct answers.

A. Configure kube-scheduler to enforce security policies.
B. Use OpenID Connect (OIDC) to integrate with an external identity provider for authentication.

C. Disable RBAC and rely solely on network policies for security.
D. Allow unauthenticated access for all internal services to simplify management.
E. Use Admission Controllers to enforce security policies on incoming requests.

Answer: B,E

Question: 397

You need to manage a complex stateful application that requires frequent configuration changes and automated updates. Which Kubernetes resource allows you to define and manage such an application as a custom resource?

A. CustomResourceDefinition (CRD)
B. PodSecurityPolicy
C. Service
D. Deployment

Answer: A

Question: 398

You have a Kubernetes cluster running a web application. The application needs to be updated with a new version, but it's critical to ensure that the application is always available to users during the update. Which approach would you use to perform this update?

A. Blue/Green Deployment
B. Recreate Deployment
C. Rolling Update Deployment
D. Job Deployment

Answer: C

Question: 399

You have been notified about upcoming deprecations in the Kubernetes API versions you are currently using in your production environment. You need to ensure your applications continue to operate smoothly with the next cluster upgrade. What is the most effective way to prepare your applications for the deprecated API versions?

A. Update your applications to use the latest stable API versions before upgrading the cluster.
B. Use a Kubernetes feature gate to enable the deprecated APIs post-upgrade.
C. Ignore the deprecation warnings and proceed with the cluster upgrade.
D. Only update the Kubernetes client libraries in your application.

Answer: A

Question: 400

You have a Kubernetes cluster with a deployment for appA in the default namespace. appA depends on an external database service hosted outside the cluster. The database service is accessible via the DNS name db.external.com. How can you configure appA to communicate with the external database service?

A. Add the db.external.com entry in the /etc/hosts file of appA pods.
B. Use a NodePort Service to expose db.external.com inside the cluster.
C. Create an ExternalName Service in the default namespace pointing to db.external.com.

D. Configure a headless Service for db.external.com in the default namespace.

Answer: C

Question: 401

You are troubleshooting an issue in your Kubernetes cluster where the application is failing intermittently. You suspect the issue is related to a specific pod. Which command will help you fetch the logs from the suspected pod?

A. kubectl describe pod [pod-name]
B. kubectl get events
C. kubectl logs -f [pod-name]
D. kubectl exec [pod-name] -- tail -f /var/log/app.log

Answer: C

Question: 402

You need to ensure that only authorized users can access a specific Kubernetes namespace to deploy applications. Which Kubernetes resource should you configure to restrict access to this namespace?

A. RoleBinding
B. NetworkPolicy
C. ResourceQuota
D. ServiceAccount

Answer: A

Question: 403

You need to expose your Kubernetes service to external traffic so that it is accessible from outside the cluster. Which type of service should you use to achieve this?

A. Use a Headless service.
B. Use a LoadBalancer service.
C. Use a NodePort service.
D. Use a ClusterIP service.

Answer: B

Question: 404

You have a Kubernetes cluster with multiple applications running in different namespaces. One of your applications, running in the production namespace, is experiencing high memory usage and frequent OutOfMemory (OOM) errors. Which of the following steps would you take to diagnose and resolve this issue?

A. Modify the service type from ClusterIP to NodePort.
B. Decrease the CPU limit for the application pod.
C. Increase the memory limit for the application pod.
D. Change the namespace of the application to default.

Answer: C

Question: 405

You have deployed a multi-tier web application in a Kubernetes cluster with separate services for the frontend, backend, and database. The frontend service is exposed using an Ingress resource. Despite configuring the Ingress, users cannot access the frontend application from the internet. What could be the problem?

A. The Ingress resource is using an invalid hostname.
B. The backend service is not correctly configured.
C. The Ingress resource is missing the path definition.
D. The Ingress controller is not installed in the cluster. **Answer: D**

Question: 406

What is the primary benefit of using a Deployment in Kubernetes?

A. To perform automatic backups of pod data.
B. To provide a persistent storage solution for pods.
C. To ensure a specified number of pod replicas are running at all times.
D. To manage network policies and security rules.

Answer: C

Question: 407

You have deployed an application in Kubernetes using a Deployment with multiple replicas. The application needs to store session data in a Redis instance running as a StatefulSet. You notice that the application pods are unable to connect to the Redis service. What could be the problem?

A. The Redis StatefulSet is not using persistent storage.
B. The Redis StatefulSet is using a ClusterIP service instead of a headless service.
C. The Redis StatefulSet pods are not labeled correctly.
D. The application Deployment and Redis StatefulSet are in different namespaces without the correct DNS resolution.

Answer: D

Question: 408

You are deploying a multi-service application where one of the services, 'Service-X', requires access to an external database. The database credentials should be kept secure and should not be hardcoded in the application's source code. How can you securely provide these credentials to 'Service-X' in Kubernetes? Select two correct answers.

A. Include the database credentials in the Pod specification as environment variables.
B. Use an external secret management system and retrieve the credentials dynamically from within the application code.
C. Store the database credentials in a ConfigMap and reference it in the Pod specification.
D. Store the database credentials in a Kubernetes Secret and reference the Secret in the Pod specification as environment variables.

E. Store the database credentials in a Kubernetes Secret and mount the Secret as a volume in the Pod.

Answer: D,E

Question: 409

You need to deploy a stateless application with three replicas and ensure that traffic is evenly distributed across all replicas. Which Kubernetes object would you use, and what key configurations would ensure proper deployment and traffic distribution?

A. Use a Pod directly with a Service of type ClusterIP.
B. Use a StatefulSet with a Service of type NodePort.
C. Use a ReplicaSet with a Service of type LoadBalancer.
D. Use a Deployment with a Service of type ClusterIP.

Answer: D

Question: 410

When Kubernetes announces the deprecation of an API version, what is the typical duration before it is removed in a future release?

A. One release cycle
B. Immediately
C. Two release cycles
D. Ten release cycles

Answer: C

Question: 411

You have a Kubernetes pod that needs to authenticate with an external service using a ServiceAccount token. How can you ensure that the pod uses the ServiceAccount token correctly?

A. Use a PersistentVolume to store the token and mount it in the pod.
B. Mount the ServiceAccount token as a volume and configure the application to use it.
C. Store the token in a ConfigMap and reference it in the pod.
D. Hardcode the ServiceAccount token in the application's configuration file.

Answer: B

Question: 412

You have a Kubernetes application composed of multiple microservices. One of the microservices, service-X, is intermittently failing to connect to the database. You need to troubleshoot and resolve this issue. Which of the following steps should you take? Select two correct answers.

A. Verify the database credentials in the configuration of service-X.
B. Ensure the service-X pod is scheduled on a node with sufficient resources.
C. Check the resource limits set on the database deployment.
D. Use kubectl logs <service-X-pod> to check the logs of the service-X pod for connection errors.
E. Increase the CPU limits for the service-X pod.

Answer: A,D

Question: 413

You need to store temporary data for a batch processing application that does not require persistence after the Pod completes its job. Which type of volume should you use?

A. PersistentVolumeClaim (PVC)
B. Secret
C. EmptyDir
D. ConfigMap

Answer: C

Question: 414

You are managing a Kubernetes cluster for an e-commerce application. The payment service needs to securely access an external payment gateway, requiring sensitive API keys. You must ensure these API keys are securely stored and consumed by the payment service without exposing them in the container image or environment variables. How should you configure this?

A. Use a PersistentVolume to store the API keys and mount it in the payment service pods.
B. Create a Kubernetes Secret to store the API keys and mount it as a volume in the payment service pods.
C. Store the API keys in a ConfigMap and mount it as an environment variable in the payment service pods.
D. Deploy the payment service with a fixed number of replicas to manage API key access.

Answer: B

Question: 415

Your application requires persistent storage, high availability, and the ability to automatically scale up during peak hours. Additionally, you need to ensure that the application can be updated without downtime. Which Kubernetes resources and configurations should you use?

A. Deployment with RollingUpdate strategy, Cluster Autoscaler, and Secret
B. Deployment with Recreate strategy, HPA, and PersistentVolumeClaims
C. StatefulSet with RollingUpdate strategy, HPA, and PersistentVolumeClaims
D. StatefulSet with Recreate strategy, Cluster Autoscaler, and ConfigMap

Answer: C

Question: 416

You are responsible for monitoring a Kubernetes application and need to ensure that you have visibility into the performance and health of the application. Which two of the following commands will help you achieve this using built-in CLI tools?

A. kubectl label pod <pod-name> env=production
B. kubectl port-forward <pod-name> 8080:80
C. kubectl top pod
D. kubectl get events
E. kubectl delete node <node-name>

Answer: C,D

Question: 417

You are responsible for managing a Kubernetes cluster hosting a multi-tier application. The application is experiencing performance issues, and you suspect that some pods are consuming more CPU and memory than expected, impacting the performance of other services. How can you ensure that each pod gets a fair share of resources without starving other pods in the namespace?

A. Use Network Policies to control the resource consumption of each pod.
B. Set resource requests and limits for CPU and memory on each pod in the application.
C. Enable HorizontalPodAutoscaler (HPA) to automatically adjust the number of pods based on resource usage.
D. Deploy the application with a high priority class to ensure it always gets the required resources.

Answer: B