

Kubernetes Certificate Expire Causes Cluster Wide Communication Halt

Contents

[Introduction](#)

[Problem](#)

[Solution](#)

Introduction

This document describes a possible outage issue that the customers might face when they have a Kubernetes based system that has been installed for more than 365 days. Furthermore, it goes through the steps needed to fix the situation and get the Kubernetes based system back up and running.

Problem

After one year of time of default installed Kubernetes cluster, the client certificates expire. You will not be able to access Cisco CloudCentre Suite (CCS). Although it will still appear up, you will not be able to log in. If you navigate to the kubectl CLI, you will see this error, "Unable to connect to the server: x509: certificate has expired or is not yet valid."

You can run this bash script to see the expiration date of their certificates:

```
for crt in /etc/kubernetes/pki/*.crt; do
    printf '%s: %s\n' \
        "$(date --date="$(openssl x509 -enddate -noout -in "$crt"|cut -d= -f 2)" --iso-8601)" \
        "$crt"
done | sort
```

You can also find an opensource workflow for Action Orchestrator that will monitor this daily and alert them to issues.

https://github.com/cisco-cx-workflows/cx-ao-shared-workflows/tree/master/CCSCheckKubernetesExpiration_definition_workflow_01E01VIRWZDE24mWlsHrqCGB9xUix0f9ZxG

Solution

New certificates have to be re-issued via Kubeadm across the cluster and then you need to re-join the worker nodes to the masters.

1. Log in to a master node.
2. Get its IP address **via ip address show**.

```
[root@cx-ccs-prod-master-d7f34f25-f524-4f90-9037-7286202ed13a3 kubernetes]# ip address show
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
inet 127.0.0.1/8 scope host lo
valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 8920 qdisc pfifo_fast state UP group default
qlen 1000
link/ether fa:16:3e:19:63:a2 brd ff:ff:ff:ff:ff:ff
inet 192.168.1.20/24 brd 192.168.1.255 scope global dynamic eth0
valid_lft 37806sec preferred_lft 37806sec
3: docker0: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc noqueue state DOWN group
default
link/ether 02:42:d0:29:ce:5e brd ff:ff:ff:ff:ff:ff
inet 172.17.0.1/16 scope global docker0
valid_lft forever preferred_lft forever
13: tunl0@NONE: <NOARP,UP,LOWER_UP> mtu 1430 qdisc noqueue state UNKNOWN group default qlen
1000
link/ipip 0.0.0.0 brd 0.0.0.0
inet 172.16.176.128/32 brd 172.16.176.128 scope global tunl0
valid_lft forever preferred_lft forever
14: cali65453a0219d@if4: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1430 qdisc noqueue state UP
group default
link/ether ee:ee:ee:ee:ee:ee brd ff:ff:ff:ff:ff:ff link-netnsid 4
```

3. Navigate to the Kubernetes directory via **cd /etc/kubernetes**.

4. Create a file called **kubeadmCERT.yaml** via **vi kubeadmCERT.yaml**.

5. The file should look like this:

```
apiVersion: kubeadm.k8s.io/v1alpha1
kind: MasterConfiguration
api:
  advertiseAddress: <IP ADDRESS FROM STEP 2>
kubernetesVersion: v1.11.6
#NOTE: If the customer is running a load balancer VM then you must add these lines after...
#apiServerCertSANS:
#- <load balancer IP>
```

6. Backup your old certificates and keys. This is not required but recommended. Make a backup directory and copy these files to it.

```
#Files
#apiserver.crt
#apiserver.key
#apiserver-kubelet-client.crt
#apiserver-kubelet-client.key
#front-proxy-client.crt
#front-proxy-client.key

#ie
cd /etc/kubernetes/pki
mkdir backup
mv apiserver.key backup/apiserver.key.bak
```

7. If you skipped Step 6., you can just delete the previously mentioned files via **rm** command like **rm apiserver.crt**.

8. Navigate back to where your **kubeadmCERT.yaml** file is located. Generate a new apiserver cert via **kubeadm --config kubeadmCERT.yaml alpha phase certs apiserver**.

```
[root@cx-ccs-prod-master-d7f34f25-f524-4f90-9037-7286202ed13a3 kubernetes]# kubeadm --
config kubeadmCERT.yaml alpha phase certs apiserver
[certificates] Generated apiserver certificate and key.
[certificates] apiserver serving cert is signed for DNS names [cx-ccs-prod-master-d7f34f25-
f524-4f90-9037-7286202ed13a3 kubernetes kubernetes.default kubernetes.default.svc
kubernetes.default.svc.cluster.local] and IPs [10.96.0.1 192.168.1.20]
```

9. Generate new apiserver kubelet cert via **kubeadm --config kubeadmCERT.yaml alpha phase certs apiserver-kubelet-client**.
10. Generate new front-proxy-client cert via **kubeadm --config kubeadmCERT.yaml alpha phase certs front-proxy-client**.
11. In the **/etc/kubernetes** folder, backup the **.conf** files. Not required but recommended. You should have **kubelet.conf**, **controller-manager.conf**, **scheduler.conf**, and possibly **admin.conf**. You can delete them if you don't want to back them up.
12. Generate new configuration files via **kubeadm --config kubeadmCERT.yaml alpha phase kubeconfig all**.

```
[root@cx-ccs-prod-master-d7f34f25-f524-4f90-9037-7286202ed13a3 kubernetes]# kubeadm --
config kubeadmCERT.yaml alpha phase kubeconfig all
[kubeconfig] Wrote KubeConfig file to disk: "/etc/kubernetes/admin.conf"
[kubeconfig] Wrote KubeConfig file to disk: "/etc/kubernetes/kubelet.conf"
[kubeconfig] Wrote KubeConfig file to disk: "/etc/kubernetes/controller-manager.conf"
[kubeconfig] Wrote KubeConfig file to disk: "/etc/kubernetes/scheduler.conf"
```

13. Export your new **admin.conf** file to your host.

```
cp -i /etc/kubernetes/admin.conf $HOME/.kube/config
chown $(id -u):$(id -g) $HOME/.kube/config
chmod 777 $HOME/.kube/config
export KUBECONFIG=.kube/config
```

14. Reboot the master node via **shutdown -r now**.
15. Once the master is backed up, check whether kubelet is running via **systemctl status kubelet**.
16. Verify Kubernetes via **kubectl get nodes**.

```
[root@cx-ccs-prod-master-d7f34f25-f524-4f90-9037-7286202ed13a3 ~]# kubectl get nodes
NAME                                     STATUS    ROLES    AGE
VERSION
cx-ccs-prod-master-d7f34f25-f524-4f90-9037-7286202ed13a1   Ready    master   1y
v1.11.6
cx-ccs-prod-master-d7f34f25-f524-4f90-9037-7286202ed13a2   Ready    master   1y
v1.11.6
cx-ccs-prod-master-d7f34f25-f524-4f90-9037-7286202ed13a3   Ready    master   1y
v1.11.6
```

```

cx-ccs-prod-worker-d7f34f25-f524-4f90-9037-7286202ed13a1   NotReady   <none>     1y
v1.11.6
cx-ccs-prod-worker-d7f34f25-f524-4f90-9037-7286202ed13a2   NotReady   <none>     1y
v1.11.6
cx-ccs-prod-worker-d7f34f25-f524-4f90-9037-7286202ed13a3   NotReady   <none>     1y
v1.11.6
cx-ccs-prod-worker-d7f34f25-f524-4f90-9037-7286202ed13a4   NotReady   <none>     1y
v1.11.6
cx-ccs-prod-worker-d7f34f25-f524-4f90-9037-7286202ed13a5   NotReady   <none>     1y
v1.11.6
cx-ccs-prod-worker-d7f34f25-f524-4f90-9037-7286202ed13a6   NotReady   <none>     1y
v1.11.6

```

17. Repeat steps 1. to 16. for each master node.

18. On one master, generate a new join token via **kubeadm token create --print-join-command**. Copy that command for later use.

```

[root@cx-ccs-prod-master-d7f34f25-f524-4f90-9037-7286202ed13a1 k8s-mgmt]# kubeadm token
create
--print-join-command kubeadm join 192.168.1.14:6443 --token mlynvj.f4n3et3poki88ry4
--discovery-token-ca-cert-hash
sha256:4d0c569985c1d460ef74dc01c85740285e4af2c2369ff833eed1ba86e1167575

```

19. Get the IPs of your workers via **kubectl get nodes -o wide**.

20. Log in to a worker like **ssh -i /home/cloud-user/keys/gen3-ao-prod.key cloud-user@192.168.1.17** and navigate to root access.

21. Stop the kubelet service via **systemctl stop kubelet**.

22. Remove the old configuration files, these include **ca.crt**, **kubelet.conf**, and **bootstrap-kubelet.conf**.

```

rm /etc/kubernetes/pki/ca.crt
rm /etc/kubernetes/kubelet.conf
rm /etc/kubernetes/bootstrap-kubelet.conf

```

23. Grab the name of the node from Step 19.

24. Issue the command for the worker in order to rejoin the cluster. Use the command from 18., but add **--node-name <name of node>** to the end.

```

[root@cx-ccs-prod-worker-d7f34f25-f524-4f90-9037-7286202ed13a1 kubernetes]# kubeadm join
192.168.1.14:6443 --token mlynvj.f4n3et3poki88ry4 --discovery-token-ca-cert-hash
sha256:4d0c569985c1d460ef74dc01c85740285e4af2c2369ff833eed1ba86e1167575 --node-name cx-
ccs-prod-worker-d7f34f25-f524-4f90-9037-7286202ed13a1
[preflight] running pre-flight checks
[WARNING RequiredIPVSKernelModulesAvailable]: the IPVS proxier will not be used,
because the following required kernel modules are not loaded: [ip_vs_rr ip_vs_wrr
ip_vs_sh] or no builtin kernel ipvs support: map[ip_vs:{}] ip_vs_rr:{}] ip_vs_wrr:{}]
ip_vs_sh:{}] nf_conntrack_ipv4:{}]
you can solve this problem with following methods:

```

1. Run 'modprobe -- ' to load missing kernel modules;
2. Provide the missing builtin kernel ipvs support

```
I0226 17:59:52.644282    19170 kernel_validator.go:81] Validating kernel version
I0226 17:59:52.644421    19170 kernel_validator.go:96] Validating kernel config
[discovery] Trying to connect to API Server "192.168.1.14:6443"
[discovery] Created cluster-info discovery client, requesting info from
"https://192.168.1.14:6443"
[discovery] Requesting info from "https://192.168.1.14:6443" again to validate TLS against
the pinned public key
[discovery] Cluster info signature and contents are valid and TLS certificate validates
against pinned roots, will use API Server "192.168.1.14:6443"
[discovery] Successfully established connection with API Server "192.168.1.14:6443"
[kubelet] Downloading configuration for the kubelet from the "kubelet-config-1.11"
ConfigMap in the kube-system namespace
[kubelet] Writing kubelet configuration to file "/var/lib/kubelet/config.yaml"
[kubelet] Writing kubelet environment file with flags to file "/var/lib/kubelet/kubeadm-
flags.env"
[preflight] Activating the kubelet service
[tlsbootstrap] Waiting for the kubelet to perform the TLS Bootstrap...
[patchnode] Uploading the CRI Socket information "/var/run/dockershim.sock" to the Node
API object "cx-ccs-prod-worker-d7f34f25-f524-4f90-9037-7286202ed13a1" as an annotation
```

This node has joined the cluster:

- * Certificate signing request was sent to master and a response was received.
- * The Kubelet was informed of the new secure connection details.

Run 'kubectl get nodes' on the master to see this node join the cluster.

25. Exit the worker and check the status on a master via **kubectl get nodes**. It should be in Ready status.
26. Repeat steps 20. to 25. for each worker.
27. The last **kubectl get nodes** should show that all nodes are in "Ready" status, back online, and joined to the cluster.

```
[root@cx-ccs-prod-master-d7f34f25-f524-4f90-9037-7286202ed13a1 ~]# kubectl get nodes
NAME                                STATUS    ROLES    AGE
VERSION
cx-ccs-prod-master-d7f34f25-f524-4f90-9037-7286202ed13a1  Ready    master   1y
v1.11.6
cx-ccs-prod-master-d7f34f25-f524-4f90-9037-7286202ed13a2  Ready    master   1y
v1.11.6
cx-ccs-prod-master-d7f34f25-f524-4f90-9037-7286202ed13a3  Ready    master   1y
v1.11.6
cx-ccs-prod-worker-d7f34f25-f524-4f90-9037-7286202ed13a1  Ready    <none>   1y
v1.11.6
cx-ccs-prod-worker-d7f34f25-f524-4f90-9037-7286202ed13a2  Ready    <none>   1y
v1.11.6
cx-ccs-prod-worker-d7f34f25-f524-4f90-9037-7286202ed13a3  Ready    <none>   1y
v1.11.6
cx-ccs-prod-worker-d7f34f25-f524-4f90-9037-7286202ed13a4  Ready    <none>   1y
v1.11.6
cx-ccs-prod-worker-d7f34f25-f524-4f90-9037-7286202ed13a5  Ready    <none>   1y
v1.11.6
cx-ccs-prod-worker-d7f34f25-f524-4f90-9037-7286202ed13a6  Ready    <none>   1y
v1.11.6
```