

Open Infrastructure - Task #7122

Setup production etcd cluster in place6

09/09/2019 01:19 AM - Nico Schottelius

Status:	Closed	Start date:	09/09/2019
Priority:	Normal	Due date:	
Assignee:	Nico Schottelius	% Done:	80%
Category:		Estimated time:	0.00 hour
Target version:			
PM Check date:			
Description			
<ul style="list-style-type: none"><li>• 3 nodes</li><li>• Ensure that sufficient permissions are used to secure access to etcd</li><li>• Include hourly backup to place5<ul style="list-style-type: none"><li>◦ Check whether we need to make a dump or can backup the data directory directly</li></ul></li><li>• Probably include letsencrypt (?) for CAs / encryption<ul style="list-style-type: none"><li>◦ Or private CA</li></ul></li></ul>			

History

#1 - 09/19/2019 01:31 PM - Ahmed Bilal

- Status changed from New to Seen

#2 - 10/02/2019 08:39 AM - Ahmed Bilal

- To Dump ETCD <https://www.npmjs.com/package/etcd-dump> (Not Working Correctly)
- To Create Snapshot <https://github.com/etcd-io/etcd/blob/master/Documentation/op-guide/recovery.md>
- To Create a Certificate Authority to issue certificates <https://coreos.com/os/docs/latest/generate-self-signed-certificates.html>
- Common Name as username <https://github.com/etcd-io/etcd/blob/master/Documentation/op-guide/authentication.md#using-tls-common-name>

#3 - 10/10/2019 02:03 PM - Ahmed Bilal

ca-config.json

```
{
  "signing": {
    "default": {
      "expiry": "43800h"
    },
    "profiles": {
      "server": {
        "expiry": "43800h",
        "usages": [
          "signing",
          "key encipherment",
          "server auth",
          "client auth"
        ]
      },
      "client": {
        "expiry": "43800h",
        "usages": [
          "signing",
          "key encipherment",
          "client auth"
        ]
      },
      "peer": {
        "expiry": "43800h",
        "usages": [
          "signing",
          "key encipherment",
          "server auth",
```

```

        "client auth"
      ]
    }
  }
}

```

#### ca-csr.json

```

{
  "CN": "ungleich",
  "key": {
    "algo": "rsa",
    "size": 2048
  },
  "names": [
    {
      "C": "CH",
      "ST": "Glarus"
    }
  ]
}

```

```
cfssl gencert -initca ca-csr.json | cfssljson -bare ca -
```

#### etcd1.json

```

{
  "CN": "etcd1",
  "hosts": [
    "2a0a:e5c0:0:5:0:78ff:fe11:d761",
    "etcd1"
  ],
  "key": {
    "algo": "ecdsa",
    "size": 256
  },
  "names": [
    {
      "C": "CH",
      "ST": "Glarus"
    }
  ]
}

```

```
cfssl gencert -ca=ca.pem -ca-key=ca-key.pem -config=ca-config.json -profile=server etcd1.json | cfssljson -bare etcd1
```

#### etcd2.json

```

{
  "CN": "etcd2",
  "hosts": [
    "2a0a:e5c0:0:5:0:78ff:fe11:d762",
    "etcd2"
  ],
  "key": {
    "algo": "ecdsa",
    "size": 256
  },
  "names": [
    {
      "C": "CH",
      "ST": "Glarus"
    }
  ]
}

```

```
cfssl gencert -ca=ca.pem -ca-key=ca-key.pem -config=ca-config.json -profile=server etcd2.json | cfssljson -bare etcd2
```

#### etcd3.json

```
{
  "CN": "etcd3",
  "hosts": [
    "2a0a:e5c0:0:5:0:78ff:fe11:d763",
    "etcd3"
  ],
  "key": {
    "algo": "ecdsa",
    "size": 256
  },
  "names": [
    {
      "C": "CH",
      "ST": "Glarus"
    }
  ]
}
```

```
cfssl gencert -ca=ca.pem -ca-key=ca-key.pem -config=ca-config.json -profile=server etcd3.json | cfssljson -bare etcd3
```

#### client.json

```
{
  "CN": "client",
  "hosts": [""],
  "key": {
    "algo": "ecdsa",
    "size": 256
  },
  "names": [
    {
      "C": "CH",
      "ST": "Glarus"
    }
  ]
}
```

```
cfssl gencert -ca=ca.pem -ca-key=ca-key.pem -config=ca-config.json -profile=client client.json | cfssljson -bare client
```

#### root.json

```
{
  "CN": "root",
  "hosts": [""],
  "key": {
    "algo": "ecdsa",
    "size": 256
  },
  "names": [
    {
      "C": "CH",
      "ST": "Glarus"
    }
  ]
}
```

```
cfssl gencert -ca=ca.pem -ca-key=ca-key.pem -config=ca-config.json -profile=client root.json | cfssljson -bare root
```

#### developer.json

```
{
  "CN": "developer",
  "hosts": [""],
  "key": {
    "algo": "ecdsa",
    "size": 256
  },
  "names": [
    {
      "C": "CH",
      "ST": "Glarus"
    }
  ]
}
```

```
]
}
```

```
cfssl gencert -ca=ca.pem -ca-key=ca-key.pem -config=ca-config.json -profile=client developer.json | cfssljson -bare developer
```

#### #4 - 10/10/2019 06:07 PM - Ahmed Bilal

##### To Start Fresh

```
rm -rf /var/lib/etcd /root/etcd-data/
```

# Run the following command on first node

```
etcd --name etcd1 --cert-file=/root/cert/etcd1.pem --key-file=/root/cert/etcd1-key.pem \
--peer-client-cert-auth --peer-trusted-ca-file=/root/cert/ca.pem \
--peer-cert-file=/root/cert/etcd1.pem --peer-key-file=/root/cert/etcd1-key.pem \
--client-cert-auth --trusted-ca-file=/root/cert/ca.pem \
--advertise-client-urls=https://[::]:2379 --listen-client-urls=https://[::]:2379 \
--initial-advertise-peer-urls=https://[::]:2380 --listen-peer-urls=https://[::]:2380 \
--initial-cluster etcd1=https://[::]:2380,etcd2=https://[2a0a:e5c0:0:5:0:78ff:fe11:d762]:2380,etcd3=
https://[2a0a:e5c0:0:5:0:78ff:fe11:d763]:2380 \
--initial-cluster-state new --initial-cluster-token etcd-cluster-1 --data-dir etcd-data
```

# Run the following command on second node

```
etcd --name etcd2 --cert-file=/root/cert/etcd2.pem --key-file=/root/cert/etcd2-key.pem \
--peer-client-cert-auth --peer-trusted-ca-file=/root/cert/ca.pem \
--peer-cert-file=/root/cert/etcd2.pem --peer-key-file=/root/cert/etcd2-key.pem \
--client-cert-auth --trusted-ca-file=/root/cert/ca.pem \
--advertise-client-urls=https://[::]:2379 --listen-client-urls=https://[::]:2379 \
--initial-advertise-peer-urls=https://[::]:2380 --listen-peer-urls=https://[::]:2380 \
--initial-cluster etcd1=https://[2a0a:e5c0:0:5:0:78ff:fe11:d761]:2380,etcd2=
https://[2a0a:e5c0:0:5:0:78ff:fe11:d762]:2380,etcd3=https://[2a0a:e5c0:0:5:0:78ff:fe11:d763]:2380 \
--initial-cluster-state new --initial-cluster-token etcd-cluster-1 --data-dir etcd-data
```

# Run the following command on third node

```
etcd --name etcd3 --cert-file=/root/cert/etcd3.pem --key-file=/root/cert/etcd3-key.pem \
--peer-client-cert-auth --peer-trusted-ca-file=/root/cert/ca.pem \
--peer-cert-file=/root/cert/etcd3.pem --peer-key-file=/root/cert/etcd3-key.pem \
--client-cert-auth --trusted-ca-file=/root/cert/ca.pem \
--advertise-client-urls=https://[::]:2379 --listen-client-urls=https://[::]:2379 \
--initial-advertise-peer-urls=https://[2a0a:e5c0:0:5:0:78ff:fe11:d763]:2380 --listen-peer-urls=
https://[2a0a:e5c0:0:5:0:78ff:fe11:d763]:2380 \
--initial-cluster etcd1=https://[2a0a:e5c0:0:5:0:78ff:fe11:d761]:2380,etcd2=
https://[2a0a:e5c0:0:5:0:78ff:fe11:d762]:2380,etcd3=https://[2a0a:e5c0:0:5:0:78ff:fe11:d763]:2380 \
--initial-cluster-state new --initial-cluster-token etcd-cluster-1 --data-dir etcd-data
```

#### #5 - 10/10/2019 07:22 PM - Ahmed Bilal

- Status changed from Seen to In Progress

#### #6 - 10/11/2019 12:47 PM - Ahmed Bilal

### Correct, permissions

```
chown -R etcd:etcd /var/lib/etcd/
```

### Queries to check if things are working correctly

#### Write something

```
ETCDCTL_API=3 etcdctl --endpoints https://[2a0a:e5c0:0:5:0:78ff:fe11:d761]:2379 --cacert ca.pem --cert
root.pem --key root-key.pem put /v1 abc
```

#### Read it

```
ETCDCTL_API=3 etcdctl --endpoints https://[2a0a:e5c0:0:5:0:78ff:fe11:d761]:2379 --cacert ca.pem --cert root.pem --key root-key.pem get /v1
```

## Enable Authentication

- Create root user, grant it root role and Enable Authentication\*\*

```
ETCDCTL_API=3 etcdctl --endpoints https://[2a0a:e5c0:0:5:0:78ff:fe11:d761]:2379 --cacert ca.pem --cert root.pem --key root-key.pem user add root
```

```
ETCDCTL_API=3 etcdctl --endpoints https://[2a0a:e5c0:0:5:0:78ff:fe11:d761]:2379 --cacert ca.pem --cert root.pem --key root-key.pem user grant-role root root
```

```
ETCDCTL_API=3 etcdctl --endpoints https://[2a0a:e5c0:0:5:0:78ff:fe11:d761]:2379 --cacert ca.pem --cert root.pem --key root-key.pem auth enable
```

## Create a non-root User e.g developer

```
ETCDCTL_API=3 etcdctl --endpoints https://[2a0a:e5c0:0:5:0:78ff:fe11:d761]:2379 --cacert ca.pem --cert root.pem --key root-key.pem role add developer
```

```
ETCDCTL_API=3 etcdctl --endpoints https://[2a0a:e5c0:0:5:0:78ff:fe11:d761]:2379 --cacert ca.pem --cert root.pem --key root-key.pem role grant-permission developer --prefix=true readwrite /v1
```

```
ETCDCTL_API=3 etcdctl --endpoints https://[2a0a:e5c0:0:5:0:78ff:fe11:d761]:2379 --cacert ca.pem --cert root.pem --key root-key.pem user add developer
```

```
ETCDCTL_API=3 etcdctl --endpoints https://[2a0a:e5c0:0:5:0:78ff:fe11:d761]:2379 --cacert ca.pem --cert root.pem --key root-key.pem user grant-role developer developer
```

### #7 - 10/16/2019 12:32 PM - Ahmed Bilal

ETCD is having some issues.

Specifically, it is saying **Cluster ID mismatch** I file an issue in ETCD's Github repository. <https://github.com/etcd-io/etcd/issues/11263>

### #8 - 10/18/2019 09:44 AM - Ahmed Bilal

Issue fixed.

### #9 - 10/18/2019 07:26 PM - Ahmed Bilal

```
[meow@meow-pc cert]$ ETCDCTL_API=3 etcdctl --endpoints https://[2a0a:e5c0:2:12:0:f0ff:fea9:c43a]:2379,https://[2a0a:e5c0:2:12:0:f0ff:fea9:c43d]:2379,https://[2a0a:e5c0:2:12:0:f0ff:fea9:c442]:2379 --cacert ~/Desktop/ungleich-issues/7122/cert/ca.pem --cert ~/Desktop/ungleich-issues/7122/cert/client.pem --key ~/Desktop/ungleich-issues/7122/cert/client-key.pem endpoint health https://[2a0a:e5c0:2:12:0:f0ff:fea9:c442]:2379 is healthy: successfully committed proposal: took = 631.531912ms https://[2a0a:e5c0:2:12:0:f0ff:fea9:c43d]:2379 is healthy: successfully committed proposal: took = 633.007889ms https://[2a0a:e5c0:2:12:0:f0ff:fea9:c43a]:2379 is healthy: successfully committed proposal: took = 634.894405ms
```

### #10 - 10/18/2019 08:16 PM - Ahmed Bilal

- % Done changed from 0 to 80

cdist type is ready. Testing underway.

### #11 - 10/19/2019 01:13 PM - Ahmed Bilal

- Deployed at place6
- Authentication enabled. Only clients with valid certificate issued by ungleich's private CA authority can access the etcd.
- Even finer control is employed by setting permissions for individual user to access specific keys or key's prefixes.

Only, backup is remaining.

```
ETCDCTL_API=3 etcdctl --endpoints https://etcd1.ungleich.ch:2379,https://etcd2.ungleich.ch:2379,https://etcd3.ungleich.ch:2379 --cacert ca.pem --cert developer.pem --key developer-key.pem endpoint health
```

```
https://etcd2.ungleich.ch:2379 is healthy: successfully committed proposal: took = 823.064847ms
https://etcd1.ungleich.ch:2379 is healthy: successfully committed proposal: took = 824.459603ms
https://etcd3.ungleich.ch:2379 is healthy: successfully committed proposal: took = 850.761864ms
```

**#12 - 10/20/2019 09:53 PM - Nico Schottelius**

Is it already in cdist?

[redmine@ungleich.ch](mailto:redmine@ungleich.ch) writes:

**#13 - 10/21/2019 10:15 AM - Ahmed Bilal**

[@Nico\\_Schottelius](#) Yes, it is in etcd-cluster branch

**#14 - 10/21/2019 12:04 PM - Ahmed Bilal**

- Assignee changed from Ahmed Bilal to Dominique Roux

Handing it over to rouxdo for review and future maintaining.

**#15 - 10/21/2019 12:08 PM - Ahmed Bilal**

- Status changed from In Progress to Feedback

**#16 - 11/30/2019 05:30 PM - Dominique Roux**

- Status changed from Feedback to Resolved

This is done now

ETCD-Cluster available at:

etcd1.ungleich.ch  
etcd2.ungleich.ch  
etcd3.ungleich.ch

Currently there are some small problems with nftables (not loaded at boot).  
Will contact alpine linux dev team

**#17 - 12/03/2019 10:08 AM - Dominique Roux**

Dominique Roux wrote:

...

Currently there are some small problems with nftables (not loaded at boot).  
Will contact alpine linux dev team

nft problem is fixed now.  
Problem was: Alpine has it's own init.d script (which works ;-)). The cdist type was already updated but the submodule was not.  
The submodule is now updated too, therefore, this should not happen again in future.

**#18 - 01/19/2021 12:48 PM - Nico Schottelius**

- Assignee changed from Dominique Roux to Nico Schottelius

**#19 - 12/31/2023 06:52 PM - Nico Schottelius**

- Status changed from Resolved to Closed