

Proxmox

How do I create a cluster ?

This article describes how to create a cluster with 4 servers.

Table of Contents

- [Explanation](#)
- [1. Preparation](#)
- [2. Create Cluster](#)
 - [2.1 Important](#)
 - [2.2 Switch to unicast](#)
 - [2.3 Can I add a node with existing virtual machines?](#)
 - [2.4 Add an empty node to the cluster](#)

Explanation

In the following, the cluster is created from the following 4 servers:

f002 (AMD Opteron 3365 on this we create the cluster)
ca130 (Intel XEON E3-1270v3 Node)
srv1583 (HPE DL360 Gen10 Node)
srv10008 (HPE DL325 Gen10 Node)

1. Preparation

Unfortunately the hostnames of the servers cannot be changed afterwards, so you should think about what you want to call them.

Below I renamed "f002" to "f002-Cluster":

```
# hostnamectl set-hostname f002-Cluster
```

The hostname in `/etc/hosts` must also be changed:

```
# nano /etc/hosts
```

Before:

Proxmox

```
127.0.0.1      localhost
217.79.181.2  f002.fuchsia.fastwebserver.de  f002

# The following lines are desirable for IPv6 capable hosts
::1          localhost ip6-localhost ip6-loopback
ff02::1      ip6-allnodes
ff02::2      ip6-allrouters
```

After:

```
127.0.0.1      localhost
217.79.181.2  f002.fuchsia.fastwebserver.de  f002-Cluster

# The following lines are desirable for IPv6 capable hosts
::1          localhost ip6-localhost ip6-loopback
ff02::1      ip6-allnodes
ff02::2      ip6-allrouters
```

2. Create Cluster

Now we can create the cluster with "pvecm create <Clustername>", in our example this is called "Testcluster":

```
root@f002:~# pvecm create Testcluster
Corosync Cluster Engine Authentication key generator.
Gathering 1024 bits for key from /dev/urandom.
Writing corosync key to /etc/corosync/authkey.
Writing corosync config to /etc/pve/corosync.conf
Restart corosync and cluster filesystem
root@f002:~# █
```

2.1 Important

It is **not** possible for us to create a cluster with the standard setting with multicast without a own VLAN, for this we have to use unicast.

*** Update ***

As of Debian 10 + Proxmox 6 this setting is already automatically set to Unicast, so only Debian 8 and 9 have to be switched to Unicast.

Update

2.2 Switch to unicast

To change this, add "transport: udpu" to "/etc/pve/corosync.conf" as shown below:

```
# nano /etc/pve/corosync.conf
```

Proxmox

```
logging {
  debug: off
  to_syslog: yes
}

nodelist {
  node {
    name: f002-Cluster
    nodeid: 1
    quorum_votes: 1
    ring0_addr: 217.79.181.2
  }
}

quorum {
  provider: corosync_votequorum
}

totem {
  cluster_name: Testcluster
  config_version: 1
  interface {
    bindnetaddr: 217.79.181.2
    ringnumber: 0
  }
  ip_version: ipv4
  secauth: on
  version: 2
  transport: udpu ←
```

After restarting Corosync and executing the command "pvecm e 1", the nodes can now be added.

```
# service corosync restart
```

```
# pvecm e 1
```

2.3 Can I add a node with existing virtual machines?

No, please make sure that they have already been copied, as in the worst case data loss from the node can occur:

```
root@cal30:~# pvecm add f002.fuchsia.fastwebserver.de
detected the following error(s):
* this host already contains virtual guests
Check if node may join a cluster failed!
root@cal30:~# █
```

2.4 Add an empty node to the cluster

Log on **to a node** via SSH and add it to the cluster with "pvecm add <FQDN or IP>".

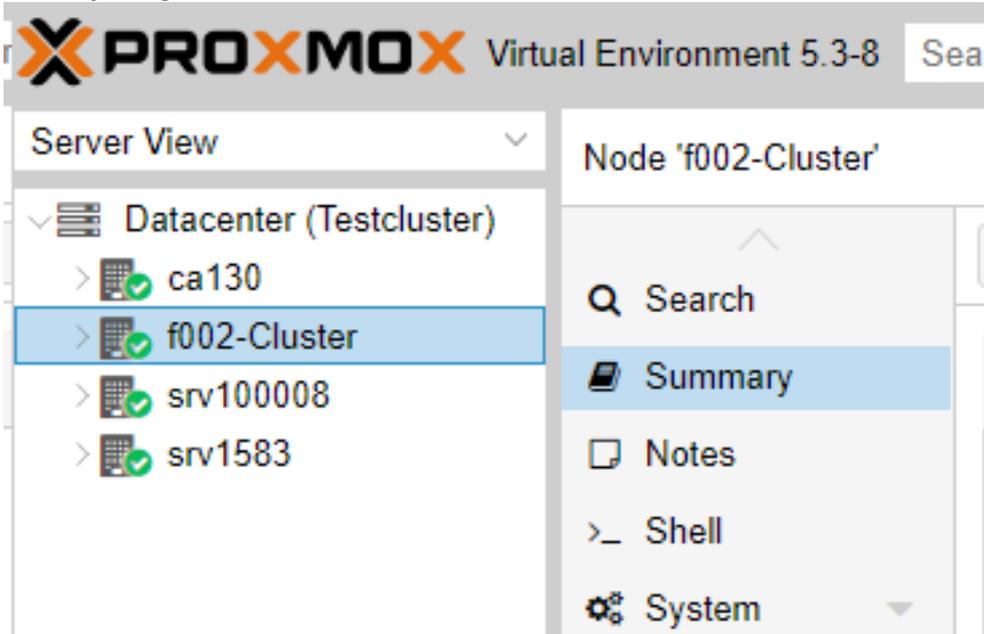
```
# pvecm add f002.fuchsia.fastwebserver.de
```

Proxmox

The output should look like this:

```
root@cal30:~# pvecm add f002.fuchsia.fastwebserver.de
Please enter superuser (root) password for 'f002.fuchsia.fastwebserver.de':
Password for root@f002.fuchsia.fastwebserver.de: *****
Establishing API connection with host 'f002.fuchsia.fastwebserver.de'
The authenticity of host 'f002.fuchsia.fastwebserver.de' can't be established.
X509 SHA256 key fingerprint is 98:24:F5:99:F3:D6:D1:6C:C8:17:54:A5:63:1E:0B:4A:8E:FC:CF:D4:C3:C6:37:68:75:E0:A5:FE:87:F2:96:E4.
Are you sure you want to continue connecting (yes/no)? yes
Login succeeded.
Request addition of this node
Join request OK, finishing setup locally
stopping pve-cluster service
backup old database to '/var/lib/pve-cluster/backup/config-1549467377.sql.gz'
waiting for quorum...OK
(re)generate node files
generate new node certificate
merge authorized SSH keys and known hosts
generated new node certificate, restart pveproxy and pvedaemon services
successfully added node 'cal30' to cluster.
root@cal30:~#
```

If everything worked, it should look like this inside the Proxmox interface:



Unique solution ID: #1478

Author: Bettina Brauer

Last update: 2021-05-21 03:34