

Proxmox - Cheatsheet

- [Proxmox Cheatsheet](#)
- [Troubleshooting](#)
 - [warning: cannot send disk \(Broken pipe\) \(I/O error\) \(already exists\)](#)

Proxmox Cheatsheet

Put Node Into Maintenance

```
ha-manager crm-command node-maintenance enable MUT7PVE201
```

Remove Node from Maintenance

```
ha-manager crm-command node-maintenance disable MUT7PVE201
```

Remove Old Nodes From HA

1. Delete the `old node`.

```
pvecm delnode MILHOSTPROX504
```

2. Navigate to the `old node` files on the `current primary` node.

```
cd /etc/pve/nodes
```

3. Remove the `old node` by name.

```
rm -rf MILHOSTPROX504
```

4. Run the following command on `all active` nodes.

```
systemctl stop pve-ha-crm.service
```

5. Run the command on the current `active` node.

```
rm -f /etc/pve/ha/manager_status
```

6. Restart `all the nodes`, starting with the node you just removed the above file from.

```
systemctl start pve-ha-crm.service
```

Wipe SSH Keys

1. Run the following command on the server that is failing to connect to the target node. Replace the below hostname with the IP address and run a second time to ensure all entries are wiped.

```
ssh-keygen -f "/etc/ssh/ssh_known_hosts" -R "MUT7PVE201"
```

2. Run the following command to test the new connection and add the new keys to the keystore. Replace the following field with the target hostname and the field below that with the relevant IP address.

```
HostKeyAlias=MUT7PVE201
```

```
root@192.168.1.75
```

```
/usr/bin/ssh -e none -o 'HostKeyAlias=MUT7PVE201' root@192.168.1.75 /bin/true
```

zfs error: cannot open /rpool/data/vm-110-disk-0: dataset does not exist

1. Verify that the disk does not exist on the node that it's sitting on.

```
zfs list | grep 110 && echo exist || echo not exist
```

2. Check other nodes using the command above to see where the disk exists.
3. Move the VM configuration file to where the disk actually resides by using the following command on the node that's currently trying to turn on the VM.

```
mv /etc/pve/qemu-server/100.conf /var/lib/vz/template/iso/template/iso/
```

4. Run the following command on the node that actually hosts the disk files.

```
mv /var/lib/vz/template/iso/template/iso/100.conf /etc/pve/qemu-server
```

5. Reenable the VM in the cluster so it can start normally using the method in the next section.
6. Once confirmed, delete the .conf file on the node where it was having problems.

VM In Error State

1. Fix the underlying issue that caused it to be in an error state in the first place, then run the following command.

```
ha-manager set vm:100 --state disabled
```

"trying to aquire lock...TASK ERROR:
can't lock file '/var/lock/qemu-
server/lock-109.conf' - got timeout"

There are several ways to fix this error. First, restarting the node. Second, restarting pve-cluster services. Third, and probably the safest, following the below steps to see whats causing the lock and potentially fixing it.

1. Run the following command to check what's locking the file.

```
lsof /var/lock/qemu-server/lock-100.conf
```

2. If necessary, run the following command to further identify what's causing the lock.

```
ps aux | grep PID
```

3. You can kill this process, or as stated above, restart services to force it to stop.

ISOs Not Showing In ISOs Share

1. Run the following command to identify shares that can be accessed via NFS from the node/s you want this share attached to.

```
pvesm scan nfs 192.168.1.1 #(NAS address)
```

2. Run the following command to add the storage to the cluster paying attention to the content types.

```
pvesm add nfs MUT7NAS-NFS-ISOs --server 192.168.20.20 --path /var/lib/vz/template/iso/ --export/volume1/MUT7PVE-ISOs --content images,iso
```

3. Verify storage was added properly in the gui and using this command.

```
cat /etc/pve/storage.cfg
```

Change HA Master Node

```
pvecm add MUT7PVE201
```

Change ZFS Pool Name While In Use

1. Deactivate the pool through the web GUI
2. Run the following command to export the pool

```
zpool export SSD_POOL_PVE2
```

3. Run the following command to import the pool under the new name.

```
zpool import SSD_POOL_PVE2 SSD_POOL_PVE1
```

local node address: cannot use IP
'10.0.40.50', not found on local node!

Check the `/etc/hosts` entry for that server and verify all entries are correct. Modify or delete any that are incorrect or stale.

Delete Replication Job

```
pvesr delete '106-4' --force
```

[Original Article](#)

Cluster Not Ready - No Quorum When Removing Nodes

```
pvecm expected 1
```

[Original Article](#)

Remove Node From Cluster

```
pvecm delnode MUT7PVE205
```

Mirror System Disk

```
sgdisk /dev/sda -R /dev/sdb  
sgdisk --randomize-guids /dev/sdb  
pve-efiboot-tool format /dev/sdb2 --force  
pve-efiboot-tool init /dev/sdb2  
zpool attach rpool /dev/disk/by-id/ata-SSDSC2KG960G8R_BTYG9284085N960CGN-part3 /dev/disk/by-id/ata-SSDSC2KG960G8R_BTYG2055022C960CGN-part3
```

(Note: checked "zpool status -v" which showed rpool as -part3 aka /dev/sda3)

[Original Article](#)

Troubleshooting

warning: cannot send disk (Broken pipe) (I/O error) (already exists)

If you see an error when trying to replicate such as these.

```
2023-11-14 14:38:08 100-0: warning: cannot send 'NVMe-R10-6D/vm-100-disk-0@__replicate_100-0_1699990683__': Broken pipe
2023-11-14 14:38:08 100-0: cannot send 'NVMe-R10-6D/vm-100-disk-0': I/O error
2023-11-14 14:38:08 100-0: command 'zfs send -Rpv -- NVMe-R10-6D/vm-100-disk-0@__replicate_100-0_1699990683__' failed: exit code 1
2023-11-14 14:38:08 100-0: [mut7pve202] volume 'NVMe-R10-6D/vm-100-disk-0' already exists
```

Verify the disk it's trying to replicate in the error and delete the snapshot replacing the disk number with the one in the error.

```
zfs destroy -r rpool/data/vm-100-disk-0
```