

@ VRTX M-630 Blades 2-4

SWITCH CONFIG NEEDED FOR THE BONDING:

@Switch – Configure Ports and Port-Group for LACP

- Locate which ports will be aggregated. At Homelab:

Blade 2 – te1/0/2 & te1/0/22

Blade 3 – te1/0/3 & te1/0/23

Blade 4 – te1/0/4 & te1/0/24

```
int [interface #1 name]
channel-group ## mode active
switchport mode trunk
description Blade # Member of Port Group ##
exit
```

```
int [interface #2 name]
channel-group ## mode active
switchport mode trunk
description Blade # Member of Port Group ##
exit
```

Blade 2 – Port Group 20

Blade 3 – Port Group 30

Blade 4 – Port Group 40

To confirm: show etherchannel summary
 Show interfaces trunk
 Show etherchannel load-balance

1. Virtual Drives:

M630-0#-01	500GB	RAID 0
M630-0#-02	1TB	RAID 0

2. Boot from ProXmoX.iso

Highlight “Install Proxmox VE (Graphical) and press E. Add “proxdebug” to the end of the “linux /boot/vmlinuz...” line.

```
GNU GRUB version 2.06-13+pmx2

setparams 'Install Proxmox VE (Graphical)'

echo      'Loading Proxmox VE Installer ...'
linux     /boot/linux26 ro ramdisk_size=16777216 rw quiet splash=silent prox\
debug_
echo      'Loading initial ramdisk ...'
initrd    /boot/initrd.img
```

Press F10 to continue.

3. Please verify the installation Target: 238GB SD Card Integrated Dell

Target Harddisk

Harddisk options

Filesystem

hdspace GB

swapspace GB

maxroot GB

minfree GB

maxvz GB

4. When presented with Network stuff... going to try assigning a “temp” mgmt. IP so the install can continue. Once rebooted:

vi /etc/network/interfaces

```
1 auto lo
2 iface lo inet loopback
3
4 # Bonded Interface (Trunk Port)
5 auto nm-bond
6 iface nm-bond inet manual
7     bond-slaves eno1 eno2
8     bond-mode 802.3ad
9     bond-miimon 100
10    bond-lacp-rate 1
11    bond-xmit-hash-policy layer3+4
12
13 # Bridge for Vlan 69 (MGMT))
14 auto vmbr69
15 iface vmbr69 inet static
16     address 10.10.69.250/24
17     gateway 10.10.69.254
18     dns-nameservers 10.10.69.193 1.1.1.1
19     bridge-ports nm-bond.69
20     bridge-stp off
21     bridge-fd 0
22     post-up ip addr 10.10.69.250/24 dev vmbr69
23
24 # Bridge for Vlan 110 (Ceph Backend)
25 auto vmbr110
26 iface vmbr110 inet static
27     address 192.168.65.250/24
28     bridge-ports nm-bond.110
29     bridge-stp off
30     bridge-fd 0
31
32
33
34
35
36 ##### BRIDGES FOR VM NETWORKS #####
37
38 # Bridge for Vlan 109 (VPN 225)
39 auto vmbr109
```

```

40 iface vmbr109 inet manual
41   bridge-ports nm-bond.109
42   bridge-stp off
43   bridge-fd 0
44
45
46
47
48 source /etc/network/interfaces.d/*

```

(leave the "source....." line at the bottom)

Systemctl restart networking.

M630-02	10.10.69.250/192.168.65.250
M630-03	10.10.69.251/192.168.65.251
M630-04	10.10.69.252/192.168.65.252

This will setup the network as I had tried for oVirt, OLVM, vCenter... etc.

5. Now at the Vlan 69 MGMT IP:8006 Webpage, ensure DNS and Hostname are setup correctly.

Edit	
Search domain	HOMELAB.CRAIN
DNS server 1	10.10.69.193
DNS server 2	1.1.1.1

```

127.0.0.1 localhost.localdomain localhost
10.10.69.250 M630-02.HOMELAB.CRAIN
192.168.65.250 M630-02.HOMELAB.CRAIN

# The following lines are desirable for IPv6 capable hosts

::1        ip6-localhost ip6-loopback
fe00::0    ip6-localnet
ff00::0    ip6-mcastprefix
ff02::1    ip6-allnodes
ff02::2    ip6-allrouters
ff02::3    ip6-allhosts

```

6. Remove the Enterprise Repo, fix Ceph Repo for free Version, Update:

```
mv /etc/apt/sources.list.d/pve-enterprise.list /etc/apt/sources.list.d/pve-enterprise.list.bak 2>/dev/null; \
mv /etc/apt/sources.list.d/ceph.list /etc/apt/sources.list.d/ceph.list.bak 2>/dev/null; \
echo "deb http://download.proxmox.com/debian/pve bookworm pve-no-subscription" >
/etc/apt/sources.list.d/pve-no-subscription.list; \
echo "deb http://download.proxmox.com/debian/ceph-quincy bookworm no-subscription" >
/etc/apt/sources.list.d/ceph.list; \
apt update && apt dist-upgrade
```

7. Install Syntax Highlighting:

```
apt install -y vim
```

Once Vim is installed, edit your **Vim config file**:

```
echo "syntax on" >> ~/.vimrc
```

Make a vi to vim Alias:

```
echo "alias vi='vim'" >> ~/.bashrc
source ~/.bashrc
```

Make it System-Wide:

```
echo "alias vi='vim'" >> /etc/profile
source /etc/profile
```

Reboot each hosts.

***** The following steps enable an auto reboot for the pve-firewall and pvestatd service that prevent the system status from being read *****

systemctl edit --full pvestatd

Find the [service] section and add or modify:

```
Restart=always
RestartSec=5s
StartLimitIntervalSec=0
```

systemctl edit --full pve-firewall

Find the [service] section and add or modify:

```
Restart=always
RestartSec=5s
StartLimitIntervalSec=0
```

Reload the Systemd and Retart the Services

```
systemctl daemon-reexec  
systemctl restart pvestatd  
systemctl restart pve-firewall
```

Enable the Services to start at boot

```
systemctl enable pvestatd  
systemctl enable pve-firewall
```

8. At Primary Hosts, from GUI, Create the Cluster (Not the Storage Cluster)

What is the Point of a Proxmox Cluster?

A Proxmox Cluster (PVE Cluster) allows multiple Proxmox nodes to work together as a single managed environment, enabling key enterprise-level features like:

-
- ✓ 1. Centralized Management (vCenter-like)
 - Instead of managing each Proxmox node separately, you get one unified web interface to manage all hosts.
 - All nodes share the same configuration, VM templates, storage, and networking.
-
- ✓ 2. Live Migration (vMotion Equivalent)
 - Move running VMs from one node to another without downtime.
 - Useful for load balancing and performing maintenance on hosts without affecting running VMs.
-
- ✓ 3. High Availability (HA)
 - If a Proxmox node crashes or reboots, VMs will automatically restart on another available node.
 - Ensures minimal downtime in case of hardware failure.
-
- ✓ 4. Shared Storage Support (Ceph Integration)
 - A cluster allows Ceph to work across multiple nodes, creating a highly available and redundant storage pool.
 - Without a cluster, each node would need its own separate storage, making Ceph useless.
-
- ✓ 5. Proxmox Backup Server (PBS) Integration
 - Cluster-wide backup management from a central interface.
 - Allows backup deduplication across multiple nodes.
-

Why Do You Need a Cluster for Ceph?

- Ceph requires multiple nodes to provide redundancy, fault tolerance, and distributed storage.
 - A Proxmox Cluster allows Ceph to use all nodes as part of a single storage system instead of managing disks separately per server.
-

```
chmod 755 /etc/corosync
```

Create Cluster

Cluster Name:

M630 Cluster

Cluster Network:

Link:

0

10.10.69.250

Link:

1

192.168.65.250

Add

Multiple links are used as failover, lower numbers have higher priority.

?

Help

Create

From the remaining hosts, “Join Cluster”. You will first need to select “Join Information” from the primary and copy to clipboard.

Cluster Join

☒ Assisted join: Paste encoded cluster join information and enter password.

Information:

Z19hZGRylpbJlEwLjEwLjY5LjI1MCIsIjE5Mi4xNjguNjUuMjUwIl0sInRvdGVtIjp7InNIY2F1dGgiOiJvbilslmNvbmZpZ192ZXJzaW9uIjoiaXBiZW50IiwidmVyc2lvbiI6IjllLCJpbmRicmZhY2UiOnsiMC16eyJsaW5rbnVtYmVyljoIMCJ9LCkxbj7ImxpbtudW1iZXliOixln19fX0=

Peer Address:

10.10.69.250

Password:

.....

Fingerprint:

09:D9:D3:CE:06:0B:87:BE:AB:13:4B:13:1F:33:39:01:58:69:9D:22:D5:EF:CC:16:8F:21:A8:AD:21:3D:1A:FF

Cluster Network:

Link: 0

10.10.69.251

peer's link address: 10.10.69.250

Link: 1

192.168.65.251

peer's link address: 192.168.65.250

?

Help

Join 'm630-cluster'

Refresh all webpages and the Datacenter should now reflect the cluster. Reboot all hosts if possible.

9. Install & Configure Ceph

Check these services first before moving on:


systemctl status pvestatd pve-firewall

pvecm status


systemctl status corosync

pvesm status

From Datacenter/Cluster Tab > Ceph > Install Ceph



Select the Installation Tab, and type "Y"



Selecting Next will take you to Configuration:

The Setup Configuration window shows the following fields and values:

- Ceph cluster configuration:**
 - Public Network IP/CIDR: 10.10.69.250/24
 - Cluster Network IP/CIDR: 192.168.65.250/24
- First Ceph monitor:**
 - Monitor node: M630-02
- Number of replicas:** 3
- Minimum replicas:** 2

A yellow callout box states: "Additional monitors are recommended. They can be created at any time in the Monitor tab."

Repeat this process on the remaining Servers...

"Create" Monitor and Manager to add the other 2 hosts.

The Monitor tab shows the following data:

Name ↑	Host	Status	Address
mon.M630-02	M630-02	running	10.10.69.250:6789/0
mon.M630-03	M630-03	running	10.10.69.251:6789/0
mon.M630-04	M630-04	running	10.10.69.252:6789/0

The Manager tab shows the following data:

Name ↑	Host	Status	Address
mgr.M630-02	M630-02	active	10.10.69.250
mgr.M630-03	M630-03	standby	10.10.69.251
mgr.M630-04	M630-04	standby	10.10.69.252

Select "OSD" > Create OSD:

Do this for each host for each drive.

default

M630-02

osd.0

osd.1

M630-03

osd.2

osd.3

M630-04

osd.4

osd.5

osd.0	hdd	bluestore	up / in	17.2.7	0.4543	1.00	0.06	465.25 GiB	0 / 0	1
osd.1	hdd	bluestore	up / in	17.2.7	0.90919	1.00	0.03	931.00 GiB	0 / 0	0
osd.2	hdd	bluestore	up / in	17.2.7	0.4543	1.00	0.06	465.25 GiB	0 / 0	0
osd.3	hdd	bluestore	up / in	17.2.7	0.90919	1.00	0.03	931.00 GiB	0 / 0	1
osd.4	hdd	bluestore	up / in	17.2.7	0.4543	1.00	0.06	465.25 GiB	0 / 0	0
osd.5	hdd	bluestore	up / in	17.2.7	0.90919	1.00	0.03	931.00 GiB	0 / 0	1


Create Storage Pools:

Virtual_Machines

VM_Templates

Efforts were made for making this CephFS for Backups were shit, and no joy. Next step is setting up the M630-01 / Server 2022 > NFS Server for backups and Log Server (NFS was the way to go, backups now mounted as NFS and sending daily)

Datacenter > Backup :



Edit: NFS

General | Backup Retention

ID: NFS-Backups Nodes: All (No restrictions) ▾

Server: 10.10.69.123 Enable: ☒

Export: /NFS-Backups

Content: VZDump backup file ▾

Preallocation: Default ▾ NFS Version: Default ▾

 ☒ Advanced

Edit Backup Job

General Retention Note Template Advanced

Node: -- All --

Storage: NFS-Backups

Schedule: 01:00

Selection mode: All

Notification mode: Default (Auto)

Send email to: robert.crain.1775@gmail.com

Send email: Always

Compression: ZSTD (fast and good)

Mode: Snapshot

Enable: ☒

Job Comment:

ID	Node	Status	Name	Type

Extras:

Datcenter > Options > Cluster Resource Scheduling: Basic & Use CRS

10. Migrate Virtual Machines from vCenter/ESXi > Proxmox Ceph Cluster:

Live Migration Attempt:

Preparation

- Before the migration, remove any guest tools specific to the old hypervisor, as it might be difficult to remove them after.
- Note down the guest network configuration, so you can manually restore it.
 - On Windows, consider removing the static network configuration, if there is any. After the migration, the network adapter will change and Windows will show a warning if you configure the same IP address on another network adapter, even if the previous one is not present anymore.
- In case of DHCP reservations, either adapt the reservation to the new MAC address of the target VM NIC, or set the MAC address on the target VM NIC manually.
- If full-disk encryption is used in the VM and the keys are stored in a virtual TPM device, consider disabling it. It is currently not possible to migrate the vTPM state to Proxmox VE from VMware. Make sure to have the manual keys to decrypt the VM available, just in case.
- Power down the source VM.

Added R630-01 as storage and was able to import no problem.

*** Update ***

Hours later Ceph seems to have crashed and in unrecoverable? The hell is going on?

PROXMOX Virtual Environment 8.3.4

Server View

Datacenter

M630-cluster

- M630-02
 - 100 (Homelab-to-Reachback-Jumpb)
 - localnetwork (M630-02)
 - NFS-Backups (M630-02)
 - R630-01 (M630-02)
 - VM_Templates (M630-02)
 - Virtual_Machines (M630-02)
 - local (M630-02)
 - local-lvm (M630-02)
- M630-03
 - 101 (Homelab-to-Reachback-Jumpb)
 - 104 (Windows-11-Hardened-Baselin)
 - localnetwork (M630-03)
 - NFS-Backups (M630-03)
 - R630-01 (M630-03)
 - VM_Templates (M630-03)
 - Virtual_Machines (M630-03)
 - local (M630-03)
 - local-lvm (M630-03)
- M630-04
 - 102 (C2HUB-200-DC-01v)
 - 103 (C2HUB-200-FS-01v)
 - 105 (Crain-Homelab-Domain-Contro)
 - localnetwork (M630-04)
 - NFS-Backups (M630-04)
 - R630-01 (M630-04)
 - VM_Templates (M630-04)
 - Virtual_Machines (M630-04)

Datacenter

- Search
- Summary
- Notes
- Cluster
- Ceph
- Options
- Storage
- Backup
- Replication
- Permissions
 - Users
 - API Tokens
 - Two Factor
 - Groups
 - Pools
 - Roles
 - Realms
- HA
- SDN
- Zones

Health

Status

Summary

No Warnings/Errors

Ceph Version:

Performance

Reads:

Status

OSDs

	In	Out
Up	0	0
Down	0	0

Total: 0

Monitors

got timeout (500)

Datacenter (m630-cluster)

- M630-02
 - 100 (Homelab-to-Reachback-Jumpb)
 - localnetwork (M630-02)
 - NFS-Backups (M630-02)
 - R630-01 (M630-02)
 - VM_Templates (M630-02)
 - Virtual_Machines (M630-02)
 - local (M630-02)
 - local-lvm (M630-02)
- M630-03
 - 101 (Homelab-to-Reachback-Jumpb)
 - 104 (Windows-11-Hardened-Baselin)
 - localnetwork (M630-03)
 - NFS-Backups (M630-03)
 - R630-01 (M630-03)
 - VM_Templates (M630-03)
 - Virtual_Machines (M630-03)
 - local (M630-03)
 - local-lvm (M630-03)
- M630-04
 - 102 (C2HUB-200-DC-01v)
 - 103 (C2HUB-200-FS-01v)
 - 105 (Crain-Homelab-Domain-Contro)
 - localnetwork (M630-04)
 - NFS-Backups (M630-04)
 - R630-01 (M630-04)
 - VM_Templates (M630-04)
 - Virtual_Machines (M630-04)

Datacenter

- Time
- System Log
- Updates
 - Repositories
- Firewall
- Disks
 - LVM
 - LVM-Thin
 - Directory
 - ZFS
- Ceph
 - Configuration
 - Monitor
 - OSD
 - CephFS
 - Pools
 - Log
- Replication
- Task History
- Subscription

Health

Status

Summary

No Warnings/Errors

Ceph Version:

Performance

Reads:

Status

OSDs

	In	Out
Up	0	0
Down	0	0

Total: 0

Services

Monitors

Server View

▼ Datacenter (m630-cluster)

▼ M630-02

100 (Homelab-to-Reachback-Jumpb

localnetwork (M630-02)

NFS-Backups (M630-02)

R630-01 (M630-02)

VM_Templates (M630-02)

Virtual_Machines (M630-02)

local (M630-02)

local-lvm (M630-02)

▼ M630-03

101 (Homelab-to-Reachback-Jumpb

104 (Windows-11-Hardened-Baselin

localnetwork (M630-03)

NFS-Backups (M630-03)

R630-01 (M630-03)

VM_Templates (M630-03)

Virtual_Machines (M630-03)

local (M630-03)

local-lvm (M630-03)

▼ M630-04

102 (C2HUB-200-DC-01v)

103 (C2HUB-200-FS-01v)

105 (Crain-Homelab-Domain-Contro

localnetwork (M630-04)

NFS-Backups (M630-04)

R630-01 (M630-04)

VM_Templates (M630-04)

Virtual_Machines (M630-04)

Node 'M630-02'

CreateEditDestroy

Pool #	Name
--------	------

▼ Open

Time

System Log

Updates

Repositories

Firewall

Disks

LVM

LVM-Thin

Directory

ZFS

Ceph

Configuration

Monitor

OSD

CephFS

Pools

Log

Replication

Task History

Subscription