Table of Contents

BIOS Tweaking

I/OAT DMA Engine	?	https://en.wikipedia.org/wiki/I/O_Acceleration_Technology
SRV/IOV Global Enable	Enabled	https://en.wikipedia.org/wiki/Single-root_input/output_virtualization
X2Apic	Enabled	https://en.wikipedia.org/wiki/Advanced_Programmable_Interrupt_Controller

IBM lab tests show that enabling the x2APIC support for Red Hat Enterprise Linux 6 guests can result in 2% to 5% throughput improvement for many I/O workloads.

NUMA Info

```
numactl --hardware
```

When possible, create smaller VMs, instead of “Monster” VMs, that fit into a single NUMA node.

Disk performance (not for SSD)

echo 100000 > /sys/block/sdX/queue/nr_requests

List CEPH Snapshot

```
rbd --pool ceph-pool snap ls vm-100-disk-1
```

```
rbd snap rm ceph-pool/vm-100-disk-1@vzdump
```

Windows Preparation

```
sysprep.exe [/oobe | /audit] [/generalize] [/reboot | /shutdown | /quit]
[/quiet] [/unattend:answerfile]
```

Import HyperV

```
qemu-img convert -O qcow2 disk.vhdx /var/lib/vz/images/112/vm-112-
disk-1.qcow2
```

BBR

/etc/sysctl.conf

```
net.ipv4.tcp_congestion_control=bbr
```

Jumbo Frames

/etc/network/interface

```
pre-up ip link set <interface name> mtu 9000
```

```
post-up ip link set eth2 mtu 9000 && ip link set eth3 mtu 9000 && ip link
set bond1 mtu 9000
```

Instantly ip link set <interface name> mtu 9000

HA Logic

shutdown: stops VMs, then move them to other nodes
reboot: stops VMs, put the into freeze state

Migrate Hyper-V to KVM

```
qemu-img check -r all disk.vhdx
qemu-img convert -O qcow2 disk.vhdx output.qcow2
```

OpenVZ / PROXMOX Tips

Proxmox “no subscription” patch
/usr/share/pve-manager/ext6/pvemanagerlib.js Modify

```
-if (data.status !== 'Active') {
+if (false) {
```

System limits Hints for Redhat Linux

Set ULIMIT

When you are using large JAVA applications, Databases
in Proxmox Hypervisor edit /etc/security/limits.conf

```
*    soft nofile 65535
*    hard nofile 65535
root soft nofile 65535
root hard nofile 65535
*    soft stack unlimited
*    hard stack unlimited
root soft stack unlimited
root hard stack unlimited
*    soft memlock unlimited
*    hard memlock unlimited
root soft memlock unlimited
root hard memlock unlimited
```

Cluster

pvecm create YOUR-CLUSTER-NAME
pvecm status
pvecm add IP-ADDRESS-CLUSTER (from new node)
pvecm nodes
pvecm delnode NAME

ext3 vs ext4

ext3 have better performance, ext4 much faster fsck.

How to edit Proxmox config without quorum

```
/etc/default/pve-cluster so that DEAMON_OPTS="-l"
```

(set local mode) and reboot
or try.

```
systemctl stop pve-cluster
pmxcfs -l
```

Change node IP address

```
/etc/network/interfaces
/etc/hosts
/etc/pve/cluster.conf
```

Update

edit **nano /etc/apt/sources.list** and add public repo

```
# PVE pve-no-subscription repository provided by proxmox.com, NOT
recommended for production use
deb http://download.proxmox.com/debian stretch pve-no-subscription
```

edit **nano /etc/apt/sources.list.d/pve-enterprise.list**
disable enterprise by hash #

apt-get update && apt-get dist-upgrade

Graphs Problems

Clean directory /var/lib/rrdcached/* then restart.

Network Bond

1) Remove vmbr0
2) Create bond0
3) Assign eth0 eth1 eth2 etc + Balance-RR (Switch support 2Gbit / up-down) or Balance-TLB (any switch 2Gb up / 1Gb down)
4 Create vmbr0 set IP and Bridge ports bond0

cat /proc/net/bonding/bond0

Root minimum size

Approx. 8GB

Speedup migration

datacenter.cfg

```
migration: network=172.24.16.1/24,type=insecure
```

KSM Sharing default at 50% used

/etc/ksmtuned.conf

KVM IO

```
Cache -> Writeback
Discard on -> LVM Thin
```

Disable the USB tablet device

Do not use balloon driver

IO Scheduler in KVM

Set noop or deadline

Mount ext4 in KVM

```
/dev/X /c ext4
defaults,noatime,nodiratime,nobh,commit=40,barrier=0,data=writeback 0 2
```

SSD / LVM-Thin

Edit harddisk setting and check “discard”

Check IOMMU Support

```
#!/bin/sh
if [ $(dmesg | grep ecap | wc -l) -eq 0 ]; then
  echo "No interrupt remapping support found"
  exit 1
fi

for i in $(dmesg | grep ecap | awk '{print $NF}'); do
  if [ $(( (0x$i & 0xf) >> 3 )) -ne 1 ]; then
    echo "Interrupt remapping not supported"
    exit 1
  fi
done
```

Single node

You can disable

```
systemctl disable zed
systemctl disable zfs
systemctl disable zfs-import-scan
systemctl disable zfs-mount
systemctl disable zfs-share
systemctl disable zfs-zed

pve-ha-crm
pve-la-hrm
pve-firewall
pvefw-logger
```

Convert QCOW2 to RAW

```
qemu-img convert -f qcow2 -O raw image.qcow2 image.raw
```