Guide d'administration

CAPS-OS v4/v5 Failover

Auteur : Franck Martinaux, franck@capsule-sa.co.za

Date: 23-Juin 2020


capsule
TECHNOLOGIES

capsule TECHNOLOGIES

Table of Contents

Principles

Mirroring of the OS HD image of the VMs to the inactive node.

Use NFSv4 for the mirroring of VM images.

Automatic Failover and Failback.

Pre-requisites on Network configuration

The two nodes must be on the same network.

The network must be have a full end to end link speed of 1Gb/s full duplex.

The latency between (measured with a standard ping) must not be higher than 5ms.

No firewall between the two servers, as the standard NFS port will be used and the TCP port 45966 is used for communication between capsvmfod services on the two servers.

Configuration files

/opt/CAPS/etc/capsvm.d/common.conf

cdrom_dir="/opt/CAPS/VMs/ISOs"
os_dir="/opt/CAPS/VMs/HDs"
templates_dir="/opt/CAPS/VMs/Templates/"
fo_dir="/opt/CAPS/VMs/FO"
fo_peer="192.168.32.64"
fo_role="0"
# fo_role 0="secondary" 1="primary"
smbios_manufacturer="Capsule Technologies (Pty) Ltd."
smbios_product="Capsule Server"

The settings to be changed are fo_peer and fo_role

Fo_peer is the IP address of the CAPSOS layer of the other server of the cluster.
Fo_role specifies the role of the current server: 1 for the primary server and 0 for secondary server

/opt/CAPS/etc/capsvm.d/enabled

This folder contains the list of the config files for various VMs

VMs can be or not part of the failover configuration.

For a VM to be part of the failover configuration, the config file of the VM must contains the following settings:

vm_failover_action="failover"

For example in a config file:

# cat vm-1
vm_profile="server"
vm_label="2019S_RECORD"

vm_memory="8192"
vm_nb_cpus="8"
vm_numa_list="0,1,2,3,4,5,6,7"
vm_backing_hdd="WIN_S2019_XPROTECT_2020R2_RECORDER.qcow2"
vm_os_hdd="WIN_S2019_XPROTECT_2020R2_RECORDER.qcow2-1"
vm_os_serial="CAPS001"
vm_smbios_serial="6749d53a"
vm_smbios_uuid="6749d53a-c48e-2900-8e21-0d5c83896aeb"
vm_autostart="no"
vm_failover_action="failover"
vm_extra_param="-device virtio-net-pci,netdev=net0,mac=52:54:00:01:02:01 -netdev type=tap,id=net0,vhost=on,script=no,downscript=no,ifname=tap0"

This must be done on both servers, config files can be different on each server (eg. CPU or Memory sizing), however vm_smbios_uuid, vm_backing_hdd and vm_os_hdd must remain identical.

vm_autostart must remain at “no”

Services and their roles

Both of the servers will run two daemons, capsvmfod and capsvmctld which have distinctive roles.

Capsvmfod

Capsvmfod is a daemon managing the NFS mount point monitoring the heart-beat between the servers, the state of the VMs and the state of the other node of the cluster.

Capsvmfod decides where and when to start or stop VMs by sending commands locally or remotely to capsvmctld daemon.

Capsvmctld

Capsvmctld is a daemon mostly calling the capsvmctl command locally. It allow the remote host to run capsvmctl commands via a network API.

Command

Capsvmfailoverctl allows to control the capsvmfod and capsvmctld daemons.

capsvmfailoverctl - Version 0.3
Usage: capsvmfailoverctl [command]

Commands:

--start	Start capsvmctld and capsvmfod daemons
--stop	Stop capsvmctld and capsvmfod daemons
--restart	Stop/Start capsvmctld and capsvmfod daemons
--status	Display current status of capsvmctld and capsvmfod
--version	Display version of this program

Log files

Capsvmfod and capsvmctld daemons are logging information into the /var/log/messages files, which is the standard syslog for centos and rhel.