Adding more hosts and shared storage

A *flexVDI* platform can easily scale up by adding new *Hosts*. The *flexVDI Manager* will balance the resources allocated to *Pools* among all the available *Hosts*, and run *Guests* in any of them as needed. *flexVDI* uses the OCFS2 clustered filesystem to grant exclusive access to *Guests*'s images by any *Host*, and to provide cluster consistency in case of communication failures. This chapter explains how to configure and use a cluster of flexVDI *Hosts*.

Adding a new Host to the cluster

Adding a new Host to a *flexVDI* cluster is as easy as following these two simple steps:

1. Install the *flexVDI* distribution in the new *Host*. Follow the Getting Started guide until you perform the Basic host configuration (**do not install a new Manager**, there is only one Manager instance per cluster). Be careful to create the same virtual bridges as in the other *Hosts*, so that they map to the same subnets.
2. Register the *Host* with your *flexVDI Manager* instance. Remember that this is done with *flexVDI Config*. In the main menu, go to "Manager" and select "Register". This time, *flexVDI Config* will ask you the *Manager* IP address first, and then the *Manager* password. After the *Host* is registered, the SSH keys will be updated in all the *Hosts*.

Once this is done, you will see the new *Host* appear in *flexVDI Dashboard*, and you can start using it right away.

Replacing a failed Host

If you are adding a *Host* that substitutes another failed one, the best way is to use **the same domain name and IP address**. In this way, the *Manager* will replace the previous *Host*. You can also remove the old *Host* with the Dashboard and then add the new one.

Configuring OCFS2

Once the *Hosts* have been registered, we have to configure the storage they are going to use. This is done with *flexVDI Config*, in the option "OCFS2" of the main menu.

OCFS2 is the file system used by *flexVDI* to manage the shared storage used by two or more *flexVDI Hosts*. OCFS2 is a cluster file system developed by Oracle, that provides high availability and performance. It is also open sourced. With OCFS2, *flexVDI* takes advantage of its *cache-coherent parallel I/O* option that provides higher performance, and uses the file system switching on error to improve availability.

OCFS2 groups *Hosts* together in storage clusters. In fact, it is not needed to put all the *flexVDI Hosts* in the same sotrage cluster. If you have many *Hosts*, you can create several (disjoint) storage clusters that match your storage architecture. In particular, at least one storage cluster is required to access shared storage devices (which will be later configured as Internal Volume in an Image Storage) even though they are used in one *flexVDI Host* only*.*

Step 1: Creating a cluster with one node

Log into one of the *Hosts* and run `flexvdi-config`. In the OCSF2 submenu, select the Create entry.

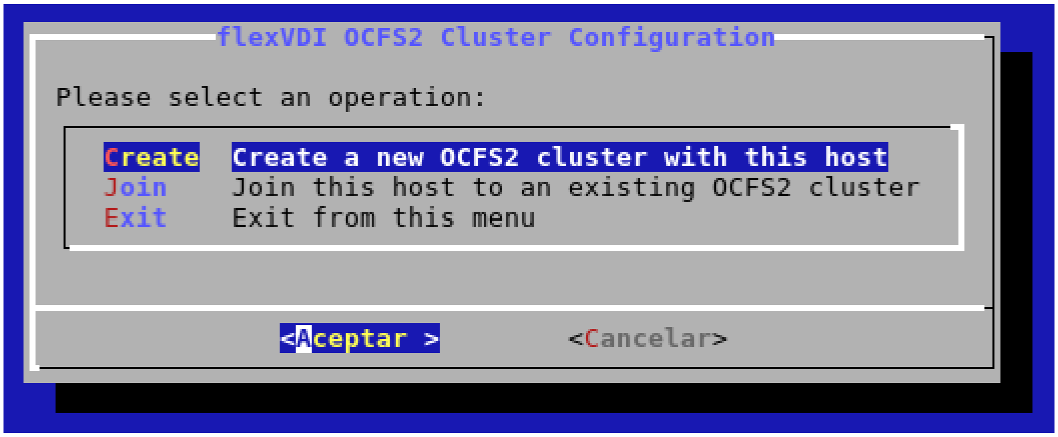

flexVDI OCFS2 Cluster Configuration

Please select an operation:

Create   Create a new OCFS2 cluster with this host
Join   Join this host to an existing OCFS2 cluster
Exit   Exit from this menu
<Aceptar >
<Cancelar>

It will prepare an OCFS2 cluster configuration with the current node as its only member. From this point, you can start using *Volumes* in an *Image Storage* with this host.

Step 2: Join other nodes to the cluster

After you have configured the first node of an OCFS2 cluster, you have to make other nodes join the cluster. After registering the *flexVDI Hosts* with the *Manager* instance, follow these steps on each of them:

1. Select OCFS2 in the main menu of `flexvdi-config`.
2. Select the option Join

3. It may ask you for the IP address and password of the Manager.
4. Enter the IP address of the first node in the cluster and press OK:

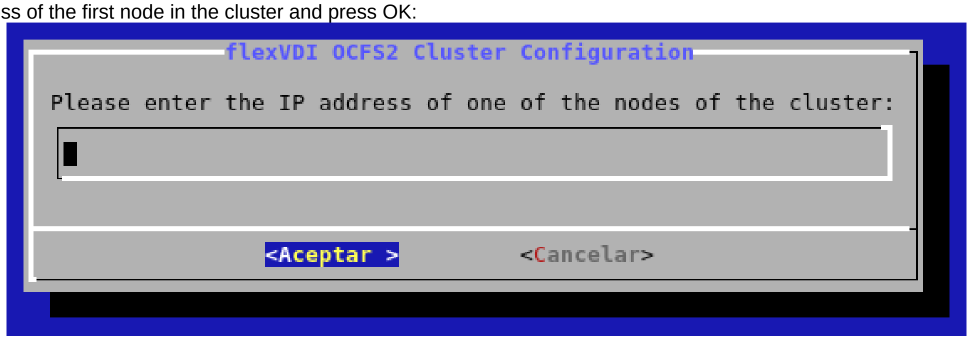

flexVDI OCFS2 Cluster Configuration

Please enter the IP address of one of the nodes of the cluster:

Do this in every Host you want to be member of the storage cluster. As said before, you can also create new storage clusters, to match your storage architecture.

Removing a Host from the cluster

This operation is a bit trickier. Removing a host from the *flexVDI* cluster can be done with the *Dashboard*. The only requirement is that no *Guest* is running in the *Host*, and that it is disabled. Then, just right-click on the desired Host and select the "Remove" option. However, the OCFS2 tools do not implement removing a node from a live storage cluster, so you must manually edit the cluster configuration and restart the OCFS2 services on each of the remaining Hosts.

> Restarting the cluster services involves unmounting and remounting all your Volumes, which results in all the running Guests being killed when their Images disappear for a moment. So, safely stop all your Guests before removing a Host from the cluster.

On each host:

1. Edit the file `/etc/ocfs2/cluster.conf`. It should look like this:

```
node:
        ip_port = 7777
        ip_address = 10.0.0.10
        number = 0
        name = flexnode01
        cluster = ocfs2

node:
        ip_port = 7777
        ip_address = 10.0.0.11
        number = 1
        name = flexnode02
        cluster = ocfs2

cluster:
        node_count = 2
        name = ocfs2
```

2. Delete the node entry of the host you are removing, adjust the node numbers of the following nodes and decrease the node count of the cluster. **Be sure to write the same configuration in all nodes**. For instance, if we removed node flexnode01 in the previous example, the resulting configuration file would be:

```
node:
        ip_port = 7777
        ip_address = 10.0.0.11
        number = 0
        name = flexnode02
        cluster = ocfs2

cluster:
        node_count = 1
        name = ocfs2
```

3. Restart the cluster services with `systemctl restart o2cb`.

Resizing an OCFS2 volume

It is possible to resize an OCFS2 volume to make it bigger (never smaller), just follow these steps:

1. Make a backup of the contents of the volume.
2. You need to unmount the volume, so:
    a. Stop all the guests with an image in that volume.
    b. Stop the flexvdi-agent service in all the hosts that share the volume. Otherwise, they will remount it as soon as they detect it is not mounted:

    ```
    # systemctl stop flexvdi-agent
    ```

    c. Unmount the volume in all the hosts.
    d. In one host, perform a filesystem check. Assuming it is in partition `/dev/sdb1`:

    ```
    # fsck.ocfs2 -fn /dev/sdb1
    ```

3. Resize the underlying device to the desired capacity. This may be a logical volume in a shared storage cluster, for instance. How you do this is out of the scope of this guide.
4. Rescan the underlying device in all your hosts. Assuming the device is `/dev/sdb`, run in all the hosts:

    ```
    # echo 1 > /sys/block/sdb/device/rescan
    ```

    If your device is part of a multipath device, rescan all the devices. Then, assuming it is called `mpatha`, run in all the hosts:

    ```
    # multipathd resize map /dev/mapper/mpatha
    ```

5. Resize the underlying device partition. Assuming the device is `/dev/sdb`, run **in one host only**:

    ```
    # parted -s /dev/sdb resizepart 0% 100%
    ```

6. Now, refresh the partition sizes in all your hosts.

    ```
    # partprobe
    ```

7. Resize the OCFS2 filesystem and check it again:

    ```
    # tunefs.ocfs2 -S /dev/sdb1
    # fsck.ocfs2 -fn /dev/sdb1
    ```

8. Finally, restart the flexvdi-agent service again in all your hosts, and they will mount the volume again in the right place.

More info:

- https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html/dm_multipath/online_device_resize
- https://www.thegeekdiary.com/how-to-resize-an-ocfs2-filesystem-on-linux/

Accessing shared storage

The following sections explain how to use a shared storage as an *Image Storage* and access its *Volumes*, and how to move your *flexVDI Manager* instance to a shared storage *Volume* to provide high availability.

- flexVDI Manager Clusterization
- Using shared storage
- Storage performance