

Frequently Asked Questions

Q: Can I install the flexVDI packages in a virtual machine with CentOS/RHEL 7?

A: Short answer, no. Long answer, flexVDI is designed to be installed on a bare metal server, and the flexvdi-config tool will complain about not finding virtualization extensions in the processor if it is run in a virtualized environment. Theoretically, you can overcome this problem if you activate nested virtualization in the outer hypervisor. However, this will come with an impact on performance and possibly other problems too. For this reason, nested virtualization is officially not supported, and you will not get any help by our support team for this configuration.

Q: What is the difference between persistent and volatile desktops?

A: A volatile desktop is a guest that can be deleted and replaced with a fresh one, and its user can keep working as if nothing happened. This is implemented in the following way:

- It must be a clone of a template, so that it can be replaced with another clone of that template.
- User's data and applications must be stored in external storage, such as a networked filesystem

A persistent desktop does not fulfill one of these conditions, so that deleting it results in not being able to recover it and/or a loss of user's data. Note that, although in flexVDI all desktops are clones by default, the guest operating system must be correctly configured so that user's data is stored in external storage. For instance, using roaming profiles in Windows or NFS home directories in Linux. The advantages of volatile desktops are:

- Templates can be updated frequently.
- Fatal errors can be easily recovered by recycling the desktop.
- Frequently recycling makes the differential images small.

Q: Which network ports should I open in my hosts' firewall?

A: You must make TCP ports 22, 443, 7777, 9443, 5800 and 5900 and up reachable from your other hosts in the platform. Of these, only TCP 443 must be reachable from the outside to communicate with the flexVDI clients.

Q: I just configured an iSCSI volume and it showed as /dev/sdx, reachable by all hosts of an Image Storage, why is it not shown in the Dashboard?

A: flexVDI Agent does not detect new disks once it is running. If you add a new device to a host, like a iSCSI volume, you have to restart the flexvdi-agent service. Then, it will appear as available to create a Volume in the Dashboard.

Q: What is the difference between an IDE controller and a Virtio controller for disk images?

A: The KVM hypervisor emulates a disk controller to which disk images are attached. It can emulate the full behaviour of an IDE controller, or it can use the virtio controller, which is virtualization-aware. In short, the virtio controller is up to 5 times faster than the IDE controller, by using direct communication buffers with the host, but it requires a special driver. This driver is present in all modern GNU/Linux distributions, but not in Windows. It is provided by our guest tools ISO, so you can use it to install Windows on a virtio disk; it is in the `viostor` directory. There is no virtio disk driver for Windows XP, you must use an IDE controller.

Q: A guest, whose disk image is in a volume that is visible by just some of the hosts of its pool, cannot be started. Why?

A: The manager has decided to start the guest in one of the hosts that cannot reach the volume, and that host's agent refuses to start the VM. flexVDI will not stop you from configuring the system in this way, and it will not try to fix it in either. So, you must be careful that **all** the hosts in a pool can reach the volumes where its VM's images are. If you have a guest's disk in a volume, and all your pools include some host that cannot reach it, create a new pool just for that guest, with the same set of hosts as the volume. **Do not** use the *preferred host* option to fix this, see next question.

Q: A guest is not started in its preferred host, even when it has enough free resources. Why?

A: The flexVDI Manager performs resource allocation at boot time, or manually when you issue a rebalance operation. It will iterate over all pools in priority order, and allocate part of the resources of each host to them until all their requirements are fulfilled. This allocation will not change unless explicitly requested, which ensures that higher priority pools always have enough resources for their guests, but leads to some curious situation. For instance, suppose there is a flexVDI cluster with two hosts, A and B, with 8 GB of RAM each. A pool is configured on both hosts which requires 8 GB of RAM. The manager decides to allocate 4 GB on each host for this pool, so there are other 4 GB free on each one. You run a first guest that requires 4 GB of RAM, and the Manager decides to start it on host A. Then, any other guest of the pool will always run on host B, even if host A is its preferred host, because all the resources of host A allocated to the pool are already being used by the first guest, even though it still has 4 GB of free RAM.

Q: My volatile desktops are suddenly freezing! what is going on?

A: You probably run out of volatile storage space. In the desktop policy, you can configure whether the differential images of volatile desktops are stored in the standard location (`/flexvdi/volatile`) or in a volume of your choice. Be careful with this decision, because differential images only store the difference with their template's image, and they can grow quite a lot. Some things that make them grow:

- Using a page file or swap space.
- Updating the OS in the volatile desktop instead of the template.
- Not recycling volatile desktops often.

If you decide to use the standard location, it is a good design choice to mount an additional SSD disk with enough space in `/flexvdi/volatile`, so that differential images have dedicated fast storage.

Q: I have hot-plugged a new disk, but flexVDI does not see the new device. What can I do?

A: Now flexVDI configuration defaults to not periodically rescan for new disk devices. After the disk is installed in the host, execute

```
systemctl restart flexvdi-agent
```

And the flexVDI Host will update its device list without disturbing end users or virtual machines. The information will be automatically propagated to flexVDI manager and flexVDI dashboard in some seconds.

Q: The kernel shows a message like "KVM: vcpu0 disabled perfctr wrmsr". Is this a problem?

A: [According to Red Hat](#), no, this warning message is of no concern. Just annoying.