

KVM Virtual Server Quick Start

This Quick Start provides a scenario for defining and operating a virtual server. It is assumed that the Linux host with KVM support has previously been installed and set up. The scenario has been tested at IBM.

Introduction

The virtual server (called a *domain* in libvirt) used in this scenario has the following resources:

- 2 virtual CPUs
- 4 GB of virtual memory
- 1 virtual disk that is backed by a zFCP-attached SCSI disk
- 1 virtual network interface that is directly connected to a bonded pair of OSA network ports.

It is assumed that the SCSI disk contains a supported Linux installation that can be IPLed, and that the resources are identified in the hypervisor as follows:

SCSI Disk: /dev/mapper/36005076305ffc1ae00000000000020d3
Bonded Interface: bond0

1 Step 1: Write a domain XML file

Use a text editor to write a domain XML file with the following content:



```
<domain type="kvm">
  <name>quickstart1</name>
  <memory unit="GiB">4</memory>
  <vcpu>2</vcpu>
  <os>
    <type>hvm</type>
  </os>
  <iotreads>1</iotreads>
  <on_crash>preserve</on_crash>
  <devices>
    <disk type="block" device="disk">
      <driver name="qemu" type="qcow2" cache="none" iotread="1" io="native"/>
      <source dev="/dev/mapper/36005076305ffc1ae00000000000020d3"/>
      <target dev="vda" bus="virtio"/>
      <boot order="1"/>
    </disk>
    <interface type="direct">
      <source dev="bond0" mode="bridge"/>
    </interface>
    <console type="pty">
      <target type="sclp"/>
    </console>
    <memballoon model="none"/>
  </devices>
</domain>
```

If you wish to run this scenario:

- Do *not* change the XML elements in the above example that are shown in **bold** typeface, which are processed by libvirt in Step 2.
- Adapt the other XML elements to suit your own environment (in the above example you will need to change the number of virtual CPUs, memory, virtual server name, and the specifications for the disk source and the interface source).
- Give the XML file the same name as the domain (for example, quickstart1.xml).

2 Step 2: Define the virtual server



To define the virtual server, enter:

```
# virsh define quickstart1.xml
```

To verify that the virtual server has been successfully defined, enter:

```
# virsh dumpxml quickstart1
<domain type="kvm">
  <name>quickstart1</name>
  <uuid>b685e7bd-f836-4a9a-a8ba-78e7f51aa564</uuid>
  <memory unit="KiB">4194304</memory>
  <currentMemory unit="KiB">4194304</currentMemory>
  <vcpu placement="static">2</vcpu>
  <iothreads>1</iothreads>
  <os>
    <type arch="s390x" machine="s390-ccw-virtio-2.12">hvm</type>
  </os>
  <clock offset="utc"/>
  <on_poweroff>destroy</on_poweroff>
  <on_reboot>restart</on_reboot>
  <on_crash>preserve</on_crash>
  <devices>
    <emulator>/usr/bin/qemu-kvm</emulator>
    <disk type="block" device="disk">
      <driver name="qemu" type="qcow2" cache="none" io="native" iotthread="1"/>
      <source dev="/dev/mapper/36005076305ffc1ae000000000000020d3"/>
      <target dev="vda" bus="virtio"/>
      <boot order="1"/>
      <address type="ccw" cssid="0xfe" ssid="0x0" devno="0x0000"/>
    </disk>
    <interface type="direct">
      <mac address="52:54:00:5a:f5:d2"/>
      <source dev="bond0" mode="bridge"/>
      <model type="virtio"/>
      <address type="ccw" cssid="0xfe" ssid="0x0" devno="0x0001"/>
    </interface>
    <console type="pty">
      <target type="sclp" port="0"/>
    </console>
    <memballoon model="none"/>
    <panic model="s390"/>
  </devices>
</domain>
```

Note: The system adds a number of elements (shown above in **bold** font) to the virtual server definition to uniquely define the virtual server and its resources. Furthermore, it replaces **machine="s390-ccw-virtio"** with **machine="s390-ccw-virtio-2.12"** (for the purposes of this scenario, the newest machine).

3 Step 3: Start, monitor, and stop the virtual server



To start the virtual server enter:

```
# virsh start quickstart1 --console
```

By specifying the `--console` option, the operating system console is displayed. You can therefore:

- View the operating system messages.
- Log in to the guest operating system once it has completed its startup.

To leave the `virsh` console, enter **Ctrl+J**.

To view the state of the virtual server, enter:

```
# virsh dominfo quickstart1
Id:                3
Name:              quickstart1
UUID:              b685e7bd-f836-4a9a-a8ba-78e7f51aa564
OS Type:           hvm
State:             running
CPU(s):            2
CPU time:          9.6s
Max memory:        4194304 KiB
Used memory:       4194304 KiB
Persistent:        yes
Autostart:         disable
Managed save:     no
Security model:    selinux
Security DOI:      0
Security label:    system_u:system_r:svirt_t:s0:c12,c1014 (enforcing)
```

To gracefully stop the virtual server, enter:

```
# virsh shutdown quickstart1
```

Some useful commands



These are some of the most useful commands that you might wish to use:

- **virsh console** - display the console of a virtual server.
- **virsh define** - define a virtual server in libvirt in state “shut off”.
- **virsh destroy** - immediately terminate a virtual server and release any resources which are being used by it.
- **virsh dominfo** - display the state of the virtual server.
- **virsh edit** - open a text editor containing the domain XML of a virtual server. All saved changes are checked for errors.
- **virsh help** - obtain online help.
- **virsh list** - browse the virtual servers that have been started. **virsh list --all** will browse all virtual servers.
- **virsh shutdown** - properly shut down a running virtual server.
- **virsh start** - start a defined virtual server that is “shut off” or “crashed”.
- **virsh undefine** - remove a virtual server definition from libvirt.

For more information about the topics described in this Quick Start, you might refer to the IBM® Knowledge Center at www.ibm.com/support/knowledgecenter/linuxonibm/liaaf/lnz_r_lib.html, or:

- *KVM Virtual Server Management*, SC34-2752.
- *Device Drivers, Features, and Commands for Linux as a KVM Guest*, SC34-2754.
- The documentation provided by your Linux distributor.

