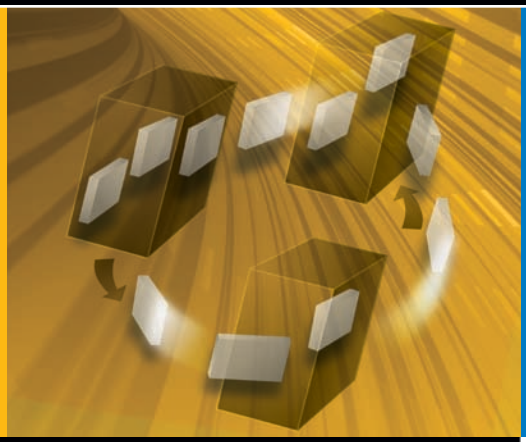

XEN VIRTUALIZATION IN RED HAT ENTERPRISE LINUX 5

BY JOSE DE LA ROSA
PUNEET DHAWAN

The open source Xen virtualization hypervisor included in the Red Hat® Enterprise Linux® 5 OS provides key components to help build dynamic, scalable virtualized environments. This article describes the basic features of Xen virtualization in Red Hat Enterprise Linux 5 and how organizations can deploy it on ninth-generation Dell™ PowerEdge™ servers.

Related Categories:

Red Hat Enterprise Linux

Virtualization

Visit ***DELL.COM/PowerSolutions*** *for the complete category index.*

Virtualization enables enterprises to consolidate multiple servers without sacrificing application isolation, scale their infrastructure as their needs grow, and increase availability through dynamic provisioning and relocation of critical systems. Combining the open source Xen hypervisor in Red Hat Enterprise Linux 5 with ninth-generation Dell PowerEdge servers and Dell storage enables enterprises to create a dynamic data center that can scale easily to meet enterprise requirements.

Xen architecture

The virtualization layer, often called the hypervisor or virtual machine (VM) monitor, abstracts underlying physical hardware to present a uniform set of hardware resources—such as processors, memory, networks, and storage blocks—to VMs. VMs running on a single system share available physical resources, with the hypervisor multiplexing key resources and maintaining isolation among different VMs. Figure 1 shows the architecture of the Xen hypervisor.

How the hypervisor abstracts the underlying physical resources defines key characteristics of the virtualization architecture. Two popular approaches are *full virtualization* and *paravirtualization*. Full virtualization presents emulated resources to VMs that mimic a standard PC architecture and standard peripheral devices. Using this approach allows operating systems to run inside a VM without modification, but may require overhead that can reduce performance.

Paravirtualization—the approach used by Xen—modifies guest operating systems to run in a virtualized environment. The VMs interface with the Xen hypervisor using hypercalls, rather than the system calls used by full virtualization. As shown in Figure 1, the real device drivers run in a special VM, or *domain,* called Domain 0 (Dom0). Rather than abstracting standard devices for the VMs, Dom0 exposes a set of class devices, such as networks and storage blocks, to the VMs. I/O data transfers to and from each VM through Xen use the XenBus memory-mapped communication channel. The VMs use paravirtualized device drivers and a paravirtualized kernel to interoperate with Dom0 and the Xen hypervisor.

Paravirtualization requires modifying guest operating systems, which is not possible for all popular OS distributions. However, Xen can take advantage of Intel® Virtualization Technology (Intel VT) and AMD Virtualization™ (AMD-V™) technology to run unmodified operating systems as well. The virtualization capabilities of Red Hat Enterprise Linux 5 coupled with ninth-generation Dell PowerEdge servers can create a flexible, powerful virtualized environment that accommodates both modified and unmodified guest operating systems.

"Combining the open source Xen hypervisor in Red Hat Enterprise Linux 5 with ninth-generation Dell PowerEdge servers and Dell storage enables enterprises to create a dynamic data center that can scale easily to meet enterprise requirements."

When Red Hat Enterprise Linux 5 is installed with virtualization capabilities, the integrated Xen hypervisor takes control of the system hardware and launches the installed Red Hat Enterprise Linux 5 distribution as Dom0. In addition to serving as the main driver domain for VMs, Dom0 also runs a set of control and management services that administrators can access through command-line interface (CLI) tools such as xm and virsh or graphical user interface (GUI) tools such as Virtual Machine Manager (virt-manager).

Virtualization support and requirements in Red Hat Enterprise Linux 5

Before creating a virtualized environment with the Xen technology in Red Hat Enterprise Linux 5, administrators should be sure they understand the support and system requirements for elements such as virtual resources, host servers and operating systems, processors, storage, packages, and the Security-Enhanced Linux (SELinux) security policy.

Virtual resources

Red Hat virtualization with Xen technology can host multiple guest operating systems, each of which runs in its own domain. Each VM handles its own applications and can only access the resources assigned to it. Assigned resources include the following:

- **Processors:** Administrators can configure a VM with multiple virtual processors, but the total number of virtual processors assigned to a VM must be less than or equal to the total number of logical processors in the host system. Virtual Machine Manager can schedule virtual processors according to the physical processor workload to help optimize available resources.
- **Memory:** Each VM is assigned a part of the host system's physical memory. Administrators should typically assign the same amount of memory to a VM as they would for the same configuration in a non-virtualized environment. Administrators can define the initial and maximum memory size when creating VMs, then increase or decrease the memory allocation at runtime without exceeding the specified maximum. The minimum amount of memory recommended for a VM is 256 MB.
- **Disk space:** Each VM is assigned a part of the host system's disk space. This disk space is unique and cannot be shared between VMs. The disk space made available to VMs can be either an image file or a disk partition.
- **Network interfaces:** Virtual network interface cards are configured with a persistent virtual Media Access Control (MAC) address. When a new VM is created, this address is selected at random from a reserved pool of over 16 million addresses, making it unlikely that any two VMs will be assigned the same one. Administrators for complex sites with a large number of VMs can allocate MAC addresses manually to help ensure that they remain unique on the network. Red Hat virtualization supports 10/100/1,000 Mbps Ethernet and 10 Gigabit Ethernet, Fibre Channel, and InfiniBand networks.

Each VM also has a virtual text console that connects to the host system. Administrators can redirect guest logins and console output to the text console, or configure VMs to use a virtual GUI console that corresponds to the physical host's standard video console. This GUI employs standard graphic adapter features such as boot messaging, graphical booting,

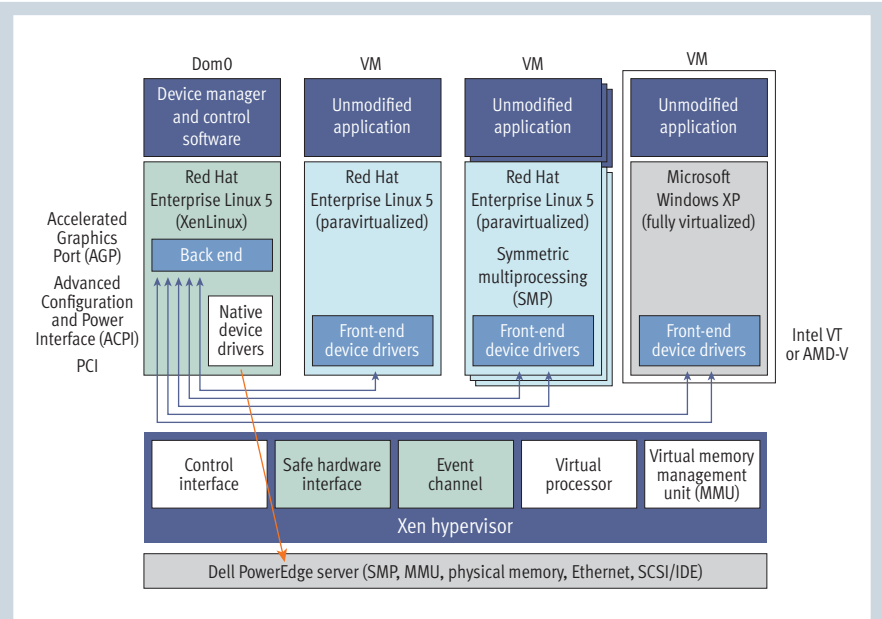

Figure 1. *Xen hypervisor architecture*

and multiple virtual terminals, and can launch the X Window System.

VMs can be identified in any of three ways:

- **Domain name:** Text string that corresponds to a VM configuration file, used to launch, identify, and control VMs
- **Domain ID:** Unique, nonpersistent number assigned to an active domain, used to identify and control VMs
- **Universally unique identifier:** Identifier controlled from the VM configuration file that helps ensure that VMs are uniquely identified by systems management tools

Host servers and operating systems

The following Dell PowerEdge servers and workstations are certified by Red Hat to run in virtualized environments with Red Hat Enterprise Linux 5 as the host OS:

- **Servers:** Dell PowerEdge server model numbers 700, 750, 800, 830, 840, 850, 860, 1650, 2600, 2650, 4600, 6600, 6650, 1750, 1800, 1850, 1855, 2800, 2850, 6800, 6850, 1900, 1950, 1955, 2900, 2950, 2970, 6950, SC440, SC1420, SC1425, SC1430, and SC1435
- **Workstations:** Dell Precision™ workstation model numbers 380n, 390n, 470n, 490n, 670n, and 690n

Red Hat Enterprise Linux 3, Red Hat Enterprise Linux 4, and Red Hat Enterprise Linux 5 can run as fully virtualized guest operating systems, but only Red Hat Enterprise Linux 4 Update 5 and Red Hat Enterprise Linux 5 can run as paravirtualized guest operating systems. Other Linux distributions, as well as Microsoft® Windows® operating systems, can run as fully virtualized guest operating systems, but are not supported by Red Hat.

As a host OS, Red Hat Enterprise Linux 5 supports up to four VMs. Red Hat Enterprise Linux 5 Advanced Platform (for four-socket servers) allows an unlimited number of VMs, enabling administrators to create as many VMs as the underlying physical hardware can support.

Note: Red Hat virtualization only fully supports matching host/guest architectures—that is, environments in which both the host and guest OS are either 32-bit or 64-bit operating systems. Running a 32-bit guest OS on a 64-bit host OS is supported only with paravirtualization, and running a 64-bit guest OS on a 32-bit host OS is not supported.

Processors

Red Hat virtualization requires that the host system's processors support Physical Address Extension (PAE) and have either Intel VT or AMD-V enabled. Administrators can determine whether a system's processors support PAE using the command `grep pae /proc/cpuinfo`. If this command returns output similar to the following, including the `pae` entry, then the system supports PAE:

```
flags : fpu vme de pse tsc
   msr pae mce cx8 apic mtrr
   pge mca cmov pat clflush
   dts acpi mmx fxsr sse sse2
   ss tm pbe nx up est tm2
```

If the command returns no output, then the system does not support PAE.

Xen uses a generic Hardware VM (HVM) layer to support both Intel and AMD™ processors. Administrators can determine whether a system's processors support Intel VT or AMD-V using the command `egrep -e 'vmx|svm' /proc/cpuinfo`. If this command returns output similar to the following, including the `vmx` entry, then the system supports Intel VT or AMD-V:

```
flags : fpu tsc msr pae mce
   cx8 apic mtrr mca cmov pat
   clflush dts acpi mmx fxsr
   sse sse2 ss ht tm pbe
   constant_tsc pni monitor
   vmx est tm2 xtpr
```

If the command returns no output, then the system does not support Intel VT or AMD-V.

For systems that support Intel VT or AMD-V, administrators must also enable that feature in the BIOS. Administrators can verify that the feature is enabled by running the command `grep hvm /sys/hypervisor/properties/capabilities`, which should return output similar to the following:

```
xen-3.0-x86_32p hvm-3.0-x86_
   32 hvm-3.0-x86_32p
```

If the command returns no output, administrators should look for a setting related to virtualization in the system's BIOS setup utility, enable that setting, save, and reboot the system.

Storage

For VM storage, Red Hat virtualization supports direct access storage devices as well as network attached storage (NAS) and storage area networks (SANs) based on Internet SCSI (iSCSI), Fibre Channel, and standard network protocols such as Network File System (NFS) and Cluster Logical Volume Manager (CLVM). Administrators can manage VM storage in multiple ways. A physical block device (hard disk partition or ISO image) on the host system can be exported to a guest domain as a virtual block device.

Packages

To run a virtualized environment, the kernel-xen kernel must be installed and running on the host system. Administrators can determine which kernel is running using the command `uname -r`; if this command does not return a kernel with the word "xen" in it, then kernel-xen is not running. If it is not running, administrators can determine whether it is installed using the command `rpm -qa | grep kernel-xen`. If this command returns no output, they must install the kernel from the installation media with the command `rpm -ivh kernel-xen`.

Next, administrators should make kernel-xen the default boot kernel, which they can do by changing the `default` parameter in /boot/grub/grub.conf to the correct number (typically 0). They should also verify that the xen, xen-libs, bridge-utils, gnome-python2-gnomekeyring, libvirt, libvirt-python, python-virtinst, virt-manager, and vnc Red Hat Package Manager

(RPM™) files are installed, or install them from the installation media if necessary.

Because of the interdependencies between these RPM packages, administrators should typically update them using the yum (Yellowdog Updater, Modified) package installer. They must first define the yum repository in /etc/yum.conf, then use the command `yum install rpm_name` to install the necessary packages. For example:

```
yum install xen
yum install virt-manager
yum install vnc
```

After installing the necessary packages, administrators should reboot the system. Finally, they should verify that the xend and xendomains daemons are running using the commands `service xend status` and `service xendomains status`. The xend daemon provides virtualization services, while the xendomains daemon allows VMs to start and stop automatically when the host system boots or shuts down. Both daemons must be running to create VMs. If they are not running, administrators should start them using the commands `service xend start` and `service xendomains start`.

SELinux security policy

The targeted SELinux security policy for Xen requires that disk images have the xen_image_t context. If a VM file image is not in that context, the host system is denied access to that image. By default, only two directories are in the xen_image_t context, as shown by the output of the command `semanage fcontext -l | grep xen_image_t`:

```
/xen(/.*)?    all files
    system_u:object_r:xen_
    image_t:s0
/var/lib/xen/
    images(/.*)?  all files
    system_u:object_r:xen_
    image_t:s0
```

If administrators want to create VM image files in other directories, they must add those directories to the SELinux policy. For example, they could use the following command to create an image in the directory /newdir:

```
semanage fcontext --add
    -t xen_image_t '/
    newdir(/.*)?'
```

They could then give this directory the appropriate context using the command `restorecon -v /newdir`.

Virtual machine creation in Red Hat Enterprise Linux 5

Administrators can take advantage of two primary tools when creating VMs: the Virtual Machine Manager GUI tool and the virt-install CLI tool.

Creating a virtual machine with Virtual Machine Manager

Virtual Machine Manager is a GUI tool provided in Red Hat Enterprise Linux 5 that administrators can use to create, pause, resume, stop, and monitor VMs. For example, administrators can use it to create a paravirtualized 32-bit Red Hat Enterprise Linux 5 guest OS on a 32-bit Red Hat Enterprise Linux 5 host OS by performing the following steps:

1. Run Virtual Machine Manager using the command `virt-manager` and connect to the local Xen host. At this point, the only domain running is Dom0.
2. To create a new VM, select File › New Machine, then click the Forward button.
3. Provide a name for the VM—for example, vm_rhel5_i386—then click the Forward button.
4. Select "Paravirtualized" as the virtualization method, then click the Forward button.
5. Provide the location of the extracted installation media files for the guest OS and, if desired, the location of a kickstart file with the system parameters already defined, then click the Forward button (see Figure 2).
6. Specify how to assign the VM disk space—a partition on the host system or an image file—and how much space to allocate, then click the Forward button (see Figure 3). If allocating the entire virtual disk now, be sure to verify that the host system has enough disk space to accommodate the specified amount.
7. Allocate the amount of memory and number of virtual processors, then click the Forward button (see Figure 4).
8. Review the specified parameters and click the Finish button to begin the VM creation. The OS installation process is the same as a non-virtual OS installation. If allocating the entire virtual disk now (as shown in Figure 3), then the system

Figure 2. *"Locating installation media" step when creating a virtual machine in Red Hat Enterprise Linux 5 Virtual Machine Manager*

Figure 3. *"Assigning storage space" step when creating a virtual machine in Red Hat Enterprise Linux 5 Virtual Machine Manager*

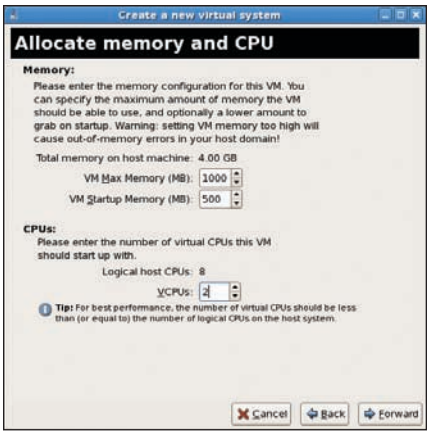

Figure 4. *"Allocate memory and CPU" step when creating a virtual machine in Red Hat Enterprise Linux 5 Virtual Machine Manager*

typically takes several minutes to create the disk space before it begins installation. When the OS installation process ends, the virtual window closes and the VM is ready to start.

9. To start the VM, enter the command `xm create vm_name`. Virtual Machine Manager should now show both Dom0 and the VM running (see Figure 5). To open a Virtual Network Computing (VNC) display window for this VM, select the VM and click the Open button (see Figure 6).

Creating a virtual machine with virt-install

Administrators can also use the virt-install CLI tool to create a paravirtualized VM. Specifying parameters in this way enables administrators to automate VM creation using shell scripts. The syntax is as follows:

```
virt-install --name=vm_name
    --ram=memory --vcpus=no_of_
    vcpus --file=vm_image_file
    --file-size=vm_disk_size
    --vnc --paravirt
    --location=OS_source_location
```

In this command, *vm_name* is the name of the VM, *memory* is the amount of memory to allocate to the VM (in megabytes), *no_of_vcpus* is the number of virtual processors to allocate to the VM, *vm_image_file* is the file to use as the disk image, *vm_disk_size* is the amount of VM disk space to allocate to the VM (in gigabytes), `--vnc` sets the VM to use VNC for graphics support, `--paravirt` denotes that the VM should be paravirtualized, and *OS_source_location* is the location of the extracted installation media files for the guest OS. This command also launches a VNC display window.

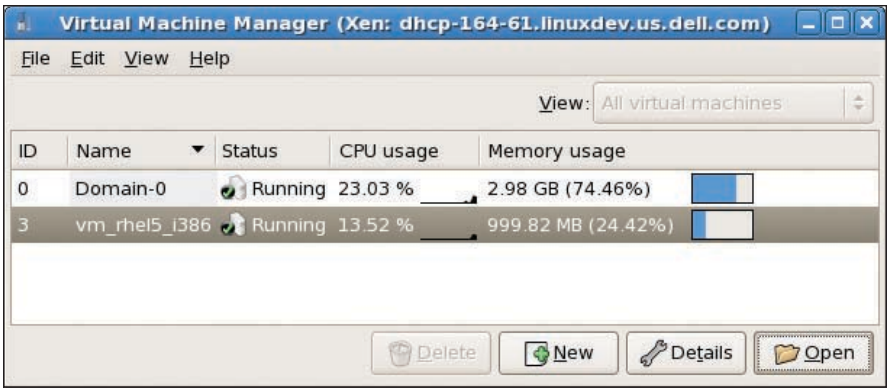

Figure 5. *Red Hat Enterprise Linux 5 Virtual Machine Manager after creating a virtual machine*

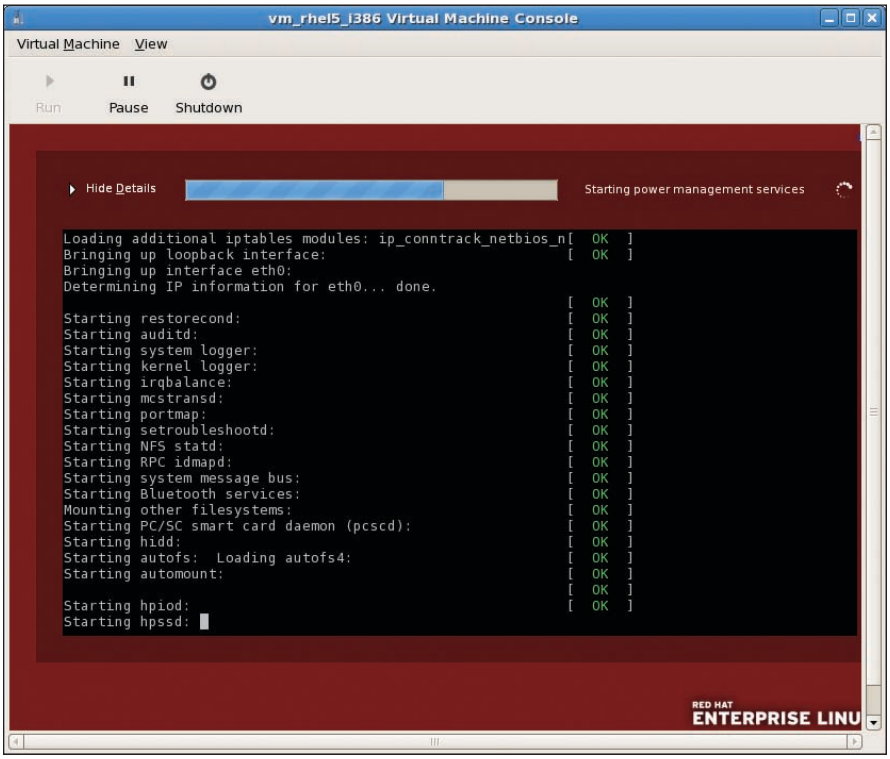

Figure 6. *Virtual machine booting in Red Hat Enterprise Linux 5*

For example, administrators could create the VM from the preceding section as follows:

```
virt-install --name=vm_rhel5_
    i386 --ram=1000 --vcpus=2
    --file=/var/lib/xen/images/
    vm_rhel5_i386 --file-
    size=10 --vnc --paravirt
    --location=http://webserver/
    pub/RHEL5/i386
```

```
# Automatically generated xen config file
name = "vm_rhel5_i386"
memory = "1000"
disk = [ 'tap:aio:/var/lib/xen/images/vm_rhel5_i386,xvda,w', ]
vif = [ 'mac=00:16:3e:0c:77:c5, bridge=xenbr0', ]
vfb = ["type=vnc,vncunused=1"]
uuid = "9847a069-51e4-8173-9d1b-2d033cf63d5b"
bootloader="/usr/bin/pygrub"
vcpus=2
on_reboot   = 'restart'
on_crash    = 'restart'

name        Specifies the name of the virtual guest.
memory      Specifies the amount of memory in megabytes.
disk        Lists the block and physical devices to export to the domain. In this case, it lists
            the virtual guest image file, the device name and the access permissions (write).
vif         Lists the randomly-assigned MAC addresses and bridges.
vfb         Specifies the virtual frame buffer. In this case, we use VNC.
uuid        Used to uniquely identify the virtual guest.
bootloader  Specifies to use the Xen bootloader.
vcpus       Specifies the numbers of virtual CPUs.
on_reboot   Specifies what action to take when the guest reboots.
on_crash    Specifies what action to take when the guest crashes.
```

Figure 7. */etc/xen/vm_rhel5_i386 configuration file*

Red Hat Enterprise Linux 5 Xen configuration and log files

The Red Hat Enterprise Linux 5 Xen configuration files are located in the /etc/xen directory. Each VM has a corresponding configuration file in /etc/xen, which is created automatically when the VMs are created and has the same name as its corresponding VM. Figure 7 shows the /etc/xen/vm_rhel5_i386 configuration file for the VM created in the "Virtual machine creation in Red Hat Enterprise Linux 5" section in this article. For more information on available configuration items, see the xmdomain.cfg man page.

The Red Hat Enterprise Linux 5 Xen log files are located in the /var/log/xen directory. The xend daemon and qemu-dm process, for example, write to multiple log files:

- **xend-debug.log:** Contains logs of event errors from xend and the virtualization subsystems (such as the frame buffer and Python scripts)
- **xend.log:** Contains data collected by the xend daemon, including system events, administrator actions, and VM operations such as create, shutdown, and destroy; this log is typically the first place administrators should look when troubleshooting event- or performance-related problems
- **xen-hotplug.log:** Contains data from hot-plug events, including events when a device or network script does not come online
- **qemu-dm.*pid*.log**: Created by the qemu-dm process for each fully virtualized guest (where *pid* is the process identifier)

Virtual machine management commands

Administrators can use the xm and virsh command-line interface tools to create, manage, and troubleshoot VMs (see Figure 8). Some commands require additional arguments; for more information, see the xm and virsh man pages.

	Command	Description
Starting and stopping VMs	xm create	Creates a domain based on a configuration file
	xm destroy	Terminates a domain
	xm pause	Pauses execution of a domain
	xm reboot	Reboots a domain
	xm shutdown	Shuts down a domain
	xm save	Saves a domain state to restore later
	xm restore	Restores a domain from a saved state
Status monitoring	xm uptime	Displays uptime for a domain
	xm top	Monitors a host and its domains in real time
	xm list	Displays domain information
	xm info	Displays host information
	xm vcpu-list	Lists domain virtual processors
	xm network-list	Lists domain virtual network interfaces
	virsh nodeinfo	Displays node information
	virsh vcpuinfo	Displays domain virtual processor information
Troubleshooting	xm console	Attaches to a domain console
	xm dump-core	Displays a core dump for a specific domain
	xm dmesg	Reads and/or clears the xend daemon's message buffer
	xm log	Displays the xend log
	virsh dominfo	Displays domain information
Performance tuning	xm mem-max	Sets the maximum amount of memory for a domain
	xm mem-set	Sets the current memory usage for a domain
	xm vcpu-set	Sets the number of active processors for a domain
	virsh dumpxml	Displays domain information in XML
	virsh dump	Saves a core dump for a specific domain to a file
Other	xm rename	Renames a domain
	xm sysrq	Sends a system request to a domain
	xm block-attach	Creates a new virtual block device
	xm block-detach	Destroys a domain's virtual block device
	xm block-list	Lists virtual block devices for a domain
	xm network-attach	Creates a new network device
	xm network-detach	Destroys a network device

Figure 8. *Commands for the xm and virsh tools for creating, managing, and troubleshooting virtual machines*

Dynamic, scalable virtualized environments

Red Hat Enterprise Linux 5 incorporates Xen virtualization technology to help create a robust, secure, easy-to-manage virtualization platform that can scale flexibly to meet enterprise needs. Combined with ninth-generation Dell PowerEdge servers, Red Hat virtualization can provide high-performance virtualized environments for enterprises of all sizes.

Jose De la Rosa *is a member of the Dell Linux Engineering Group. He has a bachelor's degree in Electrical Engineering and a master's degree in Computer Engineering from the University of Texas at El Paso.*

Puneet Dhawan *is a systems engineer in the Dell Virtualization Solutions Engineering Group. Puneet has a bachelor's degree in Electrical Engineering from Punjab Engineering College (PEC) and a master's degree in Computer Engineering from Texas A&M University.*