

Managing Xen HVM Guests

0.1 Domain creation or management

0.1.1 Create a new domain

To create a new domain or virtual machine we can use command

```
xm create <hvm_configuration_file> -c
```

Here, ‘-c’ is to attach console as soon as the domain is created.

0.1.2 Delete/Destroy running domain

To delete/destroy running domain (This is like force-off on VM) we can use:

```
xm destroy {<domain_id>|<domain_name>}
```

Here:

domain_id is the domain ID listed against domain in ‘`xm list`’ command output

domain_name is the name of domain that we specify with ‘`name`’ parameter in domain configuration file

This is like removing power from VM without giving it time to shutdown gradually. File systems, databases, network connections may be in inconsistent state and hence destroying a domain may lead to data loss. Unlike ‘`shutdown`’ and ‘`reboot`’ options discussed later this option is guaranteed to force off a domain and guest OS configuration cant prevent domain from getting powered off.

0.1.3 Pause a domain

To pause a domain so that it is no longer considered by Xen Hypervisor for scheduling we can use:

```
xm pause {<domain_id>|<domain_name>}
```

Note that paused domain will continue using RAM resoures. All files used for virtualization on host OS will still remain open. Only domain is not scheduled by hypervisor. If we pause domain for long time then the network connections (like TCP) may time-out and hence when we unpause the state of network connections may change. If network connections are with other VM and they are also paused then timeout should not occur.

0.1.4 Unpause a domain

To unpause a paused domain so that it is again considered by Xen Hypervisor for scheduling, we can use:

```
xm unpause {<domain_id>|<domain_name>}
```

0.1.5 Reboot a domain

We can use command:

```
xm reboot [options] {<domain_id>|<domain_name>}
```

to reboot a running domain. Two options are available while using reboot command

-w ⇒ This option can be used to indicate to ‘xm’ that we want to wait till the reboot completes. This prevents command from returning just after indicating to Xen hypervisor that mentioned domain should be rebooted. This can be useful in scripting or programming when we want following commands to be executed only after reboot is finished.

Note that by reboot finished we just mean that VM has rebooted. It does not means that OS has booted again. OS may take considerable more time to boot after VM has rebooted.

-a ⇒ This option can be used to reboot all running domains. In this case we do not need to specify any ‘domain_id’ or ‘domain_name’ as parameter.

Note this will work only if OS is configured to reboot, if reboot button is pressed on keyboard/PC. One can configure OS to give prompt whenever reboot / poweroff / suspend etc. buttons are pressed. In those cases the OS and hence domain may not reboot but provide option to user and wait for him/her to choose what should be done.

0.1.6 Shutdown a domain

We can use command

```
xm shutdown [options] {<domain_id>|<domain_name>}
```

to shutdown a running domain. Options for shutdown are same as options for reboot (‘-a’ and ‘-w’) and have similar effect. Here also it is necessary for guest OS to be configured to shutdown when shutdown button on keyboard/Cabinet is pressed. If guest OS is configured to prompt to user before some action is taken then shutdown is not guaranteed.

0.1.7 Saving a domain

We can save running domain in a state file. This effectively means that domain would be paused and its RAM contents will be stored in state file. This is some what like hibernating the guest OS but no support for hibernate is required from Guest OS. Also normally in hibernate the machine uses its own hard-disk for storing RAM contents. In this case the host OS would store RAM contents of guest OS. When we save a domain RAM resources and other host resources used by domain on host OS are freed and hence we can use the same RAM for other purposes like creating / restoring other domains.

To save a running domain in a state file we can use:

```
xm save {<domain_id>|<domain_name>} <state_file>
```

Note that:

- When we use save command on domain it gets renamed to ‘migrating-<old_name>’ and remains in running state for a while. After that it goes to ‘s’ (shut-down state) and eventually gets removed from list of domains.
- Sufficient space must be available on partition where we are trying to save the image so that entire RAM contents and little extra information can be saved.

0.1.8 Restoring a domain

We can restore a domain from state file using:

```
xm restore <state_file>
```

Note that:

- During restore we only specify state file and not the initial configuration with which domain was created. Hence apart from RAM contents state file also stores information about configuration of domain.
- Like pause here also network connections may time out if we are restoring domain after long time.
- When we try to restore a domain it will initially remain in blocked and paused states. Later it moves to blocked or running state.

- State file is not deleted / altered during restore and hence we can restore using same state file as many time as we want. This can help in cloning VMs. If we use this method to create clones then IP conflict will occur and same MAC address on both machines would cause problem. Such issues would need to be resolved.
- Care must be taken to ensure that sufficient RAM is free on host OS before trying to restore a guest from state file to avoid Guest from crashing.

0.1.9 Connecting to virtual console of running Xen domain

If we want to connect to virtual console of Xen domain during creation then we can use option '-c' while creating domain. If later we want to connect to virtual console then we can use:

```
xm console {<domain_id>|<domain_name>}
```

Note that to come out of console to get prompt back we have to use telnet escape sequence (Ctrl + ']'). 'Ctrl+C' is transparently passed to the console and hence wont help in coming out of console.

0.2 Domain information

0.2.1 List running domain

To list currently running Xen domains we can use command:

```
xm list
```

The above command also lists state of each running domain. The possible values for state and their meanings are:

State	Meaning
r	Running
b	Blocked - Means either waiting for I/O or is sleeping as nothing left to do
p	Paused - Domain is paused and hence will continue using RAM resources. All files will still remain open. Only domain is not scheduled by hypervisor. Note that if we pause domain for long time then the network connections (like TCP) may time-out and hence when we unpause the state of network connections may change. If network connections are with other VM and they are also paused then timeout should not occur.
s	Shutdown - The guest is requested to Shutdown or reboot or suspended. The process is being carried out.
c	Crashed - The domain has crashed. Usually this state can only occur if the domain has been configured not to restart on crash.
d	Dying - The domain is in process of dying but has not completely shutdown or crashed

0.2.2 Finding domain ID for particular domain name

In case we want to find domain ID for given domain name then we can either look for the ID against name in 'xm list' output or we can use:

```
xm domid <domain_name>
```

to get domain ID for given domain name. This can be very useful if we are writing programs that will manipulate domain and we want our program to be able to find domain ID based on domain name.

0.2.3 Finding domain name for particular domain ID

In case we want to find domain name for given domain ID then we can either look for the name against ID in 'xm list' output or we can use:

```
xm domname <domain_id>
```

to get domain name for given domain ID.

0.2.4 Finding all parameters used for running domain

To find all parameters being used by running domain we can use:

```
xm list --long
```

This will even display attached block devices, network interfaces, values for parameters like 'usb' etc.

Note that:

- We can also look at configuration file for finding the same information. But there is not guarantee that configuration file has not been altered after domain is created. Also we can attach/detach devices from running domain which may not be listed in configuration file. Hence, this method of finding configuration of running domains is more accurate.
- Command 'xenstore-ls' provides us with even more detailed information than 'xm list --long'

0.3 Information about Xen

0.3.1 Xen message buffer

We can see xen message buffer using:

```
xm dmesg
```

If we want to clear the xen message buffer then we can use

```
xm dmesg -c
```

Xen message buffer contains messages generated during Xen boot and also other important warnings / errors that Xen generates during running and hence is very nice tool to help with troubleshooting

0.3.2 Xen information

To print information about current Xen installation including basic information about current host OS, we can use:

```
xm info
```

0.4 Monitoring xen resources

If we want to see usage of various Xen resources like CPU, Block devices, network devices, which domains are running, etc. information in top like view then we can use command

```
xentop
```

In ‘xentop’ we can use:

- ‘d’ ⇒ To change the delay time after which information is refreshed
- ‘n’ ⇒ To include information about network devices / interfaces in output.
- ‘b’ ⇒ To include information about block devices in output
- ‘v’ ⇒ To see information about Virtual CPUs in output
- ‘s’ ⇒ To change sort order for displayed values.

0.5 Tips

Most ‘xm’ commands work in background. Hence, if we get prompt back after running a command it does not mean that old command has completed. We should ensure using ‘xm list’ command whether previous command(s) are finished or not.

0.6 Lab task

1. Perform following operations on Xen HVM domain:
 - (a) Create
 - (b) Destroy
 - (c) Shutdown
 - (d) Reboot
 - (e) Pause
 - (f) Unpause
 - (g) Save
 - (h) Restore
2. Find as much information about running Xen domain as possible.
3. Monitor xen domain performance using ‘xentop’.
4. Start download of big file on Guest with rate limiting. Pause domain before the download completes. Unpause domain once after a minute and in second experiment after long time (≥ 30 minutes). Observe the effect of pause/unpause action on domain on download. Explain the result obtained.

5. Start two different guests with Linux OS. Start 'scp' or remote copying of large file from one guest to other. Pause both domains for long time (≥ 30 minutes). Unpause both domains and observe effect of pause/unpause operation on file copying process. Explain the result obtained.
6. Extend program created in lab 3 with features learned in this lab including:
 - (a) Shutdown for guest
 - (b) Reboot guest
 - (c) Pause guest
 - (d) Unpause guest
 - (e) Save guest state to file
 - (f) Restore guest from state file
 - (g) Connect to guest console
 - (h) Find domain ID for given domain name
 - (i) Find domain name for given domain ID
 - (j) Give detailed info of parameters used by specific running guest.
 - (k) Display xen message buffer
 - (l) Display xen information

Optionally motivated students can learn about migrating domains from <http://www.virtuatopia.com> and try live migration of guest image from one host to other host.

0.7 Reference

Most of this information can be read from man page of 'xm' using 'man xm' command.