

Newton Linux User Group - Linux Virtual Servers

What are Linux Virtual Servers (aka LDS or LVS)?

Linux Virtual Servers allow you to create virtual private servers and security contexts which operate like a normal Linux server, but allow many independent servers to be run simultaneously in one box at full speed. All services, such as ssh, mail, Web, and databases, can be started on such a VPS, without modification, just like on any real server. Each virtual server has its own user account database and root password and doesn't interfere with other virtual servers.

How do we setup Linux Virtual Servers?

Begin by installing a minimal "host" OS. We used Redhat 9 as the Host OS in our example but you should be able to use any other Linux distribution. In addition to Redhat 9, it seems Redhat 9, Fedora Core 1, Debian, and Gentoo are reasonably supported as the Host OS.

Once your Host OS has been installed, you will need to compile and install a recent kernel with patches from the Linux Vserver project. We will discuss this in detail as well.

If you are able to successfully install your Host OS and a patched kernel, you will be well on your way to setting up your first "vserver". In our example, we simply create ONE configuration file for each vserver and extract ONE tarball that contains all the files required for each vserver. We will only use existing vserver images tonight. We will not discuss how to create a vserver from scratch yet.

How do we use Linux Virtual Servers?

Once you have created your vserver(s), you will need to perform an initial configuration for each vserver. Once the initial configuration is complete, you can access your vserver from the console or over the network via tools such as SSH.

Presentation References

LVS Summary:
<http://freshmeat.net/projects/vserver/>

Detailed Installation of Redhat 9 Host OS (Performed at Console):

Boot from your Redhat 9 installation diskette. Have your Redhat 9 driver disk handy.

1. Type the following command at the boot screen to force a Text Mode installation:
linux text

2. Installation Method:
HTTP

3. No Driver Found
Use driver diskette

4. Configure TCP/IP:
IP 192.168.111.15
NM 255.255.255.0
GW 192.168.111.254
NS 12.149.183.1

5. HTTP Setup:
SITE: mirror.datility.net
PATH: /pub/linux/distros/redhat/9/en/os/i386/disc1/

6. Installation Type: CUSTOM

7. Disk Partitioning Setup (via DISK DRUID)

| Path | ReqSize | ActSize | HDA |
|-----------|---------|---------|-----|
| /boot | 128 | 125 | 1 |
| / | 512 | 509 | 9 |
| /home | 512 | 509 | 8 |
| /usr | 3072 | 3074 | 2 |
| /var | 1536 | 1537 | 3 |
| /tmp | 256 | 509 | 7 |
| /opt | 512 | 509 | 6 |
| *swap* | 512 | 509 | 5 |
| /vservers | ~ | 31918 | 10 |

8. Network Config (ETH0)
NS2 12.149.183.4

9. Hostname Configuration:
vs1.datility.net

10. Firewall Configuration
High Security + CUSTOMIZE

11. Firewall Configuration - Custom
Allow Incoming SSH

12. Timezone Selection:
USE GMT / Timzone = America/Chicago

13. Root Password:
<password>

14. Package Group Selection:
* Development Tools
* Kernel Development
* Text-Based Internet

15. Create Boot Diskette:
YES (we recommend you keep your emergency boot disk in or near diskette drive)

Detailed Configuration of Redhat 9 Host OS (Performed via SSH):

1. Connect to Host OS via SSH:

```
ssh 192.168.111.15
```

2. Update Packages on Host OS using APT-GET (recommended):

```
[root@vs1 root]# cd /opt/
[root@vs1 opt]# mkdir apt-get
[root@vs1 opt]# cd apt-get/
[root@vs1 apt-get]# wget http://ftp.freshrpms.net/pub/freshrpms/redhat/9/apt/
apt-0.5.5cnc6-fr1.i386.rpm
[root@vs1 apt-get]# rpm -ivh apt-0.5.5cnc6-fr1.i386.rpm

## Update our local APT database
[root@vs1 apt-get]# apt-get update
## Upgrade our system packages
[root@vs1 apt-get]# apt-get dist-upgrade
```

3. Disable Unnecessary Services on Host OS (recommended):

```
## Display list of services configured to start on boot:
[root@vs1 root]# chkconfig --list | grep 3:on
kudzu          0:off  1:off  2:off  3:on   4:on   5:on   6:off
syslog         0:off  1:off  2:on   3:on   4:on   5:on   6:off
netfs          0:off  1:off  2:off  3:on   4:on   5:on   6:off
network       0:off  1:off  2:on   3:on   4:on   5:on   6:off
random        0:off  1:off  2:on   3:on   4:on   5:on   6:off
rawdevices    0:off  1:off  2:off  3:on   4:on   5:on   6:off
pcmcia        0:off  1:off  2:on   3:on   4:on   5:on   6:off
keytable      0:off  1:on   2:on   3:on   4:on   5:on   6:off
apmd          0:off  1:off  2:on   3:on   4:on   5:on   6:off
atd           0:off  1:off  2:off  3:on   4:on   5:on   6:off
gpm           0:off  1:off  2:on   3:on   4:on   5:on   6:off
autofs        0:off  1:off  2:off  3:on   4:on   5:on   6:off
iptables      0:off  1:off  2:on   3:on   4:on   5:on   6:off
isdn          0:off  1:off  2:on   3:on   4:on   5:on   6:off
sshd          0:off  1:off  2:on   3:on   4:on   5:on   6:off
portmap       0:off  1:off  2:off  3:on   4:on   5:on   6:off
nfslock       0:off  1:off  2:off  3:on   4:on   5:on   6:off
sendmail      0:off  1:off  2:on   3:on   4:on   5:on   6:off
rhnsd         0:off  1:off  2:off  3:on   4:on   5:on   6:off
crond         0:off  1:off  2:on   3:on   4:on   5:on   6:off
anacron       0:off  1:off  2:on   3:on   4:on   5:on   6:off
xinetd        0:off  1:off  2:off  3:on   4:on   5:on   6:off
```

```
## Disable unnecessary services
[root@vs1 root]# chkconfig netfs off
[root@vs1 root]# chkconfig pcmcia off
[root@vs1 root]# chkconfig atd off
[root@vs1 root]# chkconfig gpm off
[root@vs1 root]# chkconfig autofs off
[root@vs1 root]# chkconfig isdn off
[root@vs1 root]# chkconfig portmap off
[root@vs1 root]# chkconfig nfslock off
[root@vs1 root]# chkconfig sendmail off
[root@vs1 root]# chkconfig rhnsd off
[root@vs1 root]# chkconfig anacron off
[root@vs1 root]# chkconfig xinetd off
```

4. Configure and Install NTP on Host OS (recommended):

```
# Allow NTP traffic in firewall rulesets
[root@vs1 root]# vi /etc/sysconfig/iptables
-A RH-Lokkit-0-50-INPUT -p udp -m udp -s 12.149.183.243 --sport 123 -d 0/0 -j
ACCEPT
-A RH-Lokkit-0-50-INPUT -p udp -m udp -s 12.149.183.245 --sport 123 -d 0/0 -j
ACCEPT
[root@vs1 root]# service iptables restart

[root@vs1 root]# apt-get install ntp
[root@vs1 root]# ntpdate tick.datility.net
[root@vs1 root]# cp /etc/ntp.conf /etc/ntp.conf.orig
[root@vs1 root]# vi /etc/ntp.conf # see ntp.conf.txt
[root@vs1 root]# service ntpd restart
```

5. Configure SYSLOG to display "messages" to TTY12 (optional):

```
[root@vs1 root]# vi /etc/syslog.conf
*.info;mail.none;authpriv.none;cron.none /dev/tty12
[root@vs1 root]# service syslog restart
```

6. Configure Firewall to allow SSH connections from specific hosts (recommended):

```
[root@vs1 root]# vi /etc/sysconfig/iptables
-A RH-Lokkit-0-50-INPUT -p tcp -m tcp -s 10.10.10.0/24 --dport 22 --syn -j
ACCEPT
-A RH-Lokkit-0-50-INPUT -p tcp -m tcp -s 192.168.1.0/24 --dport 22 --syn -j
ACCEPT
[root@vs1 root]# service iptables restart
```

7. Double-check to be sure all packages on Host OS are current (optional):

```
[root@vs1 root]# apt-get update
[root@vs1 root]# apt-get upgrade
[root@vs1 root]# apt-get dist-upgrade
```

Detailed Config/Compile/Install of Kernel on Redhat 9 Host OS (Performed via SSH):

1. Download Kernel Source from Kernel.org and Vserver from Linux-Vserver.org

```
[root@vs1 root]# cd /opt/vserver/
[root@vs1 vserver]# ls
linux-2.4.26.tar.bz2      patch-2.4.26-vs1.27.diff.bz2
linux-2.4.26.tar.bz2.sign  split-2.4.26-vs1.27.tar.bz2
linux-vserver-1.27.tar.bz2  util-vserver-0.29.4.tar.bz2
```

2. Patch Kernel and Prepare for Kernel Configuration

```
## VERIFY KERNEL SOURCES
[root@vs1 vserver]# gpg -\-keyserver wwwkeys.pgp.net -\-recv-keys 0x517D0F0E
[root@vs1 vserver]# gpg --verify linux-2.4.26.tar.bz2.sign linux-2.4.26.tar.bz2

## EXTRACT KERNEL SOURCES
[root@vs1 vserver]# cd /usr/src
[root@vs1 src]# tar -jxvf /opt/vserver/linux-2.4.26.tar.bz2

## VSERVER SUPPORT PATCH
[root@vs1 src]# tar -jxvf /opt/vserver/linux-vserver-1.27.tar.bz2
[root@vs1 src]# patch -p0 < patch-2.4.26-vs1.27.diff
```

```
## CONTEXT QUOTA SUPPORT PATCH
[root@vs1 src]# cp /opt/vserver/patch-2.4.25-vs1.27-q0.14.diff.bz2 .
[root@vs1 src]# bunzip2 patch-2.4.25-vs1.27-q0.14.diff.bz2
[root@vs1 src]# ln -s linux-2.4.26/ linux-2.4.25-vs1.27
[root@vs1 src]# patch -p0 < patch-2.4.25-vs1.27-q0.14.diff
```

3. Copy Redhat 9 kernel config, configure/compile kernel for use with vserver:

```
[root@vs1 root]# cd /usr/src/linux-2.4.26/

[root@vs1 linux-2.4.26]# make mrproper
[root@vs1 linux-2.4.26]# make oldconfig

[root@vs1 linux-2.4.26]# cp /usr/src/linux-2.4/configs/linux-.config .config
[root@vs1 linux-2.4.26]# make menuconfig

* Block Devices / Virtual Root Device Support YES
* File Systems / Persistent Context ID UID24/GID24
* File Systems / Quota Support YES
* Quota / VFS v0 Quota Format Support YES

# NOTE: I tweaked the heck out of my kernel. You may need
# to change a few options if you receive errors while
# compiling your kernel. For example, you may need
# to compile EXT3 Support and Quota Support INTO the
# kernel <*> rather than as a module <M> or else you
# may get strange errors. For example...
#
# make modules_install gave me this before I compiled those
# two items INTO my kernel
#
# if [ -r System.map ]; then /sbin/depmod -ae -F System.map 2.4.24-vs1.24; fi
# depmod: *** Unresolved symbols in /lib/modules/2.4.24-
vs1.24/kernel/fs/ext3/ext3.o
# depmod: get_dqhash
# depmod: dlimit_transfer
# depmod: dqhash_transfer
# depmod: destroy_dqhash

[root@vs1 linux-2.4.26]# make dep
[root@vs1 linux-2.4.26]# make clean
[root@vs1 linux-2.4.26]# make bzImage
[root@vs1 linux-2.4.26]# make modules
[root@vs1 linux-2.4.26]# make modules_install
[root@vs1 linux-2.4.26]# make install

## NOTES REGARDING KERNEL COMPILE STEPS
# make mrproper # (cleans up older attempts not need for clean source)
# make oldconfig # (creates a basic configuration from running kernel maybe
# used before or in-place of step 1)
# make menuconfig # (or config, or xconfig)
# make dep # (2.2/4 kernels only not used with 2.6 kernels)
# make clean
# make bzImage
# make modules # (not needed for monolithic kernels)
# make modules_install # (not needed for monolithic kernels)
# make install # (takes the place of hand copying and editing)

## See this URL for more information
## http://www.redhat.com/docs/manuals/linux/RHL-9-Manual/custom-guide/s1-
custom-kernel-modularized.html
```

4. Reconfigure FSTAB and GRUB for use with new kernel:

```
## OUR NEW KERNEL DOES NOT SUPPORT VOLUME LABELS
## REPLACE VOLUME LABELS IN FSTAB WITH VOLUME DEVICES

[root@vs1 root]# df
Filesystem            1K-blocks      Used Available Use% Mounted on
/dev/hda9              505605        80374   399127   17% /
/dev/hda1             124427         9323   108680    8% /boot
/dev/hda8              505605        8239   471262    2% /home
/dev/hda6              505605       39614   439887    9% /opt
none                  111504          0   111504    0% /dev/shm
/dev/hda7              505605        9686   469815    3% /tmp
/dev/hda2             3099292     1056792  1885064   36% /usr
/dev/hda3             1548096     129796  1339584    9% /var
/dev/hda10            32171064     32828  30504028    1% /vservers

[root@vs1 root]# vi /etc/fstab
/dev/hda9              /                ext3      defaults      1 1
/dev/hda1             /boot            ext3      defaults      1 2
none                  /dev/pts         devpts    gid=5,mode=620 0 0
/dev/hda8             /home            ext3      defaults      1 2
/dev/hda6             /opt             ext3      defaults      1 2
none                  /proc            proc       defaults      0 0
none                  /dev/shm         tmpfs     defaults      0 0
/dev/hda7             /tmp             ext3      defaults      1 2
/dev/hda2             /usr             ext3      defaults      1 2
/dev/hda3             /var             ext3      defaults      1 2
/dev/hda10            /vservers        ext3      defaults      1 2
/dev/hda5             swap             swap       defaults      0 0
/dev/cdrom            /mnt/cdrom       udf,iso9660 noauto,owner,kudzu,ro 0 0
/dev/fd0              /mnt/floppy      auto      noauto,owner,kudzu 0 0

## OUR NEW KERNEL DOES NOT SUPPORT VOLUME LABELS
## REPLACE VOLUME LABELS IN GRUB.CONF WITH VOLUME DEVICES

[root@vs1 root]# vi /etc/grub.conf

default=0
timeout=5
splashimage=(hd0,0)/grub/splash.xpm.gz
title Red Hat Linux (2.4.26-vs1.27)
    root (hd0,0)
    kernel /vmlinuz-2.4.26-vs1.27 ro root=/dev/hda9
    initrd /initrd-2.4.26-vs1.27.img
title Red Hat Linux (2.4.20-8)
    root (hd0,0)
    kernel /vmlinuz-2.4.20-8 ro root=/dev/hda9
    initrd /initrd-2.4.20-8.img
```

5. Install Vserver Utilities:

```
[root@vs1 root]# cd /opt/vserver/
[root@vs1 vserver]# rpmbuild -tb util-vserver-0.29.4.tar.bz2
error: Failed build dependencies:
    e2fsprogs-devel is needed by util-vserver-0.29.4-0
[root@vs1 vserver]# apt-get install e2fsprogs-devel
[root@vs1 vserver]# rpmbuild -tb util-vserver-0.29.4.tar.bz2
[root@vs1 vserver]# rpm -ivh /usr/src/redhat/RPMS/i386/util-vserver-0.29.4-0.i386.rpm
```

6. Configure Vserver and Prepare Host OS:

```
## VSERVER SHOULD START IN BACKGROUND (SEE TTY9)
[root@vs1 root]# vi /etc/vservers.conf
BACKGROUND=yes

## CONFIGURE PERMISSIONS ON VSERVICES DIRECTORY
[root@vs1 /]# chmod 000 /vservers/

## Network Services on Host OS MUST bind to specific IP.  If any services bind
## to 0.0.0.0, they must be reconfigured!

## SSH is the only service listening.  It must be reconfigured.
[root@vs1 root]# netstat -an | grep LISTEN
tcp        0      0 0.0.0.0:22          0.0.0.0:*           LISTEN

## Configure SSH to bind to primary IP only!
[root@vs1 root]# vi /etc/ssh/sshd_config
ListenAddress 192.168.111.15

## SSH is the only service listening.  It is now configured properly.
[root@vs1 root]# netstat -an | grep LISTEN
tcp        0      0 192.168.111.15:22  0.0.0.0:*           LISTEN

## Create directory for individual vserver configurations
[root@vs1 root]# mkdir /etc/vservers
```

7. Create your first Virtual Server:

```
## CREATE CONFIG FILE FOR VSERVER

[root@vs1 root]# cd /etc/vservers
[root@vs1 vservers]# vi redhat9.conf
S_CONTEXT=100
IPROOT=192.168.111.16
S_HOSTNAME=redhat9.vs1.datility.net
ONBOOT="no"

## CREATE SCRIPT FOR VSERVER (RUN AT VSERVER STARTUP / STOP)

[root@vs1 vservers]# vi redhat9.sh
#!/bin/sh

case $1 in
    pre-start)
        echo "Hi!  Vserver $2 is going to start!"
        ;;
    post-start)
        echo "Hi!  Vserver $2 is done starting!"
        ;;
    pre-stop)
        echo "Hi!  Vserver $2 is going to stop!"
        ;;
    post-stop)
        echo "Hi!  Vserver $2 is done stopping!"
        ;;
esac

[root@vs1 vservers]# chmod 755 redhat9.sh
```

```
## Configure IP address on Host OS:

[root@vs1 root]# cd /etc/sysconfig/network-scripts/
[root@vs1 network-scripts]# cp ifcfg-eth0 ifcfg-eth0:0
[root@vs1 network-scripts]# vi ifcfg-eth0:0
    DEVICE=eth0:0
    BOOTPROTO=static
    IPADDR=192.168.111.16
    NETMASK=255.255.255.0
    ONBOOT=yes

[root@vs1 network-scripts]# service network restart

## Extract Filesystem for Virtual Server and Enter Virtual Server:
[root@vs1 vservers]# cd /vservers/
[root@vs1 vservers]# mkdir redhat9
[root@vs1 vservers]# cd redhat9
[root@vs1 redhat9]# tar -jxvf /vservers/templates/redhat90.tar.bz2

[root@vs1 root]# vserver redhat9 start
[root@vs1 root]# vserver redhat9 enter

## Configure Virtual Server
[root@vserver:redhat9 /]# vi /etc/resolv.conf
    nameserver 12.149.183.1
    nameserver 12.149.183.4

[root@vserver:redhat9 /]# vi /etc/hosts
    127.0.0.1      redhat9.vsl.datility.net

[root@vserver:redhat9 /]# chkconfig httpd on
[root@vserver:redhat9 /]# chkconfig syslog on
[root@vserver:redhat9 /]# chkconfig crond on
[root@vserver:redhat9 /]# exit

## Restart Virtual Server
[root@vs1 network-scripts]# vserver redhat9 stop
[root@vs1 root]# vserver redhat9 start
[root@vs1 root]# vserver redhat9 enter

## Upgrade Packages in Virtual Server:
[root@vserver:redhat9 /]# apt-get update
[root@vserver:redhat9 /]# apt-get dist-upgrade

## Configure Timezone in Virtual Server:
[root@vs1 root]# cp /usr/share/zoneinfo/America/Chicago
/vservers/redhat9/etc/localtime
```

Linux Vserver References:

Linux Vserver Website: <http://www.linux-vserver.org/>
Linux Vserver Downloads: <http://www.13thfloor.at/vserver/project/>
Jacques Gélinas Vserver FAQ: <http://www.solucorp.qc.ca/howto.hc?projet=vserver>

<plug class="shameless">
Linux Vserver Hosting @ DNI: <http://www.datility.net/>
Presentation Downloads: <http://mirror.datility.net/pub/linux/vserver/>
</plug>