VYATTA, INC. | Vyatta System

Quick Start Guide

Vyatta
Suite 200
1301 Shoreway Road
Belmont, CA 94002
vyatta.com
650 413 7200
1 888 VYATTA 1 (US and Canada)

COPYRIGHT

Copyright © 2005–2010 Vyatta, Inc. All rights reserved.

Vyatta reserves the right to make changes to software, hardware, and documentation without notice. For the most recent version of documentation, visit the Vyatta web site at vyatta.com.

PROPRIETARY NOTICES

Vyatta is a registered trademark of Vyatta, Inc.

VMware, VMware ESXi, and VMware Server are trademarks of VMware, Inc.

XenServer and XenCenter are trademarks of Citrix Systems, Inc.

All other trademarks are the property of their respective owners.

ISSUE DATE: August 2010

DOCUMENT REVISION. R6.1 v02

RELEASED WITH: R6.1.0

PART NO. A0-0090-10-0021

Contents

Welcome **1**

This Guide 1

Additional Resources 2

Chapter 1 Deploying the Vyatta System **3**

Deployment Options 3

About LiveCDs 4

Creating a LiveCD 5

Running from LiveCD 5

Image-Based Install 6

Disk-Based Install 7

Sample Disk-Based Install Session 8

Specifying the Boot Device in the System BIOS 9

Installing on VMware 9

Installing on XenServer 11

Installing XenServer and XenCenter 11

Installing the Vyatta Template 11

Creating a Vyatta Virtual Machine 12

Operating a Vyatta Virtual Machine 13

Testing Your Installation 13

Verify Release and System Type 13

Verify Connectivity 14

Chapter 2 User Interface Options **16**

The Command-Line Interface (CLI) 16

The Vyatta CLI and the System Shell 17

Command Modes 17

CLI Help 18

Command Completion 19

Viewing Long Output (“More”) 20
Showing Support Information 21

The Graphical User Interface (GUI) 21
Enabling Access to the Web GUI 22
Logging in to the Web GUI 22
Logging out of the Web GUI 24
Layout of the Web GUI Screen 24
Navigational Aids 25
Interpreting CLI Documentation for the Web GUI 27

Chapter 3 Configuration 30

Configuration Basics 30
Configuration Hierarchy 31
Adding and Modifying Configuration 31
Deleting Configuration 32
Committing Configuration Changes 33
Discarding Configuration Changes 34
Saving Configuration 35
Loading Configuration 36
Changing the Default Configuration File 37

Scenario: Basic System Configuration 37
Logging On 38
Entering Configuration Mode 39
Setting the Host Name 39
Setting the Domain Name 40
Changing Passwords 40
Configuring Interfaces 40
Configuring Access to a DNS server 41
Specifying a Default Gateway 42
Performing a Package Upgrade to the Latest Version 42

Scenario: Internet Gateway 42
Configuring Interfaces 43
Enabling SSH Access 44
Configuring DHCP Server 45
Configuring NAT 46
Configuring Firewall 46

Welcome

Thank you for choosing the Vyatta system.

Vyatta has changed the networking world by developing the first commercially supported, open-source router/firewall/VPN solution to provide an alternative to over-priced, inflexible products from proprietary vendors. Vyatta solutions offer industry-standard routing and management protocols, support for most commonly used network interfaces, and configuration via command-line interface (CLI) or graphical user interface (GUI).

Vyatta delivers the features, performance, and reliability of an enterprise-class secure router with the added benefits of flexible deployment options—x86 hardware, blade servers, virtualization—freedom to integrate applications, and the economic advantages of commodity hardware and components.

This Guide

This document is intended to:

- Provide an overview of the deployment options, and help you determine the best way to deploy the Vyatta System in your environment
- Provide an overview of the user interface options available on the system
- Walk you through basic configuration of the system based on a sample scenario

Additional Resources

More information about the Vyatta system is available in the Vyatta technical library, and on www.vyatta.com and www.vyatta.org.

Full product documentation is provided in the Vyatta technical library. To see what documentation is available for your release, see the *Guide to Vyatta Documentation*. This guide is posted with every release of Vyatta software and provides a great starting point for finding the information you need.

Chapter 1: Deploying the Vyatta System

The Vyatta System supports a number of flexible deployment options. This chapter describes methods for performing a new install of the Vyatta system.

In this chapter:

- Deployment Options
- About LiveCDs
- Image-Based Install
- Disk-Based Install
- Sample Disk-Based Install Session
- Specifying the Boot Device in the System BIOS
- Installing on VMware
- Installing on XenServer
- Testing Your Installation

Deployment Options

The Vyatta system supports a number of flexible deployment options. This section describes each option so that you can determine the best option for your needs and environment and explains how to deploy in each case:

- **Run on LiveCD.** You can run the Vyatta system directly from a LiveCD. A LiveCD runs the Vyatta software on a RAM disk on the host machine, without disturbing or changing the existing operating system on the host machine. This method of deployment is best suited for evaluation and test scenarios though a LiveCD is also required as a first step in installing to a persistent hardware device like a hard disk. You can read more about LiveCDs here.
- **Install on physical hardware.** The Vyatta system can be installed and run on most standard x86 servers and PCs. The system can be installed onto a variety of persistent devices, including hard drive, USB memory stick, and compact Flash.

  NOTE *Not all hardware supports the ability to boot from USB device or Flash; check the BIOS of your hardware to see if yours does.*

  There are two types of installations that can be performed on a persistent device:

  — **Image-based install.** The simplest, most flexible, and most powerful way to install a Vyatta system is using a binary system image. With this method, you can install multiple versions of the Vyatta system as images and switch between the images simply and easily. You install the image from a LiveCD then you reboot your system and it runs the image. You can read about this installation method here.

  — **Disk-based install.** A disk-based install also installs from a LiveCD onto a persistent device such as a hard disk partition. However, unlike an image-based install, a disk-based install uses a traditional layout of files on the disk. Additional system images may be added at a later time to a system created using a disk-based install. You can read about this installation method here.
- **Install into a virtualized environment or cloud.** Alternatively, the Vyatta system can be run as a software appliance in a virtual or cloud computing environment, allowing you to virtualize your network. The Vyatta software appliance has been optimized for XenServer and VMware virtual environments. These platforms provide a great deal of flexibility as to how the virtual machine can be configured, especially with respect to memory and Ethernet interfaces. Multiple Vyatta systems can be run simultaneously on a single hardware platform configured for multiple virtual machines.

  — **VMware**. Instructions for installing the Vyatta system on VMware can be found in here.

  — **XenServer**. Instructions for installing the Vyatta system on XenServer can be found in here.

About LiveCDs

A LiveCD runs the Vyatta software on a RAM disk on the host machine. The system uses the RAM disk for writeable sections of the file system, and uses the floppy drive or a TFTP server for saving configuration files.

The LiveCD can run on a machine with an existing operating system without disturbing or changing the previously installed operating system. Configuration is not stored on the system, but you can save configuration to a persistent device such as a USB memory stick. Also, you cannot upgrade an installation run from LiveCD; each upgrade requires a fresh LiveCD. While you are running the system from a LiveCD, you will not be able to access other applications or programs on your machine.

A LiveCD is also used to install the Vyatta system to a persistent device such as a hard disk using either an image-based install (described in Image-Based Install) or a disk-based install (described in Disk-Based Install).

Creating a LiveCD

To create a LiveCD, download one of the ISO software images directly from the Vyatta web site.

- Vyatta subscription customers use: http://www2.vyatta.com/support/software/.
- Vyatta community customers use: http://www.vyatta.com/downloads/.

Then use CD-burning software to create a bootable ISO image. Note that:

- The CD must be an ISO image: it won't work to just copy files onto the CD.
- The CD must be bootable. See the documentation for your CD-burning utility for information on how to burn an ISO image.

Running from LiveCD

▶ To get ready

- Make sure your BIOS is configured to boot from CD-ROM. For details, see Specifying the Boot Device in the System BIOS.

▶ To boot from LiveCD

1. With the system powered down, connect a VGA monitor and keyboard to it.
2. Insert the Vyatta LiveCD into the CD drive and close the drive.
3. Power up the system.

   NOTE *Systems with a Flash device may boot slowly from LiveCD, as the system looks for DMA that does not exist on the IDE device. This is expected behavior.*

   After the startup messages complete, the **Vyatta login** prompt appears:

   ```
   vyatta login:
   ```

   This prompt means that your system is running.

If you want to simply run the system from the LiveCD, go to the section Testing Your Installation to confirm that you have booted the correct system and can access the system from your network.

If you want to use the LiveCD to perform an image-based install, go to the section Image-Based Install.

If you want to use the LiveCD to perform a disk-based install, go to the section Disk-Based Install.

Image-Based Install

To perform an image-based install, you run the **install-image** command from a running LiveCD system. The system can be installed to any persistent device, such as a hard disk, USB stick, or Flash card.

The **install-image** command uses the LiveCD as the source of the image, formats the device that the system will be installed to, installs the system, and attempts to preserve the Vyatta configuration from a previous installation if one is found.

The installed system is run the next time the system reboots.

▶ To get ready

- Make sure you have a minimum of 1 GB of free space on your hard disk, USB stick, or Flash card for a root partition. A minimum of 2 GB of free space is recommended for a production installation.
- Find out the device name onto which the software is to be installed.

  NOTE *The installation process reformats the specified partition, so it is critical to specify the correct partition. For systems dedicated to Vyatta, the system selected defaults can be used. If you are installing onto a USB memory stick , insert the memory stick into a USB port. The system will detect that the USB device is installed and notify you; record the name of the device as displayed in the system message (for example,* ***/dev/sdb****). Specify this device when prompted during the installation procedure.*
- If you want to install onto a USB stick, check the BIOS setup of your hardware to confirm that it supports the ability to boot from a USB device.

▶ To install

1. Create and run the Vyatta LiveCD. For this information, see About LiveCDs.
2. Login as user **vyatta** with password **vyatta** and run the **install-image** command. To see sample system feedback for this command, see Sample Disk-Based Install Session.
3. Remove the LiveCD.
4. Make sure your BIOS is configured to boot from the persistent device the system was installed to. For this information, see Specifying the Boot Device in the System BIOS.

5 Reboot the system.

6 Test the installation. For this information, see Testing Your Installation.

Disk-Based Install

To perform an disk-based install, you run the **install-system** command from a running LiveCD system. The system can be installed to any persistent device, such as a hard disk, USB stick, or Flash card.

The **install-system** command uses the LiveCD as the source of the image, formats the device that the system will be installed to, installs the system, and attempts to preserve the Vyatta configuration from a previous installation if one is found.

The installed system is run the next time the system reboots.

▶ To get ready

- Make sure you have a minimum of 1 GB of free space on your hard disk, USB stick, or Flash card for a root partition. A minimum of 2 GB of free space is recommended for a production installation.
- Find out the device name onto which the software is to be installed.

  NOTE *The installation process reformats the specified partition, so it is critical to specify the correct partition. For systems dedicated to Vyatta, the system selected defaults can be used. If you are installing onto a USB memory stick , insert the memory stick into a USB port. The system will detect that the USB device is installed and notify you; record the name of the device as displayed in the system message (for example,* ***/dev/sdb****). Specify this device when prompted during the installation procedure.*
- If you want to install onto a USB stick, check the BIOS setup of your hardware to confirm that it supports the ability to boot from a USB device.

▶ To install

1 Create and run the Vyatta LiveCD. For this information, see About LiveCDs.

2 Login as user **vyatta** with password **vyatta** and run the **install-system** command. To see sample system feedback for this command, see Sample Disk-Based Install Session.

3 Remove the LiveCD.

4 Make sure your BIOS is configured to boot from the persistent device the system was installed to. For this information, see Specifying the Boot Device in the System BIOS.

5 Reboot the system.

6 Test the installation. For this information, see Testing Your Installation.

Sample Disk-Based Install Session

The following is a sample disk-based installation session using **install-system**. Image-based installations using **install-image** are virtually identical.

Example 1-1 Sample install session

```
Last login: Mon Oct 29 15:37:50 2007
Linux vyatta 2.6.20 #1 SMP Fri Sep 21 02:22:08 PDT 2007 i686
Welcome to Vyatta.
This system is open-source software. The exact distribution terms
for
each module comprising the full system are described in the
individual
files in /usr/share/doc/*/copyright.
vyatta:~# install-system
Welcome to the Vyatta install program.  This script
will walk you through the process of installing the
Vyatta image to a local hard drive.

Would you like to continue? (Yes/No) [Yes]: <Enter>
Probing drives: OK
Looking for pre-existing RAID groups...none found.
The Vyatta image will require a minimum 1000MB root.
Would you like me to try to partition a drive automatically
or would you rather partition it manually with parted?  If
you have already setup your partitions, you may skip this step.

Partition (Auto/Parted/Skip) [Auto]: <Enter>

I found the following drives on your system:
 sda    1074MB


Install the image on? [sda] <Enter>

This will destroy all data on /dev/sda.
Continue? (Yes/No) [No]: Yes

How big of a root partition should I create? (1000MB - 1074MB)
[1074]MB: <Enter>

Creating filesystem on /dev/sda1: OK
Mounting /dev/sda1
Copying system image files to /dev/sda1:OK
I found the following configuration files
/opt/vyatta/etc/config/config.boot
```

```
Which one should I copy to sda?
[/opt/vyatta/etc/config/config.boot]: <Enter>

Enter password for administrator account
Enter vyatta password:vyatta
Retype vyatta password:vyatta
I need to install the GRUB boot loader.
I found the following drives on your system:
 sda    1074MB

Which drive should GRUB modify the boot partition on? [sda]: <Enter>

Setting up grub: OK
Done!
vyatta:~#
```

Specifying the Boot Device in the System BIOS

Prior to booting, you should make sure that the system BIOS is set to boot from the device that you expect it to.

▶ To change the boot device

1. During the boot sequence press the appropriate key sequence (for example, <F2>) to interrupt the boot sequence and enter your system's BIOS setup program.
2. In the boot sequence menu arrange the boot devices such that the device name (for example, "USB HDD" for a USB stick or "IDE" or "SATA" for a compact Flash) is first in the list so that the system will attempt to boot from it first.
3. Save the settings and reboot the system. When the system restarts, it will boot from the device specified.

Installing on VMware

Currently, the Vyatta system supports running on VMware ESX and ESXi.

Like other virtualization platforms, VMware products provide the ability to run multiple virtual systems on a single hardware platform. VMware ESX and ESXi are virtualization platforms that run directly on system hardware in a 64-bit environment.

Vyatta provides a prebuilt VMware virtual appliances that can be run as a virtual machine on VMware ESX or ESXi. You can obtain the Vyatta VMware virtual appliances from the Vyatta web site.

- Vyatta subscription customers use: http://www2.vyatta.com/support/software/.

- Vyatta community customers use: http://www.vyatta.com/downloads/.

Before You Begin

Before installing the Vyatta system, VMware must be installed on a server and be operational. VMware ESX or ESXi and accompanying documentation can be obtained from VMware at http://www.vmware.com.

The installation procedures provided assume an operational VMware system.

Installing on VMware ESX/ESXi 4

▶ To install the Vyatta VMware virtual appliance on VMware ESX/ESXi 4

1. Copy the URL for the ESX 4 OVF file from the Vyatta web site.
   - Vyatta Subscription Edition customers use http://www2.vyatta.com/support/software/.
   - Vyatta community users use http://www.vyatta.com/downloads/.
2. Open the vSphere Client GUI.
3. From the File menu, select **Deploy OVF Template** > **Deploy from URL**.
4. Paste the URL for the OVF file in the field provided.
5. Follow the deployment dialog customizing for the local site. The virtual appliance can be run within VMware ESX/ESXi 4 like any other virtual machine.
6. Start the virtual machine.
7. Test the installation.

▶ I need more information on

Testing Your Installation

Installing on VMware ESX/ESXi 3

▶ To install the Vyatta VMware virtual appliance on VMware ESX/ESXi 3

1. Copy the URL for the ESX 3 OVF file from the Vyatta web site.
   - Vyatta Subscription Edition customers use http://www2.vyatta.com/support/software/.
   - Vyatta community users use http://www.vyatta.com/downloads/.
2. Open the VMware Infrastructure Client GUI.

3. From the File menu, select **Virtual Appliance** > **Import** > **Import from URL**.
4. Paste the URL for the OVF file in the field provided.
5. Follow the deployment dialog customizing for the local site. The virtual appliance can be run within VMware ESX/ESXi 3 like any other virtual machine.
6. Start the virtual machine.
7. Test the installation.

▶ **I need more information on**

Testing Your Installation

Installing on XenServer

Citrix XenServer is a server virtualization platform. Like other virtualization platforms, XenServer provides the ability to run multiple virtual machines on a single hardware platform.

XenCenter is the management system for XenServer. XenCenter allows you to manage multiple physical servers running XenServer, in addition to all the virtual machines running on each of the physical servers.

Virtual machines are created on XenServer using virtual machine templates. For each release of its software, Vyatta provides a XenServer virtual machine template that can be used to create virtual machines running the Vyatta system.

The Vyatta template is designed to work with XenServer 5.5 and XenCenter 5.5. To use the template, you must:

- Download the template.
- Import the template into XenServer.
- Create a Vyatta virtual machine from the template.

Installing XenServer and XenCenter

Before installing the Vyatta template, XenServer 5.5 must be installed on a server and XenCenter 5.5 must be installed on a management PC. Instructions for downloading and installing these products can be found at http://www.citrix.com.

Installing the Vyatta Template

In order to allow Vyatta system virtual machines to run on XenServer 5.5, you must install the Vyatta XenServer template.

▶ To install the Vyatta XenServer template

1. Download the compressed Vyatta XenServer template file from the Vyatta web site to a location accessible by your XenCenter machine.

- Vyatta subscription customers use: http://www2.vyatta.com/support/software/.
- Vyatta community customers use: http://www.vyatta.com/downloads/.

2. Extract the Vyatta XenServer template file from the compressed download file.
3. Within XenCenter, select **Templates** > **Import Template**.
4. On the Import Source screen, browse to the extracted template file and select **Exported Template** as the **Import Type**. Click Next; the Home Server screen opens.
5. On the Home Server screen, select the server to be used as the home server for any virtual machines created from the template. Click Next; the Storage screen opens.
6. On the Storage screen, select the storage repository where virtual disks for the template will be stored. Click Import and wait for the template to be imported. The Network screen opens.
7. On the Network screen, set up any required network interfaces. Click Next; the Finish screen opens.
8. On the Finish screen, click Finish to complete the import process. The imported Vyatta template appears in the left pane of XenCenter.

Creating a Vyatta Virtual Machine

Once the Vyatta template is loaded on the XenServer machine, you can create a virtual machine from the template.

▶ To create a Vyatta virtual machine

1. Within XenCenter, select the XenServer that you wish to add the new virtual machine to.
2. Click the **New VM** button. The Template screen opens.
3. From the Template screen, select the Vyatta template that imported. Click Next; the Name screen opens.
4. On the Name screen, enter a name and a description for the new virtual machine. Click Next; the Location screen opens.
5. On the Location screen, select **Empty DVD Drive**. Click Next; the CPU & Memory screen opens.
6. On the CPU & Memory screen, specify the number of CPUs and the amount of memory that the virtual machine is to use. Click Next; the Virtual Disks screen opens.
7. On the Virtual Disks screen, specify a virtual disk for the Vyatta virtual machine. Click Next; the Virtual Interfaces screen opens.

8 On the Virtual Interfaces screen, specify the virtual network interfaces to be installed on the Vyatta virtual machine. Click Next; the Finish screen opens.

9 On the Finish screen, select Start VM automatically if you would like the Vyatta virtual machine to start when the XenServer machine starts. Click Finish. The new virtual machine appears in the left pane of XenCenter.

Operating a Vyatta Virtual Machine

Once a Vyatta virtual machine has been created on XenServer, the virtual machine can be operated from XenCenter in the same way as any other virtual machine on XenServer.

▶ To operate the Vyatta virtual machine

1 Select the Vyatta virtual machine in the left pane of XenCenter. Click the **Start** button.

2 Select the Console tab to access the Vyatta console.

3 Test the installation. For this information, see Testing Your Installation.

Testing Your Installation

Once the system has successfully booted you will see the **vyatta login:** prompt. This indicates that the system is operational. You should:

- Verify Release and System Type
- Verify Connectivity

Verify Release and System Type

The next step is to confirm that the correct release is running and it is running on the device that you expect.

▶ To verify release and system type

1 Login as user **vyatta** with password **vyatta** (default login ID).

2 Run the **show version** command, as in the following example.

Example 1-2 Displaying version information

```
vyatta@vyatta:~$ show version
Version   :    999.kenwoodse.02200016
Copyright:     2006-2010 Vyatta, Inc.
Built by :     autobuild@vyatta.com
Built on :     Sat Feb 20 08:16:24 UTC 2010
Build ID :     1002200816-658521f
```

```
Boot via :    image
Uptime   :    01:25:29 up 9 days,  4:29,  2 users,  load average:
0.00, 0.00, 0.00
vyatta@vyatta:~$
```

The **Version:** line shows the version number of the running system. The **Boot via:** line shows the type of system that is running:

- **livecd** — The system booted from a LiveCD.
- **image** — The system booted from a persistent device and is running as an image-based system
- **disk** — The system booted from a persistent device and is running as an disk-based system.

Verify Connectivity

The final step is to confirm that the Vyatta system can be accessed on the local network. A quick and easy way to do this is to configure an Ethernet interface on the system and then "ping" the interface from another host on the network.

NOTE *Make sure that the system is physically connected to the network first.*

▶ To test system connectivity

1. At the command prompt, enter the commands shown in the example, substituting an IP address on your existing subnet. In our example:
   - The subnet is 192.168.1.0/24
   - The IP address of the interface is 192.168.1.81

   Make the appropriate substitutions for your network, as in the following example.

Example 1-3 Configuring a test Ethernet interface

```
vyatta@vyatta:~$ configure
[edit]
vyatta@vyatta# set interfaces ethernet eth0 address 192.168.1.81/24
[edit]
vyatta@vyatta# commit
[edit]
vyatta@vyatta# exit
exit
vyatta@vyatta:~$
```

2. From another host on the same subnet, ping the interface to ensure that it is up.
3. From a Linux or Windows command prompt, enter the following command (substituting the IP address you assigned to the interface):

```
ping 192.168.1.81
```

If the Vyatta system is reachable, you will see replies from it in response to the pings. If so, your system is installed and accessible on your network.

Chapter 2: User Interface Options

There are two ways you can interface with the Vyatta system: a command-line interface (CLI) and a web-based graphical user interface (GUI).

- The CLI is similar to those found on closed-source routers you may be familiar with.
- The web GUI provides an easy-to-use alternative for those who prefer a GUI to a CLI.

This chapter provides a brief introduction to both interfaces.

The remainder of this guide provides examples using only the CLI.

This chaptersection presents the following topics:

- The Command-Line Interface (CLI)
- The Graphical User Interface (GUI)

The Command-Line Interface (CLI)

This section presents the following topics:

- The Vyatta CLI and the System Shell
- Command Modes
- CLI Help
- Command Completion

- The “?” key, like the <Tab> key, will provide help strings for commands.
- Showing Support Information

The Vyatta CLI and the System Shell

The CLI of the Vyatta system includes two families of commands:

- Vyatta-specific commands for operating and configuring the Vyatta system.
- Commands provided by the Linux operating system shell in which the Vyatta CLI operates.

All Vyatta users have access to operating system commands as well as Vyatta CLI commands. The amount of access to the operating system commands varies with the privilege level of the user.

The operating system commands and constructs are available from any point within the Vyatta CLI.

Command Modes

There are two command modes in the Vyatta CLI: operational mode and configuration mode.

- Operational mode provides access to operational commands for showing and clearing information and enabling or disabling debugging, as well as commands for configuring terminal settings, loading and saving configuration, and restarting the system.
- Configuration provides access to commands for creating, modifying, deleting, committing and showing configuration information, as well as commands for navigating through the configuration hierarchy.

When you log on to the system, the system is in operational mode.

- To enter configuration mode from operational mode, issue the **configure** command.
- To return to operational mode from configuration mode, issue the **exit** command. If there are uncommitted configuration changes, you must either commit the changes using the **commit** command, or enter **exit discard** to discard the changes before you can exit to operational mode.

Issuing the **exit** command in operational mode logs you out of the system.

▶ Try it

▶ **Enter configuration mode**

In configuration mode you can **set**, **delete**, and **show** information. Enter configuration mode by typing **configure** at the command prompt in operational mode.

```
vyatta@vyatta:~$ configure
[edit]
vyatta@vyatta#
```

Notice how the command prompt changes to remind you what mode you are in.

▶ Try it

▶ **Exit configuration mode**

Exiting configuration mode returns you to operational mode.

```
vyatta@vyatta# exit
exit
vyatta@vyatta:~$
```

CLI Help

You can get CLI help by entering any of the following at the command prompt:

Type this:	To see this:
help	Displays available system shell commands
help *command*	Help and usage guidelines for the specified shell command.
<Tab> ?	For non-root users, displays all available Vyatta commands and provides command completion. For the root user,<Tab> displays all available Vyatta and shell commands and provides command completion, however, “?” will not work.

▶ Try it

▶ **Find available commands**

At the command prompt in operational mode, press the <Tab> key or the “?” key.

```
vyatta@vyatta:~$ <Tab>
clear           init-floppy     reboot          terminal
configure       install-system  set             traceroute
debug           no              show            undebug
delete          ping            telnet          vpn
vyatta@vyatta:~$
```

Command Completion

Unlike some proprietary router CLIs that accept an unambiguous prefix, the Vyatta system requires that you type the full command name. Thus, command completion is very important for usability. As an example, Cisco allows you to type “sh config”. The Vyatta system would require you to type “show configuration”, but you can get the same effect in the Vyatta system with “**sh**<Tab>**config**<Tab>” as the <Tab> completes the unambiguous command.

Pressing the <Tab> key provides command completion. You can use the <Tab> to show:

- All available commands
- All valid completions for a command partially typed in
- The next available set of options for a command. You can use the <Tab> key iteratively in this way to discover complete command syntax.

Try it

Use command completion on an unambiguous command

The following example requests command completion for the typed string **sh**. In this example, the command to be completed is unambiguous.

```
vyatta@vyatta:~$ sh<Tab>
vyatta@vyatta:~$ show
```

Try it

Use command completion on an ambiguous command

The following example requests command completion for the typed string **s**. In this case, there is more than one command that could complete the entry and the system lists all valid completions.

```
vyatta@vyatta:~$ s<Tab>
set       show
vyatta@vyatta:~$ s
```

Try it

Display help strings for commands

The following example requests command completion for the typed string **s** a second time. In this case, help strings for the possible completions are displayed.

```
vyatta@vyatta:~$ s<Tab>
Possible completions:
   set       Set system or shell options
   show      Show system information

vyatta@vyatta:~$ s
```

NOTE *The "?" key, like the <Tab> key, will provide help strings for commands.*

Viewing Long Output ("More")

If the information being displayed is too long for your screen, the screen will show a line number indication where the information breaks. There are lots of commands for controlling information in a "More" display; a few important ones are shown below.

To do this	Press this
Exit "More"	q
Scroll down one whole screen.	<Space>
Scroll up one whole screen	b
Scroll down one line.	<Enter>
Scroll up one line.	<Up Arrow>

▶ Try it

▶ **Show system configuration**

Entering **show** in configuration mode shows information you've explicitly set. Entering show -all shows information you've set plus all default information.

Enter configuration mode and enter **show -all** at the command prompt. The full default system displays and the "More" prompt displays. (Some lines are left out of the example to save space.)

```
vyatta@vyatta:~$ configure
[edit]
vyatta@vyatta# show -all
...
   ethernet eth3 {
      duplex auto
      hw-id 00:14:bf:5a:84:f9
      speed auto
   }
...
:
```

▶ Try it

▶ **Exit from a "More display"**

Within the "More" display, enter **q**. You are returned to the command prompt

```
   ethernet eth3 {
      address 192.168.1.85/24
```

```
        duplex auto
        hw-id 00:14:bf:5a:84:f9
        speed auto
    }
: q
[edit]
vyatta@vyatta#
```

Showing Support Information

If you need to report a bug or request support, you will need to supply version information for your software. You can do this in operational mode.

▶ Try it

▶ **Show support information**

If you are in configuration mode, return to operational mode.

```
vyatta@vyatta# exit
exit
vyatta@vyatta:~$
```

Use the **show version** command to display version information.

```
vyatta@vyatta:~$ show version
Version  :    3.0.2
Copyright:    2006-2008 Vyatta, Inc.
Built by :    autobuild@vyatta.com
Built on :    Wed Apr 16 08:26:33 UTC 2008
Build ID :    080416082620a705a
Boot via :    livecd
Uptime   :    14:22:45 up 35 min,  2 users,  load average: 0.00,
0.00, 0.00
vyatta@vyatta:~$
```

The Graphical User Interface (GUI)

The GUI is an alternative user interface for interacting with the Vyatta system. The GUI is turned off by default, for security reasons. If you want to use the GUI, you must enable it using the CLI.

Any operation that can be performed through the CLI (except enabling the GUI) can also be performed through the GUI. The GUI essentially reflects the structure of the CLI; in particular, the command hierarchy in the GUI follows the basic CLI configuration structure. If you are familiar with the CLI, the structure of the GUI should be straightforward to understand.

This section presents the following topics:

- Enabling Access to the Web GUI
- Logging in to the Web GUI
- Logging out of the Web GUI
- Layout of the Web GUI Screen
- Navigational Aids
- Interpreting CLI Documentation for the Web GUI

Enabling Access to the Web GUI

All you need to do to access the GUI is to enable the HTTPS service. (Note that enabling HTTPS access only allows access to the web GUI; it does not provide general SSL access.)

To access the GUI from the network, you will also need to configure an Ethernet interface with an IP address. You can read how to do this in “Chapter 3: Configuration.” In our examples, we assume you have configured an Ethernet interface with an IP address.

▶ Try it

▶ **Enable web GUI access**

1 In configuration mode, enable HTTPS access to the web GUI on the Vyatta system and commit the change.

```
vyatta@R1# set service https
[edit]
vyatta@R1# commit
[edit]
vyatta@R1#
```

Logging in to the Web GUI

▶ Try it

▶ **Access and Log in to the GUI**

1 Point your web browser at the IP address defined for the Ethernet interface. (Prefixing the IP address with **https://** is optional.) The login screen opens in your browser.

NOTE *Supported browsers include Firefox 3, Internet Explorer 7 and 8, and Google Chrome 5.*

NOTE *For security purposes, all communications to the web GUI is over HTTPS, the secure version of HTTP, using a self-signed certificate. As with any self-signed certificate, when you initially connect with a web browser you may receive a warning that the certificate is not trusted. Since this certificate is generated on the Vyatta system specifically for browser access, this message may be safely ignored and the certificate stored for future access.*

2 Log in. Any non-root user may log on to the router through the GUI. Root login through the GUI is not supported for security reasons. The default non-root user ID and password are as follows:

```
User name: vyatta
Password: vyatta
```

Passwords are independent of the means of access: the CLI and the GUI use the same login credentials.

NOTE *For security reasons, it is strongly recommened that each user accessing the system using a web browser have an individual user account.*

Once you have logged on, the main GUI screen appears as follows:

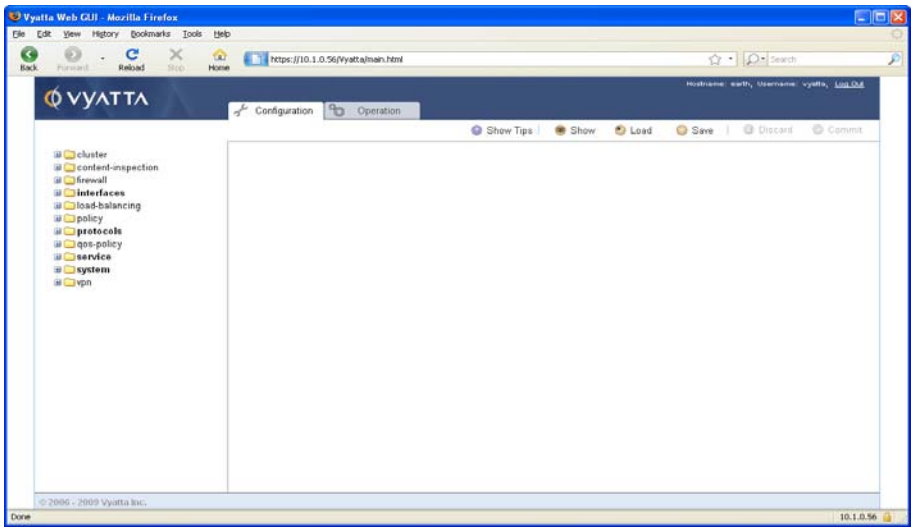

After 30 minutes of inactivity, your session is disconnected and you will be required to log on again. In addition, any uncommitted configuration changes will be discarded on session timeout.

Logging out of the Web GUI

▶ Try it

▶ Log out of the GUI

1 Log out by clicking the “Log Out” link in the top right corner of the web GUI header.

NOTE *After logout, changes that are not committed will remain available to subsequent logins for 30 minutes before being discarded.*

Layout of the Web GUI Screen

The web GUI screen is divided into four areas: **Header**, **Navigation**, **Command Buttons**, and **Input/Output**.

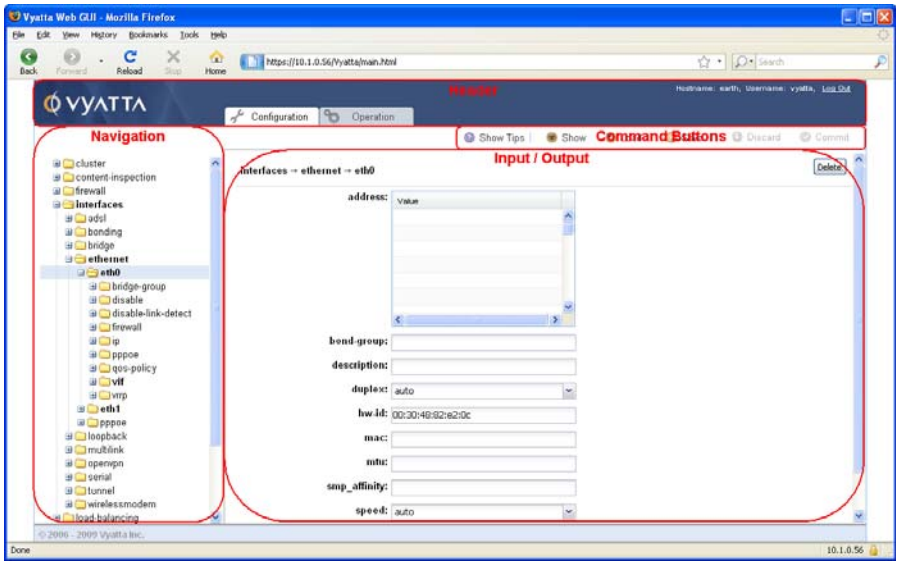

The Header Area

The Header area is located along the top of the screen. It displays the Vyatta logo, the hostname of the system, the user ID of the user who is currently logged in, the Log Out button, and the two main functional tabs: Configuration and Operation.

The Navigation Area

The Navigation area is at the left of the screen. It is essentially a hierarchical menu mirroring the command structure found in the CLI. Most commands available in the CLI are available in the Navigation area (exceptions include: configure, telnet, terminal, install-system, update

webproxy, clear screen, clear https, reboot at, shutdown at, and set). The configuration hierarchy appears when the Configuration tab is selected, and the operational command hierarchy appears when the Operation tab is selected.

The Command Buttons Area

The Command Buttons area is located between the Header area and the Input/Output area. It displays command buttons that are applicable to the selected functional tab. Within the Configuration tab the following command buttons are available:

- **Show Tips / Hide Tips**. Turns help text on and off.
- **Show**. Displays the current system configuration file.
- **Load**. Opens a dialog box to load a saved configuration file.
- **Save**. Opens a dialog box to save the current configuration. Only configurations with no uncommitted changes and no errors can be saved.
- **Discard**. Discards configuration changes that were made since the last load or commit.
- **Commit**. Commits (makes active) configuration changes that were made since the last load or commit.

Within the Operation tab, the following command button is available:

- **Show Tips / Hide Tips**. Shows or hides embedded help text.

The Input/Output Area

The Input/Output area is to the right of the Navigation area. All system input and output displays here.

The background of any fields that have been modified but not “set” will be highlighted. Configuration is “set” in the same way as when you issue a **set** command in the CLI. The equivalent of a set command is executed if you do either of the following:

- Press <Enter>
- Click the **Set** button

Navigational Aids

The web GUI provides a number of devices to aid navigation.

- Hierarchy navigation. There are two ways to navigate the hierarchy:
  - By clicking the node name. When you click the node name for the level you wish to navigate to, the next level of the hierarchy appears in the **Navigation** area and any fields applicable to the selected level appear in the **Input/Output** area.

— By clicking the “+/-” signs. When you click the “+” for the level you wish to navigate to, the next level of the hierarchy appears in the Navigation area, but no additional fields appear in the Input/Output area. This makes navigation a little quicker. Clicking the “-” sign collapses the hierarchy to the selected level.

- **Bold/Plain text in the Configuration hierarchy**. In the Configuration hierarchy, bold text identifies nodes that have been committed to the system configuration. Plain text identifies nodes that are not committed; these can be added by creating the node and clicking the Commit button.
- **Red asterix to the left of a field.** In the Input/Output area, a red asterix to the left of a field indicates that the field is a mandatory field. Mandatory fields must be filled in in order to successfully commit a change.
- **Yellow dots**. Yellow dots appear in the Navigation area and in the Input/Output area to flag uncommitted changes. These yellow dots will appear at each level of the hierarchy to guide you to fields that have been changed but not committed.

  — Modifications are indicated by plain yellow dots.

  — Additions are indicated by yellow dots with an embedded “+” sign.

  — Deletions are indicated by yellow dots with an embedded “-” sign.

  The yellow dots can be removed either by clicking the Discard button (which discards your changes), or by clicking the Commit button (which commits your changes).

NOTE *All uncommitted changes persist beyond a login session and can be seen by any logged in user. This behavior is different from the behavior of the CLI.*

- **Red dots**. Red dots appear in the Input/Output area to indicate invalid entires (e.g., entering something other than a valid IP address in a field that requires an IP address). These indicators will disappear once the problem is corrected or if you navigate to a different configuration node.

▶ Try it

▶ Create a configuration node

Creating a node in the GUI (and committing the change) adds it to the configuration file.

1. In the Navigation area, navigate to the plain text node and select the node.
2. In the Input/Output area, do either of the following:
   - For system nodes: click the Create button that appears.
   - For user-created nodes: Enter the name of the node in the text box that appears (for user-created nodes). Pressing <Enter> (or clicking the Set button) creates the node.

   The node is created, and all levels of the hierarchy for the created node appear in bold text. Because the created node represents an uncommitted change, each item in the hierarchy has a yellow dot with an embedded “+” beside it to indicate that it has been added. Once the change is committed, the yellow dots disappear and the node continues to appear in bold text.

▶ Try it

▶ Delete a configuration node

Deleting a node in the GUI (and committing the change) removes it from the configuration file.

1. In the Navigation area, navigate to the bold text node and select the node.
2. In the Input/Output area, click the Delete button.

   The node is deleted (as well as all its child nodes), but all levels of the hierarchy for the deleted node remain in bold text. Because the node deletion represents an uncommitted change, each item in the hierarchy has a yellow dot with an embedded "-" beside it to indicate that it has been deleted. Once the change is committed, the yellow dots disappear and system nodes display in plain text; user-created nodes are removed from the hierarchy.

Interpreting CLI Documentation for the Web GUI

The majority of Vyatta documentation is targeted for the CLI rather than the GUI; however, it is very easy to execute the CLI examples in the context of the web GUI, since the same functionality is available in both interfaces.

Any configuration mode commands can be executed within the Configuration tab, and any operational mode commands can be executed within the Operation tab. Clicking on the node in the Navigation area either provides an additional level of hierarchy in the Navigation area or a place to enter data within the Input/Output area.

▶ Try it

▶ Add configuration from the web GUI

Consider the following command executed in configuration mode in the CLI:

```
vyatta@R1# set interfaces ethernet eth1 address 192.168.1.61/24
[edit]
vyatta@R1# commit
[edit]
vyatta@R1#
```

To execute this sequence in the web GUI, perform the following steps:

1. Select the Configuration tab.
2. Navigate to interfaces > ethernet > eth1 in the Navigation area and click eth1.
3. Enter the address 192.168.1.61/24 in the address box in the Input/Output area.

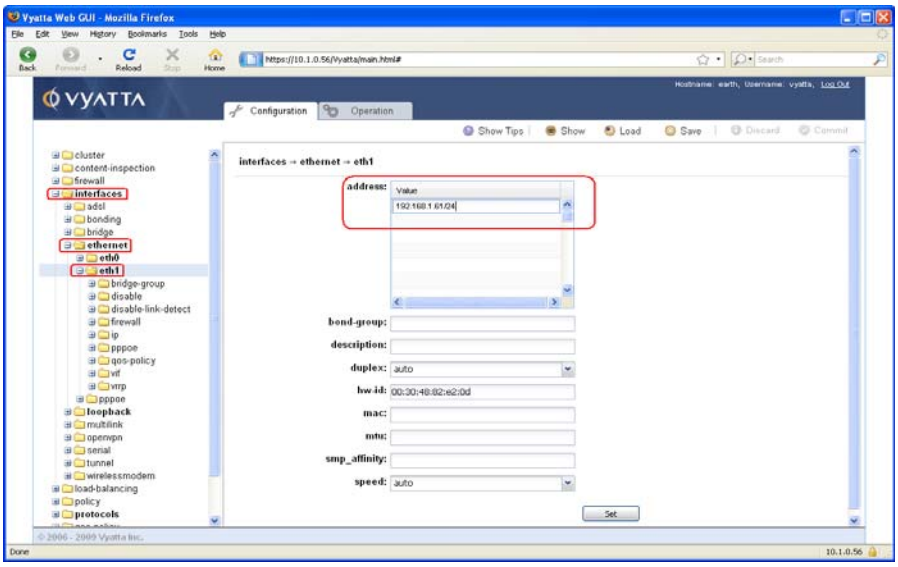

4 Clicking the Set button or pressing <Enter> executes the command and yellow dots appear to indicate the change.

5 Click Commit to commit the change. The yellow dots disappear to indicate that there are no longer any uncommitted changes.

6 You can also click the Show button to see the new **address** node in the configuration file.

▶ Try it

▶ Execute an operational command in the web GUI

Consider the following operational command:

```
vyatta@R1:~$ show interfaces ethernet eth1
eth1:  mtu 1500 qdisc pfifo_fast state UP qlen 100
    link/ether 00:30:48:82:e2:0d brd ff:ff:ff:ff:ff:ff
    inet 10.1.0.56/24 brd 10.1.0.255 scope global eth1
    inet6 fe80::230:48ff:fe82:e20d/64 scope link
       valid_lft forever preferred_lft forever

    RX:  bytes    packets     errors    dropped    overrun      mcast
       9952158      35278          0          0          0        108
    TX:  bytes    packets     errors    dropped    carrier collisions
      26597825      35122          0          0          0          0
```

To execute this sequence in the web GUI, perform the following steps:

1 Select the **Operation** tab.

2 Navigate to **show** > **interfaces** > **ethernet** > **<value>** in the **Navigation** area.

3 Select **eth1** from the **show interfaces ethernet** drop-down menu.

4 Click the **Run** button. The output appears in the **Input/Output** area.

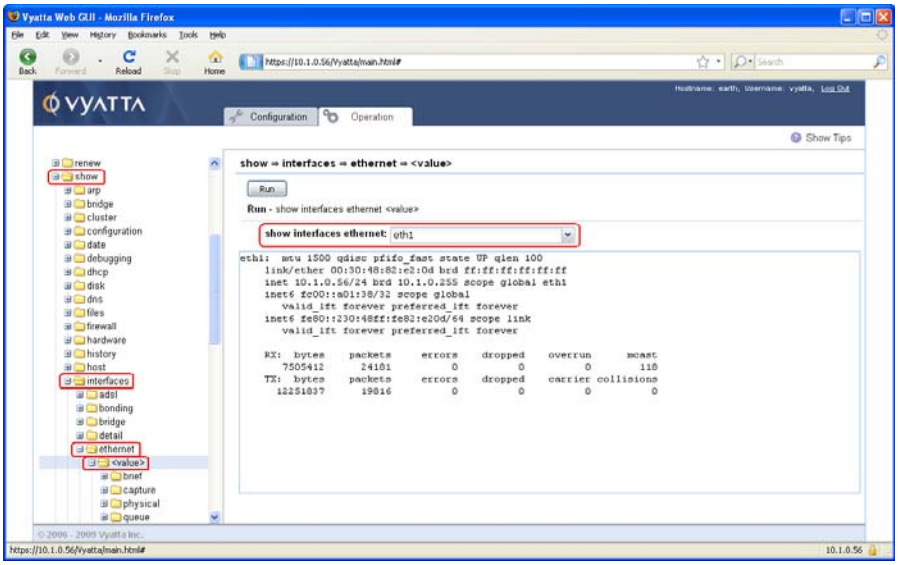

Chapter 3: Configuration

This chapter introduces you to basic configuration concepts, and then steps you through two configuration scenarios:

- Basic system configuration tasks
- A scenario with the Vyatta system acting as an Internet gateway.

The examples in this chapter use the Vyatta CLI and assume that the Vyatta system is installed on a hard disk.

This chapter presents the following topics:

- Configuration Basics
- Scenario: Basic System Configuration
- Scenario: Internet Gateway

Configuration Basics

This section presents the following topics

- Configuration Hierarchy
- Adding and Modifying Configuration
- Deleting Configuration
- Committing Configuration Changes

- Discarding Configuration Changes
- Saving Configuration
- Loading Configuration

Configuration Hierarchy

From the system's point of view, a configuration node is different from a simple configuration attribute statement. A configuration *attribute statement* takes the form *attribute value*, as in the following example.

```
protocol-version v2
```

A configuration *node* always has an enclosing pair of braces, which may be empty, as in the following example:

```
loopback lo {
}
```

or non-empty, as in the following example:

```
ssh {
   port 22
   protocol-version v2
}
```

Adding and Modifying Configuration

Add new configuration by creating a configuration node, using the **set** command in configuration mode. Modify existing configuration also by using the **set** command in configuration mode.

▶ Try it

▶ **Add configuration**

In configuration mode, set the IP address of Ethernet eth0 interface using the **set** command.

```
vyatta@vyatta# set interfaces ethernet eth0 address 192.0.2.21/24
[edit]
vyatta@vyatta#
```

Note that the configuration node for interface eth0 will already exist, assuming that your system has at least one Ethernet. That's because the system automatically discovers physical interfaces on startup and creates configuration nodes for them. For the same reason, the hardware ID (MAC address) of interface eth0 will also be known to the system.

Now use the **show** command to see the addition.

```
vyatta@vyatta# show interfaces ethernet eth0
+address 192.0.2.21/24
 hw-id 00:13:46:e6:f6:87
[edit]
vyatta@vyatta#
```

Note the "+" in front of the new statement. This shows that this statement has been added to the configuration but the change is not yet committed. The change does not take effect until configuration is committed using the **commit** command.

The configuration tree is nearly empty when you first start up, except for a few automatically configured nodes. You must create a node for any functionality you want to configure on the system. When a node is created, any default values that exist for its attributes are applied to the node.

▶ Try it

▶ **Modify configuration**

For the most part, modifying configuration is the same as adding configuration by using the **set** command. This works for identifiers of nodes containing a single instance. To change the identifier of a node for which there can be multiple instances (a "multi-node"), such as a DNS server or an IP address for an interface, you must delete the node and recreate it with the correct identifier.

You can modify configuration from the root of the configuration tree or use the **edit** command to navigate to the part of the tree where you want to change or add. This can speed up editing.

Deleting Configuration

You delete configuration statements, or complete configuration nodes, using the **delete** command.

▶ Try it

▶ **Delete configuration**

Delete a configuration node:

```
vyatta@vyatta# delete interfaces ethernet eth0 address
192.0.2.21/24
[edit]
vyatta@vyatta#
```

Now use the **show** command to see the deletion.

```
vyatta@vyatta# show interfaces ethernet eth0
-address 192.0.2.21/24
 hw-id 00:13:46:e6:f6:87
[edit]
vyatta@vyatta#
```

Note the "-" in front of the deleted statement. This shows that this statement has been deleted from the configuration but the change is not yet committed. The change does not take effect until configuration is committed using the **commit** command.

Some configuration nodes and statements are mandatory; these nodes or statements cannot be deleted. Some configuration statements are mandatory but have default values; if you delete one of these statements, the default value is restored.

Committing Configuration Changes

In the Vyatta System, configuration changes do not take effect until you commit them.

Uncommitted changes are flagged with a plus sign (for additions), a greater-than sign (for modifications), or a minus sign (for deletions).

```
vyatta@vyatta# show interfaces ethernet eth0
-address 192.08.2.21/24
 hw-id 00:13:46:e6:f6:87
[edit]
vyatta@vyatta#
```

▶ Try it

▶ **Commit configuration changes**

Commit any uncommitted changes by issuing the **commit** command in configuration mode.

Once you commit the changes, the indicator disappears. Also note that the non-mandatory configuration node (**address**) is removed from the configuration.

```
vyatta@vyatta# commit
[edit]
vyatta@vyatta# show interfaces ethernet eth0
 hw-id 00:13:46:e6:f6:87
[edit]
vyatta@vyatta#
```

Discarding Configuration Changes

Instead of deleting many specific changes, you can abandon all changes made within a configuration session by using the **discard** command or by using the **exit** command with the **discard** option.

▶ Try it

▶ **Discard configuration changes using "discard"**

Using the **discard** command is the same as deleting all changes made within a configuration session and then committing the changes. Also notice that after the **discard** command has completed you stay in configuration mode.

```
vyatta@vyatta# show interfaces ethernet eth0
+address 192.0.2.21/24
+description "This is a test"
 hw-id 00:13:46:e6:f6:87
[edit]
vyatta@vyatta# discard
Changes have been discarded
[edit]
vyatta@vyatta# show interfaces ethernet eth0
 hw-id 00:13:46:e6:f6:87
[edit]
vyatta@vyatta#
```

You can't exit from configuration mode with uncommitted configuration changes; you must either commit the changes or discard them. If you don't want to commit the changes, you can discard them using the **exit** command with the **discard** option.

▶ Try it

▶ **Discard configuration changes using "exit discard"**

Try exiting from configuration mode with uncommitted configuration changes; you won't be able to. Discard the changes by issuing the **exit discard** command.

```
vyatta@vyatta# exit
Cannot exit: configuration modified.
Use 'exit discard' to discard the changes and exit.
[edit]
vyatta@vyatta# exit discard
exit
vyatta@vyatta:~$
```

Saving Configuration

The running configuration can be saved using the **save** command in configuration mode. By default, configuration is saved to the file **config.boot** in the standard configuration directory.

- For hard disk installs the configuration directory is **/opt/vyatta/etc/config**
- For installations running off LiveCD, the configuration directory is **/media/floppy/config**.

NOTE *The **save** command writes only committed changes. If you try to save uncommitted changes the system warns you that it is saving only the committed changes.*

NOTE *Unless you save your configuration changes to the default configuration file they will not persist when the system is restarted. On restart the configuration is loaded from the **config.boot** file.*

▶ Try it

▶ Save configuration to the default configuration file

To save to the **config.boot** file in the default configuration directory, just enter **save** in configuration mode.

```
vyatta@vyatta# save
Saving configuration to '/opt/vyatta/etc/config/config.boot'...
Done
[edit]
vyatta@vyatta#
```

▶ Try it

▶ Saving configuration to another file

Save to a different file name in the default directory by specifying a different file name.

```
vyatta@vyatta# save testconfig
Saving configuration to '/opt/vyatta/etc/config/testconfig'...
Done
[edit]
vyatta@vyatta#
```

This provides the ability to have multiple configuration files for different situations (for example, test and production).

You can also save a configuration file to a location path other than the standard configuration directory **/opt/vyatta/etc/config**, by specifying a different path. You can save to a hard drive, compact Flash, or USB device by including the directory the device is mounted on in the path.

If you are running the system from LiveCD, configuration can be saved only to floppy disk (**/media/floppy/config/config.boot**). If you do not save your running configuration to floppy disk, any changes are lost on reboot.

Before saving to floppy disk, you must initialize the floppy disk for use on the system.

▶ Try it

▶ **Initialize a floppy disk for configuration files**

The floppy drive is referred to as **/dev/fd0**. It is automatically mounted in the directory **/media/floppy**.

1. Insert a blank floppy disk into the floppy disk drive.
2. At the system command prompt, enter the following:

```
vyatta@vyatta:~$ init-floppy
```

The system prepares the floppy to receive configuration files. It also saves a copy of the current configuration to **/media/floppy/config/config.boot**.

▶ Try it

▶ **Save LiveCD configuration to the default location on floppy disk**

If you want to be able to boot from the configuration file, save configuration to **/media/floppy/config/config.boot**.

```
vyatta@vyatta# save
Saving configuration to '/media/floppy/config/config.boot'...
Done
[edit]
vyatta@vyatta#
```

▶ Try it

Save LiveCD configuration to another file on floppy disk

If you just want to save a record of configuration, save to a different file name in **/media/floppy/config**.

```
vyatta@vyatta# save testconfig1
Saving configuration to '/media/floppy/config/testconfig1'...
Done
[edit]
vyatta@vyatta#
```

Loading Configuration

A configuration can be loaded using the **load** command in configuration mode. You can only load a file that has first been saved using the **save** command.

▶ Try it

▶ **Load configuration from the default directory**

Load a configuration file from the default directory by specifying only the file name.

```
vyatta@vyatta# load testconfig
```

```
Loading configuration file /opt/vyatta/etc/config/testconfig...
No configuration changes to commit
Done
[edit]
vyatta@vyatta#
```

To load from a directory other than the default directory the entire path must be specified.

Changing the Default Configuration File

To change the default configuration file to one that you have previously saved you use both the **load** and **save** commands.

▶ Try it

▶ **Load a previously saved configuration and save it as the default**

Load the previously saved configuration file (**testconfig**) from the default directory by specifying only the file name and then save it to the default (**config.boot**).

```
vyatta@vyatta# load testconfig
Loading configuration file /opt/vyatta/etc/config/testconfig...
No configuration changes to commit
Done
[edit]
vyatta@vyatta# save
Saving configuration to '/opt/vyatta/etc/config/config.boot'...
Done
[edit]
vyatta@vyatta#
```

At this point the configurations in **testconfig** and **config.boot** are the same. The currently active configuration is the same as what will be loaded when the system restarts.

Scenario: Basic System Configuration

This section steps you through initial system configuration tasks. These are tasks that are required for almost any scenario in which you might use the Vyatta system. These include:

- Logging On
- Entering Configuration Mode
- Setting the Host Name
- Setting the Domain Name

- Changing Passwords
- Configuring Interfaces
- Configuring Access to a DNS server
- Specifying a Default Gateway
- Performing a Package Upgrade to the Latest Version

Figure 3-1 shows a network diagram of the basic system configuration scenario.

Figure 3-1 Scenario: Basic System Configuration

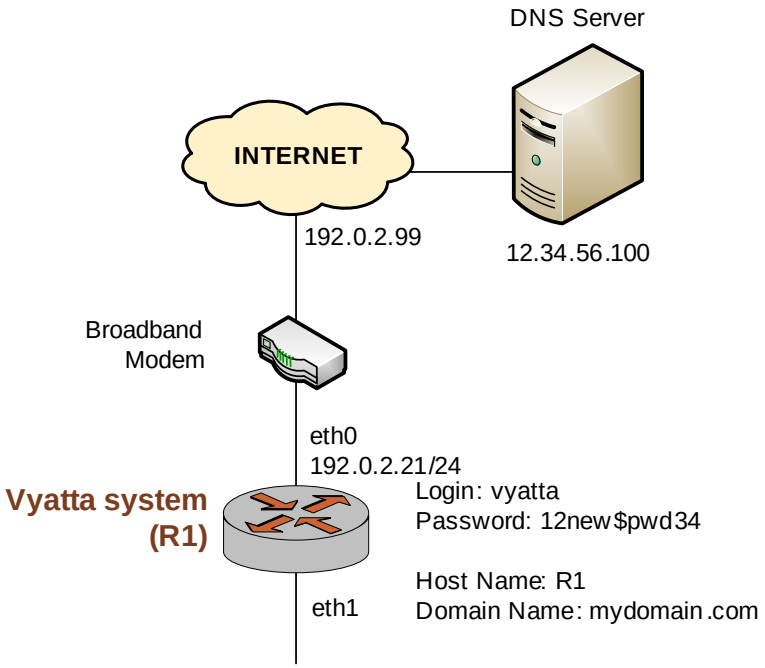

Logging On

The first step is to log on. Our examples use the predefined user **vyatta**.

▶ Try it

▶ **Log on**

Log on as user **vyatta**. The default password for this user is **vyatta**. The password is not echoed onto the screen.

```
Welcome to Vyatta - vyatta tty1
vyatta login: vyatta
Password:
Linux vyatta 2.6.20 #1 SMP Fri Sep 21 02:22:08 PDT 2007 i686
```

```
Welcome to Vyatta.
This system is open-source software. The exact distribution terms
for each module comprising the full system are described in the
individual files in /usr/share/doc/*/copyright.
Last login: Sat Nov 10 16:48:48 2007 on tty1
vyatta@vyatta:~$
```

Entering Configuration Mode

When you log on, you are in operational mode. To configure the system, you must enter configuration mode.

▶ Try it

▶ **Enter configuration mode**

Enter configuration mode by entering **configure**.

```
vyatta@vyatta:~$ configure
[edit]
vyatta@vyatta#
```

Notice how the command prompt changes to mark the move from operational mode (“:~$”) and configuration mode (“#”).

Setting the Host Name

The default host name for a Vyatta device is **vyatta**. You can change this to fit in with your environment. In our example we use a host name of R1.

▶ Try it

▶ **Set the host name**

Change the host name to R1 using the **set system host-name** command. Remember to commit all configuration changes.

```
vyatta@vyatta# set system host-name R1
[edit]
vyatta@vyatta# commit
[edit]
vyatta@vyatta#
```

The command prompt will change to reflect the new host name the next time you log in.

Setting the Domain Name

In addition to changing the host name, you must specify the system's domain name. In our examples we use **mydomain.com** as the domain name.

▶ Try it

▶ **Set the domain name**

Set the domain name using the **set system domain-name** command.

```
vyatta@R1# set system domain-name mydomain.com
[edit]
vyatta@R1# commit
[edit]
vyatta@R1#
```

Changing Passwords

The Vyatta system has one pre-defined user account:

```
user ID: vyatta    password: vyatta
```

To secure your system, you should change the password for this user account.

▶ Try it

▶ **Change your password**

Change the password of user **vyatta** to **12new$pwd34** using the **set system login user** command, as follows:

```
vyatta@R1# set system login user vyatta authentication
plaintext-password 12new$pwd34
[edit]
vyatta@R1# commit
[edit]
vyatta@R1#
```

Configuring Interfaces

The kind and number of interfaces you configure will depend on your physical device and the topology of your network. However, almost every topology will require that at least one Ethernet interface configured.

The Vyatta System automatically discovers all physical interfaces on startup and creates configuration nodes for them. In this basic scenario, we'll configure the Ethernet interface eth0 as an Internet-facing interface.

The loopback interface, which is a software interface automatically created on startup is preconfigured to IP address 127.0.0.1/8. The loopback interface will always be available as long as the device is reachable at all. This makes the loopback interface particularly useful for mapping to the system host name, as a router ID in routing protocols such as BGP and OSPF, or as a peer ID for internal BGP peers.

▶ Try it

▶ Configure an Internet-facing Ethernet interface

Configure interface eth0 with an IP address of 192.0.2.21 and a prefix length of 24. This will allow the system to reach the DNS server and default gateway configured in subsequent steps.

```
vyatta@R1# set interfaces ethernet eth0 address 192.0.2.21/24
[edit]
vyatta@R1# commit
[edit]
vyatta@R1#
```

NOTE *If the ISP in this example provided a dynamic IP address, rather than the static one shown, the command to use would be* ***set interfaces ethernet eth0 address dhcp.***

To view the configuration we use the show command:

```
vyatta@R1# show interfaces
ethernet eth0 {
   address 192.0.2.21/24
   hw-id 00:40:63:e2:e4:00
}
loopback lo {
}
[edit]
vyatta@R1#
```

Configuring Access to a DNS server

In order to be able to translate host names (such as www.vyatta.com) to IP addresses (such as 76.74.103.45), the system must be able to access a DNS server.

▶ Try it

▶ Specify a DNS server

In our example, the DNS server is at IP address 12.34.56.100. Add the DNS server using the **set system name-server** command.

```
vyatta@R1# set system name-server 12.34.56.100
[edit]
vyatta@R1# commit
```

```
[edit]
vyatta@R1#
```

Specifying a Default Gateway

When the local system does not know what route to use for a given destination, it forwards packets to the "default gateway." In our example, the ISP's gateway at 192.0.2.99 acts as the default gateway.

▶ Try it

▶ **Specify the default gateway**

Add the default gateway using the **set system gateway-address** command.

```
vyatta@R1# set system gateway-address 192.0.2.99
[edit]
vyatta@R1# commit
[edit]
vyatta@R1#
```

Performing a Package Upgrade to the Latest Version

Before you start using your system in earnest, it is advisable to perform a package upgrade to the latest version of the Vyatta software. This will provide you with any bug fixes and security patches that have been released since your installation CD image was created. Please refer to the upgrade section in the *Vyatta System Installation and Upgrade Guide* for information on how to do this.

This completes basic system configuration.

Scenario: Internet Gateway

This configuration scenario builds off of the Basic System Configuration scenario above and steps through the process of configuring the system as a basic Internet Gateway. The goal is for the system to provide the following:

1. The ability to route traffic between the Office LAN and the Internet.
2. The ability for users to access the system from the local network using SSH.
3. DHCP capability for providing dynamic IP addresses to internal devices.
4. NAT capability for translating multiple internal addresses to a single external address.
5. Firewall capability for preventing system access from the Internet.

The example assumes a system with two Ethernet interfaces: one interface facing the office LAN and the other facing the Internet.

Figure 3-2 shows a network diagram of this sample scenario.

Figure 3-2 The Vyatta system configured as an Internet gateway

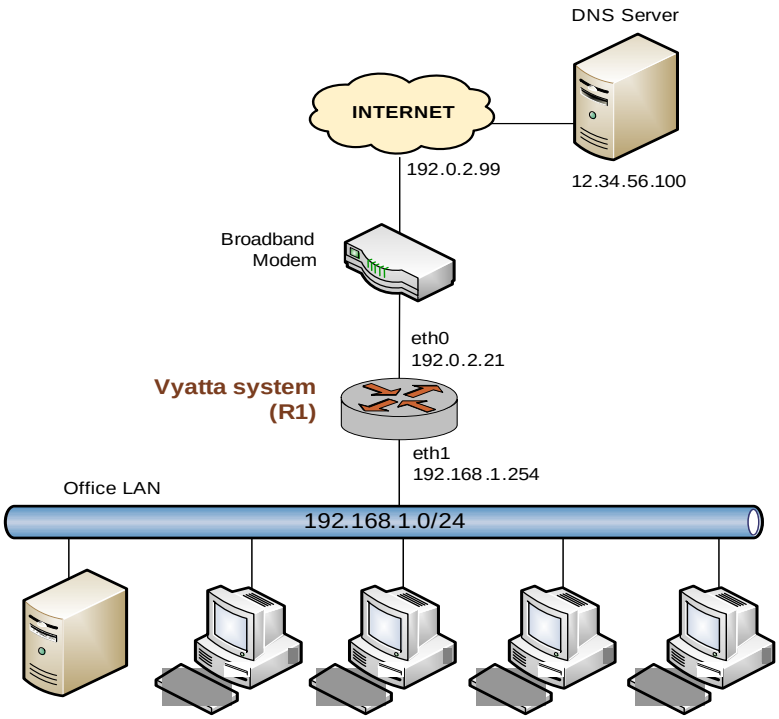

The examples in this scenario assume that you have completed the configuration shown in the basic scenario (see page 37).

This section presents the following topics:

- Configuring Interfaces
- Enabling SSH Access
- Configuring DHCP Server
- Configuring NAT
- Configuring Firewall

Configuring Interfaces

In the basic scenario, we configured an Internet-facing Ethernet interface. To act as an Internet gateway, the system will need an additional Ethernet interface facing the Office LAN. We'll use interface eth1 for this.

▶ Try it

▶ **Configure the Office LAN–facing Ethernet interface**

Assign IP address 192.168.1.254 to interface eth1 with a prefix length of 24, which is the prefix length of the office subnet.

```
vyatta@R1# set interfaces ethernet eth1 address 192.168.1.254/24
[edit]
vyatta@R1# commit
[edit]
vyatta@R1#
```

Let's look at the interfaces configured now:

```
vyatta@R1# show interfaces
ethernet eth0 {
   address 192.0.2.21/24
   hw-id 00:40:63:e2:e4:00
}
ethernet eth1 {
   address 192.168.1.254/24
   hw-id 00:13:46:e6:f6:87
}
loopback lo {
}
[edit]
vyatta@R1#
```

Enabling SSH Access

The gateway should allow users to access the SSH service from the office LAN, but not from the Internet. In this step, we enable the SSH service on the system, which will be accessed by an SSH client (Preventing access from the Internet will be done later, using the Vyatta system's firewall.)

▶ Try it

▶ **Enable SSH access**

Tip: *Setting up Telnetis done the same way: by issuing* ***set service telnet****.*

To allow SSH access, you enable the SSH service on the system. By default SSH uses the more secure "v2" protocol version.

```
vyatta@R1# set service ssh
[edit]
vyatta@R1# commit
[edit]
vyatta@R1#
```

Configuring DHCP Server

DHCP provides dynamic IP addresses to hosts on a specified subnet. In our scenario, the DHCP server provides addresses to hosts on the Office LAN (attached to interface eth1).

▶ Try it

▶ Set up DHCP Server

For the DHCP server, define an address pool from 192.168.1.100 to 192.168.1.199 to dynamically assign addresses to hosts on the Office LAN. Also, set the default router and DNS server to the values that will be assigned to hosts on the Office LAN. The default router for these devices will be the LAN-facing interface of the Internet gateway.

```
vyatta@R1# set service dhcp-server shared-network-name ETH1_POOL
subnet 192.168.1.0/24 start 192.168.1.100 stop 192.168.1.199
[edit]
vyatta@R1# set service dhcp-server shared-network-name ETH1_POOL
subnet 192.168.1.0/24 default-router 192.168.1.254
[edit]
vyatta@R1# set service dhcp-server shared-network-name ETH1_POOL
subnet 192.168.1.0/24 dns-server 12.34.56.100
[edit]
vyatta@R1# commit
[edit]
vyatta@R1#
```

Let’s view this configuration.

```
vyatta@R1# show service dhcp-server
   shared-network-name ETH1_POOL {
      subnet 192.168.1.0/24 {
         start 192.168.1.100 {
            stop 192.168.1.199
         }
         dns-server 12.34.56.100
         default-router 192.168.1.254
      }
   }

[edit]
vyatta@R1#
```

Configuring NAT

The Internet gateway should send outbound traffic from the Office LAN out through the Internet-facing interface and translate all internal private IP addresses to a single public address. This is done by defining a Network Address Translation (NAT) rule.

▶ Try it

▶ Define a NAT rule

Define a rule that allows traffic from network 192.168.1.0/24 to proceed to the Internet through interface eth0, and translates any internal addresses to eth0's IP address. (This is called "masquerade" translation.)

```
vyatta@R1# set service nat rule 1 source address 192.168.1.0/24
[edit]
vyatta@R1# set service nat rule 1 outbound-interface eth0
[edit]
vyatta@R1# set service nat rule 1 type masquerade
[edit]
vyatta@R1# commit
[edit]
vyatta@R1#
```

Let's view this configuration.

```
vyatta@R1# show service nat
   rule 1 {
      type masquerade
      outbound-interface eth0
      source {
         address 192.168.1.0/24
      }
   }

[edit]
vyatta@R1#
```

Configuring Firewall

As it is shipped the Vyatta System does not restrict traffic flow through it. That is, unless a firewall rule is applied to an interface it will allow all traffic through it. The firewall functionality provides packet filtering which enables great flexibility in restricting traffic as required. In this simple scenario, the Internet gateway should allow hosts on the local network and services on the gateway itself to initiate traffic to the Internet, but it should drop all traffic that is initiated from the Internet. This section sets up a basic firewall configuration to do this.

Essentially, this sequence defines a firewall rule set allowing traffic initiated from, or passing through, the gateway to the Internet. All other packets are implicitly denied, because there is an implicit **deny all** rule at the end of any rule set.

In general, to configure a firewall on an interface:

1. You define a number of named firewall rule sets, each of which contains one or more firewall rules.
2. You apply the each of the named rule sets to an interface as a filter. You can apply one named rule set as each of the following on an interface:
   - **in.** If you apply the rule set to an interface as **in**, the rule set filters packets entering the interface.
   - **out.** If you apply the rule set to an interface as **out**, the rule set filters packets leaving the interface.
   - **local.** If you apply the rule set to an interface as **local**, the rule set filters packets destined for the system itself.

*__Tip:__ When applying a firewall rule set, keep in mind that after the final user-defined rule, an implicit rule of **deny all** takes effect.*

▶ Try it

▶ Define a firewall rule set

For our simple example, the natural inclination is to simply create a rule to deny all inbound traffic (i.e. from any source network to any destination network) on the internet-facing interface. The problem with this approach is that outbound connections will not complete properly because response packets required to complete these connections will be denied as well. To circumvent this issue we must explicitly allow only these response packets as shown in the following example. This can be interpreted as “accept packets from established connections only” (where “established connections” includes responses to new connections). Because the final (implicit) rule in the rule set is **deny all**, this rule set will deny all other traffic on the interface destination (i.e. **in**, **out**, or **local**) that it is applied to.

```
vyatta@R1# set firewall name ALLOW_ESTABLISHED
[edit]
vyatta@R1# set firewall name ALLOW_ESTABLISHED rule 10
[edit]
vyatta@R1# set firewall name ALLOW_ESTABLISHED rule 10 action accept
[edit]
vyatta@R1# set firewall name ALLOW_ESTABLISHED rule 10 state
established enable
[edit]
vyatta@R1# commit
[edit]
vyatta@R1#
```

▶ Try it

▶ Apply the rule set to an interface

Now that we have the rule set, we need to apply it as **in** and **local** on the internet-facing interface (eth0 in our example) so that connections can only be established from these locations to the Internet.

```
vyatta@R1# set interfaces ethernet eth0 firewall in name
ALLOW_ESTABLISHED
[edit]
vyatta@R1# set interfaces ethernet eth0 firewall local name
ALLOW_ESTABLISHED
[edit]
vyatta@R1# commit
[edit]
vyatta@R1#
```

Let's view the firewall rule set we created:

```
vyatta@R1# show firewall
name ALLOW_ESTABLISHED {
   rule 10 {
      action accept
      state {
         established enable
      }
   }
}

[edit]
vyatta@R1#
```

Now let's see this rule set applied as a filter to **in** and **local** on interface eth0:

```
vyatta@R1# show interfaces ethernet
ethernet eth0 {
   address 192.0.2.21/24
   firewall {
      in {
         name ALLOW_ESTABLISHED
      }
      local {
         name ALLOW_ESTABLISHED
      }
   }
   hw-id 00:40:63:e2:e4:00
}
```

```
ethernet eth1 {
   address 192.168.1.254/24
   hw-id 00:13:46:e6:f6:87
}
```

This completes configuration of a basic Internet Gateway.