Answers

If you are really stuck and the HOWTOs yield no good result, why not write in? Our resident experts will answer even your most complicated problems!

OUR EXPERTS

Whatever your question is, we can find an expert to answer it – from installation and modem woes to network administrations, we can find the answer for you – just fire off a letter or email and it'll all be taken care of.

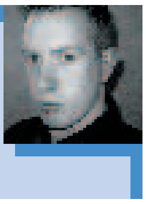

LXF Answers guy **David Coulson** is a networking and security guru with plenty of sysadmin experience to boot.

Nick Veitch is the editor of the magazine, and answers your easy questions! Or indeed anything to do with *Grub, LILO, netatalk, vi...*

Hans Huberland is Rackspace Managed Hosting's Linux expert. Send any Linux system admin questions to **sysadminqa@rackspace.co.uk**

Gentoo concerns

Q Thank you very much for including Gentoo 2004.2 with your last issue – what an amazing distribution, and what amazing documentation! I have an x86 machine connected to the net through a Draytek Vigor 2600G ADSL Modem/Router. This is also a four-port 10/100 switch with wireless capability and a firewall – quite expensive but well worth it!

The ADSL service is PPP over ATM (PPPoA) and my interface is an Intel Pro/1000 MT Desktop Adaptor, referenced as eth0. By the way, I statically compiled e1000 support into my kernel. In light of the fact that Gentoo also provides a separate e1000 module/package, was this a good decision? I have yet to find any issues with my setup.

Safe in the knowledge that I was protected by this comprehensive firewall, I've only just begun to look at IPTables, and here's my problem.

First of all, I'm a little lost as to how to configure my kernel (linux-2.4.26-gentoo-r6) for IPTables support. There seem to be several incompatible options here that I can't fathom. Secondly, if you compare some of your previous FAQs on your help pages, such as IP security, firewalls and Linux and the Internet, with some documentation I found at http://gentoo-wiki.com/HOWTO_Iptables_for_newbies, you'll notice a little difference in the number of rules and amount of detail given. I hope you're not as lost as me upon viewing the latter!

Presently, I'm fearful of tinkering before understanding things more, so please help!

James Thompson

A You should be able to get going with IPTables simply by running **iptables –nvL** from the command line. This will list the three basic 'filter' tables that you can configure to block traffic. Gentoo's kernel comes with IPTables support as default, although if you've compiled your own kernel with support for the Intel EEPro 1000 NIC, you may want to compile IPTables into the kernel rather than using modules.

Generally, it's a good idea simply to compile all of the options into the kernel because it can be very frustrating to have to reboot a firewall simply to add support for a particular IPTables feature. The documentation from Gentoo contains a very complete firewall configuration, which is beyond the needs of the vast majority of users.

The script is useful because it allows for easy modifications to permit access to and from specific ports, making it a great starting point for anyone building a complex firewall.

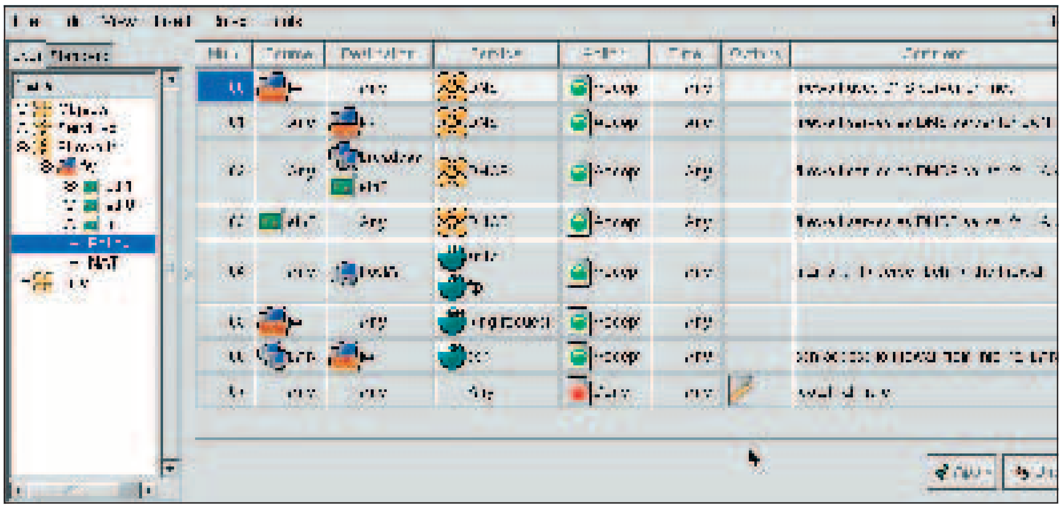

IPTables allows a complex firewall to be constructed. However, for beginners, a simple firewall is much easier to debug.

Booting Mandrake

Q I just received your December issue with the Mandrake 10.1 disks. I've installed the system, but for some reason it just won't boot.

My computer is home assembled with an AMD CPU, an ATI Rage 128 video board and a two-button mouse on COM2. I have two hard drives: one has SUSE 9.0 and Windows ME in a dual-boot and another that's used for experiments, such as the Mandrake installation. The experimental OSes are run from /dev/hdb.

I tried Mandrake 9.2, but the video was distorted and it ran like treacle. I have a feeling that the ATI board may be the problem visually.

I tried your Fedora Core 2, but apparently it couldn't stomach a mouse on COM2 (at least, it never found it) so I couldn't use that either, and I never found out how to tell it where the mouse was. Now Mandrake 10.1, when selected from GRUB, simply returns to reboot.

I managed to get something to 'take' by using the SUSE vmlinuz and initrd files instead, but there were too many errors for it to finish. Obviously they're incompatible, but the Mandrake ones seem to contain errors. Any ideas please? All offerings gratefully received.

Bill Hart

A A reboot immediately after boot generally indicates a kernel issue, and since you have an AMD CPU, it may have complaints with the i686 compiled kernel from Mandrake 10.1. Using the SUSE kernel will allow the box to boot. However, as each vendor

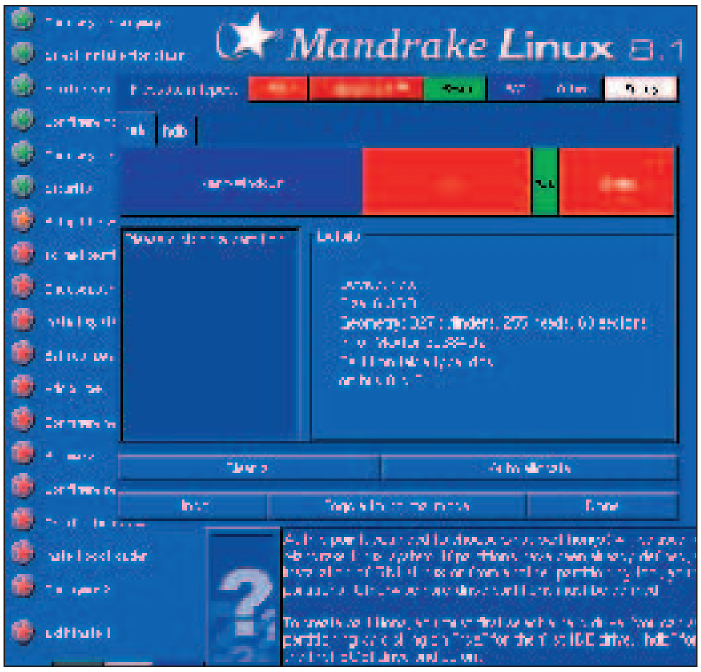

Linux can have a very flexible filesystem structure, and modifying existing disks can be done without destroying data.

has such different kernels, as you've found, it's not always successful.

Fedora Core 2 should be able to function with a mouse on COM2, or ttyS1 in Linux language. In the worst case, you can simply modify your /etc/X11/XF86Config-4 and point it to /dev/ttyS1 rather than /dev/psaux and ensure that the mouse type is set correctly. You may want to purchase a PS/2 or USB mouse with extra buttons because Linux really likes that middle mouse button, and having to click both at the same time gets tiresome very quickly.

How to get FAT

Q I note in *LXF60* (page 84) that you recommend a FAT partition, which I think can be up to 4GB, for easy read/write access from both Windows and Linux. This would be of great use to me, but I haven't been able to set this up. I run an Evesham (May 2003) with Windows XP Pro and SUSE 8.2 Pro mounted on separate 80GB Hard Discs.

Is it possible to repartition either hard disc to provide such a 4GB FAT partition without having to reload either of the operating systems and thus losing my settings? If so, how?

When I loaded SUSE 8.2, I used the recommended single partition. Now, a little wiser, I'd like to repartition that hard disk anyway for Linux use, with /home separate so I can try new Linux distros and so on, without losing my tried and trusted system. I seem to remember reading a few months ago that an easy way to partition a disc is to start loading Mandrake and stop once partitioning has been done. Does this overwrite everything already on that disc?

Please advise on the best and safest way to repartition, with a FAT partition at the end of one of the hard discs that's recognised by both Windows and Linux.

A related problem I have is that when I'm using a 32MB USB pen drive to transfer between the two systems, or indeed to other PCs, writing to the pen drive in Linux results in case changes to file names. The only way to correct this is to then read the files into Windows and then write them back from Windows. The correctly cased file names are then read by any Linux or Windows PC. Is this inherent or is it a driver problem, and do I need to load a specific driver rather than rely on a default?

My final question is this: why does the chap in the Rackspace ad (*LXF60*, page 11) always have a 12-inch/30cm rule in his shirt pocket? Is it sawn off? It must be very uncomfortable if he sits down, or does he use it for a quick scratch?

Roger Gibson, via email

A You can either use *Partition Magic* or an Open Source tool such as *GNU PartImage* to repartition a disk without wiping it. Both of these will adjust the filesystems prior to modifying the partition structure, allowing for the modifications to be made without destroying data. You can then carve out a partition on the disk and build a FAT filesystem on it. Using FAT32, you'll be able to create a partition far greater than 4GB, or alternatively you could simply mount your Windows XP filesystem and access a specific directory on the disk.

Mounting a filesystem, either disk-based or USB, using 'fat' will result in naming issues and problems with long filenames. Using 'vfat', you can ensure that information is preserved and will allow for the easy exchange of data between Linux and DOS.

A 12-inch ruler is really useful for keeping people in line, including the *LXF* editorial staff. You've got to be careful not to draw blood though.

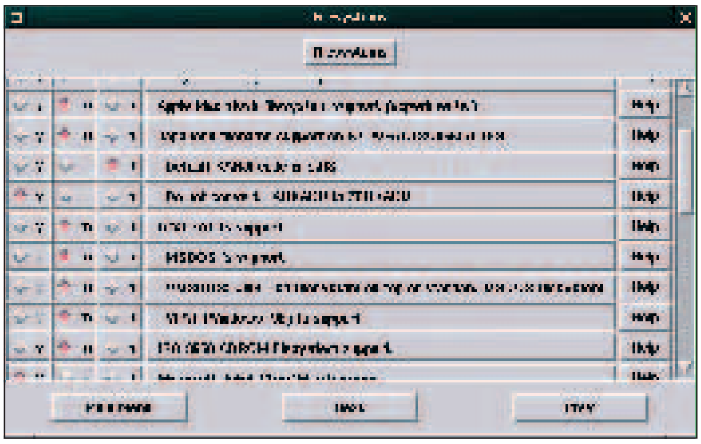

VFAT allows Windows filesystems to be accessed without losing support for large disks and long filenames.

Your answers...

Q Your 'cd $1 ; ls' answer to RS Clymo (*LXF60*, page 103, 'Bashed Up') missed the solution! Although neither alias nor script will work as intended, there's always the *bash* function command. Add this onto the end of ~/.bashrc

```
function cdls
{
      cd "$1"
      ls
}
export function cdls
```

and then give the command:

```
. ~/.bashrc
```

and you have the 'cdls' (or 'cs') as required, both now and at any future logins.

Martin, www.nottingham.lug.org.uk

It is possible to achieve what RS Clymo wants with this code:

```
alias cs='. /path/to/cs.sh'
cs.sh reads:

--BEGIN--
cd $1
ls
--END--
```

Note that it's important not to put '#!/bin/bash' on the beginning of this. Good luck!

Lukasz Wojtow

A These are both great solutions to the problem. Storing functions in *bash* is an excellent alternative to lots of scripts and aliases for those of us who tend to type the same commands over and over.

Installing KDE

Q I bought *LXF59* because it had a free version of KDE 3.3 and I wanted to try Linux as my OS. I have a 40GB hard drive partitioned to two 20GB segments, C: and D:. I run Windows XP Pro on drive C: and at some point I hoped to try a Linux system. This seemed ideal. However, I'm not a techie and I couldn't find a way of installing the program (or even getting to my D: drive for that matter!) and wonder if you could give me some guidance, either to purchase a system from PC World or learn more first!

I strongly object to the Windows stranglehold and I use alternatives where possible, like *Mozilla* for browsing and *Fastmail.fm* as my mail provider.

I have my own website at www.land-to-buy.co.uk, which was built by a university student who has since gone back to France, but I manage to keep it going, and that's »

Gaim (http://gaim.sf.net) supports almost every IM protocol used on the Internet, and a few no one will ever use.

« **done in PHP and uses *MySQL* databases. You now know the height of my ability! Could you give me any advice?**

Barry Dawson, barryhalldawson@fastmail.fm

A The KDE 3.3 discs you have are for people who already have Linux installed – you can't install Linux from them. Instead, you need a full distro such as Mandrake 10.1, which was bundled in *LXF60* (the issue after the KDE discs), or this issue, which has Fedora Core 3 on. Both of these will give you the easy headstart you're looking for. When installing Linux, pop the install CD/DVD in your disc drive and reboot, then remove the D: partition from the disk and let Linux install there. Your distro will enable you to easily repartition the disk and build the appropriate structure for your new Linux installation.

Trying on Red Hat

Q **I recently bought an Intel 865 desktop board. I have a Seagate 120GB SATA hard drive. I tried installing Red Hat 9 on it, but with no success. Can you tell me which Linux flavour I should use? Will Mandrake 10.1 detect my Seagate SATA hard disk?**

Phil Fry

A Red Hat 9.0 lacks support for the SATA chipsets that are used by the current motherboards, so a more recent Linux distribution will be necessary. If you want to stick to the Red Hat line of distribution, Fedora Core is a great choice and continues to use the RPM packages that anyone who has used Red Hat will be used to. As another option, Mandrake or SUSE will also work with SATA if a current release is used.

Yahoo and Linux

Q **After finally getting my Internet to work in Mandrake 9, which I'm happy about, I want to install Yahoo Messenger. However, I ran into a problem. On the Unix site (http://messenger.yahoo.com/messenger/download/unix.html). I'm not sure which option to choose because there isn't a Mandrake one.**

I know Mandrake was built on Red Hat but Mandrake has probably changed a lot since then and I don't have a clue where to start as I'm a newbie! Could someone please help me with this?

Clint

A A great way to use Yahoo with Linux is with Gaim (**http://gaim.sf.net**), which provides access to Yahoo, MSN, AIM and other instant messaging protocols.

Mandrake is so very far separated from Red Hat at this point that the only common feature between the two is the use of RPMs. As such, it's rare for packages for Red Hat to work with Mandrake due to the differences in libraries.

Hell breaking loose...

Q **I recently installed Oracle on my fresh Linux OS. The database was successful, although there were problems with my DBCA. In the process of fixing the above problem, someone suggested that I put this line at the top of my .java_wrapper in the jre directory of JRE: "LD_PRELOAD=/etc/libcwait.so".**

That was when all hell broke loose! My system came back with this message: "/etc/libcwait.so: cannot open shared object No such file or directory".

I decided to take the line out or try to find libcwait.so and put it in the right directory, but my system wouldn't allow me to do this. I then decided to logout and reboot. Big mistake! During the reboot, the system froze with this message: "init: error while loading shared libraries. /etc/libcwait.so: cannot open shared object No such file or directory. Kernel panic: attempted to kill init". It wouldn't go any further after this.

Could anyone help? How do I load Linux or do safe-mode loading so I can take this offending LD_PRELOAD=/etc/libcwait.so out of the .java_wrapper?

John Watson

A Booting the system from a rescue disk will allow the root filesystem to be mounted and /etc/ld.so.preload to be removed to avoid the system attempting to load /etc/libcwait.so. /etc/libcwait.so is a strange place for a library, so verifying the documentation from Oracle to ensure that the path is correct would be a great first step to solving the problem.

It will most likely be in /lib or /opt rather than /etc, although running a 'find' across the disk would find the exact library path quickly.

Hard disk hassle

Q **I've put a hard disk in my old computer. It did have a damaged hard disk but now I've reformatted it and partitioned it to a primary DOS partition. However, Windows 95 is old and I can't do anything with it because it's a new computer.**

How would it be possible to install another operating system, ever Windows ME or Linux, onto the old drive that's been re-partitioned, or am I going to have to install Windows XP instead?

Edward Handy

A You could very easily install Fedora Core or Mandrake onto the disk, or optionally do an install of both Linux

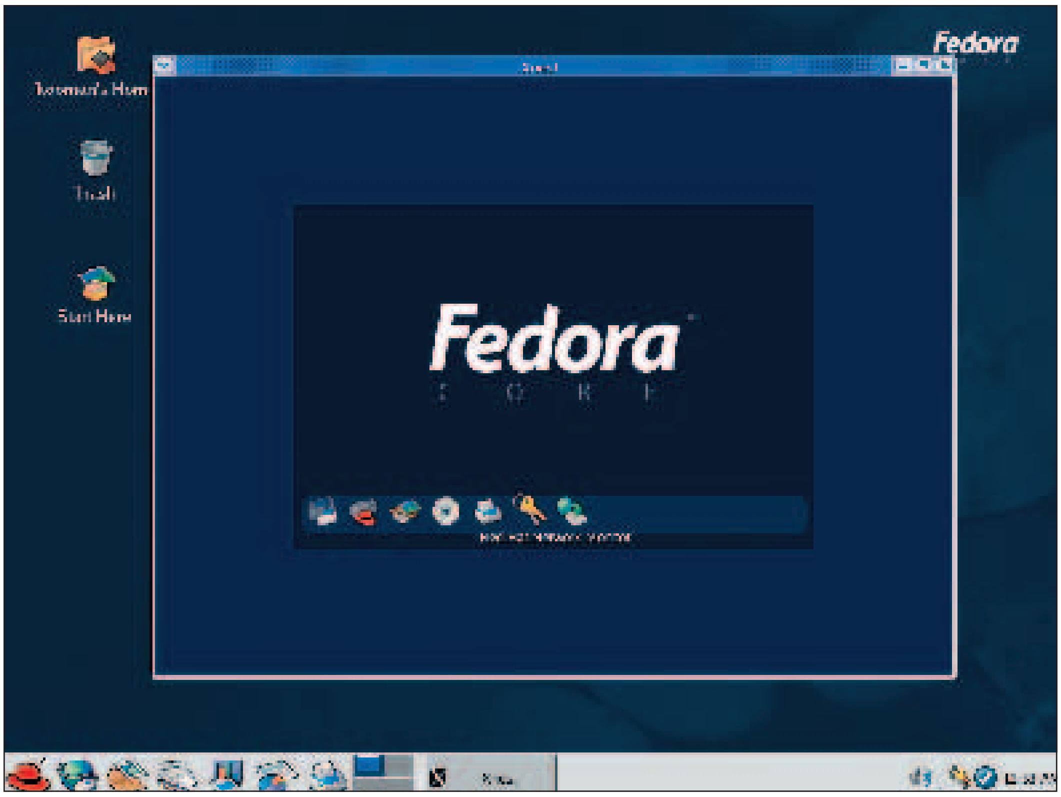

Fedora Core supports the latest hardware and gives the familiar feel of Red Hat, without the cost of Red Hat Advanced Server.

and Windows XP onto the disk. Both will repartition the disk when you install them, removing the old Windows 95 filesystem.

Drive time

Q I have a Linux PC, running Red Hat 9.2. I want to add an additional disk drive. I know this sounds like the most basic of tasks, but having only done this with Windows, I don't really know what to expect. After I've added the hardware and rebooted, what do I do next? I assume that I need to format the drive but where would I complete this task?

Am I right in thinking that Linux will automatically recognise the addition of the drive? Will I see it as an additional drive or just continuous disk space? Any clues that you could give me would be a real help.

Michael T

A When adding an extra disk to Linux, you'll have to partition it using *fdisk* and then build filesystems on the partitions you create. Once created, you can mount them in the appropriate location and use them.

To maintain a mount across a reboot, adding an entry to /etc/fstab for the new filesystem will ensure that it's mounted in the correct location when the system comes back up.

If your new disk is hdc, you can do:

```
fdisk /dev/hdc
mke2fs –j /dev/hdc1
mount /dev/hdc /home2
```

You could also copy the contents of /home onto /home2 using 'cp –fra /home/* /home2' once it's mounted, then modify /etc/fstab to mount /dev/hdc1 onto /home at boot time.

Head banging

Q I've been banging my head against this one for weeks now. Four years ago I managed to get a machine to DNAT and now I can't do it at all! At the most basic level, I'm trying this code:

```
Internet external ip on firewall =
10.x.x.5
Machine on inside of firewall =
192.168.1.2
```

The firewall can access the http server on the internal machine via port 80 without any problems, so I tried this:

```
insmod iptable_nat
iptables -F INPUT
iptables -F OUTPUT
iptables -F FORWARD
iptables -P INPUT DROP
iptables -P OUTPUT ACCEPT
iptables -P FORWARD ACCEPT

echo 1 > /proc/sys/net/ip_forward

iptables -t nat -A PREROUTING -d
10.x.x.5 -p tcp --dport 80 -j DNAT --
to 192.168.1.5:80
```

And nothing happens. I've tried many variations of source IP, interfaces and so on, but none of them seem to work. Can you tell me how to get things working?

Rob

A The first stage in any DNAT configuration is to ensure that the IP configuration on the firewall is correct, and in this case, 10.x.x.5 should be bound to the outside interface on the firewall as either an interface or an alias. Opening up ICMP traffic on the firewall and pinging the outside IP from a system will help in ensuring that the IP layer is happy.

Of course, because the outside address is in the 10.0.0.0/8 range, it won't be available from the other side of the Internet, in which case the appropriate routable address should be used.

The simplest way to debug any DNAT problem is to run 'tcpdump' on the outside interface of the firewall and review the packets that are dumped from the connections from the outside host. This will ensure that packets are being routed back and forth properly, and if a packet is seen going into the firewall but not back out again, you can work through the firewall configuration.

Your information detailed the inside address as 192.168.1.2. However, you were DNATing to 192.168.1.5. Hopefully this is just a typo, although it's always a good idea to double-check all of the firewall rules to ensure that the IP addresses are correct.

»

A QUICK REFERENCE TO: Zebra

Having a large network quickly becomes complex to administrate, particularly when it is further subdivided into subnets for various groups or individuals. Each subnet, having its own IP block, requires a separate route on gateways and, of course, it doesn't take long to have hundreds of routing entries on multiple gateways. If one of these routes is missing or incorrect, then the particular section of the network which is being referring to would become unreachable and effectively disconnected from the rest of the network.

To avoid user error, there are a number of protocols commonly used to distribute routing information over a network. The most common protocol is RIP, which is a very simple routing protocol. Any route on a router is broadcast to all other routers on the network, so there is no need to set up specific routes on gateways as they will gather all the routing information from other gateways on the network.

RIP is a distance-vector-based protocol, so each route out of a router is just as good as any other. RIP decides on the best route to use based upon the number of hops between the router and that network. Unfortunately, this can mean that you may end up with a network with five hops to a network over 100Mbit, or two hops over a 56K modem.

An alternative to RIP is Open Shortest Path First, or OSPF. Each route has a metric assigned to it giving it a 'cost', which may be available bandwidth or actual financial cost of bandwidth over that link. A popular use of OSPF is for a redundant network, so if one link fails, another takes over.

While these are standard protocols on most routers, Linux requires a service to be running in order to make use of these. Originally, gated was used on Linux systems. However, that's now a commercial product and the source code is no longer distributed. Instead, you have to use Zebra, which can handle RIP, OSPF and BGP for both IPv4 and IPv6 networks.

Installing and configuring Zebra for RIP is fairly straightforward, although using OSPF or BGP requires knowledge of how routers interact and how the protocol works. The main advantage of Zebra is that it has a configuration frontend very simple to IOS on Cisco boxes, so anyone familiar with using IOS to configure their routers can easily switch to Zebra without having to learn most of it all over again.

Docs and downloads for Zebra can be found at www.zebra.org.

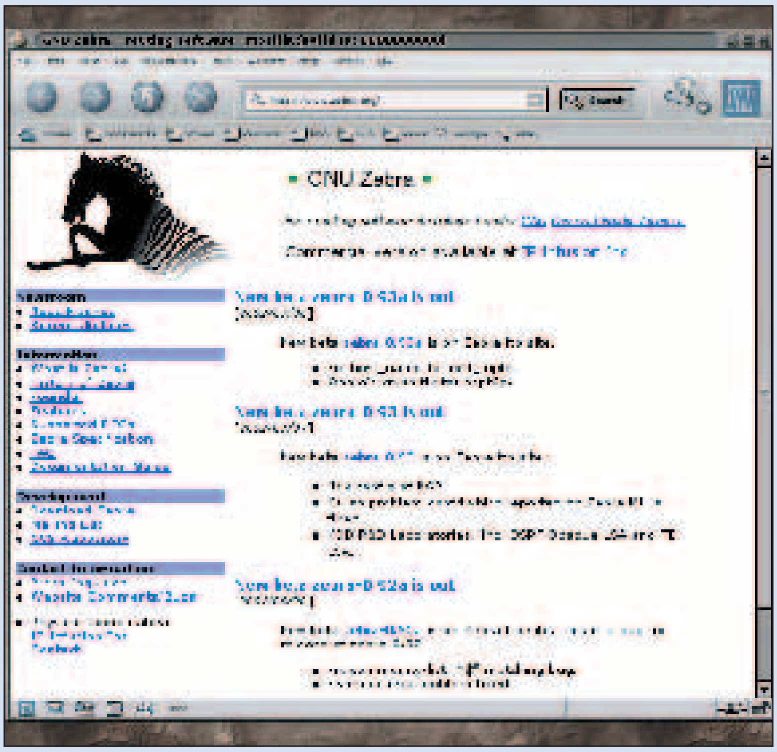

Zebra handles all of the major routing protocols that are commonly used by IPv4 and IPv6 networks.

LXF Sysadmin Answers – in association with Rackspace Managed Hosting

SYSTEM ADMINISTRATION

Ask an expert

Hans Huberland is Rackspace Managed Hosting's Linux guru. Email all your system administration and network questions and problems for our real-life system administrator to: sysadminqa@rackspace.co.uk Every published query on these system admin pages wins a Rackspace T-shirt!

ALL ABOUT RACKSPACE

When it comes to managed hosting Rackspace is unique. We are the only hosting company to guarantee 100 per cent network uptime and we will even pledge to replace faulty hardware (let's face it, it does happen) within an hour. As you would expect from Red Hat's only Advanced Hosting partner you will receive Fanatical Support™ and instant emergency response from fully qualified level three technicians, available 24-hours a day, 365 days a year. No answering machines, no diversions, no silly on-hold messages. All our customer platforms are housed in state-of-the-art, secure data centres and can be fully customised to meet any requirement. A dedicated account manager is allocated to you at the outset and will remain on your account. All Rackspace employees are committed to ensuring customer satisfaction, and will not rest until a client problem is resolved. Since Rackspace was established in 1998, it has specialised in managed hosting – nothing else. We're not an ISP, we don't sell domain names, and we won't offer website design – all we do is managed hosting. As a result, we do it very well and 97 per cent of our customers would happily recommend us. For information on how we can improve your web hosting, please contact us at www.rackspace.co.uk and ask any questions you may have about the ultimate managed hosting service.

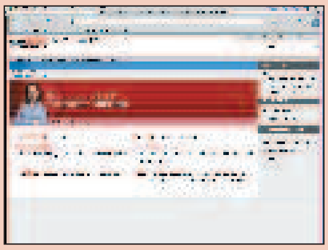

Members only

Q I hope my question here isn't too simple for your magazine column. I've been using *Apache* on my web server for some time. I must admit I found it quite difficult to configure from the command line but I eventually got it done, thanks to the help of a lot of kind-hearted Linux folk on the Internet.

I now need to add a secure area because our developers have made a members-only section. They want this to be SSL encrypted and I need to get an SSL certificate. I'm not sure how to proceed from here though. I've had a look on Google and I can't find a guide that's on a basic enough level for me.

Everything I want to do should be standard – I don't need to know about all the options and that's where I think I'm getting confused. Thanks in advance.

Nick (newbie!)

A Setting up an SSL-enabled website isn't nearly as complex as it seems at first. This can be divided into two tasks: getting the SSL certificate and configuring *Apache*. To set up the SSL certificate, you first need to generate a private key. Once generated, make sure you keep this key in a safe place because you'll need it if you ever need to regenerate your certificate or move your site to another server.

```
# cd  /etc/httpd/conf
# /usr/bin/openssl genrsa  1024 > ssl.key/mydomain-com.key
```

With this key you can generate a Certificate Signing Request (CSR). This needs to be sent to an SSL certificate provider (Thawte, Verisign and so on).

The following command will generate the CSR:

```
# /usr/bin/openssl req -new -key ssl.key/mydomain-com.key > ssl.csr/mydomain-com.csr
```

Enter your details as appropriate, taking special care to enter your domain name exactly as it will appear in your URL for the 'Common Name' – in other words, secure.mydomain.com or www.mydomain.com. Also, be sure to leave the 'Challenge password' blank. If you enter a password here, you'll need to enter this each time *Apache* starts up.

You can now head over to Verisign/Thawte and purchase a certificate. Be sure to enter the details you give them exactly as you entered them for the CSR you just generated. It will take them some time to verify your company and get back to you with your actual certificate.

When you receive your certificate, save it to you server under /etc/httpd/conf/ssl.crt/mydomain-com.crt.

Lastly, we need to tell *Apache* that this certificate exists and how to use it. Every certificate will require a dedicated IP address to listen on. Make sure that *Apache* is configured to use this IP address and is listed on port 443, then add a new Virtual Host block for your secure site.

Simply copy the details from the non-secure block and change the IP and port and add the following lines:

```
SSLEngine On
SSLCertificateFile /etc/httpd/conf/ssl.crt/mydomain-com.crt
SSLCertificateKeyFile /etc/httpd/conf/ssl.key/mydomain-com.key
```

At this stage, restarting *Apache* should bring your SSL site up. Verify this at **https://mydomain.com** by looking for the secure padlock icon in your browser.

Apache migration

Q We're currently migrating some websites from a 2.1ES server onto a new 3.0ES server. The main problem seems to be that 3.0 is using *Apache 2.0* rather than 1.3. Our websites all are all PHP based and receive substantial amounts of traffic. On the PHP website, there's a page that suggests you shouldn't really be using *Apache 2* and PHP in a production environment: http://uk.php.net/manual/en/faq.installation.php#faq.installation.apache2

So, my question is, what are my options? I presume I'm going to have to downgrade to version 1.3, but what are the consequences of doing this with regards to the up2date program?

Glen (from the Rackspace Forums)

A Rasmus's comments here are rather dated and meant more for when *Apache 2* was still brand new, less stable and had less (and less stable) module support. Also, the MPM model Red Hat uses is the default Prefork MPM, which is an order of magnitude more stable than the powerful but unstable worker MPM module.

That being said, you can be further reassured knowing that there are thousands of Red Hat Enterprise 3 production web servers running httpd-2.0 with very active PHP sites, usually with dynamic content from backend *MySQL* too. It runs perfectly.

As you're a Rackspace customer, if you have any specific code compatibility needs, please contact your support team about code compatibility issues between versions, or to ask about our code migration services, as well as bleeding edge options such as PHP5 and *MySQL4.x*. We've pre-built and tested packages that we frequently customise and install for customers.

If your situation does for some reason demand running *Apache 1.3*, this can be done because binaries are available in RPM format or can be compiled from source. You're quite correct in being concerned about up2date though – you'll need to add *Apache* to the package ignore list or it will be upgraded back to 2.0 as soon as up2date is run.

Learning to share

Q My problem is with mounting a network share from a Windows file server that I connect to from my Red Hat Enterprise Linux 3 system.

I mount the share as root.

```
[root@office root]# mount -t smbfs -o username=user,password=password \\\\fileserver.domain.com\\public /
```

WIN an Archos AV140 Video Recorder WITH RACKSPACE MANAGED HOSTING

www.rackspace.co.uk

Every month, the best question related to Systems Administration that a *LXF* reader sends in wins a prize: this month, you have the chance to win the awesome Archos AV140 Video Recorder.

The AV140 combines a modular MP4 video player and recorder, MP3 music player and recorder, digital camera and camcorder, digital photo wallet and data storage in a compact device that fits in the palm of your hand and is compatible with Linux. Based on a 40GB hard disk that is completely compatible with Linux, Windows and Mac OS, all you have to do is connect it to your computer and it will be mounted as an additional hard disk – on which you can store any type of data. USB 2.0 auto-connection ensures fast data transfer to and from your PC.

You can now enjoy 160 hours of near-DVD quality MP4 video with MP3 sound wherever you are. You can record and playback to your TV, VCR or camcorder directly in MP4 format. Playback is an AVI format file with MPEG4-SP VHS quality video and MP3 stereo sound, with the ability to record up to 2,000 hours of voice audio. For more specs, see **www.archos.com/ products/prw_ 500432_specs.html.**

```
mnt/fileserver/

[root@office mnt]# pwd
/mnt
[root@office mnt]# ll
total 16
drwxr-xr-x 2 root root 4096 Aug 15 19:31 cdrom
drwxr-xr-x 2 root root 4096 Aug 15 19:31 floppy
drwxr-xr-x 1 root root 4096 Nov 29 08:01 fileserver
```

I can view, edit and delete anything as root, but as an user on the system I can't do those options as I just get this message:

```
[dennis@office office]$ cd fileserver/
bash: cd: fileserver/: Permission denied
```

I've changed the group and the permissions of the directory, with no luck. If you have any suggestions, they would be much appreciated!
Dennis

A Dennis, all the credentials required to log onto the fileserver come from the command line you're using to do the mount. The permissions you have in place should be sufficient to allow the user to at least get a directory listing.

One thing I can pick up from the information you've given me is that you're trying to cd into the fileserver directory from the office directory and not from the /mnt directory:

```
[dennis@office office]$ cd fileserver/
```

Try the command again after running cd /mnt. If you're still having trouble, try getting a newer version of *Samba* – it's updated quite regularly on **www.samba.org**. There are binaries for Red Hat 9 that are fully compatible with EL3. You'll need to remove *Samba* and samab-common from the RPM database and install the single *Samba* RPM from ther site.

★ Star Question – AV140 winner!

This issue's lucky winner is ***Jay*** – your new portable multimedia player/recorder will be with you shortly!

Upload risks

Q **Hello. My Rackspace server hosts about 50 websites for a number of my customers. Most of them have some form of dynamic content, usually PHP based, while some use phpnuke and phpbb. I'm quite an experienced system administrator, if I say so myself, but I'm not a programmer. What level of risk is my server at by enabling my customers to upload their own PHP pages? Is there anything I can do to get better security from this?**
Jay

A Jay, because *Apache* doesn't run as root, your system shouldn't be wide open. However, if you have some bad code on your system, an attacker could still get a shell access and run commands, albeit without any privilege. Usually when someone has some exploitable PHP/CGI code, it enables you to import your own snippet of code by using remote URL execution (called 'fopen' in PHP).

Typically, this is fairly obvious when you manage to find the hack because there will be a backdoor process running as *Apache*. This will be listening on a high port and the binary will often still be left in /tmp or /var/tmp. Running through the access logs, you'll see where the hits were made and what commands they ran, usually wget'ing some C file and compiling it, then running it, thus spawning a backdoor.

You could really bolt down PHP to not allow much command execution at all, but this may be counter-productive. Many PHP-based applications, such as phpnuke, phpbb and so on, will require some loosening of restrictions to work. Ideally, sysadmins are supposed to keep an eye out for outdated software being loaded onto their servers, such as exploitable phpnuke or phpbb. This doesn't scale very well though, and as you get more users, this can become more difficult.

An alternative option is to set up a custom partitioning scheme where /tmp is a 'noexec' mounted partition, thereby preventing scripts from being executed when downloaded to /tmp. This can be implemented using a /tmp loopback file too (with /var/tmp symlinked) and it works really well. The only potential issues here are that tmp can fill up more easily since it doesn't have the full space allocation of the whole drive (this may be a feature though!), but if you start it at around 1GB, this should be large enough. Also, if /tmp is done as a mounted loopback file, the file size (partition) could be expanded to whatever size is necessary and then remounted.

Locking down apps like phpnuke can often cause more problems than it solves.

Secure SSH?

Q At my workplace we have a server running the usual Linux, *Apache* and *MySQL* combination, acting as a development and testing server for around about 100 sites we're building or have built. The sever is only open to access from the internal network, apart from SSH access to the outside world.

I now have to do some work from home but this needs to be done over a secure connection and SSH tunnelling seems like a very sensible method.

The problem is that the *Apache* server uses mod_rewrite to route http requests to the relevant site directory, but as I'd be connecting to the server through an SSH tunnel, I can't access the server through different hostnames.

Is anyone aware of a method I could use to see any of the sites without changing the server setup too drastically?

Ian Roberts

A Using SSH, you would port forward tcp/80 from the web server onto a port on the local system, such as 127.0.0.1:8080. Hosts can be maintained by modifying /etc/hosts and adding the appropriate sites and pointing them to 127.0.0.1.

An alternative to SSH would be to use IPSec, which would allow for the same DNS configuration. However, the firewall would have to allow IPSec tunnels to be established and the appropriate rules constructed. Applications such as *Vtun* and *OpenVPN* provide a similar capability using a user-space tool, although access to a system on the border of the network would be required.

There are several options available to provide secure remote access.

Rebooting bother

Q After installing Mandrake 10.1, eth0 is running well. However after reboot, I get this message: "Bringing up eth0: FAILED". Help!

```
%cat /etc/resolv.conf
search nsw.optushome.com.au
nameserver 203.2.75.132
nameserver 198.142.0.51

%lspci | grep Ethernet
00:0b.0 Ethernet controller: Realtek
Semiconductor Co., Ltd. RTL –
8129/8139C/8139C+ (rev10)
00:0c.0 Ethernet controller: Realtek
Semiconductor Co., Ltd. RTL –
8129/8139C/8139C+ (rev10)

%ifconfig eth0 192.168.0.11

%ping 192.168.0.5
PING 192.168.0.5 (192.168.0.5)
56(84) bytes of data.
From 192.168.0.11 icmp_seq=1
Destination Host Unreachable
From 192.168.0.11 icmp_seq=2
Destination Host Unreachable
From 192.168.0.11 icmp_seq=3
Destination Host Unreachable

--- 192.168.0.5 ping statistics ---
5 packets transmitted, 0 received,
+3 errors, 100% packet loss, time
3998ms
, pipe 3
%ping 192.168.0.11
PING 192.168.0.11 (192.168.0.11)
56(84) bytes of data.
64 bytes from 192.168.0.11: icmp_
seq=1 ttl=64 time=0.065 ms

--- 192.168.0.11 ping statistics ---
1 packets transmitted, 1 received,
0% packet loss, time 0ms
rtt min/avg/max/mdev =
0.065/0.065/0.065/0.000 ms

%ifconfig
eth0 Link encap:Ethernet HWaddr
00:02:44:11:DD:24
inet6 addr: fe80::202:44ff:fe11:
dd24/64 Scope:Link
UP BROADCAST RUNNING
MULTICAST MTU:1500 Metric:1
RX packets:0 errors:0 dropped:0
overruns:0 frame:0
TX packets:195 errors:0 dropped:0
overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:0 (0.0 b) TX bytes:33386
(32.6 Kb)
Interrupt:9 Base address:0x9f00

eth0:9 Link encap:Ethernet HWaddr
00:02:44:11:DD:24
inet addr:127.255.255.255
Bcast:127.255.255.255
Mask:255.0.0.0
UP BROADCAST RUNNING
MULTICAST MTU:1500 Metric:1
Interrupt:9 Base address:0x9f00

eth1 Link encap:Ethernet HWaddr
00:50:22:E9:8E:A4
inet6 addr: fe80::250:22ff:
fee9:8ea4/64 Scope:Link
UP BROADCAST MULTICAST
MTU:1500 Metric:1
RX packets:0 errors:0 dropped:0
overruns:0 frame:0
TX packets:23 errors:0 dropped:0
overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:0 (0.0 b) TX bytes:2538
(2.4 Kb)
Interrupt:11 Base address:0xae00

lo Link encap:Local Loopback
inet addr:127.0.0.1 Mask:255.0.0.0
inet6 addr: ::1/128 Scope:Host
UP LOOPBACK RUNNING
MTU:16436 Metric:1
RX packets:243 errors:0 dropped:0
overruns:0 frame:0
TX packets:243 errors:0 dropped:0
overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:20570 (20.0 Kb) TX
bytes:20570 (20.0 Kb)

%ifup eth0

Determining IP information for
eth0... done.

/sbin/ifup: line 433: 7771 Hangup /
etc/init.d/tmdns reload >/dev/null
2>&1

% /etc/init.d/network status
Configured devices:
lo eth0
Currently active devices:
lo eth1

%time /etc/init.d/network restart
Shutting down interface eth0: [ OK ]
Shutting down loopback interface: [
OK ]
Setting network parameters: [ OK ]
Bringing up loopback interface: [ OK
]
Bringing up interface eth0: [ OK ]
1.90user 0.66system 1:38.44elapsed
2%CPU (0avgtext+0avgdata
0maxresident)k
0inputs+0outputs (0major+64810mi
nor)pagefaults 0swaps

%ping 192.168.0.11
connect: Network is unreachable

%cat /etc/sysconfig/network-scripts/
ifcfg-eth0

DEVICE=eth0
BOOTPROTO=dhcp
ONBOOT=yes
MII_NOT_SUPPORTED=yes
NEEDHOSTNAME=yes
```

Adam

A For both eth0 and eth1, there are packets being transmitted but nothing being received. This suggests that DHCP requests are sent out but not responded to. Running dhclient from the command line to manually request an IP address for eth0 will output useful information, such as link failure or errors.

Why there is a 127.255.255.255 address on eth0:9 is anyone's guess. Checking out /etc/sysconfig/network-scripts and removing ifcfg-eth0:9 should remove this because it will interact with traffic on the loopback address of 127.0.0.1 since the netmask 255.0.0.0 includes this IP.

While the kernel does detect the Ethernet card, it doesn't suggest that it's working correctly. Running dmesg will show any kernel messages, indicating a timeout or other driver failure causing issues with DHCP. LXF

Submission advice

We are happy to answer all sorts of Linux related questions. If we don't know the answer, we'll find out for you! But in order to give you the best service, it helps a lot if you read the following submission advice.

- Please be sure to include any relevant details of your system. "I can't get X to work" doesn't really mean anything to us if we don't know things like what version of X you are trying to run, what hardware you are running on.
- Be specific about your problem. Things like 'it doesn't work' or 'I get an error' aren't all that helpful. In what way does something not work? What were you expecting to happen? What does the error message actually say?
- Please remember that the people who write this magazine are NOT the authors or developers of Linux, any particular package or distro. Sometimes the people responsible for software have more information available on websites etc. Try reading the documentation!

We will try and answer all questions. If we don't answer yours specifically, you'll probably find we've answered one just like it. We can't really give personal replies to all your questions.

Write to us at:
Linux Format, Future Publishing,
30 Monmouth Street, Bath BA1 2BW
or email: **lxf.answers@futurenet.co.uk**