A Convert to the Fold

Replacing Windows™ infrastructure with Linux™

David Newall

rebel.net.au

<davidn@rebel.net.au>
http://www.rebel.net.au/~davidn

ABSTRACT

A good friend and colleague asked me to help move his infrastructure from Windows to Linux. Ten out of thirteen machines running Windows, on a network with ADSL, Cable and 802.11b wireless, were converted to Linux. A firewall/router was installed, which in addition to improving security, provided traffic management policies which resulted in savings in cost of data. Over 150 email accounts across twelve virtual domains were transferred with minimal disruption. A pleasant yet unanticipated side-benefit was a freeing up of computing infrastructure, permitting functional expansion of his facilities.

1. Acknowledgement

This paper could never have been written without the full and frank assistance of my colleague, who has chosen to remain anonymous. I thank him for his trust and time, and for choosing to make the switch to Linux, which has to be a hard decision for anybody.

2. Introduction

A good friend and colleague asked me to help him move his infrastructure from Windows 2000 to Linux. He made this decision after Microsoft took six months to repair a serious security vulnerability in Windows [AD20040210][AD20040210-2][TA04-041A], and then when he applied the patch[MS04-007] his machines suddenly started sending large amounts of data to Russia. He had thirteen machines, ran his own software development and computer consultancy, and hosted twelve virtual domains (including e-shops) for clients. This was a major change in infrastructure, to put it mildly.

Over the years I had shown him unix a few times, and each time he was impressed by the incremental improvements in ease of use and aesthetics. Eventually he agreed it was as easy to install as Windows. He had discovered Open Source and SourceForge.net[SOURCEFORGE] and had a realistic idea of the effort involved in changing platform. He also had a number of new ideas that he wanted to develop, and pursuing them using an open source framework had an obvious financial appeal.

3. Designing a new system

Starting with an almost clean slate, we sat down to design the new system. Goals to be achieved were:

- Get away from Windows;
- Open solutions;
- Security;
- Ease of administration;
- Traffic management on external internet links; and
- Minimise surprise for customers.

3.1 New Network

The old facility had two separate networks, one comprised of Windows machines on an ADSL connected network with IP addresses assigned by the ISP, and the other comprised of Windows machines connected by 802.11b wireless to a home-grade Cable modem. The wireless access point provided DHCP service and performed NAT over the cable, and was principally used for web browsing. The fixed-address machines were a mixture of hosting servers, local servers and personal workstations. One machine, my colleague's personal, development workstation, (unwisely) crossed both networks.

We designed a new network comprising a central routing firewall with four segments. One segment is for official servers and is maintained at high-security. The second segment is for non-servers and is a medium-security segment. The final two segments are external internet links,

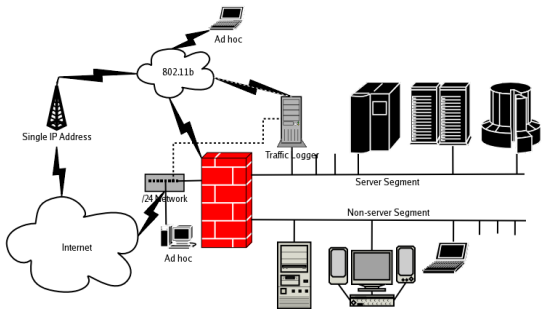

Figure 1: New Network Topology

one ADSL and one wireless link to cable; both are considered low-security.

3.1.1 External Links and Ad-hoc computers (Games)

The external links, almost by definition, have no security, and are suitable for ad-hoc networks for games, public access browsing and so on. Machines on the ADSL switch are given dynamic addresses from a small pool of public IP addresses, and route all of their packets via the ADSL modem. Machines on the wireless network are given dynamic addresses from a pool of non-routable addresses, and all of their packets are NATed via the cable modem.

3.1.2 Server Segment

The Server segment is the most restrictive environment. Access is granted according to specific type of services on specific machines. Machines are given fixed, public IP addresses. New network connections are preferentially routed via the ADSL link, but fall back to a NATed connection via the cable modem, if necessary.

3.1.3 Development Segment

The development segment has less restrictive protection. As a general rule, incoming connections are not permitted, but no restriction is made on connections to external machines. Most machines on the development network are dynamically assigned addresses from a pool of non-routable addresses. These machines have some limited trust relation with machines on the Server Segment, in that the server machines can be trust that the data really did come from a local machine.

3.1.4 Routing

Requests from the server segment are considered important, and so all connections are routed via the ADSL link if it's up, and otherwise fall back to a NATed connection via the cable modem. The development segment is considered less important, and for cost and load balancing purposes, it's web traffic is routed via the cable using a single NATed address. Other types of data are routed via ADSL, either using the machine's assigned public IP address, or by NATing the non-routable address using the router's own public address.

3.1.5 Devil-Linux

Prior to my arrival, my colleague spent time researching different Linux distributions, and found Devil-Linux[DEVIL] What he liked about it was that it runs directly from CD so there's no chance of programs being modified by intruders. (This concern might seem paranoid until you consider *why* he was replacing his infrastructure.) Devil stores configuration details on a floppy or USB disk and so when a machine dies, as they all eventually will, restoring service can be as simple as inserting the CD and configuration disks into a new machine and turning on power.

I particularly like Devil's configuration data being stored on a separate disk. It permits use of RCS to store different versions, and use of diff to examine changes prior to committing them. Bearing in mind that my colleague was somewhat new to unix, my being able to easily see his changes using one diff provided me with a sense of control and security that I might otherwise not have enjoyed.

I was concerned that the system would be too slow running off CD, but that hasn't been the case. The system is slower to initially load a program, but after first use it remains in cache and subsequent invocations are "immediate." I wouldn't recommend running unix from CD for a general workstation, but it does seem to work well for a dedicated purpose appliance.

3.1.6 fwbuilder

My colleague chose Fwbuilder [FWBUILDER] to maintain his firewall rules, and using a simple awk script, he also uses it to generate his DHCP and DNS tables. Fwbuilder is a simple GUI tool which you use to describe your network, with the ultimate aim of automatically generating a script for your firewall. In addition to Linux's *iptables*, it can emit rules for *Cisco PIX*, *ipf*, *ipfilter*, and *Berkeley packet filter*. Fwbuilder saves your data in an XML format, and this made it easy to write a script to extract MAC and IP addresses and generate DHCP and DNS tables for the complete network. As a nice touch, when you select "Install", fwbuilder calls our installation script which automatically builds new iptable rules, Bind DNS and ISC DHCP tables, transfers them to the firewall using SSH, and re-loads the appropriate services.

3.1.7 Network Logger

We built a logger for the two external links using a PC with extra network cards. Packets are logged using Ethereal[ETHEREAL] with log files rotated hourly. This is discussed in the interview, at the end of the paper.

4. Migrate Services

4.1 Email

The business provides email-account hosting for a number of domains, and we felt it was important to choose the email server wisely. Sendmail[SENDMAIL] has been the mainstay of mail servers for more than twenty years. It is powerful, flexible, and popular. Although the numbers are a year old, [CREDENTIA03] reported that Sendmail is still the leading mail transport, with more than twice as many servers as the second place getter, although it is losing market share[CREDENTIA02][BERNSTEIN]. Sendmail used to be my choice, and I have no doubt that it could have easily met my colleague's needs, but I advised against it. A 30 second glimpse at a Sendmail configuration file might help explain why.

```
#
# Copyright (c) 1998-2003 Sendmail, Inc. and its suppliers.
#       All rights reserved.
# Copyright (c) 1983, 1995 Eric P. Allman.  All rights reserved.
# Copyright (c) 1988, 1993
#       The Regents of the University of California.  All rights reserved.
#
# By using this file, you agree to the terms and conditions set
# forth in the LICENSE file which can be found at the top level of
# the sendmail distribution.
#
#

######################################################################
######################################################################
#####
#####           SENDMAIL CONFIGURATION FILE
#####
##### built by bhcompile@bugs.devel.redhat.com on Tue Oct 28 16:06:03 EST 2003
##### in /usr/src/build/320829-i386/BUILD/sendmail-8.12.10/cf/cf
##### using ../ as configuration include directory
#####
######################################################################
#####
#####   DO NOT EDIT THIS FILE!  Only edit the source .mc file.
#####
######################################################################
######################################################################

#####  $Id: cfhead.m4,v 8.108.2.3 2003/04/03 17:51:51 ca Exp $  #####
#####  $Id: cf.m4,v 8.32 1999/02/07 07:26:14 gshapiro Exp $  #####
#####  setup for Red Hat Linux  #####
#####  $Id: linux.m4,v 8.13 2000/09/17 17:30:00 gshapiro Exp $  #####



#####  $Id: local_procmail.m4,v 8.21.42.1 2002/11/17 04:25:07 ca Exp $  #####


#####  $Id: no_default_msa.m4,v 8.2 2001/02/14 05:03:22 gshapiro Exp $  #####

[...]
```

Table 1: 235 out of 1795 lines from /etc/mail/sendmail.cf

```
#####  $Id: proto.m4,v 8.649.2.24 2003/08/04 21:14:26 ca Exp $  #####

# level 10 config file format
V10/Berkeley

# override file safeties - setting this option compromises system security,
# addressing the actual file configuration problem is preferred
# need to set this before any file actions are encountered in the cf file
#O DontBlameSendmail=safe

# default LDAP map specification
# need to set this now before any LDAP maps are defined
#O LDAPDefaultSpec=-h localhost

##################
#   local info   #
##################

# my LDAP cluster
# need to set this before any LDAP lookups are done (including classes)
#D{sendmailMTACluster}$m

Cwlocalhost
# file containing names of hosts for which we receive email
Fw/etc/mail/local-host-names

[...]
```

```
###############
#   Options   #
###############

# strip message body to 7 bits on input?
O SevenBitInput=False

# 8-bit data handling
#O EightBitMode=pass8

# wait for alias file rebuild (default units: minutes)
O AliasWait=10

# location of alias file
O AliasFile=/etc/aliases

# minimum number of free blocks on filesystem
O MinFreeBlocks=100

[...]
```

Table 1: 235 out of 1795 lines from /etc/mail/sendmail.cf

```
############################
#   Message precedences    #
############################

Pfirst-class=0
Pspecial-delivery=100
Plist=-30
Pbulk=-60
Pjunk=-100

####################
#   Trusted users    #
####################

# this is equivalent to setting class "t"
Ft/etc/mail/trusted-users
Troot
Tdaemon
Tuucp

########################
#   Format of headers    #
########################

H?P?Return-Path: <$g>
HReceived: $?sfrom $s $.$?_($?s$|from $.$_)
        $.$?{auth_type}(authenticated$?{auth_ssf} bits=${auth_ssf}$.)
        $.by $j ($v/$Z)$?r with $r$. id $i$?{tls_version}
        (version=${tls_version} cipher=${cipher} bits=${cipher_bits} verify=${verify})$.$?u
        for $u; $|;
        $.$b
H?D?Resent-Date: $a
H?D?Date: $a
H?F?Resent-From: $?x$x <$g>$|$g$.
H?F?From: $?x$x <$g>$|$g$.
H?x?Full-Name: $x
# HPosted-Date: $a
# H?l?Received-Date: $b
H?M?Resent-Message-Id: <$t.$i@$j>
H?M?Message-Id: <$t.$i@$j>

#
######################################################################
######################################################################
#####
#####                   REWRITING RULES
#####
######################################################################
######################################################################

############################################
###  Ruleset 3 -- Name Canonicalization  ###
############################################
Scanonify=3

# handle null input (translate to <@> special case)
R$@                     $@ <@>

# strip group: syntax (not inside angle brackets!) and trailing semicolon
R$*                     $: $1 <@>                       mark addresses
R$* < $* > $* <@>       $: $1 < $2 > $3                 unmark <addr>
R@ $* <@>               $: @ $1                         unmark @host:...
R$* [ IPv6 : $+ ] <@>   $: $1 [ IPv6 : $2 ]             unmark IPv6 addr
R$* :: $* <@>           $: $1 :: $2                     unmark node::addr
R:include: $* <@>       $: :include: $1                 unmark :include:...
R$* : $* [ $* ]         $: $1 : $2 [ $3 ] <@>           remark if leading colon

[...]
```

```
# handle virtual users
R$+                     $: <!> $1               Mark for lookup
R<!> $+ < @ $={VirtHost} . >    $: < $(virtuser $1 @ $2 $@ $1 $: @ $) > $1 < @ $2 . >
R<!> $+ < @ $=w . >     $: < $(virtuser $1 @ $2 $@ $1 $: @ $) > $1 < @ $2 . >
R<@> $+ + $+ < @ $* . >
                        $: < $(virtuser $1 + + @ $3 $@ $1 $@ $2 $@ +$2 $: @ $) > $1 + $2 < @ $3 . >
R<@> $+ + $* < @ $* . >
                        $: < $(virtuser $1 + * @ $3 $@ $1 $@ $2 $@ +$2 $: @ $) > $1 + $2 < @ $3 . >
R<@> $+ + $* < @ $* . >

[...]
```

Table 1: 235 out of 1795 lines from /etc/mail/sendmail.cf

```
# handle local:user@host syntax -- ignore host part
R< $+ @ $+ > $* < @ $* >        $: < $1 > $3 < @ $4 >

# handle local:user syntax
R< $+ > $* <@ $* > $*           $#local $@ $2@$3 $: $1
R< $+ > $*                      $#local $@ $2    $: $1

###################################################################
###  Ruleset 93 -- convert header names to masqueraded form     ###
###################################################################

SMasqHdr=93


# do not masquerade anything in class N
R$* < @ $* $=N . >      $@ $1 < @ $2 $3 . >

R$* < @ *LOCAL* >       $@ $1 < @ $j . >

[...]
```

```
# check sender address: user@address, user@, address
R<$+> $+ < @ $* >       $: @<$1> <$2 < @ $3 >> $| <F:$2@$3> <U:$2@> <D:$3>
R<$+> $+                $: @<$1> <$2> $| <U:$2@>
R@ <$+> <$*> $| <$+>    $: <@> <$1> <$2> $| $>SearchList <+ From> $| <$3> <>
R<@> <$+> <$*> $| <$*>  $: <$3> <$1> <$2>               reverse result
# retransform for further use
R<?> <$+> <$*>          $: <$1> $2      no match
R<$+> <$+> <$*>         $: <$1> $3      relevant result, keep it

[...]
```

```
################################################
###   Local and Program Mailer specification   ###
################################################

#####  $Id: local.m4,v 8.58 2000/10/26 01:58:29 ca Exp $  #####

#
#  Envelope sender rewriting
#
SEnvFromL
R<@>                    $n                      errors to mailer-daemon
R@ <@ $*>               $n                      temporarily bypass Sun bogosity
R$+                     $: $>AddDomain $1       add local domain if needed
R$*                     $: $>MasqEnv $1         do masquerading

#
#  Envelope recipient rewriting
#
SEnvToL
R$+ < @ $* >            $: $1                   strip host part

#
#  Header sender rewriting
#
SHdrFromL
R<@>                    $n                      errors to mailer-daemon
R@ <@ $*>               $n                      temporarily bypass Sun bogosity
R$+                     $: $>AddDomain $1       add local domain if needed
R$*                     $: $>MasqHdr $1         do masquerading

#
#  Header recipient rewriting
#
SHdrToL
R$+                     $: $>AddDomain $1       add local domain if needed
R$* < @ *LOCAL* > $*    $: $1 < @ $j . > $2

#
#  Common code to add local domain name (only if always-add-domain)
#
SAddDomain
R$* < @ $* > $*         $@ $1 < @ $2 > $3       already fully qualified

R$+                     $@ $1 < @ *LOCAL* >     add local qualification

Mlocal,         P=/usr/bin/procmail, F=lsDFMAw5:/|@qSPfhn9, S=EnvFromL/HdrFromL, R=EnvToL/HdrToL,
                T=DNS/RFC822/X-Unix,
                A=procmail -t -Y -a $h -d $u
Mprog,          P=/usr/sbin/smrsh, F=lsDFMoqeu9, S=EnvFromL/HdrFromL, R=EnvToL/HdrToL, D=$z:/,
                T=X-Unix/X-Unix/X-Unix,
                A=smrsh -c $u
```

Table 1: 235 out of 1795 lines from /etc/mail/sendmail.cf

For someone switching from Windows, I think I can sum it up with one, simple statement: *Sendmail is complex*. It's true that most people configure Sendmail using the ".mc" files, and

that it can and does do everything you're likely ever to need a mail system to do, but it is also takes a lot of learning. There have also been a fair number of security incidents during it's long history. We looked at alternative mail servers and felt that Courier Mail Server[COURIER] would be better.

4.1.1 Courier

My colleague's needs were modest. His clients use Microsoft Outlook and Outlook Express, but don't use non-email functions, such as calendering. Therefore it was sufficient to provide SMTP, and POP3 or IMAP.

Courier is well documented and supports an extensive array of features. In addition to SMTP, IMAP and POP3 services, Courier includes Webmail, a minimal but functional web interface for sending and reading mail, which is a nice extra. Courier stores mail in Maildir format [BERNSTEIN2], which is much faster than traditional unix mailboxes (mbox). It supports virtual email addresses, quotas, LDAP, comes with a mailing list software, a rich filtering facility and can send and receive faxes. Users can be authenticated using a wide variety of mechanisms, and it's easy to write new mechanisms.

4.2 Web

My colleague originally had a number of virtual domains which he hosted, using a mixture of ASP and static pages. The plan was to move the static pages to Apache, and indeed Apache was loaded onto the new server, and to install a new IIS machine for the pages which couldn't be moved to Apache. Due to a late change in direction, my colleague instead chose to retire some applications and move others to third party servers. Apache is still running but is waiting for pages to serve!

5. Post Project Interview

Four months after completing the project I interviewed my colleague to hear his thoughts about the change. Here is an edited transcript of that interview.

5.1 Overall Impression

Q. What's the best thing about moving to unix?

A. The best thing? I think the thing that strikes me most in my mind, is the one to know that there isn't any code in any system that I'm running that hasn't been looked at by a number of pairs of eyes, which means there aren't secret spies in the code; and if there was one found it would be so devastating for the people who tried to get away with it and the companies who were associated with them, and anybody associated with it. The code is there so you can see it.

So it's the number of eyes, I'm really happy about the number of eyes that have looked at this thing before me so that I don't have to be as concerned.

Q. What's the worst thing about moving to unix?

A. I suppose it's the unknown. Like you know your backyard but there's always that little corner that you don't really know. Well, in linux, the corners are a lot bigger than in Windows. I know it works, I'm happy to install the upgrades, but I have to dedicate months of my life to some sub-component to really get to know it.

Probably the only other thing that would give me a bit of angst is: I've reset my clock once. When does the same amount of threat that Microsoft is enduring now, when does the same sort of threat raise it's ugly head at Linux? Has Linux got five years before it suffers the same threat as Microsoft or has it got 20 years?

Q. What was the hardest thing about moving to unix?

A. The short answer is that there's a lot there. It's the whole back yard thing.

5.2 Current Use

Q. What machines are you running now?

A. Well right now the only ones I've got running are my new firewall, my web and email server, my samba server, and my windows game machine. And a notebook, well I've got three notebooks, and my windows development machine. There's another eight machines I don't have any use for at the moment. I've got some ideas for them like backup servers.

Q. How many Windows servers do you still run?

A. None. I've changed my entire focus. Now I'm focusing on developing a whole lot of ideas I've had over the years.

With Windows letting me down so harshly, I had an awakening of whether or not I really wanted to provide services to people. It's hard enough for Telstra and major corporations to do these things well, when you get flaky this and flake that that it becomes very, very painful, so I pretty much made the active decision to leave serving for others. The only things I'll be serving in the future are my own things. I'll serve my own web sites and my own email and my own this and my own that, so the only person who can be damaged by any sort of failure will be myself and maybe immediate family, no external company. What that new focus effectively means is that a lot of the stuff I was doing is no longer

going to be necessary because it all came from the fact that I was putting together custom serving environments for customers.

Q. How central has unix been to your decision to change focus?

A. Well the ideas I've got are going to build on top of a lot of tools, and all of those tools are available in unix for free. I couldn't have afforded all those tools on a Windows system.

Q. What services did you shut down?

A. Web serving. There were thirteen domains, nine were burned to CD and handed back, and the other four have been moved somewhere else. In fact at the time I was going to put in a Windows server in to run the Windows domains, and the others, I could have done them on Apache. I'd still consider putting in an IIS server, but just make sure it was very, very tight.

5.3 Samba

The big one for me was Samba, because otherwise it's too hard to move stuff around. It's got to be easy to move stuff around otherwise it's too hard to play with. Primarily Samba lets all of my machines, based on who's using them and where they are, share data. If it wasn't for Samba I'd be running two networks: A windows network and a unix network and never the two shall meet. It's the glue that binds them together (other than whatever application is on them.)

Q. How do you share printers?

I don't print directly from linux at this stage because the bulk of my printing comes from windows.

5.4 Network Use and Security

Q. How has your traffic use changed?

A. I use about 1.2 to 1.8G a month. It turns out that before, about half of the traffic was the outside world trying to break in. The firewall keeps that out. My usage has about halved. I still do the same email, my usage there hasn't changed much. Maybe it's gone up a little as they put on a new staff member but it hasn't changed much.

Q. Have there been any security incidents since you converted, and if so, have you been affected by them?

A. There have been some. I've got my firewall turned down so tight, even to the point of having special permission to contact the Windows patch server. Nothing goes on now unless I allow it.

What I do find is that when something stops working I was far more used to pulling out a windows machine and checking its network or monitor cable. Now there are different levels of routers and firewalls, and so finding a fault, there's a lot more things to look at to get the final verdict.

Q. Are you happy to have a level of control that detailed?

A. No, I wish I could do my job without it.

Q. Tell me about your network logger.

Over the last couple of months I've had a few times when an email wasn't getting through for one of my clients, so for a couple of hours I captured every packet and I narrowed it down to the particular IP address and and no, I'm perfectly fine, it's their end that's got the problem. Maybe seven or eight cases like that where the only way I could really truly know was by looking at the traffic.

A couple of weeks ago there was another Microsoft thing when all of a sudden the traffic crept up a bit on my only three Windows machines and it was just the people out there trying it on, but the machines behind the linux firewall were unaffected. I also looked at packets coming to the firewall that were trying to break in to windows machines which I keep on the outside of the firewall and don't have the benefit of the security.

No doubt there's many tools like it, but the graphic interface, where you can colourise it, you can look through huge amounts of stuff and sort it. It's very, very powerful. I used to have my notebook doing nothing but sniffing packets so that I could take it to clients. A thousand dollar notebook and Linux used for nothing else.

The beautiful part of it is my logger just copies hour long chunks of data, and even that alone, you can run your eye down it and see what's going on, and if you're really keen you can find out exactly what's going on. Up to this point in my life I've never had everything before me. You used to feel like you could look and touch and maybe think you knew what was going on but you never really did know what was going on. The way that Ethereal breaks the packets down and understands the packets, it's not just saying here's a packet and it's from this address to that address. There's just huge amounts of data which it gives.

Half the battle, of course, is security. With my early Windows environments, I could get in, but have each machine unto itself and I'd only do what was required on each machine. The problem was that the patches and the vulnerabilities for windows were coming out so quickly, and trying to keep ten machines unbelievably, perfectly aligned with no gaps. Somehow I got a gap in there and someone got

in and re-jigged me so that it was no longer a single virus but it was my security had been compromised and once you've got that you might as well reformat and start again. But why do it with windows?

5.5 Final Thoughts

Q. Is there anything you'd have done differently?

A. No. The only thing that might have changed was because Sourceforge had a certain set of software available at that particular time. At another time it might have been a different choice, or perhaps something might not have been available, but apart from that I'd have done it the same.

Q. What advice would you give someone contemplating making the same switch?

A. Go to Sourceforge. Spend about a month looking at what's there, and what you need, because everything's there and you've got multiple choices for everything. It doesn't really matter which one you choose, so just pick one and go with it.

References

[AD20040210]
eEye Digital Security Advisory AD20040210, *Microsoft ASN.1 Library Length Overflow Heap Corruption*, http://www.eeye.com/html/Research/Advisories/AD20040210.html

[AD20040210-2]
eEye Digital Security Advisory AD20040210-2, *Microsoft ASN.1 Library Bit String Heap Corruption*, http://www.eeye.com/html/Research/Advisories/AD20040210-2.html

[BERNSTEIN]
D. J. Bernstein, *Internet host SMTP server survey*, http://cr.yp.to/surveys/smtpsoftware6.txt

[BERNSTEIN2]
D. J. Bernstein, *Using maildir format*, http://cr.yp.to/proto/maildir.html

[COURIER]
Double Precision, Inc, *Courier Mail Server*, http://http://www.courier-mta.org

[CREDENTIA02]
Credentia, *E-Mail Server Survey Results*, http://www.credentia.cc/surveys/smtp/200212/

[CREDENTIA03]
Credentia, *E-Mail Server Survey Results for April 2003*, http://www.credentia.cc/surveys/smtp/200304/

[DEVIL]
Devil-Linux, http://www.devil-linux.org

[ETHEREAL]
Ethereal, http://www.ethereal.com

[FEDORA]
Fedora™ Project, *Fedora Core 1*, http://fedora.redhat.com

[FWBUILDER]
Firewall Builder, http://www.fwbuilder.org

[MS04-007]
Microsoft Security Bulletin MS04-007, *ASN.1 Vulnerability Could Allow Code Execution (828028)*, http://microsoft.com/technet/security/bulletin/MS04-007.asp

[SENDMAIL]
The Sendmail Consortium, *Sendmail*, http://sendmail.org

[SOURCEFORGE]
SourceForge.net, http://sourceforge.net

[TA04-041A]
US-CERT Technical Cyber Security Alert, TA04-041A, *Multiple Vulnerabilities in Microsoft ASN.1 Library*, http://www.us-cert.gov/cas/techalerts/TA04-041A.html