

The Apprentice's Notes

Jurjen Bokma

The Apprentice's Notes

Jurjen Bokma

Table of Contents

Introduction	x
I. 2007	1
1. June 2007	4
On writing DocBook documentation	4
On writing WebPlatform Documentation from Linux	4
On creating ERDs for PostgreSQL	5
On installing Eclipse under Debian	5
On running a Subversion server	5
On Using the G: drive under Linux	7
On using iPrint printers from Linux	8
On using SystemImager	8
Enabling X11 forwarding for SSH by default	9
HTML redirect	9
2. July 2007	10
Initial RAMdisks in SystemImager	10
On the need for a proper unattended partitioning tool	10
On creating a Debian client for the IWI	10
On update-alternatives under Debian	10
Dissecting RPM packages	11
IBM Tivoli on Debian	11
On the Bootleg host imager	15
PXE boot	15
The bootleg boot script	16
Drawing a bundle	17
How to determine what a host boots	17
On moving DHCP from an endangered host	17
On moving anti-spam and viruschecking from an endangered host	18
On using Badblocks with ReiserFS	20
Finding the bad area using badblocks	20
Notifying the ReiserFS of the bad area	21
On the Syslog-NG log server	23
Creating a Syslog-NG server	24
Creating a Syslog-NG client	24
On configuring CfEngine	25
The server part	25
The client part	27
Creating a yum repository under Debian	30
On using Procmail and Vacation	31
3. August 2007	32
Screen rotation in X with NVidia	32
Firewalling with FWBuilder	32
Resetting IPtables	32
Firewall Rules	32
4. September 2007	33
Configuring CUPS clients	33
Configuring NTP clients	33
RedHat Certification	34
Using dmesg to stop log cluttering console	35
Installing new machines for the IWI in 2007	35
Setting global key bindings in Emacs	36
OWL	36
Links	36
Using NCPMOUNT to access central storage	39

5. November 2007	40
CPU buying advice	40
The problem at hand	40
Theoretical solutions	40
A poor but practical solution	40
Some more useful CPU bookmarks	41
Trusted/Treacherous Computing	41
Stopping the IWI.net mail service	41
The current situation	41
The new situation at a glance	42
The Transition	42
Mounting USB devices from udev with permissions of the user at the console	50
Tracking down the authoritative nameserver	52
Installing Mathematica 6.0 under Linux	52
The problem	52
The solution	53
Centralizing the DHCP service	53
The problem	53
The Implementation	53
No Remote Login possible at the IWI	54
The Problem	54
Investigation and solution	54
II. 2008	55
6. January 2008	58
A Configuration Repository idea	58
Introduction	58
What <i>is</i> configuration?	58
What is <i>not</i> configuration?	58
Storing Configuration and the Unit of Configuration	58
Parametrization and its implications	60
Cups Command Line Options	60
No Sound on the student PCs (unresolved)	60
In search of a proper Keyboard	60
Imaging Linux boxes with Zenworks imaging	61
Can ZENworks imaging be used to clone a Linux machine (and if so, how)?	61
What are the restrictions ZENworks imaging puts on the way a machine is partitioned?	62
What limitations does ZENworks imaging impose on the filesystems used?	62
Can we use ZENworks imaging to put Linux in a designated space on a harddisk without damaging any other OSes or data already present on the disk?	62
Does ZENworks imaging support RAID? LVM?	62
Can we safely use ZENworks imaging on machines with multiple disks?	62
No SSH to my server possible	63
Creating a parser with Bisonc++ and flex	63
7. February 2008	64
Slapd takes 100% CPU on sched_yield()	64
SSH tunneling	64
Using Bacula for backup	64
debmirror on Ubuntu	65
Newest OpenSSL and BIND on a 64-bit Debian machine	65
Multipath Fibrechannel interface to SAN under Ubuntu	65
X access from under sudo	67
Quick source NAT with IPTables	68
8. April 2008	69
Installing the SuSE iPrint client under Debian	69
Converting the Novell iPrint client to a Debian package	70
9. May 2008	75
Installing OpenBSD on a Soekris Net5501-70	75

10. June 2008	83
Installing Linux over Windows without BIOS access	83
Turning nVidia driver on on machines that have the libraries and an NVidia card ..	84
A Firewall Install Script	84
Fixing the NIS port	86
Transferring the IWI printers to IPrint	86
Remote Firefox acutally remote	87
11. July 2008	88
Installing SpaceWalk (using a remote database)	88
Installing CentOS unattendedly	92
Remote access to Windows XP from Linux	93
Creating a SpaceWalk Channel	94
12. August 2008	96
Booting Ubuntu Hardy unattendedly using preseed	96
Cloning NTFS partitions at the file level (and booting them)	96
Introduction	96
Expectations	97
The experiment	97
Conclusion	98
Using DHCP-initialized PXELinux under VMWare	99
Labelling partitions during Linux unattended install	99
Introduction	99
Debian/Ubuntu	99
Reverse Engineering the ERD of an Oracle database	99
Using WebDav to connect to the so-called Y:-drive	100
XML versus web template engines	100
Installing 64-bit Matlab on Linux	101
Installing 64-bit Maple 11 under Linux	102
Installing 64-bits Mathematica on Linux	108
13.	112
Unable to mount USB disk under Debian	112
III. Appendices	113
A. Indices	115

List of Figures

5.1. Sketch of current mail flow at the IWI	42
5.2. Sketch of mail flow at the IWI with all forwards in place	45
5.3. Sketch of mail flow at the IWI with MX records redirected	47
5.4. Sketch of mail flow at the IWI with users talking to CIT servers	48
5.5. Sketch of mail flow at the IWI with iwi200 off	49
12.1. Screenshot of Maple Install: Introduction	103
12.2. Screenshot of Maple Install: Choose Install Folder	103
12.3. Screenshot of Maple Install: Choose Type of Licensing	104
12.4. Screenshot of Maple Install: Pointing out the License Server	105
12.5. Screenshot of Maple Install: Pre-Installation Summary	106
12.6. Screenshot of Maple Install: Installation Complete	107

List of Tables

2.1. Units of disk space used by programs involved in a ReiserFS badblocks detection	23
6.1. Keyboard Features	60

List of Examples

2.1. A Tivoli wrapper script	13
2.2. A Tivoli user options file	14
2.3. A Tivoli system options file	14
2.4. Bootleg intervention	16
2.5. Example of sethostip usage	17
2.6. Log of bad sectors on iwi202	20
2.7. SCSI errors in the log	23
2.8. Example cfsvr.d.conf	28
2.9. Example update.conf	28
2.10. Example cfagent.conf	29
4.1. Snippet from ntp.conf showing “iburst” flag to “server” statement	33

Introduction

These are the notes I take while I work. If you come across them, you are free to use them to your advantage. I do not, however, write them for anyone but myself. So you may find old, deprecated, clumsy, stupid and just plain incorrect statements and procedures here.

Warning

Your fault if you take my notes for granted.

Part I. 2007

Table of Contents

1. June 2007	4
On writing DocBook documentation	4
On writing WebPlatform Documentation from Linux	4
On creating ERDs for PostgreSQL	5
On installing Eclipse under Debian	5
On running a Subversion server	5
On Using the G: drive under Linux	7
On using iPrint printers from Linux	8
On using SystemImager	8
Enabling X11 forwarding for SSH by default	9
HTML redirect	9
2. July 2007	10
Initial RAMdisks in SystemImager	10
On the need for a proper unattended partitioning tool	10
On creating a Debian client for the IWI	10
On update-alternatives under Debian	10
Dissecting RPM packages	11
IBM Tivoli on Debian	11
On the Bootleg host imager	15
PXE boot	15
The bootleg boot script	16
Drawing a bundle	17
How to determine what a host boots	17
On moving DHCP from an endangered host	17
On moving anti-spam and viruschecking from an endangered host	18
On using Badblocks with ReiserFS	20
Finding the bad area using badblocks	20
Notifying the ReiserFS of the bad area	21
On the Syslog-NG log server	23
Creating a Syslog-NG server	24
Creating a Syslog-NG client	24
On configuring CfEngine	25
The server part	25
The client part	27
Creating a yum repository under Debian	30
On using Procmail and Vacation	31
3. August 2007	32
Screen rotation in X with NVidia	32
Firewalling with FWBuilder	32
Resetting IPtables	32
Firewall Rules	32
4. September 2007	33
Configuring CUPS clients	33
Configuring NTP clients	33
RedHat Certification	34
Using dmesg to stop log cluttering console	35
Installing new machines for the IWI in 2007	35
Setting global key bindings in Emacs	36
OWL	36
Links	36
Using NCPMOUNT to access central storage	39
5. November 2007	40
CPU buying advice	40

The problem at hand	40
Theoretical solutions	40
A poor but practical solution	40
Some more useful CPU bookmarks	41
Trusted/Treacherous Computing	41
Stopping the IWInet mail service	41
The current situation	41
The new situation at a glance	42
The Transition	42
Mounting USB devices from udev with permissions of the user at the console	50
Tracking down the authoritative nameserver	52
Installing Mathematica 6.0 under Linux	52
The problem	52
The solution	53
Centralizing the DHCP service	53
The problem	53
The Implementation	53
No Remote Login possible at the IWI	54
The Problem	54
Investigation and solution	54

Chapter 1. June 2007

On writing DocBook documentation

DocBook can be edited using various editors [<http://wiki.docbook.org/topic/DocBookAuthoringTools>] We use the nXML mode package [<http://www.thaiopensource.com/nxml-mode/>] (also here [<http://ourcomments.org/Emacs/nXhtml/doc/nxhtml.html>]) to author DocBook under Emacs [<http://infohost.nmt.edu/tcc/help/pubs/nxml/index.html>]. See also my DocBook-under-Debian-HOWTO [<http://www.cs.rug.nl/~jurjen/iwi-howtos/linux/DocBook-under-Debian-HOWTO.html>] for a more in-depth explanation. A few handy shortcuts in nXML-mode:

- **c-enter**: tab-completion
- **c-c c-n**: goto next error
- **c-c c-s c-t**: choose schema when starting a new file
- **c-c c-f**: match last tag we are in with end-tag
- **ESC TAB**: finish current tag (or show possible options)

A good reference is DocBook: The Definitive Guide [<http://www.docbook.org/tdg/en/html/docbook.html>]. A list of templates would be handy. DocBook XSL: The Complete Guide [<http://www.sagehill.net/docbookxsl/index.html>] is another excellent piece of documentation, as are the docs that come with the “docbook-xsl-doc” package under Debian.

Conversion to XHTML is done with

```
xsltproc \  
--stringparam base.dir $(@D)/\  
--stringparam use.id.as.filename 1 \  
--stringparam root.filename "" \  
--stringparam chunker.output.encoding UTF-8 \  
/usr/share/xml/docbook/stylesheet/nwalsh/xhtml/onechunk.xsl \  
file.xml
```

in which case the id of the top level tag is the basename of the file that is generated ¹

On writing WebPlatform Documentation from Linux

One may use VMWare, install Windows and use Internet Explorer to write docs in the online Xopus editor. The LWP pages can be found at <http://www.rug.nl/rc/doelgroepen/medewerkers/LINUXwerkplek> [<http://www.rug.nl/cit/doelgroepen/medewerkers/LINUXwerkplek>]. Recursive download using wget:

¹ It also gets an extension: “html” in this case.

`wget -r --no-passive-ftp --ftp-user=p012345 --ftp-password=passwd -Dwebplatform.rug.nl -no-parent -S --directory-prefix=${HOME}/downloaded/WebPlatform-mirror/ftp://webplatform.rug.nl/webroot/dev/rc/doelgroepen/medewerkers/LINUXwerkplek/`. Single items can be picked using plain `ftp`. The WebPlatform uses an ftp server connected to an Oracle database. XML files are validated upon upload, and must pass validation in order to be stored. The XML schema used is home-made specifically for the WebPlatform. Editors like `oxygen` can do validation on them, but only partially. Unfortunately, the schema is stored in multiple files on the server, and only the top-level file is referenced in the header of the XML documents. The top-level schema uses local references to point to lower level files, and these will not be found when validating off-server. I see false negatives (local validation invalid while file is accepted by server). False positives may also occur. This makes off-server editing of files unfeasible. There is no password-less upload, although key-based login is being worked on.

The HTML generated from DocBook can be inserted into an XML template, and this will display properly. However, some features² will not work, and I consider this a kludge.

On creating ERDs for PostgreSQL

A good, free ERD editor for PostgreSQL is called “clay” [<http://www.azzurri.jp/en/software/clay/index.jsp>] and made by “Azzurri” [<http://www.azzurri.jp/en/index.jsp>]. It works under eclipse [<http://www.eclipse.org/downloads>]. It is distributed in zip-files which can be unzipped into the eclipse directory, and it needs the GEF (Graphical Editor Framework) [<http://www.eclipse.org/gef>], which can be downloaded from the Eclipse site.

On installing Eclipse under Debian

Eclipse is installed using just `apt-get install eclipse`. However, the SUN Java [<http://java.sun.com/j2se/1.5.0/download.jsp>] it needs has to be installed by downloading it separately, and using “java-package” and fakeroot to create a package, like this: `fakeroot make-jpkg jre-1_5_0_03-linux-i586.bin`. The created package can then be installed. This is all documented nicely here [<http://www.debian-administration.org/articles/142>]. Under 64-bit linux, using Debian Etch, once Eclipse starts, it doesn't find the java vm because it doesn't search in the `/usr/lib64` tree. Add a line reading `/usr/lib64/j2re1.5-sun` to `/etc/eclipse/java_home` and it will be found.

On running a Subversion server

Subversion [<http://subversion.tigris.org>] is the Version Control System of choice for most people in the environment I work in (Computing Science students and systems administrators). A good source of documentation is the Red Bean Book [<http://svnbook.red-bean.com/nightly/en/index.html>]. Subversion is easily installed by `apt-get install subversion`.

Procedure 1.1. Creating a SVN repository

1. Create the repository on the server: `svnadmin create /var/lib/svn/grid`
2. Import the layout of the repository: `svn import template/ file:///var/lib/svn/grid -m "initial import"`

²xref-tags for example

```

Adding    template/workernode
Adding    template/workernode/trunk
Adding    template/workernode/branches
Adding    template/workernode/tags
Adding    template/iserv
Adding    template/iserv/trunk
Adding    template/iserv/branches
Adding    template/iserv/tags
Adding    template/master
Adding    template/master/trunk
Adding    template/master/branches
Adding    template/master/tags

```

Committed revision 1.

Procedure 1.2. Creating a Subversion server using Apache(version 2)

1. Create a separate system group 'svn': **addgroup --system --group svn**
2. Change ownership of created repositories: **chgrp -R svn:svn /var/lib/svn/**
3. Add the apache user to the 'svn' group³: **adduser www-data svn**
4. Add some more users to the 'svn' group: **adduser username svn**⁴
5. Install mod_dav by **apt-get install libapache2-svn**.
6. Edit /etc/apache2/mods-available/svn.conf to contain

```

<Location /svn>
DAV svn

# any "/svn/foo" URL will map to a repository /usr/local/svn/foo
SVNParentPath /var/lib/svn

AuthType Basic
AuthName "Subversion repository on myhost.com"
AuthUserFile /etc/subversion/passwd

<LimitExcept GET PROPFIND OPTIONS REPORT>
Require valid-user
</LimitExcept>

Order deny,allow
Deny from all
# Localhost
Allow from 127.0.0.1

```

³ We are assuming here that Apache is installed in /usr/local.

⁴ I think membership of this group is only important if it is required by Apache asking for it in a 'require' statement. When a user logs in via the web server, writing of the repositories is done by the web server. But it is nice anyway in case a user wants to do file-based access instead of via the web server.

7.
 - a. Enable the configuration: `cd /etc/apache2/mods-enabled && ln -s ../mods-available/svn.conf ./`
 - b. And reload the Apache configuration: `/etc/init.d/apache reload`
8. Create a new passwd file for Apache and put a username/passwd in it. `htpasswd -cs /etc/subversion/passwd username`

Procedure 1.3. Using the Apache SVN server

1. Check out ⁵ an (empty) working copy with `cd / && svn co http://myserver.com/svn/repository/project/trunk/ ./`
2. Add files to subversion control using `svn add /etc`
3. Commit it like this: `svn commit -m"Adding /etc"`

Procedure 1.4. Enabling ssh access to the Subversion repository

- Nothing needs to be done. Just link `/svn` to `/var/lib/svn` for convenience and use it saying `svn co svn+ssh://myhost.com/svn/repository/project/trunk ./`

On Using the G: drive under Linux

The G:-drive is an ncp volume. In order to mount is, ncpfs is needed: `apt-get install ncpfs`. Then mount it saying: `/usr/bin/ncpmount -S Cluster02_usr04_server -U p012345.staff.rug.nl -P PassWd -V /Usr04/Acc/P012345 -t 900 -r 30 -A cluster02-usr04.staff.rug.nl /home/username/g`^{6,7}

⁵ I tried to use the `svn import` command to do the initial fill of the project, but it failed giving a 405 (resource not allowed) error on the PROPFIND request in the Apache logs. The checkout and commit works fine so I won't spend time on it.

⁶ Both the `-S` and `-A` options need to be specified, with arguments.

⁷ The argument of the `-S` and `-V` options can readily be derived from the comments about the Server Volume Name (following the Drive letter in Windows Explorer. The argument to the `-A` option takes some black magic if these comments are well formed and informative, or may have to be sniffed or pried out of a DOS box.

On using iPrint printers from Linux

The University has a site about printing with iPrint [<http://www.rug.nl/fwn/voorzieningen/ictbeheer/werkplek/PrintViaiPrint>] which leads you to a list of printers [<http://iprint-02.id.rug.nl/ipp>].

Note

Please note that the IP numbers of printers do change every now and then.

Note

Please also note that printing straight to the printer -although is is the *only* suitable way- is frowned upon by the maintainers of the iPrint system.

⁸ Under the 'Information' icon at the end of each line, find all information you need to add this printer to e.g. CUPS [<http://www.cups.org>].

On using SystemImager

SystemImager [<http://sourceforge.net/projects/systemimager/>] has a Wiki [http://wiki.systemimager.org/index.php/Main_Page].

Procedure 1.5. Installing the server

1. **apt-get install systemimager-server systemimager-server-flamethrower** installs the server.
2. In `/etc/systemimager/systemimager.conf` the variables `DEFAULT_IMAGE_DIR` and `TFTP_DIR` must be set to sensible values. `NET_BOOT_DEFAULT` Needs to be set to "LOCAL" to make the client install once and boot from disk later.
3. **si_mkbootserver** `--interface=eth2` `--localdhcp=y` `-kernel=/var/lib/systemimager/boot/i386/glite-WN/kernel` `-initrd=/var/lib/systemimager/boot/i386/glite-WN/initrd.img` `--tftpdire=/var/lib/tftpboot` `-pxelinux=/usr/lib/syslinux/pxelinux.0` to create DHCP config
4. **si_mkclientnetboot** `--verbose` `--netboot` `--clients "node116"` to create `/var/lib/tftpboot/pxelinux.cfg/0A000374`
5. **si_getimage --golden-client 10.0.3.116 --image node116**
6. **si_addclients**

Procedure 1.6. Creating the golden client

⁸ Unfortunately, this URL might not be accurate at all times. As of this writing, another one, with a "5" instead of a "2" [<http://iprint-05.id.rug.nl/ipp>] is more up-to-date, but not linked to.

1. **apt-get install systemimager-client** installs the client on a Debian host. In other distributions this may not be so straightforward.
2. **pushd /usr/share/systemimager/boot/ && ln -s i386 x86_64 && si_prepareclient --server 10.0.3.1 -yes**
3. **si_prepareclient --server 10.0.3.1**

Enabling X11 forwarding for SSH by default

Put the following in `~/.ssh/config`:

```
ForwardX11 yes
```

According to the manual, options can also be set on a per-host basis, like this:

```
<some general settings>
```

```
Host host1
ForwardX11 yes
<some more host-specific settings>
```

```
Host host2
ForwardX11 no
```

HTML redirect

```
<!DOCTYPE html PUBLIC "-//W3C//DTD HTML 4.01 Transitional//EN"
"http://www.w3.org/TR/html4/loose.dtd">
<html>
<meta http-equiv="refresh" content="3; URL=http://www.cs.rug.nl/~jurjen/Apprenti
<head>
<title> Redirecting to http://www.cs.rug.nl/~jurjen/ApprenticesNotes.html </titl
</head>
<body>

The <a href="http://www.cs.rug.nl/~jurjen/ApprenticesNotes.html">Apprentice's No

<Redirect / http://www.cs.rug.nl/~jurjen/ApprenticesNotes.html>

</body>
</html>
```

Chapter 2. July 2007

Initial RAMdisks in SystemImager

Documentation at IBM [<http://www.ibm.com/developerworks/linux/library/l-initrd.html>] indicated that initial RAMdisks used to be created using an actual loop-mounted RAMdisk and `gzip`. However nowadays they are made using an ordinary directory and `cpio`, like this: `find ./ | cpio -H newc -o > /tmp/new-initrd.cpio`. Still it is hard to get systemimager to create a (kernel,initrd) pair that will boot an amd64 Etch box properly. With "BOEL" kernel it fails complaining about `insmod.old` (which belongs to the 2.4 line of kernels), with UYOK it complains "No init found" (which might be a module problem). To add to this, the systemimager site is down today, and the systemimager people have deemed it necessary to create their source packages without actual source. They just download from the systemimager site, which is down. So we cannot work. Furthermore, the Debian maintainer just dropped the package, so even if the site comes up, we still haven't got an amd64 package. So much for the systemimager. I have my hopes back on Harm.

On the need for a proper unattended partitioning tool

In most unattended installation tools, partitioning is problematic. Most can only rewrite complete partition tables, although some¹ do preserve partitions. The Windows installer only has a very limited understanding of partitioning. And always one must carefully point out which partition to preserve.

There is a need for a tool that allows us to list criteria for partitions to be preserved, and criteria for partitions to be created. The tool should probably use prioritization and rules to create a situation as close to what we want as it can.

It would be nice if this tool were a distribution-independent script, so we could load it from any installer as a pre-installation hook, and use it to prepare the disk(s) for the limited abilities of the installation programs. This would also free us from the need to place the machine into different PXE groups during different phases of its installation.

Could PartEd [<http://www.gnu.org/software/parted/manual/parted.html>] do the trick?

On creating a Debian client for the IWI

Some parts still missing here... `apt-get install -t etch-backports openoffice.org openoffice.org-style-industrial openoffice.org-style-tango openoffice.org-style-hicontrast openoffice.org-officebean openclipart-openoffice.org odbc-postgresql openoffice.org-help-en-gb`

On update-alternatives under Debian

Under Debian, `/etc/alternatives` holds links to programs for which several alternatives are available and installed. These links can be maintained automatically or by hand. It may happen (as with `sun-java6-bin` from Etch-Backports) that a program creates links there, but when it is removed, the links remain and dangle. This can be remedied by finding the dangling links and running `update-alternatives --auto` on them, like this: `find -L /etc/alternatives/ -type l -lname *java-6-sun* -exec basename {} \;|xargs -i update-alternatives --auto {}`

¹most notably FAI, perhaps KickStart as well?

Dissecting RPM packages

RPMs can be unpacked in the current directory using the command: `rpm2cpio package.rpm|cpio -idmv --no-absolute-filenames`

IBM Tivoli on Debian

IBM's Tivoli is a storage solution that the RuG uses for backup/restore. It can be downloaded here [<ftp://service.boulder.ibm.com/storage/tivoli-storage-management/maintenance/client/v5r4/>]. Documentation for the client is supposed to be here [http://www.tivoli.com/support/public/Prodman/public_manuals/td/StorageManagerClient5.4.html], but that takes you to a generic entry point where you have to search for the correct route to the actual docs first. I find this [<http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/topic/com.ibm.itsmc.doc/clients.html>] a good starting point.

Procedure 2.1. Log of Tivoli installation

1. Preparing the system

- a. The java we're going to use needs `libstdc++5`: **apt-get install libstdc++5**
- b. The java-based `dsmj` needs the Korn shell, so we install that: **apt-get install ksh**
- c. **Installing Java v1.4**

The `dsmj` executable supports only java 1.4 and no higher. SUN doesn't offer version 1.4 for amd64. So we try Blackdown java:

- i. Download:
ftp://ftp.easynet.be/blackdown/JDK-1.4.2/amd64/rc1/j2re-1.4.2-rc1-linux-amd64.bin
- ii.
- iii. As an ordinary user, create the package: **fakeroot make-jpkg /tmp/j2re-1.4.2-rc1-linux-amd64.bin**
- iv. Install the package: **dpkg -i ~jurjen/blackdown-j2re1.4_1.4.2+rc1_amd64.deb**

2. Getting the software

- a. We fetch the software ²:

```
wget ftp://service.boulder.ibm.com/storage/tivoli-storage-management/maintenance/client/v5r4/
wget ftp://service.boulder.ibm.com/storage/tivoli-storage-management/maintenance/client/v5r4/
```

² Do fetch them both. Some libs are only in the former

- b. Unpack the tarfiles with:

```
for TAR in `ls -1 *.tar` ; do
DIR=`basename $TAR .tar`
mkdir $DIR &&
pushd $DIR &&
tar xf ../$TAR &&
popd
done
```

- alien** Fails miserably, so we unpack all the rpms with:

```
for DIR in `ls -d1 *LinuxX86` ; do
UNPACK=${DIR}/unpack \
if mkdir ${UNPACK} && pushd ${UNPACK} ; then
for RPM in `ls -1 ../*.rpm` ; do
rpm2cpio ${RPM}|cpio -idmv --no-absolute-filenames
done
popd
fi
done
```

3

- c. Copy the `./opt` files into place:

```
cp -ru 5.4.0.0-TIV-TSMBAC-LinuxX86/unpack/opt/tivoli /opt/
cp -ru 5.4.1.0-TIV-TSMBAC-LinuxX86/unpack/opt/tivoli /opt/
```

- d. Not all the programs in `{dsmc, dsmj, dsagent}` react identically to setting the `DSM_DIR` environment variable. So in order to avoid using that variable, we run all these programs from the directory they reside in. They then still search for message files in locations where they are not, so we fool them with a symlinks:

```
pushd /opt/tivoli/tsm/client/ba/bin/ && \
ln ../lang/en_US/ ./
```

- e. We also want our config in `/etc/`, so we create `/etc/tivoli` and link into it:

³ This puts files in the `unpack/opt` and `unpack/usr` directories, of which the latter contain just symlinks.

```
pushd /opt/tivoli/tsm/client/ba/bin/ && \  
ln -s /etc/tivoli/dsm.sys ./ && \  
ln -s /etc/tivoli/dsm.opt ./
```

- f. Link a xerces library sought by dsmagent into place:

```
pushd /opt/tivoli/tsm/client/api/bin && \  
ln -s libxerces-c1_6_0.so libtsm541xerces-c1_6_0.so
```

3. Configuring the software

- a. We can now run the **dsmc** binary with the following script:

Example 2.1. A Tivoli wrapper script

```
#!/bin/bash  
  
# The executable to run  
TIVOLI=/opt/tivoli/tsm/client/ba/bin/dsmc  
#TIVOLI=/opt/tivoli/tsm/client/ba/bin/dsmagent  
#TIVOLI=/opt/tivoli/tsm/client/ba/bin/dsmj  
  
export PATH="/usr/lib/j2re1.4-blackdown/bin:$PATH"  
  
HOST=`hostname`  
  
# I don't know why we set these  
export LANG='en_US.iso88591'  
export -n LC_CTYPE  
# Now I do: we want files with diacritics in their names to be backed up t  
export LANG=en_US  
export LC_TYPE=en_US  
export LC_ALL=en_US  
  
export DSM_LOG='/var/log/tivoli/'  
  
# Tivoli is not properly packaged for Debian, and brings its own libraries  
# so we set LD_LIBRARY_PATH accordingly  
export LD_LIBRARY_PATH="/opt/tivoli/tsm/client/api/bin;\ \  
/opt/tivoli/tsm/client/api/bin64;\ \  
/opt/tivoli/tsm/client/ba/bin/plugins;\ \  
/opt/tivoli/tsm/client/hsm/bin;\ \  
/opt/tivoli/tsm/client/icc32/icc/icclib;\ \  
/opt/tivoli/tsm/client/icc64/icc/icclib"
```

```
pushd `dirname ${TIVOLI}`  
exec ./`basename ${TIVOLI}` $@
```

- b. We edit `/etc/tivoli.dsm.opt` to contain:

Example 2.2. A Tivoli user options file

```
COMPRESSALWAYS YES  
ARCHSYMLINKASFILE NO  
DATEFORMAT 3  
NUMBERFORMAT 4  
  
DOMAIN / /boot /home /opt /spareboot /usr /usr/local /var /srv/si /srv/ftp
```

- c. And we put in `/etc/tivoli/dsm.sys`:

Example 2.3. A Tivoli system options file

```
SERVERNAME ADSM  
COMMMethod TCPip  
TCPPOINT 1500  
TCPSEVERADDRESS your.backupserver.com  
  
NODENAME provided.bythebackupadmin  
PASSWORDACCESS GENERATE  
RUNASSERVICE YES  
COMMMethod TCPip  
  
SCHEDMODE PROMPTED  
TCPBUFFSIZE 32  
TCPWINDOWSIZE 64  
Largecommbuffers yes
```

In the above `dsm.sys`, we temporarily set `RUNASSERVICE` to `NO`. Then we run the `dsmc` program once and issue the command **query schedule** so it starts a session with the server and provides for an automatically updated password in `/etc/adsm/TSM.PWD`. And we set it back to `YES`.

4. Running the software

We can now start `dsmd` by slightly modifying Example 2.1, “ A Tivoli wrapper script ” It will still give us errors after some time, but it runs long enough to exclude some ftp mirrors from being backed up.

In order to be backed up nightly, the server needs to have the scheduler running, which is started by `dsmd sched`. We create an `initrc` from `/etc/init.d.skeleton` and make it start and stop saying `update-rc.d tivoli defaults`.

On the Bootleg host imager

PXE boot

A computer can be configured to boot from its NIC using PXE, as described in Booting from PXE

Procedure 2.2. Booting from PXE

1. Computer is powered up and boots into BIOS ROM
2. BIOS boots into PXE ROM
3.
 - a. PXE does DHCP request
 - b. DHCP server responds with
 - IP number
 - TFTP server to contact
 - (possibly) directory to fetch from
 - and PXE executable to fetch there (i.e.`pxelinux.0`)⁴
4. PXE boots into PXE executable
5. PXE executable fetches configuration from TFTP server^{5 6}
6. PXE executable determines from the configuration file
 - which kernel to fetch
 - which boot parameters to set (if any)
 - which RAMdisk to fetch (if any)

⁴ It fetches the kernel and the RAMdisk, and boots into them.
⁴ Typically from `/var/lib/tftpboot`, which is often aliased to `/tftpboot`.

⁵ Typically from `/tftpboot/pxelinux.cfg/`.

⁶ The name of the file that is fetched is by default the IP number in hexadecimal. But this can be altered by the DHCP server.

The system has been booted.

The bootleg boot script

Harm copied a RAMdisk from Ubuntu-7.0.4 and it it changed just a few files:

```
find -mtime -7 -exec ls -ld {} \;
drwxr-xr-x 11 root root 4096 2007-07-10 12:23 .
drwxr-xr-x 13 root root 4096 2007-07-10 12:17 ./scripts
-rw-r--r-- 1 root root 2726 2007-07-05 09:22 ./scripts/local
-rwxr-xr-x 1 root root 5009 2007-07-05 17:35 ./scripts/start_bootleg
-rwxr-xr-x 1 root root 5009 2007-07-05 17:35 ./scripts/start_bootleg.iwi
-rwxr-xr-x 1 root root 5008 2007-07-05 17:36 ./scripts/start_bootleg.wing
```

When in `/conf/initramfs.conf` `BOOT` is set to `local`, `/scripts/local` is run by **init**. At the end of `/scripts/local`, a Example 2.4, “Bootleg intervention” is added which starts the script `/scripts/start_bootleg`.

Example 2.4. Bootleg intervention

```
#
# Bootleg intervention : By Harm@cs.rug.nl
# Initial : Mon May 21 16:53:33 CEST 2007
# This is the new bootleg intervention for 2.6.20 kernels and up
# Using cpio initrd's
#
#
# Unmount the root fs from the initrd script
umount /root

echo "Starting bootleg"
sleep 1      # calm down a little bit

/scripts/start_bootleg

# Mount the ubuntu root fs again
mount ${roflag} -t ${FSTYPE} ${ROOTFLAGS} ${ROOT} ${rootmnt}

# End Bootleg intervention
```

This `start_bootleg` script mounts an NFS share from a server ⁷ and starts a secondary script that resides on the mount ⁸ Once the secondary script is running, the boot is considered to be in its “2nd boot

⁷ Which server is configured in the `start_bootleg` script itself. Typically this is the file server for the domain.

phase”.

Drawing a bundle

Taking a fully installed client, and rebooting it from the NIC in such a way that its partitions are copied into tarfiles which are then stored on the server is called “drawing a bundle”.

How to determine what a host boots

Assuming that a client host is configured (in its BIOS) to boot from the NIC, the first point at which we can influence from the server how the client boots is in the DHCP table. A hosts place in the DHCP tables typically changes when it moves from one subnet to another, or when it changes from Windows-only to Linux-only. Generally, in order to use PXE for booting, it must be in a scope that offers the settings we described in Booting from PXE .

The second point at which we can influence the client from the server is at the PXE configuration file. At the IWI, this is done with the **sethostip** command, which takes two parameters: first name of a file already present in /tftpboot/pxelinux.cfg, then the hostname. See Example 2.5, “ Example of **sethostip** usage ”.

Example 2.5. Example of sethostip usage

The command **sethostip -t bootpxe-sarge.sda iwi101** would make the host iwi101 boot as specified in the pxe configuration file /tftpboot/pxelinux.cfg/bootpxe-sarge.sda.

To be continued with exports, hosts.client, /etc/bundle, gen_domainprof, overwrites, addons, NEWDIST, MAKEPARENTDIST etc.etc.

On moving DHCP from an endangered host

iwi202 Was the DHCP server and it also checked mail for viruses and spam. Now it is becoming old: it has one bad sector already. The machine is out of service and I have a week left before my vacation: no time to buy new hardware or do serious migration of users. I'm moving mailchecking and DHCP to iwi2.

Procedure 2.3. Things to do on the new server

1. Bring the machine up to date: **apt-get update && apt-get dist-upgrade**
2. Configure the IP number statically. Modify /etc/network/interfaces to contain a stanza:

```
#iface eth0 inet dhcp
iface eth0 inet static
address nnn.nnn.nnn.nnn
netmask 255.255.255.0
```

⁸ Which secondary script is also configured in the start_bootleg script. As of this writing, bootleg is started.

3. Make sure it finds the nameserver when the NIC comes up. It seems this thing runs BIND already, it is its own nameserver:

```
search mydomain.com
nameserver 127.0.0.1
```

4. install the DHCP daemon: **apt-get install dhcp3-server**
5. Enable the DHCP server by editing `/etc/default/dhcp3-server`:

```
INTERFACES="eth0"
```

6. Set aside the original dhcpd config: **cd /etc/dhcp3 && mv dhcpd.conf dhcpd.conf.dist**
7. Copy the DHCP config from the old system, `iwi202`:

```
cd /etc/dhcp3 && \
scp root@iwi202:/etc/dhcp3/dhcpd.conf ./ && \
scp -r root@iwi202:/etc/dhcp3/IWI-NET ./
```

8. Start the DHCP server: **/etc/init.d/dhcp3-server restart**
9. Copy `/usr/local/bin/dhcprestart` from another host for convenience.
10. Put the DHCP daemon in the startup scripts: **update-rc.d dhcp3-server defaults**
11. And start the DHCP daemon: **/etc/init.d/dhcp3-server restart**

Procedure 2.4. Things to do on the old server

1. Remove the DHCP server from the startup scripts: **update-rc.d -f dhcp3-server remove**
2. And stop the DHCP daemon: **/etc/init.d/dhcp3-server stop**

On moving anti-spam and viruschecking from an endangered host

Procedure 2.5. Things to do on the new server

1. Install some packages: **apt-get install spamassassin spame spamoracle spampd spamprobe clamav clamav-base clamav-daemon clamav-freshclam postfix postfix-tls**
2. Leave `/etc/spamassassin` and `/etc/clamav` as they are (for now)
3. Alter `/etc/default/spamassassin`:

```
# Change to one to enable spamd
ENABLED=1
#...
#OPTIONS="--create-prefs --max-children 5 --helper-home-dir"
OPTIONS="-c -m 10 -H -i my.own.ip.number -p 783 -u spamass"
```

4. Add a SpamAssassin-specific system user: **adduser --system spamass**
5. Start SpamAssassin: **/etc/init.d.spamassassin restart**
6. Edit `/etc/default/spampd`:

```
AUTOWHITELIST=1
#...
LOCALONLY=0
```

7. Start the spampd daemon: **/etc/init.d/spampd start**
8. Start the ClamAv daemons:

```
/etc/init.d/clamav-daemon start && \  
/etc/init.d/clamav-daemon clamav-freshclam
```

9. Edit `/etc/amavis/conf.d/05-domain_id`:

```
@local_domains_acl = ( ".$mydomain", "my.first.domain.com", "my.second.domain"
```

10. Edit `/etc/amavis/conf.d/20-debian-defaults` (the last two lines are the modification) in order to grant access from other machines than localhost:

```
$inet_socket_port = 10024; # default listening socket
@inet_acl = ( '127/8', 'my.ip.nnn/24' ); #This needed to grant access to mail
$inet_socket_bind = undef; #This needed too to grant access to
```

11. Restart the amavis daemon: `/etc/init.d/amavis restart`

Procedure 2.6. Things to do on the mail server

- On the mail server, all we have to do is to point PostFix to the new mails scanner via a small snippet in `/etc/postfix/main.cf`⁹:

```
content_filter = smtp-amavis:[ip.of.new.filterhost]:10024
```

On using Badblocks with ReiserFS

Finding the bad area using badblocks

We have error messages on server `iwi202` that look like shown in Example 2.6, “Log of bad sectors on `iwi202`”. The problem repeats twice in quick succession (8 seconds between occurrences) about every twelve minutes, but doesn’t stick to fixed post-the-hour times, so we don’t believe a cron job causes it. The machine does react more slowly than usual. I will move important processes off the machine, but some minor items may stay on it, and I want to see if I can get rid of the problem by making the ReiserFS stop using the single block that is causing errors.

Example 2.6. Log of bad sectors on `iwi202`

```
Jul 17 09:11:48 src@iwinnn kernel: scsi1: ERROR on channel 0, id 0, lun 0, CDB: 0x28 00 0a b5 f0 fa 00
Jul 17 09:11:48 src@iwinnn kernel: Info fld=0xab5f0fd, Current sd08:09: sns = f0 3
Jul 17 09:11:48 src@iwinnn kernel: ASC=11 ASCQ= 0
Jul 17 09:11:48 src@iwinnn kernel: Raw sense data:0xf0 0x00 0x03 0x0a 0xb5 0xf0 0xfd 0x0a 0x00 0x00
Jul 17 09:11:48 src@iwinnn kernel: I/O error: dev 08:09, sector 104206368
Jul 17 09:11:56 src@iwinnn kernel: scsi1: ERROR on channel 0, id 0, lun 0, CDB: 0x28 00 0a b5 f0 fa 00
```

⁹ This assumes that in `/etc/postfix/master.cf` you already have a snippet like this to enable it to receive from the mails scanner: `10025 inet n - n - - smtpd -o content_filter= -o local_recipient_maps= -o relay_recipient_maps= -o smtpd_restriction_classes= -o smtpd_client_restrictions= -o smtpd_helo_restrictions= -o smtpd_sender_restrictions= -o smtpd_recipient_restrictions=permit_mynetworks,reject -o mynetworks_style=host -o mynetworks=my.first.ip.range/no_bitsmasked,my.second.ip.range/no_bitsmasked -o strict_rfc821_envelopes=yes`

```
Jul 17 09:11:56 src@iwinnn kernel: Info fld=0xab5f0fd, Current sd08:09: sns = f0 3
Jul 17 09:11:56 src@iwinnn kernel: ASC=11 ASCQ= 0
Jul 17 09:11:56 src@iwinnn kernel: Raw sense data:0xf0 0x00 0x03 0x0a 0xb5 0xf0 0xfd 0x0a 0x00 0x00
Jul 17 09:11:56 src@iwinnn kernel: I/O error: dev 08:09, sector 104206368
```

The sector that causes errors -104206368- is located in `/dev/sda9`, which is mounted as `/var`. I could run `badblocks` on the entire disk if I put the machine in single-user mode and unmounted `/var`, but I'd rather be as unobtrusive as possible, as I'll see notifications of bad sectors turning up in the logs anyway. According to the `badblocks` manual, I can say: **`badblocks -c<blocks-at-a-time> <device> <end-block> <start-block> -i <former-badblocks-report>`** Badblock counts in blocks of 1024 bytes, whereas we know the location of the bad sector in 512-byte sectors. So we compute the location of the sector in blocks: **`echo -e "104206368\n2\n\n"|dc`**, which yields `52103184`. Then we issue the command to check the partition: **`badblocks -c64 /dev/sda9 52103222 52103152 |tee ~/bad_blocks.dev.sda9`** A few blocks after our culprit appear to be bad as well: `52103184 52103185 52103186 52103187`

Notifying the ReiserFS of the bad area

Note

It appears that `reiserfstune` may not be run on a mounted file system. So we must unmount the file system after all. The only advantage we created is that this can be (hopefully) a quick operation now.

Note

It also appears that `reiserfstune` cannot not accept the output of `badblocks`. Their units of disk space differ (see Table 2.1, “Units of disk space used by programs involved in a ReiserFS `badblocks` detection”) so we have to convert: **`for n in `cat bad_blocks.dev.sda9` ; do echo -e "${n}\n8\n\n"|dc ; done > converted-badblocks-file`**

```
6512898
6512898
6512898
6512898
```

When we try to notify the filesystem of its bad blocks (**`reiserfsck --add-badblocks <converted-badblocks-file> --fix-fixable /dev/sda9`**), the command returns an error and the message that the block under consideration is already in use, and please use `reiserfsck` to repair.

Warning

This `reiserfsck` then fails with a segmentation fault, and we are glad to escape with our filesystem intact. There is *no way* that I will use **`reiserfsck --rebuild-tree`** on an already populated filesystem.

Note


I think we'd better stay away from ReiserFS from now on.

We use another route, and do a **find /var/ -type f -exec cat {} \;>/dev/null** on the affected filesystem. This fails with the message

```
find /var/ -type f -exec cat {} \;>/dev/null
cat: /var/lib/postgresql/8.1/main/base/16629/16667: Input/output error
```

and since we know that only a single sector is affected, this must be the file that causes the messages in our logs. So we'll do the following:

1. Stop all services that use the filesystem by switching to single-user mode:

```
init 1 
```

2. Stop `syslog-ng` too, as it uses `/var` and is still active in runlevel 1

```
/etc/init.d/syslog-ng stop 
```

3. unmount the filesystem¹⁰: **umount /var**
4. mount it somewhere else: **mkdir /mnt/sda9 && mount /dev/sda9 /mnt/sda9**
5. Move the file that lies on the bad sector to another filesystem: **dd if=/mnt/sda9/lib/postgresql/8.1/main/base/16629/16667 of=/home/16667 conv=noerror**¹¹
6. Unmount the partition: **umount /dev/sda9**
7.
 - a. Notify the filesystem of its bad blocks using `reiserfstune`: **reiserfstune --add-badblocks ~/bad_blocks.dev.sda9.base4096 /dev/sda9**
 - b. When this fails with the already-in-use message, we try `/sbin/reiserfsck --badblocks` **~/bad_blocks.dev.sda9.base4096 /dev/sda9**

Note

This failed during an earlier try, but it succeeded this time. YMMV.

8. Remount the partition at the alternative mount point: **mount /dev/sda9 /mnt/sda9**
9. Copy the file back in place: **mv /home/16667 /mnt/sda9/lib/postgresql/8.1/main/base/16629/**

~~10. Unmount the partition from its alternative mount point: **umount /dev/sda9**~~

¹⁰Of course, first we have to unmount (possibly remote) file systems mounted on subdirectories of `/var`, like f.i. `/var/mail`.

¹¹We cannot use `cp`, as it will stop when it encounters the bad sector, and copy only part of the file.

11. Mount it in its usual place: **mount /dev/sda9**
12. Start syslog again:

```
/etc/init.d/syslog-ng start ❶
```

13. Return to multi-user mode:

```
init 5 ❶
```

❶ A long list of scsi-driver errors (as in Example 2.7, “SCSI errors in the log”) is still in the kernel ringbuffer . During start/stop/reload of `syslog-ng` they will scroll across the console. This may look disturbing, but it is *not* an indication that the bad part of the disk is still being accessed.

After all this is done, we see no more SCSI errors in the logs, and **debugreiserfs -B /tmp/bad /dev/sda9 && cat /tmp/bad** confirms that block 6512898 is bad.

Example 2.7. SCSI errors in the log

```
Raw sense data:0xf0 0x00 0x03 0x0a 0xb5 0xf0 0xfd 0x0a 0x00 0x00 0x00 0x00 0x11
I/O error: dev 08:09, sector 104206368
scsil: ERROR on channel 0, id 0, lun 0, CDB: 0x28 00 0a b5 f0 fa 00 00 08 00
Info fld=0xab5f0fd, Current sd08:09: sns = f0 3
ASC=11 ASCQ= 0
```

Table 2.1. Units of disk space used by programs involved in a ReiserFS badblocks detection

program	unit	size
kernel sata driver	sectors	512 bytes
badblocks	block	1024 bytes
ReiserFS	block	4096 bytes

On the Syslog-NG log server

Note

ToDo: log over TCP instead of UDP, and encrypt communication between client and server (using a tunnel?)

This document shows plain-text logging over UDP. While this is simple, it is hardly bandwidth-efficient, and certainly not secure.

Creating a Syslog-NG server

Syslog-NG [<http://www.balabit.com/network-security/syslog-ng/>] is an improvement upon `syslog` with regard to configurability. We followed the SysLog-NG Administrator Guide [http://www.balabit.com/dl/html/syslog-ng-admin-guide_en.html/bk01-toc.html], in which `syslog-ng` is documented well.

Procedure 2.7. Creating a `syslog-ng` server

1. Install `syslog-ng`: **`apt-get install syslog-ng`**
2. In `/etc/syslog-ng/syslog-ng.conf.dist`, configure the server to listen to incoming logs:

```
source s_all {
    # message generated by Syslog-NG
    internal();
    # standard Linux log source (this is the default place for the
    # function to send logs to)
    unix-stream("/dev/log");
    # messages from the kernel
    file("/proc/kmsg" log_prefix("kernel: "));
    # use the following line if you want to receive remote UDP logs
    # (this is equivalent to the "-r" syslogd flag)
    # enabled --JB 20070718
    udp(); ❶
};
```

❶ Uncomment this to make `Syslog-NG` listen on udp port 514

3. Restart the daemon: **`/etc/init.d/syslog-ng restart`**

Now you can see the daemon listen on udp port 514: **`netstat -ltn|grep syslog`**

```
udp        0          0 0.0.0.0:514          0.0.0.0:*
unix  2      [ ACC ]     STREAM  LISTENING   7009
```

Creating a Syslog-NG client

Procedure 2.8. Configuring a Syslog-NG client

1. Install the software: **apt-get install syslog-ng**
2. Modify `/etc/syslog-ng/syslog-ng.conf`:
 - a. Declare the name of the syslog server:

```
destination syslog_server { udp(log.server.my.com); }
```

- b. Make the system log to the syslog server:

```
log { source(src); destination(syslog_server); };
```



- Make sure to put this at the top of the list of “log” statements, *before* any statements that carry “final” flags.

3. Restart the daemon: **/etc/init.d/syslog-ng restart**
4. If you have a server running already, you can now test the configuration by logging on to the client and saying something like: **echo "JohnDoe testing Syslog-NG" | logger**

You should see the message turn up in the log files of the server.

On configuring CfEngine

In a freshly installed Debian Etch, the package `cfengine2` [<http://www.cfengine.org/>] comes with

- `/var/lib/cfengine2/inputs` a symlink to `/etc/cfengine`, and
- `/etc/default/cfengine2` looking like

```
RUN_CFSERVD=0
RUN_CFEEXEC=0
RUN_CFENVD=0

CFSERVD_ARGS=" "
```

The server part

1. First of all, the server needs to run. Modify `/etc/default/cfengine2`:

```
RUN_CFSERVD=1 ❶  
RUN_CFEEXECD=0 ❷  
RUN_CFENVD=0  
CFSERVD_ARGS= " "
```

- ❶ The server must run
 - ❷ If the server isn't a client of another server, best leave this off. Circular dependencies don't seem like a good idea at this time.
2. Then, we need configuration files in `/etc/cfengine`. CfEngine looks for them in `/var/lib/cfengine2/inputs`, but that is a symlink to `/etc/cfengine`. To get started, we need at least the following:

```
server:/etc/cfengine# find /etc/cfengine/ -type f  
/etc/cfengine/cfrun.hosts  
/etc/cfengine/cfservd.conf  
/etc/cfengine/masterfiles/update.conf  
/etc/cfengine/masterfiles/cfagent.conf
```

Note

In my case, these files are generated using `iserv`¹². Whenever **iserv-update** is run, the templates in `/etc/iserv/templates/etc/cfengine/` are used to generate instances in `/var/lib/iserv/generated/etc/cfengine/`, which are then copied to `/etc/cfengine/`. The Makefile (`/etc/iserv/Makefile`, and its inclusion `/etc/iserv/copy-targets`) direct this behaviour.

Note

`/etc/cfengine/cfrun.hosts` *should* be created by `iserv`, but isn't. This is on the To-Do-list for `iserv`.

- a. `/etc/cfengine/cfrun.hosts` should simply contain the FQDNs of all hosts this server is ever going to service, one per line. List them all and be done with it.
- b. `/etc/cfengine/cfservd.conf` can be fairly simple: see Example 2.8, “Example `cf-servd.conf`”
- c. `/etc/cfengine/masterfiles/update.conf` should be tampered with as little as possible. Its sole responsibility is to keep the configuration files for CfEngine on the client identical to those on the server. If the master copy /

¹² `iserv` Is a suite I made. It consists of a couple of scripts that generate configuration files for various services using a central stash of known variables like network addresses of interfaces, serviced domains, etc. There is no use in looking for it on the Web, as it is not past the kludge phase, and I haven't open-sourced it yet. If anyone who reads this knows of an existing system that can do this, I would be grateful for a quick notification.

`etc/cfengine/masterfiles/update.conf` becomes incorrect, the client copies also become incorrect at the next run of `cfagent`, and this cannot be repaired, as this is the file that should do the repair, and it has become incorrect.

Warning

Do not mess with `/etc/cfengine/masterfiles/update.conf` unless you know exactly what you are doing, *and* you have redundant mechanisms in place to assure that a newer version does not depend on CfEngine to spread, *and* you are wearing clean underwear.

That being said, take a look at Example 2.9, “Example `update.conf`”

- d. All the previous configuration just serves to get `/etc/cfengine/masterfiles/cfagent.conf` into place. This is the file that does all the work. We present a minimal version in Example 2.10, “Example `cfagent.conf`”, but it should greatly be elaborated upon. For help, please refer to the `cfagent` reference [<http://www.cfengine.org/docs/cfengine-Reference.html#Cfagent-reference>].
3. We are now ready to restart the `cfserverd` daemon: **`/etc/init.d.cfengine2 restart`**

Note

In order to make the `update.conf` and `cfagent.conf` available for clients which don't have a configured `cfagent` yet, `iserv` also makes them available through `ftp`, by placing a copy in e.g. `srv/ftp/pub/local/os/linux/distributions/sl/installer/scripts/`¹³

The client part

1. We start by installing the package: **`apt-get install cfengine2`**
2. Then we enable the `cfagent` daemon by editing `/etc/default.cfengine2`:

```
RUN_CFSERVD=0
RUN_CFEEXECD=1 ❶
RUN_CFENVD=0

CFSERVD_ARGS= " "
```

- ❶ `cfexecd` is a wrapper around `cfagent`, so this enables it. For redundancy, we could also put a cronjob in `/etc/cron.hourly` that runs `cfagent`. `Cfagent` should then restore this setting in `/etc/default/cfengine2`.
3. Then we should download an `update.conf`: **`rm /var/lib/cfengine2/inputs/update.conf && wget -nc -P /var/lib/cfengine2/inputs ftp://master.grid.rug.nl/pub/local/os/linux/distributions/sl/installer/scripts/update.conf`**

¹³ This location is not the only one that carries a copy, and subject to rapid change over time.

4. ... and run cfagent: **cfagent -vqK**

Example 2.8. Example cfserverd.conf

```
control:
    domain = ( mynet )
    TrustKeysFrom = ( 10.0.3.0/24 ) ❶

    MaxConnections = ( 50 )
#####
admit: # or grant:
    /var/lib/cfengine2/inputs/masterfiles *.mynet ❷
    /etc/cfengine/masterfiles *.mynet ❷
    /etc/cfengine/trees *.mynet
```

- ❶ With this line, clients that don't have identification keys may generate them and the server will trust them the first time they connect. Without it, new keys will have to be transferred either from server to client or the other way around by other means than CfEngine.
- ❷ Access needs to be granted to `/var/lib/cfengine2/inputs/masterfiles`, but also to `/etc/cfengine/masterfiles`, because `/var/lib/cfengine2/inputs` is a link to `/etc/cfengine/masterfiles`

Example 2.9. Example update.conf

```
control:
    actionsequence = ( copy links processes tidy ) # Keep this simple and const
    domain         = ( mynet ) # Needed for remote copy
    policyhost     = ( master.mynet ) # This is the host part of where our conf
    master_cfinput = ( /var/lib/cfengine2/inputs/masterfiles ) # This is the dir
    AddInstallable = ( new_cfenvd new_cfserverd )
    #
    # Workdir is not identical on all clients
    #
    workdir        = ( /var/cfengine )
    linux::
```

```
workdir          = ( /var/lib/cfengine )
debian::
workdir          = ( /var/lib/cfengine2 )
scientific_sl_3::
workdir          = ( /var/cfengine )
solaris::
  cf_install_dir = ( /usr/local/sbin )
linux::
  cf_install_dir = ( /usr/local/sbin )
  # Serve proxy servers first
!AllBinaryServers::
  SplayTime = ( 1 )

# This is the section that does the copy of update.conf and cfagent.conf from se
copy:
  $(master_cfinput)          dest=$(workdir)/inputs
                             r=inf
                             mode=700
                             type=binary
                             exclude=*.lst
                             exclude=*~
                             exclude=#*
                             server=$(policyhost)
                             trustkey=true

# We clean up old data
tidy:
  $(workdir)/outputs pattern=* age=7
```

Example 2.10. Example cfagent.conf

```
control:
  actionsequence = ( files copy shellcommands )
  domain         = ( mynet )
  timezone       = ( MET )
  smtpserver     = ( smtphost.mynet ) # used by cfexecd
  sysadm         = ( jurjen@cs.rug.nl ) # where to mail output
  policyhost     = ( master.mynet )
```

```
#####
copy:

any::
    /etc/cfengine/trees/any
        dest=/
        r=inf
        timestamps=preserve
        type=binary
        exclude=*~
        server=$(policyhost)
#         purge=true # Never use this, or lose _all_ data on your client!

shellcommands:

    "/etc/iserv/client/scripts.all"
        useshell=true
        background=true
        #ifelapsed=15
        #expireafter=15
```

Creating a yum repository under Debian

We have a server running Debian, but we want to serve packages to machines running Scientific Linux (a RedHat clone). We try to set up a yum repository.

Procedure 2.9. Setting up the repository

1. Creating and populating the repository is fairly simple: create any directory - /
srv/ftp/myrepos for our purposes, and copy the rpm-files into it.
2. Now we need two packages installed: **apt-get install yum createrepo**
3. Indexing the repository for the client can be done in two ways:
 - a. **createrepo /srv/ftp/myrepos** for use with newer versions of yum
 - b. **cd /srv/ftp/myrepos && yum-arch .** for use with older versions

Now we can configure the client to use our repository by putting the following stanza into /
etc/yum.conf:

```
[My Repository]
name=myrepos
baseurl=ftp://myhost.mynet.com/myrepos/
```

Usage of yum (and other package managers) is nicely described in the Naked Ape's PackageManager-

CheatSheet [<http://nakedape.cc/wiki/PackageManagerCheatsheet>].

Warning

There is a catch though: some rpm files will not be accepted by yum-arch, either with a message error `public key not available` or simply with ignoring bad rpm: `<filename>`.

`createrepo` will not complain, but whether it is passing or just silently ignoring I did not test.

Nice repositories to steal from are Dries's [<http://dries.ulyssis.org/rpm/packages/index.html>] and Dag's [<http://dag.wieers.com/rpm/packages/>].

On using Procmail and Vacation

vacation needs a file `~/ .vacation.msg`, and upon running **vacation -I** it will create a fresh database of users in `~/ .vacation.db`.

It can be called from the `~/ .procmailrc` like this:

```
:0c ❶  
| /usr/bin/vacation -a my.alias@my.com -a another.alias@mycom -a y.e.t.another@a
```

- ❶ The “c” generates a copy of every message, so messages are not lost after being sent to vacation.

Of course, procmail needs to be run in order for the `~/ .procmailrc` to be respected. If procmail is not the default delivery agent, this can still be accomplished by setting the appropriate `~/ .forward`:

```
"|IFS=' ' && exec /usr/bin/procmail -f- || exit 75 #user"
```

Chapter 3. August 2007

Screen rotation in X with NVidia

Warning

This setting is specific for the NVidia driver. I didn't get screen rotation to work on an Intel Q963/Q965 chipset with `xserver-xorg-xserver-i810` driver on Debian Etch when that was stable. I did find some reports of it working with a driver provided by Intel, but that one is still so much in development that it stays in Unstable.

Put into the "Device" section of `xorg.conf`:

```
Option "RandRRotation" "true"
```

Then use `xrandr` to rotate the screen: `xrandr -o left`

Firewalling with FWBuilder

Resetting IPTables

IPTables can not easily be turned off via the `/proc` interface or by removing modules. But messed up rules can be replaced by the following as long as the modules themselves are functioning properly.

```
echo "Stopping $Firewall"  
$IPTABLES -F  
$IPTABLES -X  
$IPTABLES -P INPUT ACCEPT  
$IPTABLES -P OUTPUT ACCEPT  
$IPTABLES -P FORWARD ACCEPT
```

Firewall Rules

Warning

Interface-specific rules are evaluated first, and thus take precedence over general rules.

Chapter 4. September 2007

Configuring CUPS clients

The CUPS daemon needs not run on CUPS clients. Under Debian, it can be turned off by `/etc/init.d/cupsys stop && update-rc.d -f cupsys remove`. Then create `/etc/cups/client.conf` by `echo "ServerName cupsserver.your.domain.com" > /etc/cups/client.conf` and try `lpq` to see whether it worked.

Note

Of course you still have to allow the client to use some printers on the CUPS server.

Configuring NTP clients

The `ntp` daemon used for timekeeping may not by default be configured well for all circumstances. In particular, the “`iburst`” and “`burst`” directives may be needed in server statements in order to let the machine under config trust its external time sources well enough to let it choose these for reference instead of the local clock. And if this is not enough, the local clock (recognizable as a “`127.127.something`” server entry) may have to be taken out of the configuration altogether.

Example 4.1. Snippet from `ntp.conf` showing “`iburst`” flag to “`server`” statement

```
#server 127.127.1.0 ❶
#fudge 127.127.1.0 stratum 10 ❷
server 0.debian.pool.ntp.org burst iburst ❸
server 1.debian.pool.ntp.org burst iburst ❸
server 2.debian.pool.ntp.org burst iburst ❸
server 3.debian.pool.ntp.org burst iburst ❸
```

- ❶ Local clock commented out, or *it* will be taken as the reference.
- ❷ “Fudging” the local clock to al (low) stratum 10 appeared not to be enough to keep it from being taken as the reference clock.
- ❸ This “(i)burst” needed here, or system clock won't be set.

This measure should result in something like the following output of `ntpq -p`:

remote	refid	st	t	when	poll	reach	delay	offset	jitter
+ntp1.hro.nl	192.87.106.2	2	u	6	64	377	7.027	4007.26	838.555
*ntp2.hro.nl	193.79.237.14	2	u	60	64	377	7.068	4067.85	577.443
-ns1.iseer.nl	193.67.79.202	2	u	12	64	377	4.415	4668.18	344.353
+news.vps.budget	193.79.237.14	2	u	24	64	377	4.595	3854.39	911.959
xRN2-R6509-RP.ne	192.36.143.150	2	u	22	64	377	0.798	4824.30	259.985

If the local system is really screwed up, removing the `drift` file may also help.

Note

It may take half an hour for the daemon to reign the local clock. If the `drift` file was removed, it will only be created again after an hour of continuous operation.

Extensive documentation on `ntp` is available as "The Network Time Protocol (NTP) Distribution" at www.eecis.udel.edu [<http://www.eecis.udel.edu/~mills/ntp/html/index.html>].

RedHat Certification

We are going to do the installation and management of a large number of Linux desktops, possibly in various configurations and in multiple subnets. We already have a basic idea of how to manage this, and we know how to do it on a small scale. Still, a course on how to do it using RedHat would be nice. Skills we need to learn:

- How to do system administration RedHat-style
- How to create rpm packages.
- How to install a RedHat box through PXE.
- How to maintain and monitor a lot of RedHat boxes pull-style. (i.e. Using a satellite server)
- How to maintain a lot of other machines (i.e. linux, but not RedHat) with the same tools.

We go to the Red Hat Training Pre-assessment Questionnaires [<https://www.redhat.com/apps/training/assess/>] in order to test our skill. Some questions are rather RedHat-specific. Anyway, we score (on 033, 133, 253): 36,36,30 (Jurjen), 40, 36, 33 (Heiko), 30, 36, 27 (Arjan). So perhaps we would like to follow

- RH253 (Red Hat Linux Networking and Security Administration) [https://www.redhat.com/training/rhce/courses/rh253_content.html] is what we can do already.
- RHS333 - Red Hat Enterprise Security: Network Services [<http://www.europe.redhat.com/training/course/RHS333>] (RHCE required)
- We probably want the RHCE: RH300 - Red Hat Certified Engineer (RHCE) [https://www.redhat.com/training/rhce/courses/rhce_content.html] or perhaps the rapid track: RH300 RHCE Rapid Track Course Outline [https://www.redhat.com/training/rhce/courses/rhce_content.html] See also
- RH336 - JBoss for Administrators [<http://www.europe.redhat.com/training/course/RH336>]
- This is what we definitely want: RH401 - Red Hat Enterprise Deployment, Virtualization, and Systems Management [<http://www.europe.redhat.com/training/course/RH401>] (do the questionnaire?)
- Perhaps RH423 - Red Hat Enterprise Directory Services and Authentication [<http://www.europe.redhat.com/training/course/RH423>]
- Later perhaps RHS333 - Red Hat Enterprise Security: Network Services [<http://www.europe.redhat.com/training/course/RHS333>]

- Later too RHS429 - Red Hat Enterprise SELinux Policy Administration [<http://www.europe.redhat.com/training/course/RHS429>]

Using dmesg to stop log cluttering console

The kernel ringbuffer is usually also echoed to the console to some extent. `dmesg` controls exactly to which extent. Use `dmesg -1` to cause only critical warnings to be echoed to the console.

Installing new machines for the IWI in 2007

Procedure 4.1. Log

1. Changing BIOS settings
 - a. Under Storage->Boot order: Put "Hard drive, integrated SATA" on top, followed by "Network Controller". Disabled all other (CD-ROM and USB)
 - b. Under Storage -> Storage Options: Disabled Removable Media Boot and Legacy Diskette Drive. Set "SATA Emulation" to "RAID"
 - c. Under Advanced -> Power-On options: Set Post Messages to Enable. Set MEBx Setup Prompt to Displayed Set After Power Loss to Previous State
 - d. Under Advanced -> Onboard Devices: Set "Diskette Controller" to Disable.
 - e. Left all other settings as they were, most notably under Advanced -> Device Options: left S5 Wake on LAN at "Enable"
 - f. Saved changes and Exit
2. Put the machine in the DHCP config
3. The NIC fails to be detected.
 - a. According to an Ubuntu forum thread [<http://ubuntuforums.org/showthread.php?t=551720>], the Intel e1000 driver is too old. Fetched a new one here [<http://downloadcenter.intel.com/confirm.aspx?httpDown=http://downloadmirror.intel.com/9180/eng/e1000-7.6.5.tar.gz&agr=N&ProductID=983&DwnldId=9180&strOSs=All&OSFullName=All%20Operating%20Systems&lang=eng>], and put it on a machine with the same kernel as the install we are running (that is demanded in the README of the driver). Unpacked, compiled.
 - b. On the install server, copied the `initrd.gz` to a temp directory, unpacked it, created a directory 'ramdisk', cd into it and dit `cpio -i < ../initrd`
 - c. Copied the new `e1000.ko` over the old `ramdisk/lib/modules/2.6.18-5-amd64/kernel/drivers/net/e1000/e1000.ko`.
 - d. In `ramdisk`, did `find . | cpio -o -H newc > ../initrd` then `cd ..` and `gzip -9 initrd`

- e. Copied the new initrd to the tftp tree, restarted the master's tftp daemon
- f. Copied the amd64 tftp subtree to the local tftp server (for speed)
- g. Pointed the installer to Updated `/etc/dhcp3/dhcpd.conf`:

```
if substring (option vendor-class-identifier, 0, 3) =
    {
        filename "ftp://master.grid.rug.nl/"
    }
```

- h. Restarted the DHCP daemon, and restarted the client

Installation starts...

4. Install does not complete

- a. Removed “DEBIAN_FRONTEND=noninteractive” from kernel parameters in `/var/lib/tftpboot/debian/etch/amd64/debian-installer/amd64/pxelinux.cfg/default`. Now we see the usual menus, only the disappear again as they are filled in by the preseed.
- b. We get the error “No root file system (is defined)”

Setting global key bindings in Emacs

`M-x global-set-key <key> <command>`

OWL

Links

- Protege [<http://protege.stanford.edu/>] is a Java-based Ontology [http://protege.stanford.edu/publications/ontology_development/ontology_101-noy-mcguinness.html] Editor
- OWL [<http://www.w3.org/TR/owl-features/>] is a second order predicate logic language built on top of DAML+OIL, which in turn is built on top of XML and RDF. There exists an accesible Wine advice server [<http://www.ksl.stanford.edu/people/dlm/webont/wineAgent/>] implementation that makes use of JTP [<http://www.ksl.stanford.edu/software/JTP/>], that is used as a web-based reasoning system.
- SPARQL [<http://www.w3.org/TR/rdf-sparql-query/>] is a query language for RDF [ht-

[tp://www.w3.org/RDF/](http://www.w3.org/RDF/)

- theFigtrees.net [<http://thefigtrees.net>] Has a SPARQL FAQ [<http://thefigtrees.net/lee/sw/sparql-faq>] with links to implementations [<http://esw.w3.org/topic/SparqlImplementations>].
- The “Semantic Reasoner” entry of Wikipedia [http://en.wikipedia.org/wiki/Semantic_Reasoner] has a list of available semantic reasoners with a (limited) feature comparison.
- W3.org has a FAQ on the Semantic Web [<http://www.w3.org/2001/sw/SW-FAQ>]
- There is an extensive list [<http://esw.w3.org/topic/SemanticWebTools>] of tools available for implementing parts of the Semantic Web at esw.w3.org.
- Chimaera [<http://ksl.stanford.edu/software/chimaera/>] might be of help in “creating and maintaining distributed ontologies on the web”. In doing so, the DAML Ontology Library [<http://www.daml.org/ontologies/>] also might come in handy.
- Joseki [<http://www.joseki.org/>] is an HTTP engine with SPARQL support, Java-based and offering an HTTP service on port 2020. Its site states nothing about OWL though, just SPARQL and RDF.
- ARQ [<http://jena.sourceforge.net/ARQ/>] is a query engine for Jena [<http://jena.sourceforge.net/>] that supports SPARQL. Jena in turn is a semantic framework for Java, and it does support OWL (as well as RDF, RDFS, and SPARQL). In addition, it is Open Source.
- Virtuoso [<http://virtuoso.openlinksw.com/wiki/main/Main/VOSSparqlProtocol>] is another Open Source HTTP SPARQL server. Not much is stated about its reasoning capacities, though. And it looks a bit convoluted, being an SQL server, a web server, a webdav server and SPARQL server all in one. Not much of “do one thing and do it well”, and nothing about OWL...
- Pellet [<http://pellet.owldl.com/>] is an Open Source OWL DL reasoner in Java. Among others, it does have a command line interface and a DIG server, but nothing is stated about HTTP.
- Minerva is part of the IBM Integrated Ontology Development Toolkit [<http://www.alphaworks.ibm.com/tech/semanticstk>], and it seems nice (Open Source and Eclipse and all that). It is reported (by IBM) to be 10-20 times slower on Apache Derby than on IBM's own DB2, but it does support SPARQL, although not sure whether we can talk HTTP to it.
- Sesame [<http://www.openrdf.org>] is an open source framework for storage, inferencing and querying of RDF data. It does have SPARQL support in its newest version. Anton Jansen provided this link. He uses Sesame in combination with OWLIM (see next item), and also with ELMO (at the Sesame site), which maps Java classes to the OWL concepts.
- OWLIM [<http://www.ontotext.com/owlim/>] is a high-performance semantic repository developed in Java (this link also from Anton). In its System Documentation [<http://www.ontotext.com/owlim/OWLIMSysDoc.pdf>] there is also a performance analysis.
- There is also this Scalability report on triple store applications. [<http://simile.mit.edu/reports/stores/index.html>]
- When looking for a stash to store ontology data in, analogous to a database, one googles for "semantic repository" [<http://www.ontotext.com/owlim/>]

].

- In Protege, if you get the An error related to DOT has occurred message, you are in need of the graphviz package. Under Debian, it can simply be installed by saying **apt-get install graphviz**.

Using Protege (4.0 alpha) as an editor, just to get acquainted with the matter, I tried to develop a little ontology all for myself. The idea behind is to have an ontology that is fit to derive configuration files from. As an example, the config file for a DHCP server will contain the IP numbers of name servers, paths to files on a TFTP server, MAC addresses of network cards of computers, etc. etc. DHCP configuration files are much like one another if we travel from server to server, yet they differ in what specific IP numbers, MAC addresses etc. are mentioned. We could easily generate such a file for our customer by asking them questions like "What is the IP number of your TFTP server?" "What are the MAC addresses and intended IP numbers of all network cards in the same VLAN the DHCP server is in, grouped by subnet and, within these groups, grouped by OS and boot method?" In order to store the answers to these questions, we could create a database.

However, by their nature some of the concepts we try to handle here are more suited to classes as used in object-oriented programming languages than to the rigidity of databases. If the ratio of number of tables to rows per table in a database is any measure of efficiency, the database doesn't give nice figures in this matter, and the effort to create it so that later data might be added, the effort to enforce business logic, and the effort to create and maintain the user interface(s) make this approach seem doomed to futility. Ontologies on the other hand are well suited to an approach that bears likeness to the classes of an object oriented programming language. It can easily handle IP numbers that stem from a list of servers no further specified and IP numbers that come from a list of network cards, handling them as IP numbers in both cases, but still distinguishing between IP numbers that we know as a property of a network card, and IP numbers that we know nothing more of.

An additional advantage of ontologies is that they offer data sharing across multiple repositories. If our customer wants to expand upon our knowledge base, they are free to do so, and we can even reimport their knowledge base back into our own, possibly designating it as read-only (see managing imports in protege-owl [<http://protege.stanford.edu/doc/owl/owl-imports.html>] and also SeparatingClassesAndInstances [<http://protege.cim3.net/cgi-bin/wiki.pl?SeparatingClassesAndInstances>]).

The above links were reached from the Protege WikiHomePage [<http://protege.cim3.net/cgi-bin/wiki.pl/>], which is a useful source of information, as is the W3C page on Semantic Web Best Practices and Deployment Working Group [<http://www.w3.org/2001/sw/BestPractices/>]. A nice graphical representation of the complexity of some semantic web tools and technologies is the Naive OWL Fragments Map [http://www.ontotext.com/inference/rdfs_rules_owl.html] (scroll to bottom of page).

The first issue I ran into is that of user-defined types. The list of available types in OWL is not as rich as that of most RDBs, and certainly not as easily extensible as object oriented programming languages. Where in my PostgreSQL databases I can define a column as being of datatype MAC-address or IP number (be it IPv4 or IPv6), I cannot do so in an OWL ontology. There is an effort to expand and extend the available data types [<http://www.w3.org/TR/2004/REC-owl-guide-20040210/#Datatypes1>] in user-defined datatypes in protege [<http://protege.stanford.edu/plugins/owl/xsp.html>], but it is of limited scale and it uses annotation properties of RDF, which are not visited by the reasoner, which makes the extended data types tag along instead of becoming part of the system. *Right now, intricate data types cannot be created in an ontology (just as they cannot be in databases), and a mapping from interface data types to ontology data types is still necessary, as is the implementation of some business logic and data restrictions in that interface.* We cannot define something as an IP number, we must store it as a simple string, use it as a string where we can, and when we need to do IP number arithmetic, we must convert it to an actual IP *number* in our interface. But as I discovered, we may not want these intricate data types anyway...

Another matter is that of unique identifiers. Protege can designate a property of a class as “functional”, which means that two instances with equal values in this property are inferred to represent the same instance. Whether this can be used in the same fashion as “primary key”s in databases needs some more reading and thought on my part. And then there are multi-column primary keys in databases, which allegedly are modeled with “Combined Inverse Functional Properties” in OWL. Haven't found a good source of information for this yet.

One more thing that is highly desirable to a systems administrator is knowing what information is present in what configuration files. As an example: if the IP number of our TFTP server changes, the DHCP configuration file will have to reflect this change. This might seem obvious, but often it isn't trivial to figure out all the places where a single change instantiates. So we would like to have a system that doesn't only generate configuration files, but also `knows' what information it needs to generate them.¹ If we are to pull this off, we might have to resort to an OWL full ontology, which cannot be handled by most reasoners, and which cannot be guaranteed to be handled by any reasoner in finite time. Or we can use one of the tricks described in the W3C draft Representing Classes As Property Values on the Semantic Web [<http://www.w3.org/TR/swbp-classes-as-values/>]. In any case, we can reason about classes as values, or we can reason about classes, but we cannot reason directly about the values of properties, and current reasoners cannot reason about instances either. So we probably shouldn't care much about whether we can do AND operations on IP numbers here. That is to be done in the interface.

Using NCPMOUNT to access central storage

`apt-get ncpfs`

```
ncpmount -t 900 -r 30 -A rugstor01.staff.rug.nl -S rugstor01 -U  
p012345.STAFF.RUG.NL /mountpoint/
```

¹ Another thing that we would like to keep track of is the interdependencies between machines, services and configuration files. E.g. what other services will stop if our DHCP config is broken. Although we can derive some of that information from a system that knows what information is needed for what config files, we cannot ultimately know what interdependencies are present in the actual system but not in the knowledge base, so we can never rely on such a system.

Chapter 5. November 2007

CPU buying advice

The problem at hand

Desktop computers are roughly divided into “low end”, “mid range” and “high end”. The boundaries between these sections of a single range remain a bit blurry. Available processors differ in cost, computations performed per second, compatibility with other hardware, power consumption (and consequentially, cooling needs), features, communication speed, and a host of other parameters. Still, most PC vendors offer some choice in the exact type of processor that goes into a PC. Our first question is: What criteria define the range of processors from which we wish to pick, and what values should demarcate its edges? And secondly, we want to know how to specify these half a year or more in advance?

Theoretical solutions

Assuming that the maximum price we are willing to pay is fixed, the price of the CPU as a percentage of the entire system would be a good candidate for a criterium, especially since that would give us a nice translation to compare against other optimization parameters such as amount of memory and CPU to memory bandwidth. However, this kind of information doesn't seem to be readily available on the Web, and it may be particularly hard to pry it from a PC vendor.

FLOPS characteristics on the other hand are indeed available, but although they may be good for comparing machines used for actual computing, they aren't very useful in benchmarking desktop machines, because they favour CPUs with good floating point performance, which is of lesser importance in office work, especially with the advent of powerful GPUs to handle eye candy. Benchmarking is better, but to know which benchmarks to pick, one must study the characteristics of the benchmarks, and have a good idea of what software is to be run on the machines.

Based on a short investigation, I come up with two candidates to obtain benchmarking data from: the infamous Tom's Hardware [http://www23.tomshardware.com/cpu_2007.html], with a choice of benchmarks to pick from, and PassMark software [<http://www.cpubenchmark.net/>] for a single benchmark that incorporates all use a desktop PC is likely to be put to. My advice would be to specify minimum requirements for any number of these benchmarks for a CPU to be considered.

Remains the question of what these minimum requirements are to be. In order to answer this question, we need to know how fast the graphs of CPU computation prowess slide by in time. If we had a list of release dates of the processors shown in the charts, we could just slide a 6 month window across the charts and we would have a good approximation of what would be reasonable to ask. However, I was unable to find such a list, and release dates of individual processors seem harder to come by on the Web than expected.¹

A poor but practical solution

So while we don't have the field experience to come up with cost of CPU as a percentage of the system or a history of CPUs used in the past, we still need to come up with an advice for the next six months or so. Our best bet in this case is probably to surf to the sites of some large vendors, e.g. HP [<http://h10010.www1.hp.com/wwpc/us/en/en/WF02d/12454-12454-64287.html>] or Dell [<http://www.dell.com>].

¹ The big manufacturers seem eager to forget what they produced months ago, and the search results on the Web in general are very much contaminated with anticipative and promotional material. And if a date is indeed found, it remains unsure whether the “release” means actual availability of the device. And if that is certain, it is usually unclear which particular model in a range of processors is referred to by the particular poetic circumscription used.

, pick the current system of choice, step up the default CPU by two, look it up in the benchmark charts, and demand that the CPUs to be chosen from half a year ahead in time all rank higher on those benchmarks. If we do not trust ourselves on that, we may watch what others recommend. The University of Pennsylvania in the US states some sensible thoughts in their Computer Recommendations Page [<http://www.upenn.edu/computing/arch/standards/>].

Some more useful CPU bookmarks

- The Big Processor Guide [<http://www.10stripe.com/featured/cpu/index.php>]
- TechReport on CPUs [<http://techreport.com/cpu/>]
- Feature Comparison Chart at Intel.com [<http://www.intel.com/support/processors/mobile/pm/sb/CS-007967.htm>]
- List of Intel Core Duo processors at Wikipedia.org [http://en.wikipedia.org/wiki/Intel_Core_2] (with some approximate release dates)
- The AMD CPU choice start page [http://www.amd.com/gb-uk/Processors/ProductInformation/0,,30_118,00.html]
- The Intel Desktop CPU comparison page [<http://compare.intel.com/pcc/default.aspx?familyID=1&culture=en-US>]

Trusted/Treacherous Computing

As a founding member of the Trusted Computing Group [http://en.wikipedia.org/wiki/Trusted_Computing_Platform_Alliance]², HP builds into its PCs a so-called TPM, a Trusted Platform Module. This device can readily be used by the OS to identify the individual machine it is built into. Depending on the implementation of the chip, it could also be used to restrict the use of a PC to purposes approved of by the hardware manufacturer³. It can certainly be used by the manufacturer of an OS to put whatever restrictions they like on the use of the PC. From a vendor lock-in point of view⁴, it seems wise to demand of a PC manufacturer that they offer hardware without this 'feature'. If hardware with such a module is favoured for other reasons, or even, perhaps, for the very existence of it, the least we want to demand is that the module can be turned *off* by the BIOS in such a way that it cannot be turned on again by the OS while the machine is running.

Stopping the IWI net mail service

The current situation

Until now, the IWI has had its own mail servers, equipped with spam detection, anti-virus measures, local delivery with `procmail` enabled, and IMAP and POP services to retrieve mail from. It also kept its own mailing lists. Unless significant effort is put into buying and configuring new servers, this situation cannot persist for much longer, as the hardware is long out of service, and the people who supported the server have been reassigned to other tasks. This document tries to plan the moving of users' mail from

² I presume "Trusted" must be an adjective pertaining to "Computing".

³ Do a Google search on ""Treacherous Computing"".

⁴ I will not dig into ethics nor politics, nor even corporate warfare here :)

the old IWI server to the present CIT servers, and lists the changes that users will experience.

The new situation at a glance

The goal of the change is a situation in which all mail for the IWI is handled at the CIT servers. Mail generated at the IWI is directed to those same servers as soon as possible. Users fetch their mail from CIT IMAP servers.

The Transition

Where to make changes

At the IWI, as can be seen in Figure 5.1, “Sketch of current mail flow at the IWI”, `iwi200` is the mail server. Its configuration is in the directory `/etc/postfix` and the files therein. Some of the lists used by `postfix` were sometimes published via NIS from `/etc/mail/primary_mx_hosts/etc/postfix` on `iwi1`, but this no longer seems to be the case. The most important files in either directory are `aliases`, and `virtual`.

Note

Please note that `aliases` often resides in `etc` instead of in `etc/postfix`. The IWI MTA, `PostFix`, uses these two tables when delivering mail it considers itself the final destination for. In the case of `virtual` addresses, it looks up the address in the left hand side of the `virtual` table, and sends it on to the address(es) on the right hand side⁵. Once the address to deliver to is local (i.e. *not* virtual), the `aliases` table is used for the same sort of lookup.

Before being delivered, mail received by `iwi200` is scanned for SPAM and viruses at `iwi202`. `iwi202` Is also the Authenticated SMTP server. Mail received by `iwi202` via ASMTTP is not scanned, as the sender is known, and only IWI employees have accounts. `iwi202` Is not going to feature prominently in this plan, since it will just become jobless when `iwi200` is taken down.

Mail delivery at `iwi200` is influenced on a per-user basis by `~/.forward`⁶, and often also by `~/.procmailrc`⁷. If delivered without `.forward` or `.procmail` intervention, it will end up in `/var/spool/mail/${USER}`⁸, but **procmail** can also access home directories, and sort mail into boxes located there.

`iwi200` Serves the mailboxes in `/var/spool/mail` via IMAP and POP. Via IMAP, the users' home directories can also be reached.


Figure 5.1. Sketch of current mail flow at the IWI

⁵It does this over and over again until the resulting address occurs in the LHS no more.

⁶if it exists

⁷if that exists and is called from the `.forward`

⁸MBOX-style



Preliminary actions

Before the CIT mail server can start taking over *iwi200*'s duties, a couple of things must be sorted out. Among others, we must know which IWI account maps to which RuG account, and we must have some replacement for the IWI mailing lists. The required steps are listed in Preliminary steps to be taken before moving the mail

Procedure 5.1. Preliminary steps to be taken before moving the mail

1. Find a CIT address for every local IWI user.
 - a. List all addresses mentioned in virtual and aliases. Drop all addresses containing a “|” (pipe), drop all addresses *iwi200* doesn't consider itself final destination for (i.e. addresses with “@”-signs, but outside the {*iwinet,cs,math*}.*rug.nl* domains).
 - b. Use the **postalias** command to peruse both lists for reducing each of the remaining addresses to a local user (i.e. an address with no “@”-sign in it).

- c. For each of the remaining addresses, find the users' RuGmail address.

Note

Most of these users' RuGmail addresses can be found simply by looking in their `.forward` or `.procmail` files, where the forward we want to accomplish later is already set. For the other users, the secretaries (`secretariaat@cs.rug.nl`, `secretariaat@math.rug.nl`) will be able to provide so-called P-numbers if they know the users' actual name (as opposed to their username). In cases where this mapping from username to actual name is unclear, `peter@cs.rug.nl` will be able to clarify.

- d. Create a lookup table with local IWI users on the left hand side, and RuGmail addresses on the right hand side. This is the `IWI-to-RuGmail` mapping

Warning

There is a category of IWI addresses that do not have RuGmail equivalents. These include a few former IWI employees as well as family members of IWI employees. To my knowledge, the University of Groningen doesn't serve IMAP to people in these categories. Inform them of their predicate, asking them to supply an address to forward to. If they provide it, put it in the list. If they don't, service to them will necessarily be dropped.

2. Move mailing lists away from `iwi200`

- a. On `iwi200`, in `aliases` and `virtual`, find all simple lists, i.e. entries with a “,” (comma) in the right-hand column, and all lists managed by **slist** or **majordomo** (which are easily recognized by the occurrence of these strings in the RHS column). For each list, find all the recipients.

Note

The recipients on `slist`-type lists can be found in `iwi200/home/slist<name-of-list>/dist` on `iwi200`. Lists maintained by the **majordomo** command have their actual list of addresses in `/var/lib/majordomo/lists/<name-of-list>` in `iwi200`.

- b. For each of the lists, contact `systembeheer@cs.rug.nl` to ask whether the list can be dropped or not.
- c. For each of the remaining lists, contact the CIT helpdesk to ask for a list with the appropriate recipients on the CIT list server. If the list is going to have a new address, send notification to all recipients on the list.

3. Ensure correct mailflow from the IWI

- a. Install and configure a very simple mail server that
 - accepts mail only from within the IWI domains⁹,
 - does no SPAM or virus checking,
 - does no local delivery, only forwarding, to the CIT mailserver,

⁹This prevents it from becoming an open relay too.

- has a mapping from local usernames to RuGmail addresses,
- b. Ensure that local administration routines include adding an entry to this lookup table when creating a local account.
- c. Make all IWI machines use this simple server as a smarthost.

Note

This has the advantage of bringing all potential local problems to a single machine, where they will be seen sooner. It also postpones the need for the IWI to use central accounts only.

Note

This smarthost is meant for machine-generated mail only. Human users should use either `iwi200` or the CIT mailer for outgoing mail.


Redirecting the bulk of mail flow

The IWI SMTP server can largely be switched out of the mail circuit by setting the MX record for all domains owned by the IWI to point to the central SMTP server instead of the IWI smtp server. The steps to be taken are in Steps to be taken in taking the IWI SMTP server out of the mail flow .

Procedure 5.2. Steps to be taken in taking the IWI SMTP server out of the mail flow

1. Now that the above Step 2 has been completed, the mailing lists in `aliases` and `virtual` are superfluous. Comment them out, put their addresses in the `relocated` table along with their CIT counterparts, and have `PostFix` reload its tables.
2. All users must be warned that mail that used to end up in their IWI mailboxes is going to be redirected to the central mail server. Particular stress must be put on the fact that they should not forward mail from central accounts to the IWI servers any more. It must also be communicated to them that server-side mail sorting will end.
3. Use the `IWI-to-RuGmail-mapping` to rewrite the right hand side of `virtual` (and `aliases`, if applicable) in such a way that no entry redirects to local users any more. Have `PostFix` reload its tables.
4. Make the users' mail files in `/var/spool/mail` read-only, and put a tail on `/var/log/mail` to see whether `PostFix` has trouble. It shouldn't. The mail flow is now like in Figure 5.2, "Sketch of mail flow at the IWI with all forwards in place".

Figure 5.2. Sketch of mail flow at the IWI with all forwards in place



5. If possible, try to make `aliases` empty by this time, with all translations done in `virtual`.
6. Offer the `virtual` table as it is now to the CIT postmasters (possibly via the helpdesk), and wait for them to
 - make the CIT mailer accept mail for the IWI domains,
 - perform upon the addresses in those domains the address translations as specified in the table.

Warning

It may be against central policy to forward mail in the same way the IWI server does for local users with forwards to addresses outside the RuG. If that is the case, their mail may have to be stored in central mailboxes instead of keeping the forward. Notify them of that.

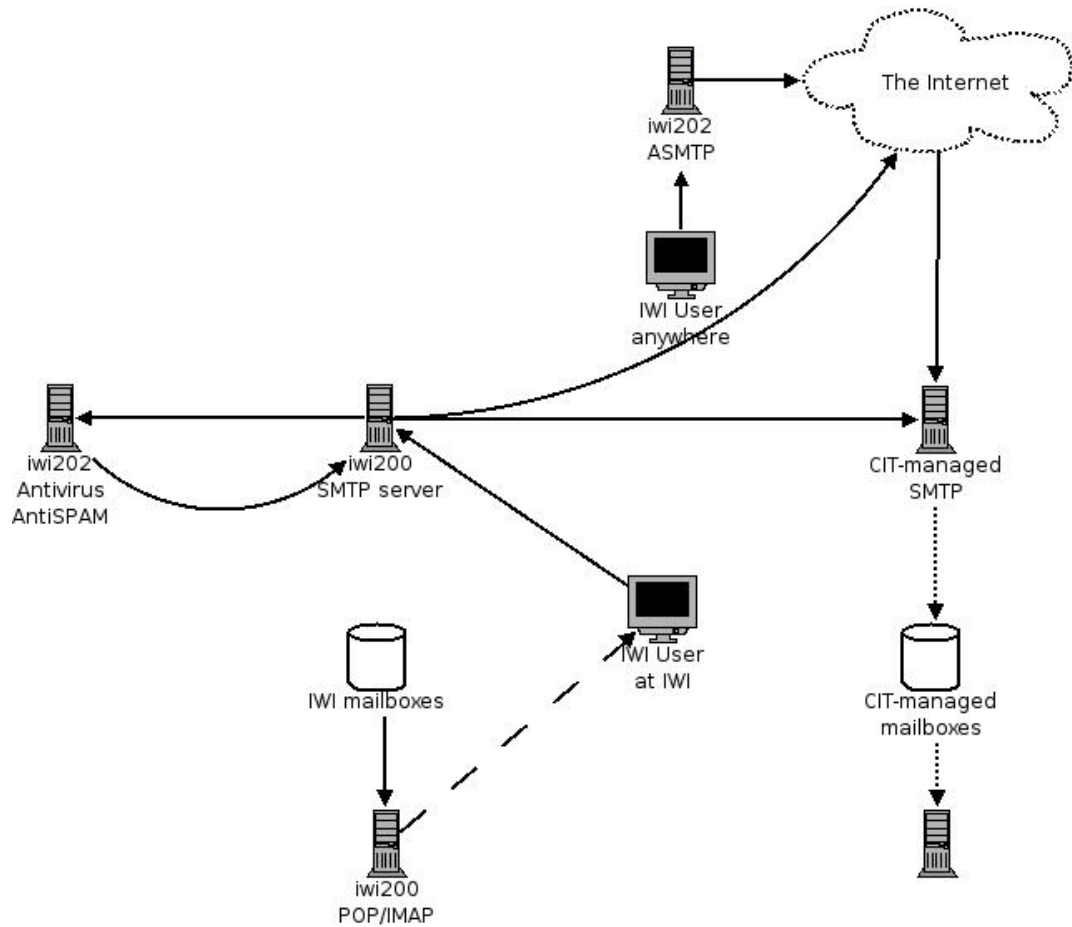
7. Test the new recipients on the CIT mailer, or have the postmasters test them, and verify that they work.
8. Once the previous steps are completed, the `MX` records of the IWI domains and all machines on them must be set to point to the central SMTP server. This will lead to a rapid decline of the mail

flow to the IWI servers, and a proportional increase of the flow to the central SMTP server.

Warning

It is not yet time to turn off `iwi200`, because many IWI users still send mail via it, as can be seen in Figure 5.3, “Sketch of mail flow at the IWI with MX records redirected”

Figure 5.3. Sketch of mail flow at the IWI with MX records redirected



Moving the users' mailboxes, and adjusting their settings

We need to take down the old POP/IMAP server as well as the SMTP server, as they are of equal age and state. So the users' mail must be moved from the IWI mailboxes to the central mailserv. This is near trivial, as modern mail programs are able to open both accounts at the same time, and entire mail folders can be dragged from the old account to the new one. The only problem lies in forcing the users to actually move their mail. Best approach is probably to state a deadline, but offer support to those who can't trust themselves with the task. While we're at it, the users must also be taught to use the CIT mailer for outgoing mail.

Procedure 5.3. Adjusting users' settings

1. Train two or three helpdesk employees to understand Linux mailers and the problem at hand.
2. Write and send out documentation telling the users that they should stop using `iwi200` for outgoing mail, and start using the CIT mailer. Offer helpdesk to those who don't understand.
3. Write and send out documentation telling the users that they should move the content of the IWI mailboxes to their RuGmail accounts. Offer helpdesk for those who don't understand.
4. Over the course of a few weeks, watch `/var/spool/mail` on `iwi200` grow thinner. Then, tar and gzip whatever remains, and send it to those who didn't clean up their mailboxes.

Warning

One could wish to ensure that no mailboxes were used in users' home directories any more either. This is done by turning off the IMAP and POP daemons, and waiting for complaints.

Cleanup of `iwi200`

Now that this is all done, we get to the situation as illustrated in Figure 5.4, “Sketch of mail flow at the IWI with users talking to CIT servers”. All that is left to be done is to watch the logs of `iwi200` to ensure there is no mail coming in any more. If need be, the local users' `@iwinet.rug.nl` addresses can be put in the `relocated` table for a while before the service is finally turned off and we end up with Figure 5.5, “Sketch of mail flow at the IWI with `iwi200` off”. Turning off the ASMTTP service at `iwi202` has much lower impact, and can be handled separately.

Figure 5.4. Sketch of mail flow at the IWI with users talking to CIT servers

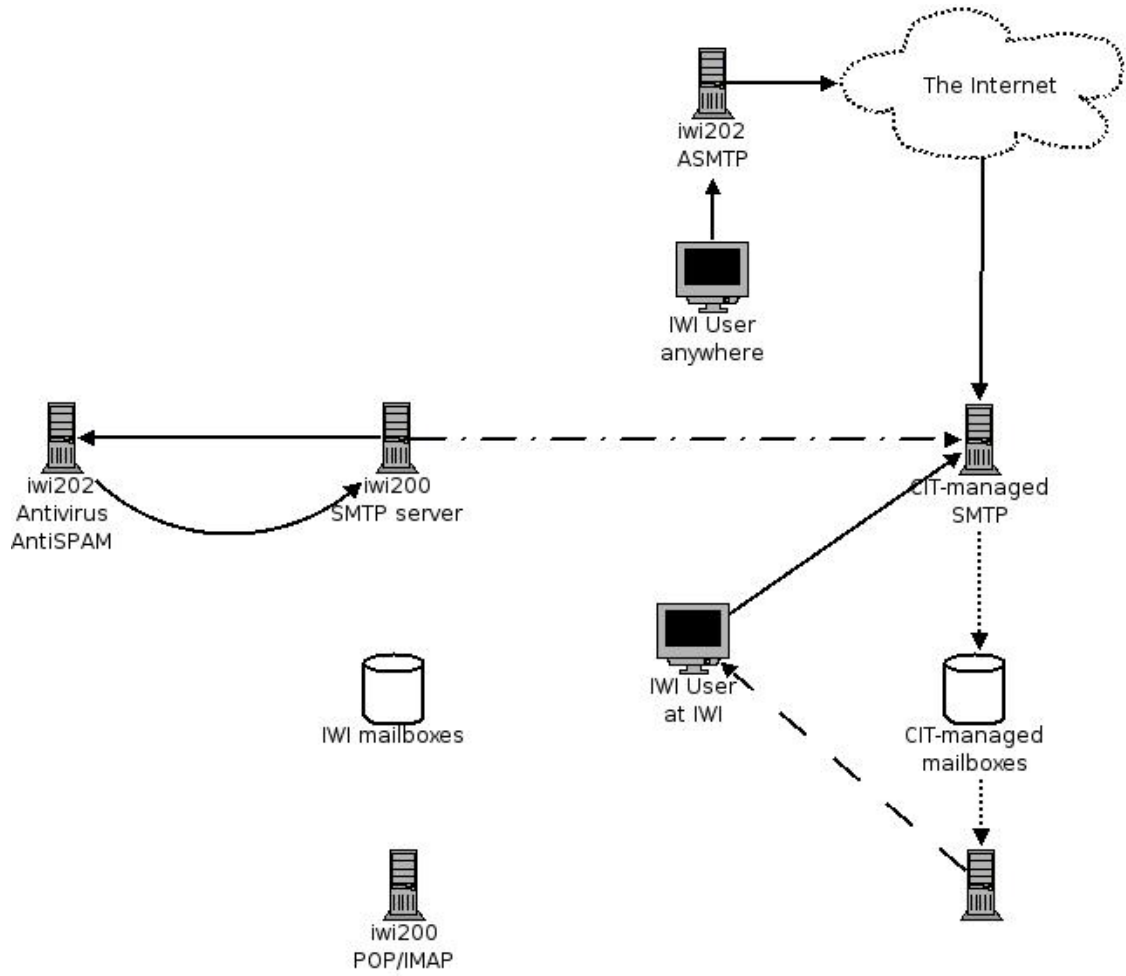



Figure 5.5. Sketch of mail flow at the IWI with iwi200 off



Mounting USB devices from udev with permissions of the user at the console

We assume that a user plugs in their USB device only after having logged in. When they do, we want the device to automatically become mounted, under the ownership of the user that controls the console.

We use Writing udev rules [http://reactivated.net/writing_udev_rules.html] for documentation, and connect the device.

Procedure 5.4. Figuring out the device characteristics and writing a matching rule

1. `lsusb` shows:

```
Bus 004 Device 014: ID 0421:0410 Nokia Mobile Phones 6630 I
```

2. `find /sys/ -type f -exec grep -li nokia {} \;` `2>/dev/null` shows:

```
/sys/devices/pci0000:00/0000:00:1d.3/usb4/4-1/product  
/sys/devices/pci0000:00/0000:00:1d.3/usb4/4-1/maker
```

3. `udevinfo -a -p /sys/devices/pci0000:00/0000:00:1d.3/usb4/4-1/` gives:

```
<snipped for brevity>  
looking at device '/devices/pci0000:00/0000:00:1d.3/usb4/4-1/  
KERNEL=="4-1"  
SUBSYSTEM=="usb"  
DRIVER=="usb"  
ATTR{product}=="Nokia 6630"  
ATTR{manufacturer}=="Nokia"  
ATTR{maxchild}=="0"  
ATTR{version}==" 2.00"  
ATTR{devnum}=="14"  
ATTR{speed}=="12"  
ATTR{bMaxPacketSize0}=="64"  
ATTR{bNumConfigurations}=="1"  
ATTR{bDeviceProtocol}=="00"  
ATTR{bDeviceSubClass}=="00"  
ATTR{bDeviceClass}=="02"  
<snipped for brevity>
```

Note

Note the “SUBSYSTEM” here and the “SUBSYSTEMS” in the udev rule.

4. So in `/etc/udev/nokia6630.rules`, we create a rule like

```
DRIVER=="usb", SUBSYSTEMS=="usb", BUS=="usb", SYSFS{manufacturer}=="Nokia", SYS
```

5. And in order to activate the rule, we **cd** to `/etc/udev/rules.d` and link to the file containing the rule: **ln -s ../nokia6630.rules 40_nokia6630.rules**.

Note

This link needs to be in a proper place in the alphabetical order, as too late may mean that other rules have been applied to the device already.

6. **udevcontrol reload_rules** to make the udev subsystem reconsider its position on our device.
7. Then we re-plug the device, and we should see some effects in the logs, and `/dev/nokia6630` appears.

Note

It is also possible to run **udevtrigger** instead of replugging.
We now have a mount point and a device owned by the console user.

Note

I saw some commands pertaining to USB subsystems still fail because they access `/dev/bus/usb` and deeper, which they weren't allowed to read (let alone write).

Tracking down the authoritative nameserver

The authoritative nameserver for host `myhost.mydomain.com` can be found with **dig my-host.mydomain.com +trace**.

Reverse lookup is with **dig -x 189.165.13.239 +trace**.

`dig` is part of the `dnsutils` package under Debian.

Installing Mathematica 6.0 under Linux

The problem

Mathematica 6.0 was substituted at short notice for 5.2, and I'm trying to install it. I tried to install it from `wingtip85`, and that worked, but now it won't start, twice complaining:

```
xset: bad font path element (#100), possible causes are:  
Directory does not exist or has wrong permissions  
Directory missing fonts.dir
```

Incorrect font server address or syntax

Wolfram has a page about font problems [http://documents.wolfram.com/mathematica/GettingStarted/SystemAdministrationGuide/UnixLinuxAdministration/Fonts.html].

The solution

1. The Mathematica fonts need to be present on the *X server*. Copy them from the X client (that has the Mathematica binaries) to the X server:

```
scp -r /opt/Mathematica-6.0/SystemFiles/Fonts/Type1 user@othermachine:~/MathFonts
```

2. On the X server, add the new font directory to your font path:

```
xset fp+ ~/MathFonts
```

3. Start Mathematica. Half of the error message will still show up, but the program will start and run.

Note

As a side note: as of version 6.0, the IWI license server has been replaced with the University wide one.

Centralizing the DHCP service

The problem

The IWI DHCP server is a single machine with less-than optimal backup and no access for helpdesk personnel. By putting the DHCP tables in a NetBeheer [https://rasp1.service.rug.nl] database (as built and maintained by Arjan Meijer) and periodically restarting the server, we ascertain access for helpdesk personnel, and by merging with other DHCP servers we reduce maintenance cost.

The config cannot be transferred as-is, because some information currently present in the DHCP tables doesn't fit in to the database without modification. Some files that are in the DHCP config tree but are not actually linked to contain information that shouldn't be lost, and some comments added in host declarations should also be put into the database.

The Implementation

1. For NetBeheer efficiency DHCP options need to be at the bottommost group level instead of higher group levels or even subnet or shared-network level. I migrated most options to that level.

Note

While doing so I also cleaned up a bit. This is bad practice in principle, but I did it anyway.

Note

There are some DHCP options that are at higher level even in NetBeheer. These options I left at their levels.

2. We are now in the process of waiting and seeing what broke after this change. A serious problem does occur (see the section called “ No Remote Login possible at the IWI ”, but I don't believe it is connected to the DHCP config.
3. Another problem occurs: Windows users cannot log in any more: the clients complain about `Tree or Server not found`. I revert to the old config. Will retry with the Windows section of the DHCP config unaltered (instead of amended with higher-level section data).

No Remote Login possible at the IWI

The Problem

I try to do my daily login at the IWI, and ssh waits a long time (three minutes), then the connection is cut off. This is symptomatic for NIS-servers not functioning. A day earlier I have seriously edited the DHCP config (see the section called “ Centralizing the DHCP service ”, and I am worried that I have mis-directed the NIS clients to the wrong NIS server, a problem which would typically surface after the clients renew their DHCP leases.

Investigation and solution

Since my own account is a NIS account, I cannot use that on the clients.

Warning

I try to use the root account, but the ssh times out on the NIS first, before even trying the root account. This is a security problem.

The servers don't accept remote root login, so I seem to have no access to the institute. However, the students' domain still has machines that will accept my username (which is an indication that DHCP might not be the problem), and via them I can get to the NIS server. From there, I can get to the clients, and NIS is indeed the problem. (I cannot `su` to my account.) However, the NIS settings on the client are correct, and restarting the NIS client doesn't solve the problem.

Restarting the NIS server *does* solve the problem. But it reduces my possibilities for finding out what the problem was in the first place.

Part II. 2008

Table of Contents

6. January 2008	58
A Configuration Repository idea	58
Introduction	58
What <i>is</i> configuration?	58
What is <i>not</i> configuration?	58
Storing Configuration and the Unit of Configuration	58
Parametrization and its implications	60
Cups Command Line Options	60
No Sound on the student PCs (unresolved)	60
In search of a proper Keyboard	60
Imaging Linux boxes with Zenworks imaging	61
Can ZENworks imaging be used to clone a Linux machine (and if so, how)?	61
What are the restrictions ZENworks imaging puts on the way a machine is partitioned?	62
What limitations does ZENworks imaging impose on the filesystems used?	62
Can we use ZENworks imaging to put Linux in a designated space on a harddisk without damaging any other OSES or data already present on the disk?	62
Does ZENworks imaging support RAID? LVM?	62
Can we safely use ZENworks imaging on machines with multiple disks?	62
No SSH to my server possible	63
Creating a parser with Bisonc++ and flex	63
7. February 2008	64
Slapd takes 100% CPU on sched_yield()	64
SSH tunneling	64
Using Bacula for backup	64
debmirror on Ubuntu	65
Newest OpenSSL and BIND on a 64-bit Debian machine	65
Multipath Fibrechannel interface to SAN under Ubuntu	65
X access from under sudo	67
Quick source NAT with IPtables	68
8. April 2008	69
Installing the SuSE iPrint client under Debian	69
Converting the Novell iPrint client to a Debian package	70
9. May 2008	75
Installing OpenBSD on a Soekris Net5501-70	75
10. June 2008	83
Installing Linux over Windows without BIOS access	83
Turning nVidia driver on on machines that have the libraries and an NVidia card ..	84
A Firewall Install Script	84
Fixing the NIS port	86
Transferring the IWI printers to IPrint	86
Remote Firefox acutally remote	87
11. July 2008	88
Installing SpaceWalk (using a remote database)	88
Installing CentOS unattendedly	92
Remote access to Windows XP from Linux	93
Creating a SpaceWalk Channel	94
12. August 2008	96
Booting Ubuntu Hardy unattendedly using preseed	96
Cloning NTFS partitions at the file level (and booting them)	96
Introduction	96
Expectations	97
The experiment	97

Conclusion	98
Using DHCP-initialized PXELinux under VMWare	99
Labelling partitions during Linux unattended install	99
Introduction	99
Debian/Ubuntu	99
Reverse Engineering the ERD of an Oracle database	99
Using WebDav to connect to the so-called Y:-drive	100
XML versus web template engines	100
Installing 64-bit Matlab on Linux	101
Installing 64-bit Maple 11 under Linux	102
Installing 64-bits Mathematica on Linux	108
13.	112
Unable to mount USB disk under Debian	112

Chapter 6. January 2008

A Configuration Repository idea

Introduction

A systems administrator produces configurations, of computers and related devices, and of groups of these. The configurations I produce are often interrelated, and sometimes differ only minimally between different installations. They are most often a mixture of copied and written material, and parts of it may be generated. I would like to have a way of storing them, and a way of effectively re-using configuration I created in the past.

What *is* configuration?

When I have to configure some machines, I start with the bare machines ¹. The hardware, e.g. architecture of the CPUs, the number of harddisks per machine, or type of network connectivity device, are the first configured items. Then I pick an OS, a particular version of it, and I boot a particular kernel, and perhaps modify some boot parameters. Usually the machines have disks that need to be partitioned, the partitions need to be assigned to tasks, networking devices must be assigned IP numbers. Then, when the machines are running, they must perform their tasks. Some service, even if it is merely the fact that the machine exists and is listening on the network, must be offered to some clients. Once the machines are running, we get to finetuning. Some services are offered to some clients only. Some values in a configuration file may have to be altered to optimize performance.

So there we have it. *Configuration consists of the choices made when putting together a system* ². Some configuration is stored in files on the configured system, but some is also stored implicitly before the system is even connected to a power supply. Some configuration items are rather independent, but other items are closely interrelated, and one doesn't make sense without the other. However, often only the configuration is stored, and the sense is in the mind of the configurator, and in some commentary on the side if we're lucky. But there is always a reason for configuration, if only that it was there by default and we didn't change it.

What is *not* configuration?

Once a system is running, it changes with time. Logs are filled, users put their data somewhere, databases are filled. This is not configuration, it is *state*. I do not want to store state, I want to store configuration. The line between the two is sometimes thin. A system may be configured to periodically fetch and apply available patches to its software. After two years of running and being patched, what is its configuration? Is it the way it was installed originally, or is it the way it was installed originally *plus* two years' worth of patching?

Storing Configuration and the Unit of Configuration

If our purpose is to recreate a system from scratch with only its hardware and its configuration, we must devise a way to store *any* possible configuration, not only that which consists of files, but also the configuration implicit in e.g. hardware choices, and the way devices are interconnected. As configuration consists of choices made, we must at least be able to store all the possible choices, with the reasons for each. But there are some cases where we choose not to choose, and that still is a choice, which must be stored. Let's try as an exercise for the mind to configure a four-machine team of web servers.

¹ Even if I choose the hardware myself, for the purpose of this article, I still start with the bare machine.

² And as such, it can be stored separate from the system. If we have the configuration, and the hardware is still available, we can rebuild the system from scratch.

Choice of hardware can be achieved by listing name and type of parts (e.g. An ASUS P5M2-E/4L motherboard, a Quad-Core Intel Xeon 3000 processor, 16 GB of RAM, etc etc), and interconnectivity can be described by listing the connectors of each device, and listing which is connected to which. Some of the hardware is chosen for a reason, e.g. availability of a certain part, or performance. If so, it makes sense to store the reason with the configuration, or even as part of it. It makes sense to me to store the configuration of one machine as a unit, and making clear that the other three have the same configuration as the first one. When that is clear, we can mention that the four are in some way connected by network wires.

The choice of OS can be simply from a list, but here we encounter the first interdependency: some OSes are not available for some hardware, or have no drivers for them. So we must refer back to choices made earlier, in order to explain the ones made here. Kernels to be booted have names, and perhaps paths. Boot parameters are just strings.

But we do get to a fundamental choice here. Some OSes have default boot parameters. We could decide to store these default boot parameters with our configuration, or we could store just what we added and what we removed. If the default has changed by the time we try to reproduce our system, we may get different results from what we got the first time if we store only the difference. So it is probably best to store the default, mark it as default (the distributor's default, not ours), and store what we finally decided was best as well.

Then the partitioning. We could store sizes as absolutes, as percentages of disk size, as combinations of these (minimum, percentage, maximum), or we could refer to an entire algorithm of computing the partition sizes. Network devices are a bit tricky -as are disks- in that they don't always appear with the same name, even across reboots of a single system. So there must be some way of circumscribing a device, e.g. "the disk that has tag so-and-so", "the networking device with MAC-address 00:11:22:33:44:55", "any disk with more than 20GB on it, which has space for another virtual ext3 fs of at least 700MB" or "any network device that can reach www.google.com".³

As for the list of software installed, there is software that we explicitly need, which must be listed, and there is software that tags along, which must also be listed, just in case we were unknowingly using it and depending on it anyway.

Now the machine is running, and we start to configure in earnest. First of all, as these are public web servers, with Debian Linux and Apache, we enable IPtables and set it to drop all traffic except on port 80 and port 22, and to allow only two specific IP numbers on port 22. This piece of configuration depends on the fact that IPtables is available and installed. It also depends on the names of the network interfaces. Furthermore, it makes not all that much sense if no HTTP or SSH server is listening, but it doesn't do any harm either. On the other hand, it *does* harm things if any other service were to function on this machine, since that would be rendered useless. All ethernet packages to it would be dropped. This information needs to be stored. Then the actual configuration is done in several files. With a GUI, an IPtables script is created, and it is stored in `/etc/firewall`. Another script, `/etc/init.d/iptables`, is created that uses the former to turn on or turn off firewalling. Due to the names and locations of these files, the configuration is fit only for Debian, or more precisely, for distributions that have their initscripts in `/etc/init.d` and don't have their own files in `/etc/firewall`.

Resuming, we can state that "configuration consists of the choices made during the installation of a system". With the entire configuration available, we can buy new hardware (if available) and bring up a new system that is for our purposes identical to the old one when that was installed. State is *not* a form of configuration. Items of configuration come with a reason, which should be stored with the configuration. They usually come with a default, which should also be stored. Stretches of configuration that exist for the same reason can be grouped together, but stretches of config that exist for different reasons must be split. It must be possible for different stretches of config that are spread over separate sections of a file, over different files or directories, or even across different systems, to be combined in a larger package.

³ This cries for plugins, but they will have dependencies of their own. Figuring out the size of a disk works differently under Debian than under Windows Vista.

Parametrization and its implications

Consider a modest Linux desktop machine I'm going to use at work. Its settings came straight out of the box, but I provide it with a local `IPtables` configuration, which allows only my old desktop and my home machine to access it via `ssh`. This is something I do often. My home machine is accessible only from work, and I have other machines that have similar restrictions. I would like to store a single version of my `IPtables` script, and for each instance just change the list of other machines that are allowed `ssh` access.

This can be accomplished using *parametrization*. We split the part of the config that remains constant: “grant access to this machine via `ssh`”, and the part that varies: the list of grantees. Then we store both parts separately, and use some program to join them together in order to recreate the config file.

Let's assume that we implement this as a `config server`, and that the `config client` downloads its configuration from this server.

Cups Command Line Options

When using a Cups printer server, command line options can be passed to `cupsdoprint` as well as `lpr` to specify how to print. The full list is at <http://www.cups.org/documentation.php/options.html>, but the following example comprises all I usually need:

```
lpr -Ppr2 -o media=A3,Transparency -o landscape -o sides=two-sided-long-edge -o fitplot my-file.jpg
```

No Sound on the student PCs (unresolved)

The students have new PCs, but no sound. The `kmix` program shows no devices to control. And the sound card isn't recognized either: `00:1b.0 Audio device: Intel Corporation Unknown device 293e (rev 02)`.

So we update the PCI specs:`update-pciids`. Now `lspci` shows the following output:`00:1b.0 Audio device: Intel Corporation 82801I (ICH9 Family) HD Audio Controller (rev 02)`.

Perhaps the hardware detection system can now handle our device... Nope. Work in progress.

In search of a proper Keyboard

I'd like to have a keyboard with the standard IBM 104 key layout my employer uses everywhere, but with an ergonomic, tilted shape, a bit like the Microsoft ergonomic keyboards used to have. But of course I'd like to have one with mechanical key switches instead of the spongy touch of a membrane. And a *low* key operation force please...

So we find Adesso (all rubber), Cherry (no tilt, no tent), Maxim, Datahand (filthy expensive, and allegedly breaks often), ErgoMagic from ComfortKeyboard (Conductive rubber), GoldTouch (rubber), Kinesis (perhaps not rubber)...

Table 6.1. Keyboard Features

Manufacturer	Type Name	Switch Type	Adjustability	Layout	Actuation force	Price	Other
Belkin							
Logitech							

Manufacturer	Type Name	Switch Type	Adjustability	Layout	Actuation force	Price	Other
Microsoft							
Adesso							
Datahand							
ComfortKey-board	ErgoMagic [http://www.comfortkey-board.com/keyboards_ergo-magic.html]	Conductive Rubber Membrane	Split, tent, turn, change places				
GoldTouch							
Kinesis	Maxim [http://www.ergocanada.com/products/keyboards/kinesis_maxim.html]	Membrane	Split, tent	52 gram	??		
Kinesis	FreeStyle [http://www.kinesis-ergo.com/freestyle-solo-features.htm]	Membrane	Completely free	??	??		

Imaging Linux boxes with Zenworks imaging

Novell Zenworks [http://www.novell.com/products/zenworks/] has so-called “preboot services”, one of which is imaging [http://www.novell.com/documentation/zlm72/lm7admin/index.html?page=/documentation/zlm72/lm7admin/data/bve6kps.html#bx5wxyg]. My colleague Arjan has set up and maintains a few imaging servers, and I’m trying to find answers to the following questions:

Can ZENworks imaging be used to clone a Linux machine (and if so, how)?

Yes, it can, but only to a very limited extent. I restored a simple setup with the root filesystem on a primary partition (/dev/hda1) and swap on a logical (/dev/hda5). I was able to include the the swap in that image, but restoring it failed. Then I took an image that excluded the swap, but when I restored that image, the swap partition was left out as well. So... if we want to restore a Linux disk that has a swap partition, we have to do some partitioning by ourselves. That may be practical for the occasional restore, but it isn’t for an enterprise grade deployment system.

What are the restrictions ZENworks imaging puts on the way a machine is partitioned?

None. But if you want to restore, and have to do any partitioning (as is already the case if you have a swap partition -see previous answer-), then all tools you have are **fdisk**, **sfdisk** and the **img** program Novell provides. There are a few things I think should be repaired in the **img** program before it becomes useful⁴:

- It cannot remove logical partitions (it can only remove the entire extended partitions).
- It refuses to create swap partitions of sensible size, only succeeding if they are *huge*.
- It sometimes goes into interactive mode when provided with command line parameters, thus making scripting hard.
- It has a curses-like GUI, but the GUI sometimes stalls, not reflecting what is happening any more.
- When restoring only a single partition from an image, it retrieves the entire image, *and* instead of just skipping over the entire lot, it checks every file in the image, and skips it if it isn't in the appropriate partition.

What limitations does ZENworks imaging impose on the filesystems used?

As far as I see, any filesystem can be used, but only FAT16, FAT32, NTFS, Extended2 and ReiserFS are handled at the file level. Other filesystems are copied byte by byte. This may mean that they cannot be restored to partitions that differ in size from the one they came from. And it is annoying that it cannot create just partitions, instead insisting on creating the filesystem too.

Can we use ZENworks imaging to put Linux in a designated space on a harddisk without damaging any other OSes or data already present on the disk?

Yes, but only if we do the partition handling ourselves via **sfdisk**, or if we put linux in an extended partition *of its own*, and do not use a swap partition.

Does ZENworks imaging support RAID? LVM?

No.

Can we safely use ZENworks imaging on machines with multiple disks?

After the other results, I didn't even try.

⁴I may be malinformed, but then the criticism in these lines goes straight back to the makers of the documentation on the program

No SSH to my server possible

Symptoms: I have a server, to which I have an OpenSWAN IPsec connection, and I also use ssh sometimes. Now my client has rebooted, and the IPsec connection remains down, while I cannot ssh to the server any more. The server is still up, and via other machines I can still access it, with ssh too, just not from this particular client.

It turns out that on the server, the IPsec daemon is so busy trying to reconnect, using the SSH port, that the client never gets its request through.

Creating a parser with Bisonc++ and flex

I used flex [http://www.gnu.org/software/flex/manual/html_node/flex_toc.html] and bisonc++ [<http://bisoncpp.sourceforge.net/bisonc++.html>] to create a rudimentary parser for syslinux config files [<http://syslinux.zytor.com/faq.php>]. I realize that a parser is overkill to extract the location of the image-to-boot and the boot parameters from such a simple file, but I also needed a way to distinguish these config files from others, and I just longed to have a go at creating a parser. These are the guidelines I set myself for the next time:

- Finish the scanner before starting work on the parser.
- Don't do in the parser what can be done in the scanner (like e.g. joining newlines and whitespace into a single token).
- Character classes in the scanner go within range brackets, like this: "[[:blank:]]".
- When trying to create a parse tree while using non-pointer semantic values, one shouldn't use a list of pointers in the node, as the location they point to will be overwritten by subsequent parsing actions.

Chapter 7. February 2008

Slapd takes 100% CPU on sched_yield()

The LDAP daemon slapd (ancient version 2.2.23) and slapcat share 90% CPU between themselves on a machine I want to use for other tasks too. My colleague Heiko says that with a version this old, this may be caused by calling slapcat while slapd is running. Slapcat will try to access the database directly instead of using the LDAP protocol, and two processes accessing the same BDB backend simultaneously is more than the backend can handle. So I kill slapcat, kill slapd, restart the slapd daemon. Now the slapd daemon takes nearly 100% CPU, and **strace** shows that it spends all that time on the “sched_yield” call. This, according to many sources on the web, is caused by a corrupt backend, which can be mended by **cd**ing to `/var/lib/ldap/databasename` and running **db4.2_recover** without parameters. This works: slapd can now be started, and takes acceptable toll on the CPU.

SSH tunneling

I never did this before, but an SSH tunnel is made like this: **ssh -f -L 1234:localhost:902 remote-host sleep 10**

The “-f” tells `ssh` to background itself just before command execution, the “-L [bindaddress:]port:localhost:hostport” tells `ssh` to encrypt and forward local connections to port on localhost to hostport on remotehost. So in this case, if `vmware` is listening on port 902 on remotehost, we can connect a local `vmware` to port 1234 on localhost, and talk to the remote server encryptedly.

Using Bacula for backup

Bacula [<http://www.bacula.org>] is a nice solution for backing up moderate levels of data in simple networked configurations. It is not as fast as commercial solutions, and more difficult to configure than a cron-tabbled `rsync`, but it will handle multiple clients, multiple storages, and multiple backup regimes as long as you don't overask. Here are some of the pitfalls I encountered the last time I reconfigured it.

1. I use Bacula with PostgreSQL, and the scripts are geared more towards MySQL. The scripts mentioned in `/etc/bacula/bacula-dir.conf` for dumping the database can easily be replaced by a single call to `pg_dump`. This saves me the hassle of understanding the included scripts, which are just wrappers anyway.
2. Second point is to get authentication right. The passwords and resource names in `bacula-dir.conf`, `bacula-sd.conf`, `bacula-fd.conf` and `bconsole.conf` must match one another as stated in this FAQ-entry [<http://www.bacula.org/en/rel-manual/faq.html#AuthorizationErrors>].
3. Volume recycling is a nice feature, but I prefer to have a new volume for every run, and to manually throw away volumes when space is running out. Must make sure that volumes have been “pruned” by that time, i.e. the database has forgotten what they contain.
4. Automatic Volume Labeling isn't completely clear to me yet. Although I have set up everything according to the examples, I still had to ‘add’ volumes to my storage once. This was a one-time event, and new volumes are automatically added, but I still don't see why it wouldn't work the first time.
5. If “Use Volume Once” is set on a pool and only the date (as opposed to only the time) is used to distinguish the file to back up to, no job can run twice a day. (Talking 'bout kicking in open doors, eh?)
- 6.

debmirror on Ubuntu

The `~/ .gnupg/trustedkeys` file used by the **debmirror** script has another name under Ubuntu. Symlinking to the expected name is enough.

Newest OpenSSL and BIND on a 64-bit Debian machine

- I try to compile the latest OpenSSL, and the compiler barfs that it cannot find the most basic of include files.

Problem: the `libc6-dev` package hasn't been installed (as it shouldn't be on a server, but that's another matter).

Solution: **apt-get install libc6-dev**

- Problem: `libc6-dev` won't install, because `libc6` is the wrong (too high) version.

Cause: somebody has installed one from another version of Debian (this is Etch), which can still be seen in the history (**dpkg -i libc6_2.7-8_amd64.deb**).

Solution: **apt-get install libc6=2.3.6.ds1-13etch5 && apt-get install libc6-dev**

Now OpenSSL does compile.

- So we do **cd /home/jurjen/src/openssl-whateverversion ; ./config --prefix=/home/jurjen/openssl**

Then we do **make && make install**

- Then we proceed to BIND: **cd /home/jurjen/bind-someversion ; ./configure -**
-prefix=/home/jurjen/bind --with-openssl=/home/jurjen/openssl¹

We proceed with **make install**.

Multipath Fibrechannel interface to SAN under Ubuntu

1. **apt-get install multipath-tools**
2. Edited `/etc/multipath.conf`:

```
defaults {
polling_interval    30
failback            immediate
no_path_retry       5
rr_min_io           100
path_checker        tur
```

¹ BIND deems some versions of OpenSSL too old, and will complain if it encounters one. It didn't, so apparently it is satisfied with the OpenSSL I offered

```
user_friendly_names yes
}

multipaths {
multipath {
wwid 360050768019101ca2800000000000031
alias IBM-2145
}
}
```

multipath -l now gives:

```
IBM-2145 (360050768019101ca2800000000000031) dm-4 IBM,2145
[size=2.0T][features=1 queue_if_no_path][hwhandler=0]
\_ round-robin 0 [prio=0][enabled]
\_ 3:0:0:0 sdb 8:16 [active][undef]
\_ round-robin 0 [prio=0][enabled]
\_ 3:0:1:0 sdc 8:32 [active][undef]
```

3. Now, `/dev/sdc`, which we used to mount, is in use by `multipathd`, and we cannot mount it any more.

/etc/init.d/multipath-tools restart

mount -t xfs /dev/disk/by-id/dm-uuid-mpath-360050768019101ca2800000000000031 /srv/linux/ runs forever with no CPU load, and cannot be killed, so I think it is waiting for I/O.

/bin/sh -f /usr/sbin/xfs_check /dev/disk/by-uuid/../../mapper/IBM-2145 also ends up waiting for I/O.

Furthermore, `/var/log/daemon.log` reads:

```
...
Feb 23 18:06:16 sil7 multipathd: 8:16: mark as failed
Feb 23 18:06:17 sil7 multipathd: 8:32: mark as failed
Feb 23 18:06:46 sil7 multipathd: 8:16: reinstated
Feb 23 18:06:46 sil7 multipathd: 8:16: mark as failed
Feb 23 18:06:47 sil7 multipathd: 8:32: reinstated
Feb 23 18:06:47 sil7 multipathd: 8:32: mark as failed
Feb 23 18:07:16 sil7 multipathd: 8:16: reinstated
Feb 23 18:07:16 sil7 multipathd: 8:16: mark as failed
Feb 23 18:07:17 sil7 multipathd: 8:32: reinstated
Feb 23 18:07:17 sil7 multipathd: 8:32: mark as failed
Feb 23 18:07:46 sil7 multipathd: 8:16: reinstated
Feb 23 18:07:46 sil7 multipathd: 8:16: mark as failed
Feb 23 18:07:47 sil7 multipathd: 8:32: reinstated
Feb 23 18:07:47 sil7 multipathd: 8:32: mark as failed
Feb 23 18:08:16 sil7 multipathd: 8:16: reinstated
Feb 23 18:08:16 sil7 multipathd: 8:16: mark as failed
```

```
Feb 23 18:08:17 sil7 multipathd: 8:32: reinstated
Feb 23 18:08:17 sil7 multipathd: 8:32: mark as failed
Feb 23 18:08:46 sil7 multipathd: 8:16: reinstated
Feb 23 18:08:46 sil7 multipathd: 8:16: mark as failed
Feb 23 18:08:47 sil7 multipathd: 8:32: reinstated
Feb 23 18:08:47 sil7 multipathd: 8:32: mark as failed
```

And **multipath -l** now gives:

```
IBM-2145 (360050768019101ca2800000000000031) dm-4 IBM,2145
[size=2.0T][features=1 queue_if_no_path][hwhandler=0]
\_ round-robin 0 [prio=0][enabled]
\_ 3:0:0:0 sdb 8:16 [failed][undef]
\_ round-robin 0 [prio=0][enabled]
\_ 3:0:1:0 sdc 8:32 [failed][undef]
```

(It seems to cycle through “failed” and “active”).

Also, after **/etc/init.d/multipath-tools stop**, I/O still fails on both `/dev/sdc` and `/dev/sdb`. And when the `multipath-tools` are running, I/O fails on `/dev/mapper/IBM-2145`.

4. Rebooting solves. We can now at least mount `/dev/mapper/IBM-2145` read-only. Still it is a Windows solution. And I don't understand what went wrong.
5. And now the problem is back: the filesystem on `/dev/mapper/IBM-2145` is mounted read-only (and without filesystem check) on `/srv/linux`, but both a **find** on the mount point and an **xfs_check** on the block device end up waiting for I/O.

And this time, rebooting ends with a machine that responds to ping, but nothing else. I have no idea whether it hangs in shutdown (the unmount?) or didn't come up properly.

Note

Arjan solved this one. It turns out that the add-in card was broken. He replaced it with a new one, and in the process replaced the entire server too.

Note

This was something of a bad start, but things got better. The rest of this adventure is in <http://www.cs.rug.nl/~jurjen/iwi-howtos/linux/Osis>.

X access from under sudo

1. Merge your `~/.Xauthority` with the `~/.Xauthority` of root:

sudo xauth merge ~*user*/.Xauthority

2. Or for a more permanent solution:

- a. Modify your `~/ .bashrc` to contain the following:

```
export XAUTHORITY=~/.Xauthority
```

- b. Modify your `/etc/sudoers` to contain the following:

```
Defaults env_keep += "DISPLAY XAUTHORITY"
```

Quick source NAT with IPtables

Note

`eth0` Is the outside world, `eth1` is the private network.

1. **echo 1 > /proc/sys/net/ipv4/ip_forward**
2. **iptables -t nat -A POSTROUTING -o eth0 -j MASQUERADE**
3. **iptables -A FORWARD -i eth0 -o eth1 -m state --state RELATED,ESTABLISHED -j ACCEPT**
4. **iptables -A FORWARD -i eth1 -o eth0 -j ACCEPT**

Chapter 8. April 2008

Installing the SuSE iPrint client under Debian

1. The printserver page [<http://myprintserver.mydomain.com/ipp>] has a link to the client to install. Use it to download the Novell rpm package.

Note

The link is a scripted thing that determines your platform, so in order to download the client for platform X, you first have to install an OS and a browser on platform X. But see the next chapter for a solution.

2. The “Cool Solution [<http://www.novell.com/communities/node/4208/using-iprint-with-ubuntu>]” suggested by Novell is to use the command **alien -d --scripts iprint...rpm** to convert the package to a .debpackage. This disregards the fact that the resulting package is somewhat in conflict with the Filesystem Hierarchy Standard, in that it puts its files in /opt, which is for application software, whereas iPrint would belong in /usr.

However, we can still do it on a system that doesn't care. We also need to install the libglitz package:

```
sudo apt-get install libglitz1
```

3. Now we install the beastie:

```
dpkg -i novell-iprint-xclient-sl_6.500.20070508-2_amd64.deb
```

4. Following a successful installation, according to the above webpage, we need to issue a few commands (mainly to set environment variables):

```
./opt/novell/iprint/bin/iprint-gnome-init  
./etc/profile.d/novell-iprint.sh
```

5. And then of course we've got FireFox working, but not IceWeasel, so we need a few more commands:

```
pushd /usr/lib/iceweasel/plugins/  
sudo cp ../../firefox/plugins/npnipp* ./
```

6. Now we can restart the browser, go to the printer lists, and add printers to our hearts' content.

Note

We must of course have rights to add printers to our local machine.

Converting the Novell iPrint client to a Debian package

1. From an iPrint servers' main page, all client rpm can be downloaded. So there is no need to install platform X in order to download the client rpm for platform X.
2. We unpack the Debian package:

```
mkdir -p unpack/DEBIAN &&\  
dpkg -X novell-iprint-xclient-sl_6.500.20070508-2_amd64.deb unpack &&\  
dpkg -e novell-iprint-xclient-sl_6.500.20070508-2_amd64.deb unpack/DEBIAN
```

3. Then we enter the unpack directory, and start moving things around:
 - a. We move the log file to a proper place:

```
cd unpack  
mkdir -p var/log/iprint && rm -rf var/opt/
```

- b. We move the manpage to a proper place:

```
mkdir -p usr/share/man/man1 &&\  
mv opt/novell/man/man1/iprintcmd.1 usr/share/man/man1 &&\  
gzip usr/share/man/man1/iprintcmd.1 &&\  
rm -rf opt/novell/man/
```

- c. We move around some libraries:


```
mv opt/novell/lib/* usr/lib  
mv opt/novell/lib64/* usr/lib64  
rm -rf opt/novell/lib*
```

- d. We push around some messages:

```
mv opt/novell/iprint/share/locale usr/share/
```

- e. We push around (and convert in the process) some icons:

```
mkdir -p usr/share/icons/gnome/32x32/apps/  
convert opt/novell/iprint/share/icons/iprint_32.gif usr/share/icons/gnome/32x32/apps/iprint.p  
mkdir -p usr/share/icons/kdeclassic/32x32/apps  
convert opt/novell/iprint/share/icons/iprint_32.gif usr/share/icons/kdeclassic/32x32/apps/iprin  
  
mkdir -p usr/share/icons/gnome/16x16/apps  
convert opt/novell/iprint/share/icons/iprint_16.gif usr/share/icons/gnome/16x16/apps/iprint.p  
mkdir -p usr/share/icons/kdeclassic/16x16/apps  
convert opt/novell/iprint/share/icons/iprint_16.gif usr/share/icons/kdeclassic/16x16/apps/iprin  
  
rm -rf opt/novell/iprint/share/
```

- f. We move around a few binaries:

```
mkdir -p usr/bin  
mv opt/novell/iprint/bin/* usr/bin/  
rm -rf opt/novell/iprint/bin
```

- g. We move around some plugins:

```
mkdir -p usr/lib/iprint
mv opt/novell/iprint/plugin/* usr/lib/iprint/
rm -rf opt/novell/iprint/
```

- h. We move around some desktop links:

```
mkdir -p usr/share/applications/kde/
mv opt/kde3/share/applications/novell-iprint.desktop usr/share/applications/kde/

mkdir usr/share/autostart
mv opt/kde3/share/autostart/novell-iprint.desktop usr/share/autostart/

mkdir -p usr/share/gnome/autostart
mv opt/gnome/share/autostart/novell-iprint.desktop usr/share/gnome/autostart
```

- i. We handle some plugins:

```
mkdir -p usr/lib/mozilla/plugins/
mkdir -p usr/lib/firefox/plugins/
mkdir -p usr/lib/iceweasel/plugins/
cd usr/lib/mozilla/plugins/
rm /*
ln -s ../../iprint/npnipp* ./
cd ../../firefox/plugins/
rm -rf /*
ln -s ../../iprint/npnipp* ./
cd ../../iceweasel/plugins/
ln -s ../../iprint/npnipp* ./

rm -rf opt
```

- j. We clean up some odds and ends that are not needed any more because we put our files in default directories:

```
rm -rf etc/ld.so.conf.d
rm -rf etc/profile.d
```

- k. We leave `etc/opt/novell/iprint/libiprint.conf` in place because we suspect it has been hardcoded here and there.
- 4. We modify some files that contain SuSE-specific path info:

```
sed -i 's%/opt/novell/iprint/bin/iprint-listener%/usr/bin/iprint-listener%g; s%/opt/novell/iprint/shar
sed -i 's%/opt/novell/iprint/bin/iprint-listener%/usr/bin/iprint-listener%g; s%/opt/novell/iprint/shar
sed -i 's%/opt/novell/iprint/bin/iprint-listener%/usr/bin/iprint-listener%g' usr/bin/iprint-gnome-init

cd usr/share/autostart/
rm novell-iprint.desktop
ln -s ../applications/kde/novell-iprint.desktop ./
```

- 5. We modify the Debian control file to make the package depend on `libglitz1`:

```
sed -i '/^Depends:/ s%$%, libglitz1%g' DEBIAN/control
```

- 6. We truncate `DEBIAN/conffiles`

```
cd etc/ && ln -s opt/novell/iprint/libiprint.conf ./
echo -e "/etc/opt/novell/iprint/libiprint.conf\n/etc/libiprint.conf" > DEBIAN/conffiles
```

- 7. We update the md5 checksums:

```
find -type f -exec md5sum {} \; |grep -v DEBIAN|sed 's%\.\./%%g' > DEBIAN/md5sums
```

- 8. We create a directory `/var/log/novell/iprint/client` and modify the postinstall script to set permissions on it

```
mkdir -p /var/log/novell/iprint/client &&
sed -i 's%\var\opt\novell\log\novell\iprint%g' DEBIAN/postinst
```

- 9. Then we build a package from it:

```
mv unpack novell-iprint-xclient-sl-amd64 \&\&
dpkg-deb -b novell-iprint-xclient-sl-amd64
```

Warning

The path to `/var/opt/novell/log` seems to be hardcoded into several binaries and not configurable in `/etc/libiprint.conf`. So in `/etc/libiprint.conf`, the line “TraceLevel off” should remain. Alternatively, we could recompile from source, which we didn't do here.

Procedure 8.1. Trying out the new package

- First we remove the package we created with `alien`, and we check what remains of the package:

```
sudo apt-get remove novell-iprint-xclient-sl libglitz1
for LINE in $(find etc usr var -not -type d) ; do if [ -e "${LINE}" ] ; then echo "${LINE} still exists
```

Chapter 9. May 2008

Installing OpenBSD on a Soekris Net5501-70

1. Connecting the device

- Connect the null modem cable to the serial ports of your net5501 [<http://www.soekris.com/net5501.htm>] box and your PC.
- Connect a UTP cable to the port labeled “eth0” on the Soekris box, and to your network.

Warning

This procedure assumes you have a DHCP and a TFTP server handy.

- Do not connect the power cord yet.

2. Configuring the terminal emulator

- a. Install `Minicom` [<http://alioth.debian.org/projects/minicom/>] on the PC.

apt-get install minicom

- b. Press **CtrlA+Z** then **O** to get into the configuration menu. Now configure `Minicom` in such a way that it emulates an ANSI terminal at 19200 baud with 8-bits-one-parity and one stop bit. And make sure it connects to the serial port your cable is plugged into.
- c. Save these settings

3. Boot into the comBIOS

- a. Now connect the power cable (both ends ;) and see your box boot.

```
POST: 012345689bcefghips1234ajklmnopqr , , , tvwxy
```

```
comBIOS ver. 1.33 20070103 Copyright (C) 2000-2007 Soekris Engineering.
```

```
net5501
```

Slot	Vend	Dev	Class	Rev	Cmd	Stat	CL	LT	HT	Base1	Base2	Int
0:01:2	1022	2082	1010	0000	0006	0220	08	00	00	A0000000	00000000	10
0:06:0	1106	3053	0200	0096	0117	0210	08	40	00	0000E101	A0004000	11
0:07:0	1106	3053	0200	0096	0117	0210	08	40	00	0000E201	A0004100	05
0:08:0	1106	3053	0200	0096	0117	0210	08	40	00	0000E301	A0004200	09
0:09:0	1106	3053	0200	0096	0117	0210	08	40	00	0000E401	A0004300	12
0:14:0	104C	AC23	0604	0002	0107	0210	08	40	01	00000000	00000000	
0:20:0	1022	2090	0601	0003	0009	02A0	08	40	80	00006001	00006101	
0:20:2	1022	209A	0101	8001	0005	02A0	08	00	00	00000000	00000000	
0:21:0	1022	2094	0C03	1002	0006	0230	08	00	80	A0005000	00000000	15
0:21:1	1022	2095	0C03	2002	0006	0230	08	00	00	A0006000	00000000	15
1:00:0	100B	0020	0200	0000	0107	0290	00	40	00	0000D001	A4000000	10
1:01:0	100B	0020	0200	0000	0107	0290	00	40	00	0000D101	A4001000	07
1:02:0	100B	0020	0200	0000	0107	0290	00	40	00	0000D201	A4002000	10
1:03:0	100B	0020	0200	0000	0107	0290	00	40	00	0000D301	A4003000	07

4 Seconds to automatic boot. Press Ctrl-P for entering Monitor.

comBIOS Monitor. Press ? for help.

> ?

comBIOS Monitor Commands

```

boot [drive][:partition] INT19 Boot
reboot                          cold boot
download                         download a file using XMODEM/CRC
flashupdate                      update flash BIOS with downloaded file
time [HH:MM:SS]                 show or set time
date [YYYY/MM/DD]              show or set date
d[b|w|d] [adr]                 dump memory bytes/words/dwords
e[b|w|d] adr value [...]       enter bytes/words/dwords
i[b|w|d] port                  input from 8/16/32-bit port
o[b|w|d] port value            output to 8/16/32-bit port
run adr                          execute code at adr
cmosread [adr]                 read CMOS RAM data
cmoswrite adr byte [...]       write CMOS RAM data
cmoschecksum                    update CMOS RAM Checksum
set parameter=value            set system parameter to value
show [parameter]               show one or all system parameters
?/help                          show this help

```

- b. Set date and time:

date 2008/05/16 22:42:15

- c. Set new connection speed

set ConSpeed=57600

Warning

The minicom settings should also be adjusted (at next boot).

4. Boot from PXE

- a. Have the following DHCP snippet and restart the DHCP service:

```
group
{ # openbsd-clients
  next-server 192.168.5.200;
  filename "pxeboot_openbsd_43";
  host soekris { hardware ethernet 00:00:24:XX:XX:XX ; fixed-add
}# end group openbsd-clients
```

- b. On the TFTP server, go into your TFTP directory and download a few files¹:

```
cd /var/lib/tftpboot
wget http://osis.service.rug.nl/pub/os/bsd/openbsd/4.3/i386/pxeboot
mv pxeboot pxeboot_openbsd_43
wget http://osis.service.rug.nl/pub/os/bsd/openbsd/4.3/i386/bsd.rd
mv bsd.rd openbsd_43.rd
mkdir etc
cat <<EOF > etc/boot.conf
set tty com0
stty com0 57600
boot openbsd_43.rd
EOF
```

- c. Boot the Soekris box into PXE:

boot f0

```
> boot f0
```

```
Intel UNDI, PXE-2.0 (build 082)
Copyright (C) 1997,1998,1999 Intel Corporation
VIA Rhine III Management Adapter v2.43 (2005/12/15)
```

¹Do use your favourite mirror

```
CLIENT MAC ADDR: 00 00 24 CA 65 D4
CLIENT IP: 192.168.5.4 MASK: 255.255.255.0 DHCP IP: 192.168.5.200
GATEWAY IP: 192.168.5.251
probing: pc0 com0 com1 pxe![2.1] mem[639K 511M a20=on]
disk: hd0+*
net: mac 00:00:24:ca:65:d4, ip 192.168.5.4, server 192.168.5.200
>> OpenBSD/i386 PXEBOOT 2.02
switching console to com0
>> OpenBSD/i386 PXEBOOT 2.02
com0: changing speed to 57600 baud in 5 seconds, change your terminal to ma

com0: 57600 baud
booting tftp:openbsd_43.rd: 4780308+874136 [52+178240+163973]=0x5b821c
entry point at 0x200120
```

```
Copyright (c) 1982, 1986, 1989, 1991, 1993
The Regents of the University of California. All rights reserved.
Copyright (c) 1995-2008 OpenBSD. All rights reserved. http://www.OpenBSD.org
```

```
OpenBSD 4.3 (RAMDISK_CD) #645: Wed Mar 12 11:31:03 MDT 2008
deraadt@i386.openbsd.org:/usr/src/sys/arch/i386/compile/RAMDISK_CD
cpu0: Geode(TM) Integrated Processor by AMD PCS ("AuthenticAMD" 586-class)
cpu0: FPU,DE,PSE,TSC,MSR,CX8,SEP,PGE,CMOV,CFLUSH,MMX
real mem = 536440832 (511MB)
avail mem = 512524288 (488MB)
mainbus0 at root
bios0 at mainbus0: AT/286+ BIOS, date 20/70/03, BIOS32 rev. 0 @ 0xfac40
pcibios0 at bios0: rev 2.0 @ 0xf0000/0x10000
pcibios0: pcibios_get_intr_routing - function not supported
pcibios0: PCI IRQ Routing information unavailable.
pcibios0: PCI bus #1 is the last bus
bios0: ROM list: 0xc8000/0xa800
cpu0 at mainbus0
pci0 at mainbus0 bus 0: configuration mode 1 (no bios)
pchb0 at pci0 dev 1 function 0 "AMD Geode LX" rev 0x31
"AMD Geode LX Crypto" rev 0x00 at pci0 dev 1 function 2 not configured
vr0 at pci0 dev 6 function 0 "VIA VT6105M RhineIII" rev 0x96: irq 11, address
ukphy0 at vr0 phy 1: Generic IEEE 802.3u media interface, rev. 3: OUI 0x004
vrl at pci0 dev 7 function 0 "VIA VT6105M RhineIII" rev 0x96: irq 5, address
ukphy1 at vrl phy 1: Generic IEEE 802.3u media interface, rev. 3: OUI 0x004
vr2 at pci0 dev 8 function 0 "VIA VT6105M RhineIII" rev 0x96: irq 9, address
ukphy2 at vr2 phy 1: Generic IEEE 802.3u media interface, rev. 3: OUI 0x004
vr3 at pci0 dev 9 function 0 "VIA VT6105M RhineIII" rev 0x96: irq 12, address
ukphy3 at vr3 phy 1: Generic IEEE 802.3u media interface, rev. 3: OUI 0x004
ppb0 at pci0 dev 14 function 0 "TI PCI2250 PCI-PCI" rev 0x02
pcil at ppb0 bus 1
sis0 at pcil dev 0 function 0 "NS DP83815 10/100" rev 0x00, DP83816A: irq 10
nsphyter0 at sis0 phy 0: DP83815 10/100 PHY, rev. 1
sis1 at pcil dev 1 function 0 "NS DP83815 10/100" rev 0x00, DP83816A: irq 7
nsphyter1 at sis1 phy 0: DP83815 10/100 PHY, rev. 1
sis2 at pcil dev 2 function 0 "NS DP83815 10/100" rev 0x00, DP83816A: irq 10
nsphyter2 at sis2 phy 0: DP83815 10/100 PHY, rev. 1
sis3 at pcil dev 3 function 0 "NS DP83815 10/100" rev 0x00, DP83816A: irq 7
nsphyter3 at sis3 phy 0: DP83815 10/100 PHY, rev. 1
glxpcib0 at pci0 dev 20 function 0 "AMD CS5536 ISA" rev 0x03: rev 0, 32-bit
```



```

pciide0 at pci0 dev 20 function 2 "AMD CS5536 IDE" rev 0x01: DMA, channel 0
wd0 at pciide0 channel 0 drive 0: <SanDisk SDCFH2-004G>
wd0: 4-sector PIO, LBA, 3919MB, 8027712 sectors
wd0(pciide0:0:0): using PIO mode 4, DMA mode 2
pciide0: channel 1 ignored (disabled)
ohci0 at pci0 dev 21 function 0 "AMD CS5536 USB" rev 0x02: irq 15, version 1
ehci0 at pci0 dev 21 function 1 "AMD CS5536 USB" rev 0x02: irq 15
usb0 at ehci0: USB revision 2.0
uhub0 at usb0 "AMD EHCI root hub" rev 2.00/1.00 addr 1
isa0 at glxpcib0
isadma0 at isa0
pckbc0 at isa0 port 0x60/5
pckbd0 at pckbc0 (kbd slot)
pckbc0: using irq 1 for kbd slot
wskbd0 at pckbd0: console keyboard
npx0 at isa0 port 0xf0/16: reported by CPUID; using exception 16
pccom0 at isa0 port 0x3f8/8 irq 4: nsl6550a, 16 byte fifo
pccom0: console
pccom1 at isa0 port 0x2f8/8 irq 3: nsl6550a, 16 byte fifo
usb1 at ohci0: USB revision 1.0
uhub1 at usb1 "AMD OHCI root hub" rev 1.00/1.00 addr 1
biomask e145 netmask ffe5 ttymask ffe7
rd0: fixed, 3800 blocks
PXE boot MAC address 00:00:24:ca:65:d4, interface vr0
root on rd0a swap on rd0b dump on rd0b
erase ^?, werase ^W, kill ^U, intr ^C, status ^T
(I)nstall, (U)pgrade or (S)hell?

```

5. Install OpenBSD

It is time to follow the steps in the installation manual [<http://www.openbsd.org/faq/faq4.html#Install>].

Procedure 9.1. Configuring the OpenBSD box

1. Add a mere mortal user

```
adduser username
```

2. Add the user to the `/etc/sudoers` file

```

<snip>
# User privilege specification
root    ALL=(ALL) SETENV: ALL
username ALL=(ALL) ALL

```

<snip>

3. Configure the packaging system

Put in `~/ .profile` a stanza

```
PKGPATH=ftp://ftp.nluug.nl/pub/OpenBSD/4.3/packages/i386
export PKGPATH
```

and re-source the file:

```
./ .profile
```

4. Install some packages

```
sudo pkg_add -v syslog-ng-1.6.8
sudo pkg_add -v isc-dhcp-server-3.1.0
```

5. Configure some network interfaces

a. Edit `/etc/hostname.vr1` to create a WAN NIC:

```
echo "dhcp NONE NONE NONE" > /etc/hostname.vr1
```

b. Edit `/etc/hostname.vr2` to create a LAN for which this box will be the DHCP server:

```
net 10.0.12.1 255.255.255.0 NONE
```

c. Edit `/etc/hostname.vr3` to create another LAN (e.g. the DMZ):

```
net 10.1.154.1 255.255.255.0 NONE
```

6. Configure SSH

Make sure the SSH daemon doesn't listen on the WAN interfaces. For now, make it listen on all LAN NICs, including the config NIC (later on we can remove all but the config NIC). Add the following lines to `/etc/ssh/sshd_config`:

```
ListenAddress 10.0.12.1
```

```
ListenAddress 10.1.154.1
ListenAddress 192.168.5.4
```

7. (Compile and) Configure DHCPD

The *package* `isc-dhcpd-3.1.0` that we installed has not replaced the DHCP daemon executable `/usr/sbin/dhcpd` that was in the *file set* `base43.tgz`. Instead, a new file `/usr/local/sbin/dhcpd` was added that contains the daemon we want to use.

Follow this tutorial [http://grommit.com/~ranga/notes/openbsd_3.7_dhcpd.html] to get it running, chrooted and all.

Fun! Now we have a perfectly chrooted DHCP server, but it won't pass the `PXElinux` options to the clients, so `PXElinux` loads the default config files. That was not the idea. This is a consequence of the patches OpenBSD applied to the daemon. So we install another instance of OpenBSD (on a virtual machine, and this time *with* the compiler on it), and fetch the source of the ISC `dhcpd` daemon. This compiles without error, and we copy just the `dhcpd` binary to the router. Now this of course doesn't support opening all files as root and then dropping privileges, so we leave the chroot out for the moment. But it does support `PXElinux` all right.

8. Compile and Configure BIND9

- a. Get the BIND source and unpack it:

```
ftp http://ftp.isc.org/isc/bind9/9.4.2/bind-9.4.2.tar.gz
tar xzf bind-9.4.2.tar.gz
```

- b. Configure the installer:

```
./configure --with-libtool --with-openssl --enable-ipv6 --with-dlz-filesystem --with-dlz-stub
make
```

2

- c. Since installation is fairly complex and I don't know how to log only the copy actions, I copy the entire tree to the target machine, and run **make install there**:

On the compiling machine:

```
tar cvzf /tmp/bind-9.4.2-compiled.tgz bind-9.4.2
scp /tmp/bind-9.4.2-compiled.tgz ordinaryuser@router:/tmp
```

² The options “`--with-dlz-postgres --with-dlz-bdb --with-dlz-mysql --with-dlz-ldap`” would've been nice too, but I'm not wasting my time on options I'm not sure I'm going to use here

On the router:

```
cd ~
tar zxvf /tmp/bind-9.4.2-compiled.tgz
cd bind-9.4.2/
sudo make install
rm -rf bind-9.4.2 /tmp/bind-9.4.2-compiled.tgz
sudo find /usr/local/sbin/ -type f -group wheel -exec chown root:bin {} \;
```

Warning

It is imperative that the path on the target machine where **sudo make install** is to run is identical to the path on the build machine where **make** has run.

Chapter 10. June 2008

Installing Linux over Windows without BIOS access

This was a fun one: we have a laptop with an unknown BIOS password, and we want to return it with the password still set, so erasing the BIOS is not an option. However, we do have Administrator access to the Windows running on the laptop, so in principle, adding another OS is possible. Turns out it is...

1. Visit Mr. Herbert's page [<http://marc.herbert.free.fr/linux/win2linstall.html>] on this subject, and read it.
2. Download the latest WinGRUB from gna.org [<http://download.gna.org/grub4dos/>], and pick from it the file `grldr`. Put it in `C:\`.
- 3.

- a. Make `C:\boot.ini` visible:

```
chattr c:\boot.ini -s -h -r
```

- b. Add the following line to `c:\boot.ini`:

```
C:\grldr="Chainload GRUB"
```

- c. Make `C:\boot.ini` invisible again:

```
chattr c:\boot.ini +s +h +r
```

4. Fetch from wherever you like a kernel and RAMdisk to boot, and put them in `C:\boot`. It's nice to fetch a network installation RAMdisk. These tend to require fewer CDs in order to work.
5. Edit `C:\boot\grub\menu.lst` to contain something along the lines of:

```
title Install Linux
kernel (hd0,0)/boot/linux kernel-and-other-options
initrd (hd0,0)/boot/initrd.gz
```

6. Optionally, use Partition Magic to make some room for Linux.
7. Reboot the laptop, and when booting pick the "Chainload GRUB" option. In the GRUB menu, pick "Install Linux".
8. You are now in a regular Linux installer.

Turning nVidia driver on on machines that have the libraries and an NVidia card

The following script can be put in `/etc/init.d` and linked to from `/etc/rcS.d` in order to modify the default Debian `vesa` video driver on machines that have an NVidia card and have the NVidia drivers installed:

```
#!/bin/bash

XORG_CONF=/etc/X11/xorg.conf
NVIDIA_LIBS_INDICATOR="/usr/lib/xorg/modules/drivers/nvidia_drv.so"
NVIDIA_CARD_INDICATOR=nVidia

if [ -f ${NVIDIA_LIBS_INDICATOR} ] && lspci|grep VGA|grep nVidia > /dev/null
then
    sed -i 's%\([ \t]*Driver[ \t]\{1,\}\)%"vesa"%\1"nvidia"%' ${XORG_CONF}
fi
```

On Debian machines, the proper links in `/etc/rcn.d` can be created with the command: **sudo update-rc.d nvidia defaults**

Note

The above script doesn't ensure that the drivers are actually suitable for the card, and it certainly doesn't finetune the `xorg.conf` for use with the NVidia card.

A Firewall Install Script

Remotely turning on a firewall always carries the risk of locking yourself out. Rather than properly stealing a script, I made my own [<http://www.cs.rug.nl/~jurjen/scripts/firewall>]. It has features attractive to me: before installing a new firewall, it checks whether I can still work when the new configuration is active. And it can often be used stand-alone (with just the binaries it needs, but no additional config) on fresh installations¹. It has a `--help` option, but basic usage is:

1. Get the script and put it in `/usr/local/bin` or the like.

```
wget -P /usr/local/bin http://www.cs.rug.nl/~jurjen/scripts/firewall \
&& chmod a+x /usr/local/bin/firewall
```

2. Create the config directory for the script:

```
mkdir /etc/firewall
```

3. Usually this is not needed, but if your `PATH` is incomplete or if the script needs to run without it, create `/etc/firewall/firewall.cfg`, with paths to binaries the script needs.

¹Yet, I am aware, it is programmed rather erratically.

```
IPTABLES=/sbin/iptables
IPTABLES_SAVE=/usr/sbin/iptables-save
IPTABLES_RESTORE=/usr/sbin/iptables-restore
MD5SUM=/usr/bin/md5sum
LN=/bin/ln
MV=/bin/mv
CP=/bin/cp
RM=/bin/rm
ECHO=/bin/echo
CAT=/bin/cat
TTY=/usr/bin/tty
AWK=/usr/bin/awk
TRUE=/bin/true
FALSE=/bin/false
EGREP=/usr/bin/egrep
GREP=/usr/bin/grep
BASENAME=/usr/bin/basename
DATE=/bin/date
SLEEP=/bin/sleep
KILL=/bin/kill
TOUCH=/usr/bin/touch
WC=/usr/bin/wc
WHICH=/usr/bin/which
```

4. Run the script to see whether it works:

```
firewall status firewall is off
```

5. Install your first firewall ruleset:

- a. Configure the firewall by some other means (e.g. by hand-typing iptables commands), and save the configuration

Note

Note that saving an empty ruleset won't succeed. You actually have to configure something, *anything*

:

```
iptables -I INPUT 1 -p tcp -s 10.10.10.10 -j DROP  
firewall save
```

- b. Once there is a configuration, you can also load earlier iptables dumps:

```
firewall -t iptables-save-output-file update
```

- c. Or you can control a script that installs a firewall:

firewall -s *firewall-generating-script* update

Note

The **firewall** script also has some diagnostics. For example, to figure out whether the firewall currently running is according to the last stored configuration: **firewall analyze**.

Warning

While the script does basic checking to make sure you can still press <ENTER> after starting the firewall, it doesn't check to see if you can log out, and ssh back to the machine under scrutiny.

Fixing the NIS port

When protecting a NIS server with IPTables, the problem arises that `yppserv` doesn't always pick the same port number to serve on, and relies on the portmapper to convey its location to the clients. This is all well, but it is moderately hard to punch holes in the firewall every time the NIS server is restarted, and close them a gain afterwards.

The solution is painfully simple: `yppserv` accepts the `-p` option, which fixes the port it is listening on.

Note

As an aside, the real solution would be for SuSE to always reconfigure the firewall when the NIS server has restarted, which it does not.

Transferring the IWI printers to IPrint

According to the IWI CUPS server [<https://iwi202.iwinet.rug.nl:631/printers/>], the following are IWI printers:

- `pr1` [<http://iwi221.iwinet.rug.nl>] (Kopieerhok 3e verdieping)
- `pr2` [<http://pr2.iwinet.rug.nl>] (Kopieerruimte 4e verdieping)
- `pr3` [<http://iwi223.iwinet.rug.nl>] (Kopieerhok 5e verdieping)
- `pr-oud` [<http://iwi220.iwinet.rug.nl>] (deprecated)
- `k3` [<http://iwi224.iwinet.rug.nl>] (Bernoulliborg 3e verdieping)
- `k4` [<http://iwi211.iwinet.rug.nl>] (Bernoulliborg 4e verdieping)
- `k5` [<http://iwi213.iwinet.rug.nl>] (Bernoulliborg 5e verdieping)
- `prico101` [<http://iwi97.iwinet.rug.nl>] (deprecated)
- `prlab` [<http://iwi215.iwinet.rug.nl>] (Practicumruimte ONT/OIC)

After creating a Debian version of the Novell IPrint package (see the section called “ Converting the Novell iPrint client to a Debian package ”), and installing the package on a client, users can add printers to the CUPS config via IPrint, but only if they are members of the `lpadmin` group. In `/etc/cups/cupsd.conf`, we can add a line

```
SystemGroup staff
```

to change the group which' membership is demanded. The proper solution would be to add all appropriate users to appropriate groups, but this 'll do.

Remote Firefox acutally remote

In order to actually start a remote firefox instead of having it connect to a locally running instance:

firefox -no-remote

Chapter 11. July 2008

Installing SpaceWalk (using a remote database)

RedHat SpaceWalk [<http://www.redhat.com/spacewalk>] is the Open Source version of RedHat's satellite software. Two machines are involved in its installation:

- A Database server, running “Oracle Database 10g Enterprise Edition Release 10.2.0.3.0 - 64bit Production”
- The Spacewalk server proper, running CentOS 5.2.

Following the SpaceWalk HowToInstall [<https://fedorahosted.org/spacewalk/wiki/HowToInstall>], I did the following:

Procedure 11.1. Installing SpaceWalk

1. Add the EPEL repositories to the SpaceWalk server

```
rpm -Uvh http://download.fedora.redhat.com/pub/epel/5/i386/epel-release-5-2.noarch.rpm
```

```
Retrieving http://download.fedora.redhat.com/pub/epel/5/i386/
warning: /var/tmp/rpm-xfer.Zy9sSY: Header V3 DSA signature: N
Preparing...                               #####
1:epel-release                             #####
```

2. Add the spacewalk repository to yum's repository list

Edit `/etc/yum.repos.d/spacewalk.repo` to contain:

```
[spacewalk]
name=Spacewalk
baseurl=http://spacewalk.redhat.com/yum/rhel/5Server/$basearch/
gpgkey=http://spacewalk.redhat.com/yum/RPM-GPG-KEY-spacewalk
enabled=1
gpgcheck=1
```

3. Install the spacewalk package

```
yum install spacewalk
```

```
<snip>
Error: rhns-app conflicts with specsपो
Error: rhns-xp conflicts with specsपो
Error: Missing Dependency: oracle-instantclient-basic = 10.2.
Error: Missing Dependency: oracle-instantclient-basic is need
Error: Missing Dependency: oracle-instantclient-basic >= 10.2
```

4. Remove specsपो and install Oracle client

rpm -e specsपो

Fetch oracle-instantclient-basic-10.2.0.4-1.i386.rpm from the Oracle website [http://www.oracle.com/technology/software/tech/oci/instantclient/htdocs/linuxsoft.html] and install it:

```
rpm -iv oracle-instantclient-basic-10.2.0.4-1.i386.rpm oracle-instantcli-
ent-jdbc-10.2.0.4-1.i386.rpm oracle-instantclient-sqlplus-10.2.0.4-1.i386.rpm oracle-in-
stantclient-devel-10.2.0.4-1.i386.rpm
```

```
Preparing packages for installation...
oracle-instantclient-basic-10.2.0.4-1
oracle-instantclient-devel-10.2.0.4-1
oracle-instantclient-jdbc-10.2.0.4-1
oracle-instantclient-sqlplus-10.2.0.4-1
```

5. Try to install the Spacewalk application once more

yum install spacewalk

```
<snip>
Transaction Summary
=====
Install      227 Package(s)
Update       0 Package(s)
Remove       0 Package(s)

Total download size: 149 M
Is this ok [y/N]: y
<snip>
warning: rpmts_HdrFromFdno: Header V3 DSA signature: NOKEY, k
Importing GPG key 0x430A1C35 "Spacewalk <spacewalk-devel@redh
Is this ok [y/N]: y
<snip>
noarch 0:1.1.3.4.0-2jpp.ep1.1.el5.1
```

Complete!

6. Configuring SpaceWalk: 1st try

```
export PATH="${PATH}:/usr/lib/oracle/10.2.0.4/client/bin"
ORACLE_HOME=/usr/lib/oracle/10.2.0.4
export LD_LIBRARY_PATH=/usr/lib/oracle/10.2.0.4/client/lib
semanage fcontext -a -t textrel_shlib_t '/usr/lib/oracle/10.2.0/client/lib/*'
restorecon -R /usr/lib/oracle/10.2.0/client/lib
```

```
[root@host ~]# spacewalk-setup --disconnected
* Loading answer file: /usr/share/spacewalk/setup/defaults.co
* Setting up environment and users.
** GPG: Initializing GPG and importing RHN key.
* Setting up database.
** Database: Setting up database connection.
DB User? username
DB Password?
DB SID? sid
DB hostname? db.host.yourdomain.com
DB port [1521]? 1522
DB protocol [TCP]?
** Database: Testing database connection.
** Database: Populating database.
sh: dbhome: command not found
*** Progress: #
* Performing initial configuration.
* Activating Satellite.
** Loading Satellite Certificate.
** Verifying certificate locally.
There was a problem activating the satellite: Certificate exp
[root@host ~]# date
Thu Feb 19 12:24:04 CET 2015
```

7. Configuring SpaceWalk: 2nd try

```
ntpdate ip-of-ntp-server
```

```
28 Jul 11:48:44 ntpdate[15293]: step time server 129.125.60.2
[root@host ~]# date
Mon Jul 28 11:48:48 CEST 2008
```

```
yum -q install usermode-gtk pyOpenSSL
cp /usr/share/spacewalk/setup/defaults.conf spacewalk-install-answers-2.conf
vi !$
spacewalk-setup --disconnected --answer-file=spacewalk-install-answers-2.conf
```

```
[root@host ~]# spacewalk-setup --disconnected --answer-file=
* Loading answer file: spacewalk-install-answers-2.conf.
* Setting up environment and users.
** GPG: Initializing GPG and importing RHN key.
* Setting up database.
** Database: Setting up database connection.
** Database: Testing database connection.
** Database: Populating database.
sh: dbhome: command not found
*** Progress: #####
* Performing initial configuration.
* Activating Satellite.
** Loading Satellite Certificate.
** Verifying certificate locally.
** Activating Satellite.
* Enabling Monitoring.
* Creating SSL certificates.
Email Address [j.bokma@cs.rug.nl]?
** SSL: Generating CA certificate.
** SSL: Deploying CA certificate.
** SSL: Generating server certificate.
** SSL: Storing SSL certificates.
Use of uninitialized value in chown at /usr/bin/rhn-generate-
Use of uninitialized value in chown at /usr/bin/rhn-generate-
* Deploying configuration files.
* Update configuration in database.
* Restarting services.
Installation complete.
Visit https://host to create the satellite administrator acco
[root@host ~]#
```

8. **Configure SpaceWalk via web-interface**

setenforce 0 To temporarily turn off `selinux` and log in to your `https://host.service.domain.com` to do the rest of the configuration.

Installing CentOS unattendedly

1. **Create a kickstart file.**

a. Install CentOS manually on a machine.

b. Install the X window system¹.

```
sudo yum groupinstall "X Window System"
```

c. Install `system-config-kickstart`. Also install `xauth` if you're about to run `system-config-kickstart` remotely, 'cause without it X11 forwarding won't work.

```
sudo yum install system-config-kickstart xauth
```

d. Run `system-config-kickstart` and specify what your unattendedly installed machine should look like.

```
system-config-kickstart
```

2. **Use the kickstart file in combination with DHCP and PXE**

a. Create a PXELinux config file like this one:

```
MENU TITLE centoslinux/5.1/i386
default desktop-001

prompt 1
timeout 600

label desktop-001
kernel images/distr/centoslinux/5.1/os/i386/isolinux/vmlinuz
ipappend 2
append initrd=images/distr/centoslinux/5.1/os/i386/isolinux/initrd
```

b. Tell the DHCP server to point the client to `pxelinux.0` with the above file as config. (I won't explain here how to do that)

c. Boot the client

¹You could probably make do with a much smaller subset, but this does ensure that remote X applications can be displayed.

Remote access to Windows XP from Linux

We have a user who wants to access her Windows computer while sitting behind her Linux computer.

Procedure 11.2. Enabling remote desktop on a Windows computer and accessing it from Linux.

1. Gathering basic data

a. Figuring out the IP number and/or FQDN of the Linux computer.

Log in on the Linux computer, open a shell, and type **hostname -f**. The answer is the FQDN of the Linux computer. Write it down for later reference.

Also type **hostname -i**. This is the IP number of the Linux computer. Write it down, too

b. Figuring out the IP number and/or FQDN of the Windows computer

Log in on the Windows computer, open a shell (click Start, Run . . . , type “cmd”, click Ok), and type “ipconfig”. The IP number of the Windows computer is now listed behind “IP Address...”. Write it down.

To translate the IP number of the Windows computer to its FQDN, log in on the Linux computer, and type **host ip-number-of-windows-machine**, with “ip-number-of-windows-machine” replaced by the actual IP number you wrote down.

2. Enabling the Remote Desktop service on the Windows computer.

Start the Windows computer, log in, and do the following:

- a. Right-click the My Computer icon on the desktop, click Properties, and then click the Remote tab.
- b. Turn on Remote Desktop by selecting the check box Allow users to remotely connect to this computer.
- c. Designate users by clicking the Select Remote Users... button. (In my case, I didn't need to add any users, as the account I logged on with was already enabled.)
- d. Click Ok

3. Actually starting the Remote Desktop service

- a. Log in on the Windows computer
- b. Right-click the My Computer icon, click services
- c. (Scroll down and) right-click Terminal Services, click properties.
- d. Set Startup type to “automatic”, and click the Start button.

- e. Click Ok.
4. **Allowing Remote Desktop connection through the Windows firewall**

Log in on the Windows computer and do the following:

 - a. Click Start, Settings, Control Panel.
 - b. Doubleclick the Windows Firewall icon, then the Exceptions tab.
 - c. Select the Remote Desktop checkbox.
 - d. Now click the Edit... button, then click Change scope....
 - e. Choose the Custom list: option, and fill in the *IP number of the Linux computer*.
 - f. Click Ok thrice.
 5. **Preparing the Windows computer for remote use**

Leave the Windows computer running, but do log off.
 6. **Connecting to the Windows machine from Linux**
 - a. Have your systems administrator install the rdesktop program for you.
 - b. Log in on the Linux computer and start a shell.
 - c. Type `rdesktop -g 1024x786 ip-number-of-windows-computer`, with “ip-number-of-windows-computer” replaced by the actual IP number of the Windows computer you wrote down. The “-g 1024x786” sets the resolution rdesktop shows Windows in. You may want to fiddle with it a bit.
 - d. Log in and use your Windows PC from behind your Linux PC.

Creating a SpaceWalk Channel

Follow the Wiki [<https://fedorahosted.org/spacewalk/wiki/UploadFedoraContent>]:

Procedure 11.3. Creating a Spacewalk channel

•

```
wget https://fedorahosted.org/spacewalk/attachment/wiki/UploadFedoraContent/create_channel.py
wget https://fedorahosted.org/spacewalk/attachment/wiki/UploadFedoraContent/extended_create_ch
```

Chapter 12. August 2008

Booting Ubuntu Hardy unattendedly using preseed

Ubuntu Hardy can be installed with preseeded answers according to Ubuntu Documentation [<https://help.ubuntu.com/8.04/installation-guide/i386/preseed-contents.html#preseed-partman>] (it can also be done using Kickstart [<https://help.ubuntu.com/8.04/installation-guide/i386/automatic-install.html>]). However, where they say “To specify the keymap as a boot parameter, use `console-setup/ask_detect=false console-setup/layoutcode=us`. The layout code is an X layout name, as would be used in the `XkbLayout` option in `/etc/X11/xorg.conf`”, it is worth noting that this particular preseeding must be done *per kernel parameter*.

Note

If you don't do this preseeding per kernel parameter, you end up with an Afghani console specified in `/etc/default/console-setup`.

Note

Even if you preseed `console-setup console-setup/*` in your preseed file, these values don't end up in the output of `debconf-get-selections` on the system being installed, and you still end up with an Afghani console.

Cloning NTFS partitions at the file level (and booting them)

Introduction

As the title of this document shows, I consider myself the Apprentice. The Guru [<http://www.cs.rug.nl/~harm>]¹ has created BootLeg, system cloning software that clones systems not at the filesystem, but at the *file* level. Please take a moment to consider the effects of this subtlety. Most disk/partition cloning software requires that the sizes of restored partitions are identical to the original. This is no problem with backup/restore operations, but for deployment of OS'es it becomes inconvenient, as different hardware soon forces the administrator to keep another image for that particular hardware.

BootLeg works at a higher level, and doesn't suffer from this drawback. As it images files into tar bundles, it allows you to restore them to any size partitions, or even to restore multiple images to a single partition, as long as the partition is large enough to hold all files in the image(s)².

This method has worked perfectly for years, producing way fewer errors than unattended installs do, and it would be a good way to install thousands of more-or-less identical machines. But it has two limitations: it works only with file systems that Linux reads and writes well, and only with bootloaders that can read filesystems, as the locations on restored partitions of files needed for boot may differ wildly from the original. For these reasons, it was particularly hard to incorporate the cloning of Windows XP into the tool, as this would involve NTFS filesystems, which Linux didn't read all that well. Even cloning an NTFS partition sector-by-sector was hard, as Windows stores information about the disk geometry in the partition's boot sector, and NTLDR uses that information.

¹Credit where credit is due

²And as long as you handle differing partition numbers -if any- after imaging.

The advent of NTFS-3G [<http://www.ntfs-3g.org/>], brought writable NTFS within Linux range, and so I investigated whether it is possible to apply the BootLeg tarbundle method to clone NTFS file systems from Linux. If it is, we have an advantage over NtfsClone [<http://man.linux-ntfs.org/ntfsclone.8.html>], which suffers the same disadvantages other *filesystem* cloners do, as it assumes too many similarities [<http://www.linux-ntfs.org/doku.php?id=ntfsclone>] when used for deployment. Although `relocntfs` [<http://www.linux-ntfs.org/doku.php?id=contrib:relocntfs>] could help, it seems to be too immature and unmaintained to be put in a production environment. `NTFSResize` [<http://www.linux-ntfs.org/doku.php?id=ntfsresize>] from the NTFSprogs suite could help, but then again the NTFS-3G developers criticize the NTFSprogs developers for breaking the FS sometimes (ToDo: citation needed).

Expectations

One would expect the cloning of NTFS to work, but ACLs and the like may lose information, if I understand the docs [<http://pagesperso-orange.fr/b.andre/security.html>] correctly. And one would expect booting into Windows to fail if the C:-drive is cloned, as the NTFS bootsector [<http://www.ntfs.com/ntfs-partition-boot-sector.htm>] holds information about the disk geometry and partition size that NTLDR needs. (See also here [<http://mirror.href.com/thestarman/asm/mbr/NTFSBR.htm>] and here [<http://bootmaster.filerecovery.biz/appnote3.html>]).

The experiment

Procedure 12.1. Cloning an NTFS partition step by step

1. Install Windows on a PC.
2. Reboot the PC into some Linux that has the NTFS-3G driver (and that *doesn't* touch the Windows partition). I used the System Rescue [http://www.sysresccd.org/Main_Page] Live Distro.

3. Mount the Windows C:-drive

```
ntfs-3g /dev/hda1 /mnt/windows
```

4. Copy the contents of the C:-drive to another machine

```
tar cvjC /mnt/windows -f - | ssh user@other.machine.com "cat > /home/user/C-disk.bz2"
```

Note

This operation took an hour and twenty minutes, which is twenty-five minutes *longer* than the Windows install itself. `bzip2` is rather slowish, and the filesize obtained with it is only marginally smaller than with `gzip`.

5. Just to make sure, copy the contents of the first 16 512-byte sectors of the system as well. Windows may have stored data there that is required for booting.

```
dd if=/dev/hda1 bs=512 count=16 | ssh user@other.machine.com "cat > /home/user/C-bootsect.dd"
```

6. Take a different PC, and boot it into your Live Distro. Prepare the system for Linux install by re-partitioning (you can use `fdisk`). Make at least a primary NTFS partition (type 7) with enough space to hold the data we just copied away. Make it a different size from the one you took the image from, and make it active, too. While you're at it, make a small other partition that will hold Grub's stage2.

7. Just to make sure we're not helped inadvertently by an old Windows that was ever on that disk, wipe out a couple of thousand bytes from the beginning of the Windows partition.

```
dd if=/dev/zero of=/dev/hda1 bs=512 count=1000
```

8. Create a filesystem on the NTFS partition.

```
mkfs.ntfs -f /dev/hda1
```

9. Mount the partition

```
ntfs-3g /dev/hda1 /mnt/windows
```

10. Restore the image onto the clean filesystem.

```
ssh user@another.machine.com "cat /home/user/C-disk.bz2" | tar xjvC /mnt/windows -f -
```

11. Install Grub in the MBR of the disk.

```
mkfs.vfat /dev/sda5
mount /dev/sda5 /mnt/boot
grub-install --root-directory=/mnt/boot /dev/sda
```

Make sure Grub has a stanza that looks like:

```
#
rootnoverify (hd0,0)
chainloader +1
```

12. Reboot the machine and try to boot into Windows via GRUB.

Conclusion

This procedure works like a charm when the receiving PC is the same one that the image was created from, even if the first 16 sectors of the Windows partition are wiped out. But if the original is a real PC with an ATA disk and the intended clone is a VMWare virtual machine with a virtual SCSI disk, booting fails with the message that “\WINNT\inf\biosinfo.inf is missing or corrupt”. In our case, it must be missing, as it is nowhere to be found on the NTFS-3G-mounted partition. If we restore the 16 salvaged sectors before creating the NTFS filesystem and then unpack the image, we get a bit further, but we still run into “STOP: 0x00000007B”.

Note

Cloning NTFS at the file level is still not feasible if the NTFS partition is to be booted from.

Using DHCP-initialized PXELinux under VM-

Ware

When trying to use PXELinux under VMWare, we run into limitations of the configurability of the VMWare DHCP daemon. Simple solution is to let a local ISC DHCP daemon bind to the `vmnet` network interfaces. When doing so, it is necessary to kill the `vmnet-dhcpd`.

Note

When the VMWare services are restarted, the ISC DHCP daemon will notice that the `vmnet` interfaces are down, but fail to see them coming up again. It is necessary to restart the DHCP daemon too.

Labelling partitions during Linux unattended install

Introduction

Some unattended Linux installers support installing in the free space on a disk. When using such installers, it is convenient to be able to remove old versions of the same installation prior to the partitioning stage. This can be done by labelling the filesystems on the partitions we create, and later remove partitions by label.

Warning

There is of course always the risk of a user naming their data partition with the same label we are using. And there is the chance of a partition losing its filesystem or filesystem label, thus getting stuck on the disk. The first is just stupidity against which there is no cure. The second will require manual intervention.

GNU `parted`, `findfs`, `mkswap` and `e2label` are not all available in the installers under scrutiny... But perhaps we can tag the filesystems with files...

Debian/Ubuntu

.

Reverse Engineering the ERD of an Oracle database

To obtain the ERD of an Oracle 10g database from the database itself, we use `Toad Data Modeler` from `ToadSoft` [http://www.toadsoft.com/toaddm/toad_data_modeler.htm]. This program runs under Windows and must be payed for, but there's no Linux tool that I know of that does the trick, so here goes:

1. Install Windows on a scrap box or a virtual machine.
2. Fill in anything they want to know on `ToadSoft`'s registration pages, download the program and in-

stall it.

3. Click File->Reverse Engineering and fill in the account data that lead to your Oracle database.
4. Click Next until satisfied.

Note

Although this got me an ERD, the layout of the thing is hideous and of no use (many tables overlapping, and autolayout doesn't help at all). The Quest employee following up on my download who called the other day advised to use one of the Computer Associates tools, but its trial version offers only 35 or so entity ERDs, from which I won't be able to judge whether the tool can handle the 300 or so SpaceWalk produces. I won't go through the hassle of trying that out.

Using WebDav to connect to the so-called Y:-drive

Under the KDE, Konqueror supports the WebDav protocol. Filling in the URL "webdav://p1234567@netstorage.id.rug.nl:80/oneNet/NetStorage/" suffices to connect to the server.

It seems that neither `davfs2` (**mount.davfs**) nor `cadaver` properly handle the NetStorage server. Sniffing reveals that they never do the PROPFIND Konqueror does. I won't investigate further now.

XML versus web template engines

When generating webpages from data, there are a host of techniques available to the clueless sysadmin. XML, XSD, XSLT and XPATH (well explained at [zvon.org](http://www.zvon.org) [<http://www.zvon.org>] and [xmlfiles.com](http://www.xmlfiles.com) [<http://www.xmlfiles.com>]) provide a generic way to store data, and transform it to whatever format desired.

Advantages

- fairly well defined
- human readable (well, with some effort)
- platform-independent
- Separating data from presentation possible. (Data in XML, structure in XSD, presentation in XSLT. However, XSLT not completely decoupled from XSD.)
- Integrity checking possible.

Disadvantages

- XSLT has ugly and cumbersome constructs for conditional and loop evaluation.

- It has the power of a programming language, but it neither looks like imperative nor logical programming.
- It isn't obvious where to steal XSLT for generating web forms and the like. Perhaps has to be written from scratch.
-

But XML c.s. have a rather steep learning curve, and if you just want to fill in a few values in your HTML pages, template engines [http://www.whenpenguinsattack.com/2006/07/19/php-template-engine-roundup] like e.g. django [http://www.djangoproject.com], smarty [http://www.smarty.net] or cheetah [http://www.cheetahtemplate.org] may be a quicker solution.

Advantages

- fairly simple to learn, looks like simplified programming,
- templates to generate CRUD webpages for databases readily available,
- extensible by writing native-language plugins.

Disadvantages

- Usually depends upon particular programming language for processing, and will generate output in same language.
- functionality tends to differ between versions,
- Integrity checks probably depend on compiler.
- Database-as-persistent-storage model with data definition in programming language not model of my choice.

Installing 64-bit Matlab on Linux

Note

The Installation Manual [http://www.mathworks.com/access/helpdesk/help/base/install/index.html] at the MathWorks site is sound. Follow the preparatory steps below, then follow that link.

Procedure 12.2. The installation of Matlab on Linux step by step

1. Preparations

- a. Locate the installation media. See the section called “ Using NCPMOUNT to access central storage ” for how to mount a Windows server. The .iso-images you're looking for are prob-

ably under `/mountpoint/ISO/RUG/APS/`.

b.
Note

This step is only needed if the installation images are not mounted on the installing machine. Do make sure the receiving machine has space for the medium.

Send the CD/DVD of your choice to the machine that is going to do the installation.

```
scp MATHWORKS_R2008A_LIN_MAC_SUN.iso user@machine.domain.com:/tmp
```

c.
Note

This step is only needed if the installed files are going to be on an NFS share.

Mount the target share `read-write` before starting the install:

```
ssh user@machine.domain.com  
sudo mount -o remount,rw /opt
```

2. **Logging in**

```
ssh user@machine.domain.com
```

3. **Mounting the installation medium**

```
sudo mkdir /mnt/loop  
sudo mount -o loop,ro /tmp/MATHWORKS_R2008A_LIN_MAC_SUN.iso /mnt/loop  
cd /mnt/loop
```

4. Follow the Installation Manual [<http://www.mathworks.com/access/helpdesk/help/base/install/index.html>] from Mathworks. It's much better than this one.

Installing 64-bit Maple 11 under Linux

Note

The preliminary steps for this installation are identical to those described under Step 1 of the section called “Installing 64-bit Matlab on Linux”.

1. **Logging in**

```
ssh user@machine.domain.com
```

2. **Mounting the installation medium**

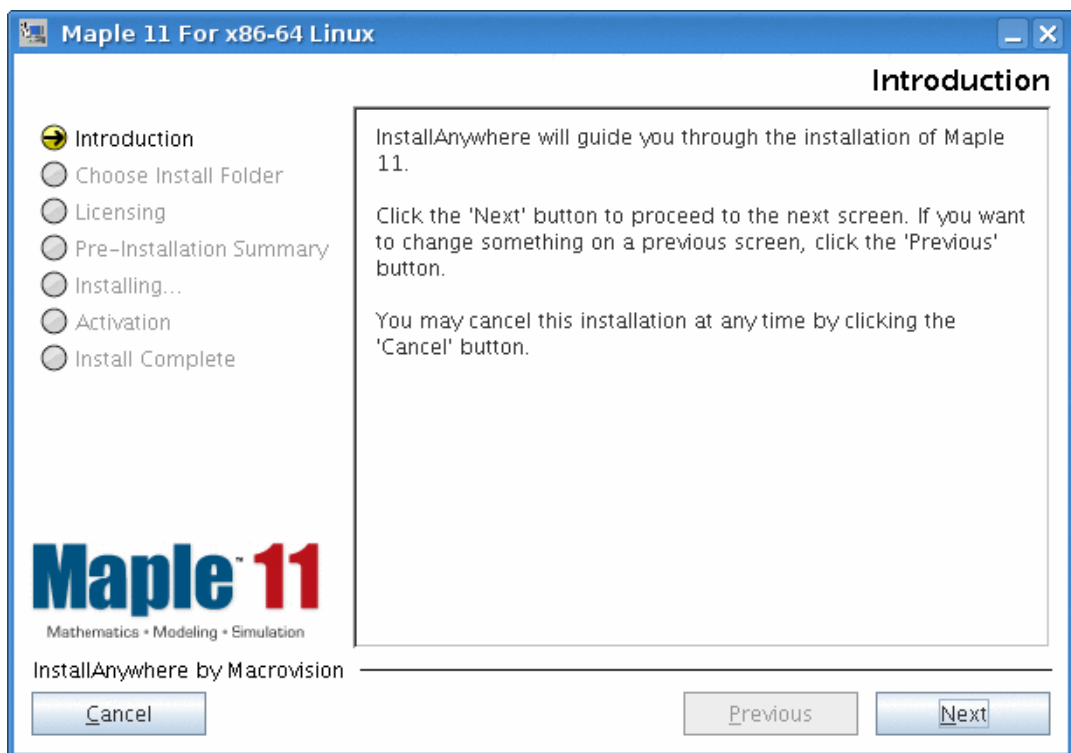
```
sudo mkdir /mnt/loop  
sudo mount -o loop,ro /home/jurjen/MAPLE11-LINUX.iso /mnt/loop/  
cd /mnt/loop
```

3. **Starting the installer**

```
sudo sh ./installMapleLinux64
```

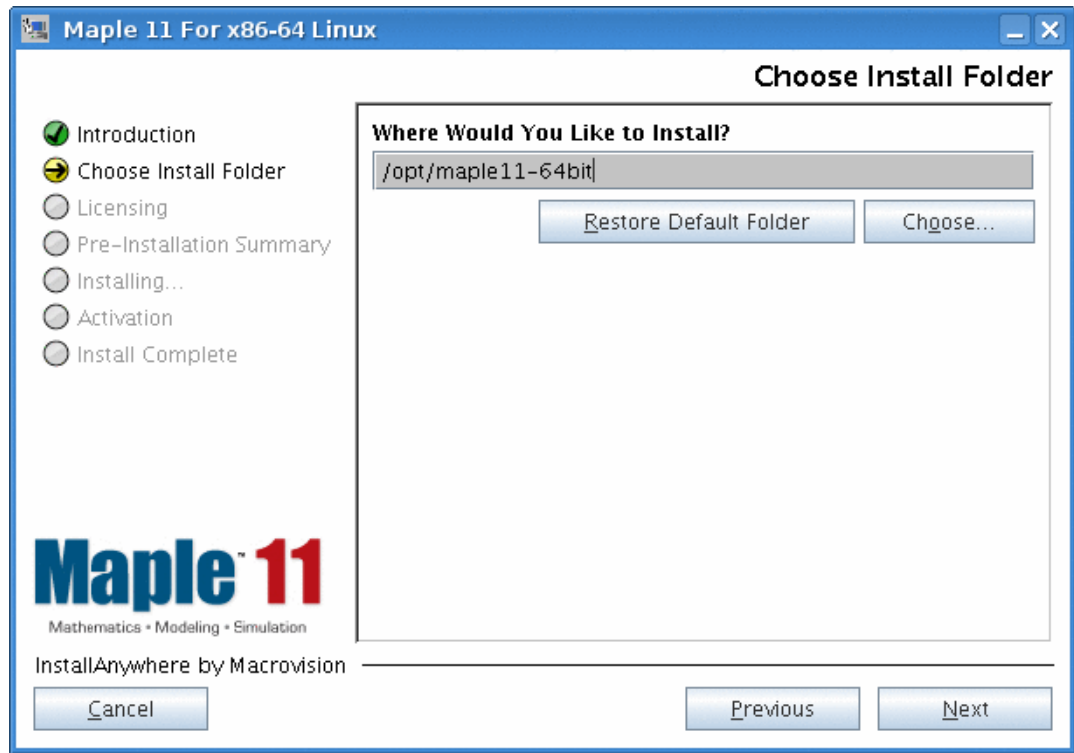
4. Click Next

Figure 12.1. Screenshot of Maple Install: Introduction



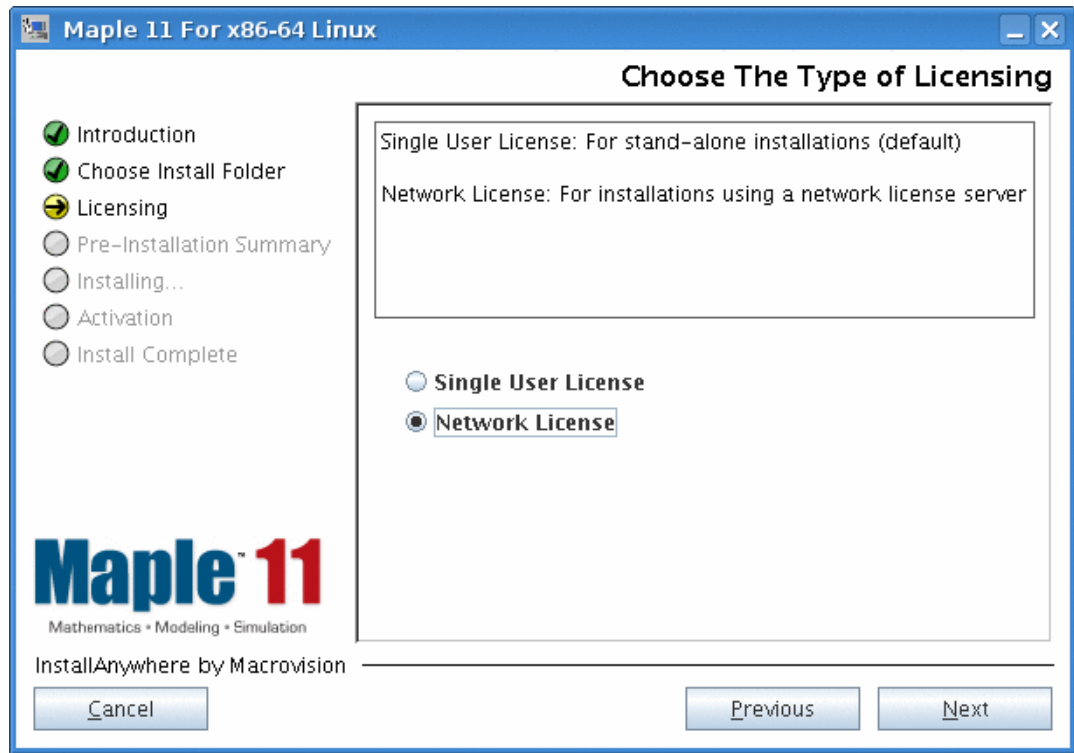
5. Type the path to the directory where you want Maple installed and click Next.

Figure 12.2. Screenshot of Maple Install: Choose Install Folder



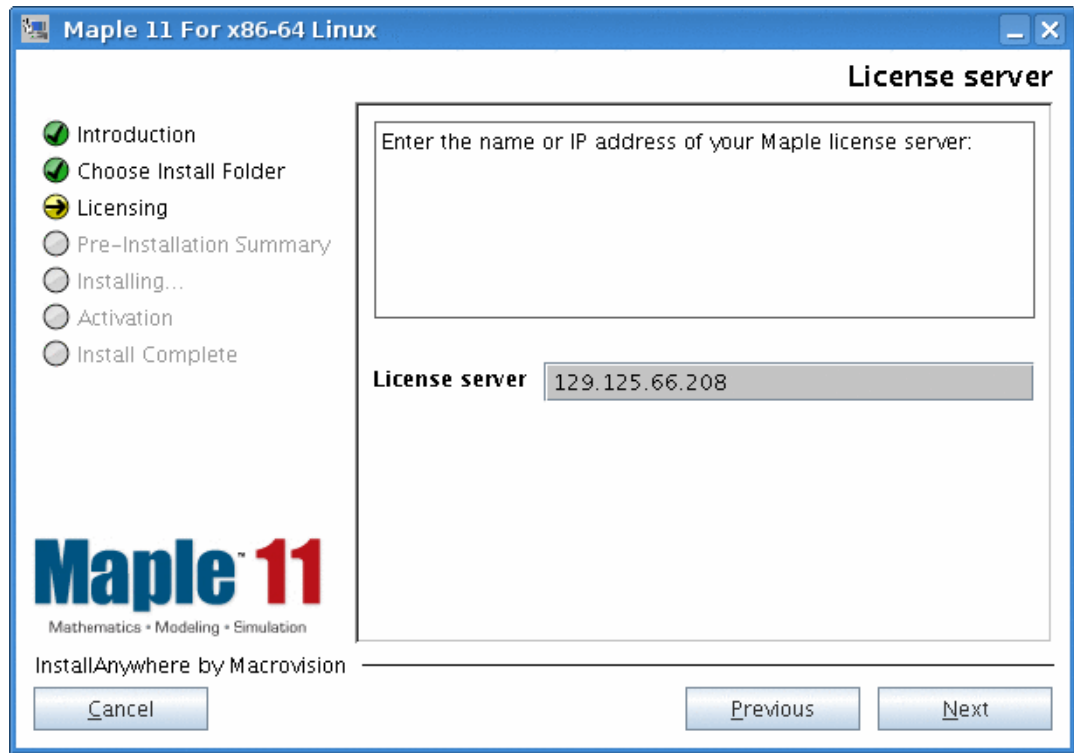
6. Choose "Network License" and click Next

Figure 12.3. Screenshot of Maple Install: Choose Type of Licensing



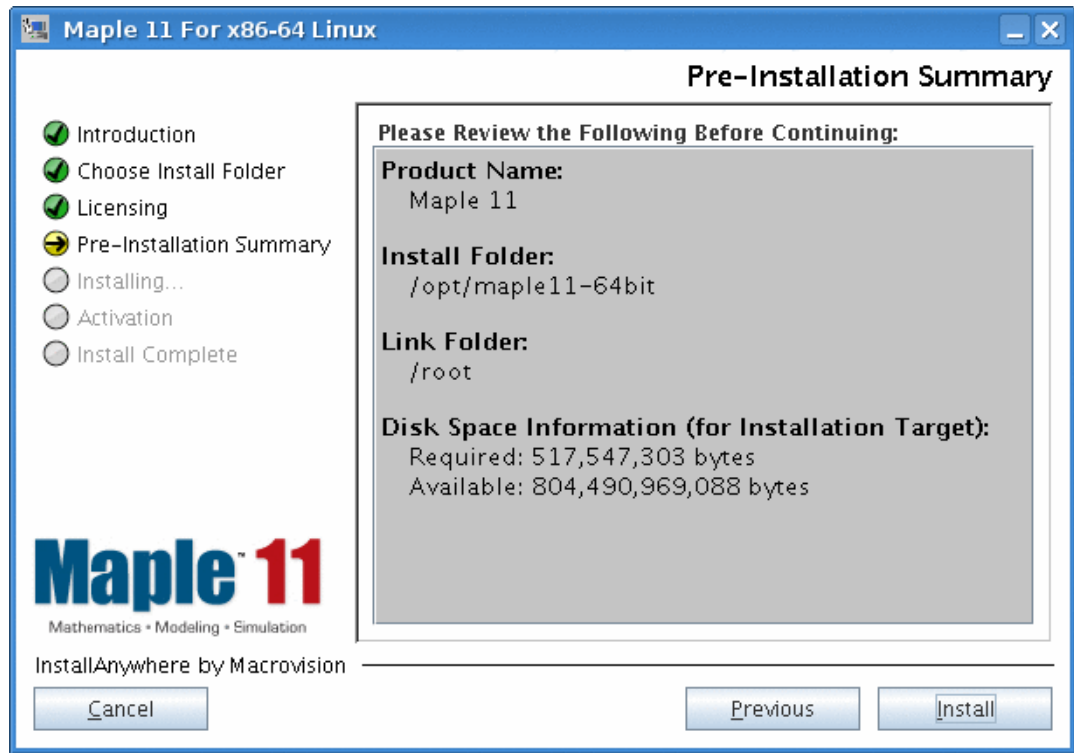
7. Type the address of the Maple License Server and click Next.

Figure 12.4. Screenshot of Maple Install: Pointing out the License Server



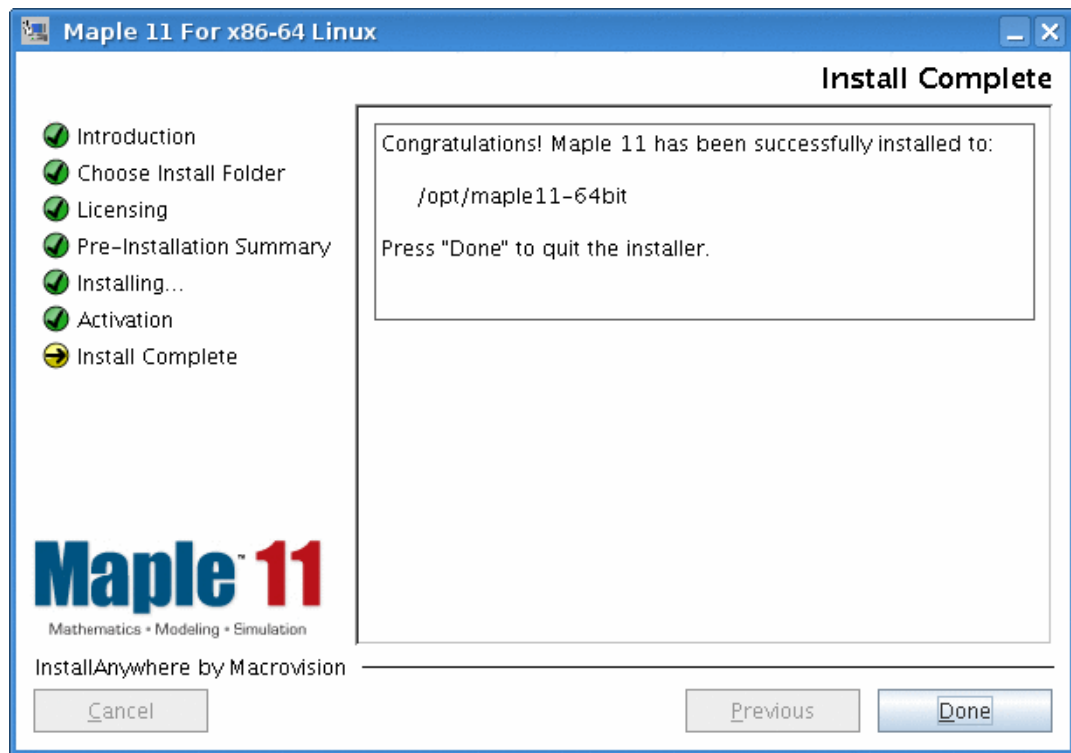
8. Overview the install data and click Next

Figure 12.5. Screenshot of Maple Install: Pre-Installation Summary



9. Click Done to get rid of the “Installation Complete” message.

Figure 12.6. Screenshot of Maple Install: Installation Complete



10. Creating links to the executables

```
sudo su -
cd /opt/local/bin
/opt/local/bin# ln -s /opt/maple11-64bit/bin/maple ./maple11.64bit
/opt/local/bin# ln -s /opt/maple11-64bit/bin/xmaple ./xmaple11.64bit
/opt/local/bin# ln -s /opt/maple11-64bit/bin/mint ./mint11.64bit
```

11. Starting Maple

```
/opt/local/bin/xmaple11.64bit
```

Note

You may want to set your path, so you can just type **Mathematica.64bits**

Installing 64-bits Mathematica on Linux

Note

The preliminary steps for this installation are identical to those described under Step 1 of the section called “Installing 64-bit Matlab on Linux”.

Procedure 12.3. The installation of Mathematica on 64-bits Linux step-by-step

1. **Logging in**

```
ssh user@machine.domain.com
```

2. **Mounting the installation medium**

```
sudo mkdir /mnt/loop
sudo mount -o loop,ro /home/jurjen/Mathematica\ -\ Unix.iso /mnt/loop
cd /mnt/loop
```

3. **Starting the installer**

```
cd /mnt/loop/Unix/Installer
sudo sh MathInstaller
```

4. **Picking the installation platform**

```
-----
Mathematica 5.2 Installer
-----
```

```
Copyright (c) 2005 Wolfram Research, Inc. All rights reserved.
```

```
WARNING: Mathematica is protected by copyright law and international
prosecuted to the maximum extent possible under law.
```

```
For which of the following platforms would you like to install
```

- (1) Linux x86 (32 and 64 bit)
- (2) IBM AIX Power (64 bit)
- (3) HP Tru64 Unix
- (4) HP-UX PA-RISC (64 bit)
- (5) SGI IRIX (64 bit)
- (6) Linux IA-64
- (7) Sun Solaris (x86 64 bit)
- (8) Sun Solaris UltraSPARC

```
Type your selection (multiple choices can be separated with space)
>> 1
```

5. Choosing the completeness of the installation

```
The following installation methods are available:
```

- (1) Full
- (2) Minimal

```
Type your selection, or press ENTER to select (1):
> 1
```

6. Choosing the installation directory

```
Enter the installation directory, or press ENTER to select /usr/local
> /opt/Mathematica-Rr2008a-64bit
```

```
Create directory (y/n)?
> y
```

```
Now installing...
```

```
[*****]
```

7. Choosing the directory for scripts

```
Type the directory path in which the Mathematica scripts will be installed
> /opt/local/bin
```

8. renaming the scripts

```
The scripts 'MathKernel', 'Mathematica', 'math', 'mathematica', 'mathkernel',
```

- (1) Overwrite

(2) Rename

```
Type your selection, or press ENTER to select (1):  
> 2
```

9. Giving an extension for the old scripts

```
Type a file name extension to be appended to the name of the  
> .32bits
```

10. renaming all scripts to their intended names

```
sudo su -  
pushd /opt/local/bin  
for i in MathKernel Mathematica math mathematica ; do mv $i $i.64bits && mv $i.32bits $i ; done
```

11. Setting the license server

```
sudo su -  
echo \!129.125.66.208 > /opt/Mathematica-Rr2008a-64bit/Configuration/Licensing/mathpass
```

12. Starting Mathematica

```
/opt/local/bin/Mathematica.64bits
```

Note

You may want to set your path, so you can just type **Mathematica.64bits**

Note

In case you get a message about missing fonts, you may want to take a look at the section called “Installing Mathematica 6.0 under Linux”.

Unable to mount USB disk under Debian

Somebody after inserting a USB stick into their PC gets the pop-up message: A security policy in place prevents this sender from sending this message to this recipient, see message bus configuration file (rejected message had interface "org.freedesktop.Hal.Device.Volume" member "Mount" error name "(unset)" destination "org .freedesktop.Hal").

Seeing that in `/etc/group` there was already a user `usb` listed, and guessing that USB disks would probably get that GID from `udev`, I added a section to `/etc/dbus-1/system.d/hal.conf` reading:

```
<!-- You can change this to a more suitable user, or make per-group
  <policy group="usb">
    <allow send_interface="org.freedesktop.Hal.Device.Volume"/>
  </policy>
```

The I added to `/etc/security/group.conf` the following:

```
mount; :0 ;*;*;usb
```

Part III. Appendices

Table of Contents

A. Indices 115

Appendix A. Indices

Index

A

- alternatives
 - /etc/alternatives, 10
- Amavis
 - enabling in PostFix, 20
- Apache server
 - subversion repository, 6
- Azzurri
 - Clay, 5

B

- backports
 - problems, 10
- backup
 - Bacula, 64
 - Tivoli, 11
- bad font path element, 52
- bad sector
 - in a PostgreSQL file, 20
 - with a file on it, 20
- bad sectors, 21
- badblocks, 20
- BIND
 - compile from source, 65
- BIOS settings
 - for booting at the IWI, 35
- bisonc++, 63
- boot parameters, 36
- booting
 - from PXE, 15
- Bootleg, 15
- bootleg
 - drawing a bundle, 17
- bundle
 - in Bootleg, 17

C

- CentOS
 - unattended install, 92
- Certification
 - Red Hat, 34
- CfEngine, 25
 - client configuration, 27
 - server configuration, 25
- Circumventing BIOS, 83
- ClamAv, 19
 - configuration, 19
 - enabling remote access, 19
- Clay

- Azzurri, 5
- cloning, 61
 - NTFS, 96
- configuration, 58
- console logging
 - setting verbosity, 35
- CPU choice, 40
- CPU usage, 64
- CUPS, 60
 - client-only, 33

D

- dangling symlinks, 10
- debian
 - packaging, 70
- Debian desktop
 - for the IWI, 10
- Debian packaging system
 - messing with -, 65
- debmirror, 65
- DHCP, 53
 - server setup, 17
 - VMWare
 - PXE, 99
- dmesg, 35
- DocBook, 4
- docbook
 - editors, 4
- dual boot
 - circumventing BIOS, 83

E

- Eclipse, 5
 - installation, 5
- emacs, 4
 - global key bindings, 36
- ERD
 - Oracle, 99
- ERDs
 - PostgreSQL, 5

F

- Fibrechannel, 65
- firefox
 - plugins, 69
 - remote, 87
- firewalling
 - firewall management script, 84
 - fwbuilder, 32
 - iptables, 84
 - portmapper, 86
- flex, 63
- font problems
 - Mathematica, 52
- fwbuilder, 32

H

- harddisk
 - failure, 21
- hardware testing
 - for BootLeg bundle compliance, 35
- headless, 75
- HTML redirect, 9

I

- iPrint, 70, 86
 - under Debian, 69
 - under Linux, 8
- IPtables
 - truning off, 32
- iptables, 84
- iserv, 26
- IWI
 - Debian desktop, 10

J

- java
 - java-package, 5
 - Sun Java, 5

K

- kernel
 - ringbuffer, 23
- kernel parameters, 36
- kernel ringbuffer, 35
- keyboards, 60
- kickstart, 92
- knowledge bases
 - OWL, 36

L

- logging to console
 - setting verbosity, 35
- lpr, 60
- LWP, 4

M

- mail scanner
 - setup, 18
- mail server
 - mail scanner
 - setup, 18
 - migrating users from a mailserv, 41
 - stopping a mailserv, 41
- Maple
 - installation, 102
- Mathematica
 - font problem, 52
 - installation, 108
- Matlab
 - installation, 101
- minicom, 75
- mirroring, 65

multi-user mode
in Knoppix, 23

N

nameserving
authoritative, 52
reverse lookup, 52

NAT
source -, 68

ncpfs, 7

ncpmount, 7, 39

Netbeheer, 53

NetFilter
turning off, 32

Netware Volumes
mounting under Linux, 7

nis, 54
fixing port, 86

Nokia, 51

NTFS
cloning, 96

NTFS-3G, 97

NTP
client
prying loose from hwclock, 33

null modem cable, 75

NVidia
detecting card presence, 84
picking driver, 84

nXML-mode package, 4

O

OpenBSD, 75
basic configuring, 79
compiling from source, 81
packaging system, 80

OpenSSL
compile from source, 65

Oracle, 88
ERD, 99

OWL
toolchain candidates, 36
Web Ontology Language, 36

P

packaging
debian, 70

parser generator, 63

partitioning
unattended, 10

plug-and-play
udev, 50

portmapper
firewalling, 86

PostFix
enabling an external virusscanner, 20

PostgreSQL

- ERDs, 5
 - with bad sector in database file, 20
- PostgreSQL, 64
- preseed, 96
- printing
 - CUPS
 - client-only, 33
- PXE
 - booting, 15

R

- RAMdisk
 - initRAMdisk, 16
 - with SystemImager, 10
- Red Hat
 - Certification, 34
- ReiserFS
 - badblocks, 20
 - do not use, 22
- reiserfsck
 - failure, 21
- repository
 - of rpms, 31
 - yum, 30
- ringbuffer, 23, 35
- rpm, 11
 - repositories to stela from, 31
 - rpm2cpio, 11

S

- SAN, 65
- Satellite, 88
- screen rotation
 - X
 - NVidia, 32
- scsi
 - kernel errors, 23
- slapd, 64
- Soekris
 - Net5501-70, 75
- sound, 60
- source NAT, 68
- Spacewalk, 88
 - channel, 94
- spamassassin
 - startup options, 19
- spampd
 - startup options, 19
- ssh, 54
 - tunnelling, 64
 - X11 forwarding, 9
- subversion
 - creating a repository, 5
 - server, 5
- subversion repository
 - Apache server, 6
- sudo
 - X access, 67

- symlinks
 - dangling, 10
- sys, 51
- Syslog-NG
 - server, 24
- syslog-ng
 - client configuration, 24
- system-config-kickstart, 92
- SystemImager, 8
 - Golden Client, 8
 - init RAMdisk, 10
 - server, 8

T

- template
 - version control repository, 5
- Template Engines, 100
- terminal emulator
 - minicom, 75
- Tivoli
 - backup, 11
- Treacherous Computing, 41
- Trusted Computing, 41
- tunnelling
 - ssh, 64

U

- Ubuntu
 - unattended install, 96
- udev
 - reload_rules, 52
 - SUBSYSTEM, 51
 - SUBSYSTEMS, 51
 - udevcontrol, 52
 - USB, 50
- udevinfo, 51
- unattended install
 - CentOS, 92
 - Ubuntu
 - preseed, 96
- update-alternatives, 10, 10
- USB
 - udev, 50

V

- vacation, 31
- version control, 5
 - repository template, 5
 - subversion, 5
- VMWare
 - DHCP
 - PXE, 99

W

- WebDAV, 100
- WebPlatform, 4
- wget, 4

Windows XP
remote desktop from Linux, 93

X

X
screen rotation
NVIDIA, 32

X11
forwarding, 9

XML, 100
schema, 5
validation, 5

Xopus, 4

xset, 52

xsltproc, 4

Y

yum
repository, 30

Z

ZenWorks, 61