



This Month's Meeting – Connecting PCs and Linux

Linux networking is not limited to TCP/IP and the common services you find on most UNIX boxes. It also provides lots of client- and server-side support for interconnectivity with PC networks, such as Netware and SMB-based networks (as used by MS LAN Manager, Windows for Workgroups/95/NT). This month, we'll take a look at Linux's support for these, and how Linux can play a role in a predominantly PC network environment.

We'll try to show as much as possible given the limited time, including connecting Windows to Linux using Samba, connecting Linux to a NetWare server as a client using ncdfs, and using Linux as a client and a server at the same time. If we have time, we'll look into intranet services with this configuration.

Please note our meeting location. This month, we'll be meeting at IBM Canada's offices in the TD Centre, at the corner of Portage and Main. We'll gather in the lobby on the main floor – please try to be there by about 7:15 PM. Steve Moffat will then take us up to the meeting room just before the meeting starts at 7:30. Don't be late, or you may not get in.

Parking is available either in the parkade behind the TD building, off Albert Street, or in the ground level lot just north of the TD building. Entrance to the lot is from Albert Street, behind the parkade. Either way, parking is a \$1.25 flat rate for the evening. You purchase your ticket from a dispenser, so make sure you've got exact change – a loonie and a quarter, or 5 quarters.

Next Meeting

The Manitoba Unix User Group will take a summer break, resuming regularly scheduled meetings as of September 9, 1997. Check your mailbox in August for details on the upcoming MUUG year!

Linux Installation Project

By Kendall G. Clark, kclark@dal284.computek.net

It all started with a simple question: 'Why don't we install Linux at all our meetings instead of at only some of them?' The North Texas Linux Users Group had been in existence for only about five months, and we wanted to make sure to spread the word in North Texas about Linux. We wanted to educate the computing public in our area about Linux, but we also wanted to let experienced computer users know that Linux could handle anything they threw at it.

After meeting at Texas Christian University for our first few meetings, we signed a contract with the DFWXchange that enabled NTLUG to meet at the Dallas Infomart. The DFWXchange is an umbrella organization that allows Dallas-Fort Worth users groups to meet at the Infomart for free, with all costs being absorbed by the many commercial vendors who also meet at the Infomart during the Super Saturday Sale. So every month between 3,000 and 5,000 computer users from the Dallas-Fort Worth Metroplex converge on the Infomart—the premier meeting facility in the Southwest devoted exclusively to computer and technology events and organizations—to participate in users groups meetings and to take advantage of some really good prices on computer-related hardware. It's a big party.

It didn't take long for the NTLUG leadership to realize that we had stumbled upon a great opportunity: we wanted to let computer users in our area know about Linux, and we were now meeting every month in a facility filled with thousands of potential Linux converts. Our solution was to start the Linux Installation Project, which we call the LIP.

The goal of LIP is simply to install Linux on as many computers as possible. Those of us who participate in this project month-to-month have discovered that the very best way to advocate the use of Linux is to make it easy and painless for the unconverted to do just that: namely, run Linux on their computer of choice. In other words, Linux is its own best advocate. After a few weeks without a crash, most people say goodbye to Windows 95 with zealous enthusiasm. We like to think of LIP as an ongoing Linux Installation Festival that allows us to convert computer users to Linux one at a time.

The first step in establishing LIP as a well-run, consistent endeavor was to find someone to lead the effort. NTLUG is fortunate to have a technologically advanced membership, and it was fairly easy to find someone to lead the LIP; in fact, we found two such people: Mike Dunn and Bill Petersen, both of whom are experienced Unix and Linux Sysadmins. Under their guidance, and due to the generosity of NTLUG members, NTLUG's LIP has solicited and organized enough computer hardware to perform up to many simultaneous installations of Linux by all the usual methods, although we've found that cdrom installations are usually the most trouble free.

The word has now spread in and around the Dallas-Fort Worth Metroplex—from schools and universities to computer vendors and other users groups—that NTLUG's LIP is the place to go for a painless installation of Linux onto PCs, laptops, servers, and even Alpha platforms. We have expanded our efforts at the LIP booth to include Linux advo-

cacy, advertisement for Linux vendors who supply us with materials, the sale of Linux CDs (thanks to Bradley Glonka at Linux Systems Labs), and even basic Linux system administration and maintenance. We also spend a lot of time explaining to the uninitiated masses what makes Linux free and what makes it so much fun.

While we have been happy with the results so far, the LIP has more work to do. We want to expand our sales efforts to include other kinds of Linux merchandise (the sales of which go to support NTLUG and LIP), and we'd also like to expand our hardware assets to enable more simultaneous installations. Finally, we also want to develop our users group assets to such an extent that we can go to other DFW-area computer events and setup Linux installation and advocacy booths. NTLUG's approach to the Linux Installation Project can be summed up in the phrase: "Linux is free. Life is good."

If you want to learn more about the North Texas Linux Users Group or our Linux Installation Project, or if you're a Linux Users Group and would like to talk about setting up your own local version of LIP, please visit the NTLUG website (<http://www.ntlug.org>) or contact me at kclark@computeek.net.

Finally, I would be guilty of ingratitude if I did not thank the following people and organizations that have made the LIP possible. Please forgive me if I've forgotten anyone. It's just about impossible not to meet great people when you work with Linux.

Copyright 1997, Kendall G. Clark
Published in Issue 17 of the Linux Gazette, May 1997

More 2 Cent Tips & Tricks

X Limitation to 8 Bit Color

From: Gary Masters
<gmasters@devcg.denver.co.us>

I read your question in Linux Gazette regarding an X limitation to 8 bit color when the system has more than 14 megs of RAM. Where did you find that information? I ask because my system has 24 megs of RAM, and I run 16 bit color all the time. One difference between our systems is that I am using a Diamond Stealth 64 video card.

The place I tell X to run in 16 bit mode is in the file /usr/X11R6/bin/startx. There is a line in this file that begins with serverargs. I get 16 bit mode by giving "-bpp 16" as an argument in this line (e.g. serverargs="-bpp 16").

One problem I did have was that the OpenLook Window Manager (olwm) did not like 16 bpp mode. I solved this by switching to the OpenLook Virtual Window Manager (olvwm). I also had success using the Tab and FV Window Managers (twm & fvwm) in 16 bpp mode. Coming from a SunOS background, I'm used to OpenLook.

Screen Blanking Under X

From: Gary Masters
<gmasters@devcg.denver.co.us>

I read your question in the Linux Gazette regarding unwanted screen blanking under X after upgrading to a newer distribution of Linux. I had the same frustration. Apparently the X servers included in the Xfree86 version distributed with current Linux distributions has screen blanking compiled as a default behavior.

This behavior can be controlled with the -s option to the server. Look in the startx script for a line that begins

with serverargs and add "-s 0". This will disable the X screen blank.

Doubleclick Internet User Profiles

From: Krage Javier Sittler
<kragen@pobox.com>

Check out the description of what doubleclick.net does at <http://www.doubleclick.net/frames/adinfo/dartset.htm>. Then decide whether you want to be added to their database of Internet user profiles. If not, you can use the script below; I run it in my /etc/rc.d/rc.inet1. It prevents any DoubleClick banners from being displayed, prevents any cookies from being set, and prevents DoubleClick from collecting any data on you.

It also does the same thing with linkexchange.com, because I find their constant banners and requests for cookies annoying. If you'd prefer, you can take out the linkexchange lines. However, this will also keep you from receiving *any* information from doubleclick or linkexchange directly... so you can't visit their web sites either.

On my machine, I put the script in /etc/rc.d/rc.doubleclick and run it from /etc/rc.d/rc.inet1 at boot time, so I'm always protected from DoubleClick.

```
# Script begins below:
#!/bin/sh # By Krage Sitaker, 21 April1997
# Prevent any packets from reaching
doubleclick.net
/sbin/route add -net 199.95.207.0 netmask
255.255.255.0 lo
/sbin/route add -net 199.95.208.0 netmask
255.255.255.0 lo
# And ad.linkexchange.com too!
/sbin/route add -net 204.71.189.0 netmask
255.255.255.0 lo
```

How to Mount/Unmount For Users

From: Kidong Lee
<kidong@shinbiro.com>

When I mount/unmount file, I have to login as root. It's not convenient for me & other users. but, I found the solution that user who is not root can do mount/

umount in mount man page.

Take a look at /etc/fstab.

```
# <device><mountpoint><filesystemtype>
<options> <dump><fsckorder>
/dev/hdb /cdrom iso9660
ro,user 0 0
```

Note "user" in options field. In options field, if you add "user", users can do mount/umount.

Using ftp Commands in Shellscript

From: Walter Harms,
<Walter.Harms@Informatik.Uni-Oldenburg.DE>

Working on several different networks means that you always need to copy your data from net to net. Most ppl use rcp but like most SysOps I found this to be a terrible security hole. So as I started this job my first business was to rewrite several scripts that were using rsh,rcp etc. I replaced them with an ftp based script ftp - <input> out 2> out.err. It's easy to see that this was not a good idea because ftp was not intended as shell-commando like cp,mv and the other guys. So I was happy to find the ftplib on a linux-CD. It's a nice lib that I used to build cmds like ftpmv, ftpcp, ftpm.. This made my scripts much slimmer and simpler. I have some terrible copy-scripts running but no problems copying on different systems like Ultrix or AIX.

Example using ftpget (from the ftplib
Author Thomas Pfau)

```
ftpget sunsite.unc.edu -r/pub/Linux
ls-lR.gz
```

This command reads the file /pub/Linux/ls-lR.gz from sunsite.unc.edu. Likewise there are other commands with the lib: ftpdir ,ftpsend, ftpm

Who needs ftplib?

Everybody tired of typing ftp... every evening to get the latest patches or whatever. Everyone who is regularly copying with ftp the same Datafiles. Why use ftplib?

Of course you can add it to you own application but more experienced users don't have to use these r-commands anymore. An ftpd is available for the majority of systems so it is easier to access more of them.

Any drawbacks?

Of course, for any ftp session you need a user/paswdr. I copy into public area using anonymous/email@ others will need to surly a password at login, what is not very useful for regular jobs or you have to use some kind of public login but still I think it's easier and better to use than the r-cmds.

Including Graphics in Linuxdoc SGML

From: Martin Michlmayr
tbm@cyrius.com

You can already include PostScript images in Linuxdoc-SGML which will get included in TeX output (and consequently in DVI and PostScript). Linuxdoc-SGML doesn't support images for HTML, however.

An example:

```
<figure> <eps file = "neuro">
<caption><label id = "decade"> Decade of
the brain </figure>
```

You can make references to the figure with <ref id="decade">

PostScript is already supported and the developer version of SGML-Tools (the successor of Linuxdoc-SGML) now supports HTML as well. You can specify a PostScript and a GIF file and depending on the output (TeX or HTML) the respective image will be included.

X Configuration Issues

From: Michael J. Hammel,
mjhammel@emass.com

If you get sufficiently tweaked by the X monitor config problems, I suggest X Inside's AcceleratedX package. Its much simpler to configure than the XFree package for both cards and monitors. I used to work for them, but haven't in over a year. I still use their package because its the easiest to handle all the video card/monitor details.

BTW, the monitor setup is menu based. If your monitor is not listed you can just use one of the multisync if single frequency generic configs. No dot clocks, but you do need to know your monitor's frequency capabilities. These should be listed in the monitor's documentation.

The package is a commercial distribution and runs about \$100 (last time I checked). They change their name to Xi Graphics recently and the domain for xinside.com might not be working right now. Try <http://www.xig.com>.

Multiple X Displays

From: Michael J. Hammel
mjhammel@emass.com

Setting up the software is probably fairly straight forward. I've never used MetroX (I use AcceleratedX instead), however. Basically you'll have two choices:

Multiple displays (host:0.0 and
host:1.0)

Multiple screens of the same display
(host:0.0 and host:0.1)

The second choice is the one you need if you want to move the mouse between the two monitors - like when the mouse goes past the right edge of the first monitor it shows up on the left edge of the second monitor. You'll have to check with Metro to find out which of these options is supported and how to configure for it.

The hardware problem is tougher. The problem lies in the fact that PC's were not originally designed with the idea that multiple display adapters would be installed. The BIOS looks for an adapter at certain locations (IRQ, I/O address) and, unless the second card is configurable to some other address, the system will find multiple cards. What happens next is in-determinant. Some systems won't boot. Some do but don't display to either monitor correctly.

The trick is to find video adapters that were designed to be used in conjunction with other video adapters. Many are not. The easiest way for you to find out is check with Metro about what combinations of video adapters they know work together. Chances are good the ones you have don't. I know X Inside had a list of cards they knew work together. You could search their web site (<http://www.xinside.com> or <http://www.xig.com>) and see if that info is still there.

Color Depths with X

From: Michael J. Hammel
mjhammel@emass.com

After fiddling with the `xf86config` file in a concerted effort to coax X into displaying 16 bit color, I was dismayed to learn that with my current hardware (16 megs RAM and a Cirrus Logic GL-5426) 16 bit color is **impossible**...not because of any hardware in-capability, but because of a certain limitation of X Windows itself...a problem with linear addressing. Seems that to have 16 bit color under X, one must have linear addressing enabled, which only works if the system has **no more than 14 megs RAM**.

Horse hockeys. 16 bit color is a limitation of the video subsystem and has nothing to do with the memory of your system. Linear addressing in the XFree86 X servers might be tied to system memory amounts, but that would be a limitation in the XFree86 server, not in X. X defines "method without policy", so such limitations just aren't built into X.

A couple of things you should note: The number of colors available under 16bit displays is actually **less** than the number available to 8bit displays. Why this is true has to do with the way 16bit display hardware works. The actual color palette for 8 bit displays can have millions of colors - it can only display 256 colors at a time, however. Frugal

use of colormaps can allow you to have nearly exactly the right colors for any given application. 16 bit displays only have a palette of 65k (roughly) colors. Once those are used up, you're outta luck.

I'm not completely clear on what makes this difference such a problem but if you visit the Gimp User's mailing list (see the Linux Graphics mini-howto: <http://www.csn.net/~mjhammel/linux/lgh.html>) and ask this question you'll get similar replies. Its been discussed quite at length on the developers list, and most of them read the User's list.

BTW, if you want to see if Linear Addressing is the real problem, try the X Inside AcceleratedX demo server and see if it works in 16 bit color for you. Generally, your video card needs at least 1M of on board RAM (not system memory - this is video memory on the video card) to run in 16Bit mode, but then you'll probably only be able to run in 640x480 or (at most) 800x600 resolution. To run at higher resolutions you'll need more video memory.

Figuring Out the Boot Process

From: David Ishee
dmi1@ra.MsState.Edu

One of the things that is confusing about Linux at first is which files Linux uses to load programs and get the system started at bootup. Once you figure out which programs are run during the boot process, which order are they run? Here is an easy solution.

On my Red Hat 4.0 system, the `/etc/rc.d` directory tree is where everything happens. There are a lot of shell scripts in this set of directories that are run when the system boots. To give yourself a little more info, add some echo statements to the files. For example:

```
edit /etc/rc.d/rc.sysinit and add the
following lines at the beginning
echo ""
echo "**** Running /etc/rc.d/rc.sysinit ****"
```

echo ""

Now when the system is booting you can see exactly when `rc.sysinit` is run, and what programs it launches. Repeat the above process for all the scripts you find.

Now if the system hangs or gives an error during bootup you have a better idea of where to look. If you don't have any problems while booting then at least you have more info about what Linux is doing.

ftping Home

From: Kaj J. Niemi, kajtzu@4u.net

I read your article about `ftping` home with dynamic IPs.. Here's something you might need if you get tired of looking at the screen every time you want to find out the IP.

```
ADDRESS='/sbin/ifconfig | awk
'BEGIN { pppok = 0 } /ppp.* / { pppok
= 1; next } {if (pppok == 1) {pppok =
0; print } }'| awk -F: '{print $2 }'| awk
'{{print $1 }}'
```

Just replace the `ppp.*` with whatever you want (if you have multiple `ppps` running). The easiest thing would be to write a script called `ftphome` (or similar) and make it first assign the address and then doing `ftp` or `ncftp` `$ADDRESS`. The snippet is originally from a local firewall, at the part where it needs to know what its' own address is. :-) A friend of mine at `mstr@ntc.nokia.com` wrote this for me.

Contact Information

To contact the MUUG board for membership information or anything else, send e-mail to `board@muug.mb.ca`. We have a Web presence as well, at <http://www.muug.mb.ca/>, where you can find all kinds of information, including details of upcoming and past meetings and presentations and references related to them. E-mail the editor at `editor@muug.mb.ca`.