

This Month's Meeting

This month we are pleased to announce we will be having Brian Pauls speaking on the topic of TCP/IP. This is a broad topic, and Brian has broken it down into the following areas, with the first few being more of an overview; more attention will be given to later areas:

A Very Brief History of TCP/IP; How computers communicate (a brief overview of the OSI); Layer 2 communication (MAC addresses etc.); Layer 3 communication (Network addresses, ARP, etc.); Internetwork Communication (Routers, routing tables, routing protocols); Network addressing (the class system and subnetting); (Anatomy of an IP address); and Troubleshooting tools.

Our meeting this month is Tuesday, the 13th of January. 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.

Notes On Last Month's Newsletter

You might be wondering why you didn't get last month's newsletter until this month, along with the January issue. The answer is: the postal strike. We felt that since we couldn't guarantee delivery of the newsletter before the December meeting that it wouldn't serve as a meeting notice, and you probably wouldn't have time to read it until after Christmas anyway.

Those of you who have provided us with an accurate email address received a meeting notification by that method, and we also posted our meeting notice (as usual) to various local newsgroups. Don't forget you can always check the MUUG web site (<http://www.muug.mb.ca/>) for the latest MUUGlines issue, which is always posted before each month's meeting.

Perfect Backup+ Personal Edition

Unisource Systems, Inc. announced today the release of the famous PerfectBACKUP+ Personal Edition, a fully functional version of their best-selling PerfectBACKUP+ V5.5. Having received continued and tremendous support from the LINUX community, and in recognition of LINUX becoming our #1 best-selling platform we are giving something back. The PerfectBACKUP+ Personal Edition is unrestricted and free to anyone. Its freely redistributable and can be use for either private or commercial use.

Information about, and the program itself can be obtained from <http://www.unisrc.com/>.

Running Multiple Instances of X on One Video/Monitor (VCs)

From: Guillermo S. Romero famrom@ran.es

Hello,

*I have tried to run multiple X servers with only one card and one monitor. Is this possible, or is it normal that the second X server does not run? I used startx display :0 the first time, and :1 the second. I have a 1024K video board (#9GXE64 PCI, S3 864), and normal config is 8 bpp, 1024*768 virtual desktop, running on a remix of RedHat 4.0, 4.1 and 4.2, with XFree86 as server. Maybe I did not understand the man page (English is not my first language). Any suggestion?*

The normal way this is done is using the form:

```
startx -- :0 &
startx -- :1 &
```

... The -- is used by startx and xinit to separate an optional set of client parameters from the set of display/server options and parameters.

If you ran the command:

```
startx xterm -e myprog -- :1 &
```

... it would start X Windows with a copy of xterm which would be running 'myprog' (whatever that might be). The remainder of the line informs the X server to use display number one (which would be VC -- virtual console -- number eight on most Linux systems).

(On my systems it would start on VC#14 -- accessed with the {Right Alt}+{F2} key combination. I routinely configure mine with 24 VC's -- the first twelve of which have "getty's"

(login prompts) and the next eleven of which are available for X (xdm's or otherwise), using 'open' commands, or for dumping status output from a process (like 'make' or 'tail -f').

Read the man pages for startx and xinit one more time. I'm pretty sure that the man pages have all been translated into Spanish -- so you might want to hunt those down.

Thanks!!!

Read the man pages for startx and xinit one more time.

Sure, and with a dictionary. ;]

I'm pretty sure that the man pages have all been translated into Spanish -- so you might want to hunt those down.

Try: man-pages-es-0.2-1.src.rpm:
ftp://ftp.redhat.com/pub/contrib/SRPMS/man-pages-es-0.2-1.src.rpm

The Spanish Howto is small, too global, sure it does not cover that. And I still have problems with my tilde chars and accents, Spanish is not supported a lot (Linux or another OS, always late and bad)... The system explained in that howto does not work (but thats another question, whose solution maybe... magic? real support?). GSR

I'm afraid I'm completely ignorant of internationalization issues with Linux. I do know that there is quite a bit of work done on Linux boxes in Japan, Germany, Italy and, naturally enough, Finland (where Linus comes from).

As bad as it seems -- Linux' support for other languages is probably the best in the world. Unfortunately I don't have the skill or resources to point you to the support and resources you need.

Since your English is clearly adequate to discuss these issues with me -

- you might consider contributing some of your time to a translation effort (get the LIGS, NAG, and SAG portions of the Linux Documentation project translated, and "beef up" (improve) the Spanish-HOWTO.

I highly recommend that you find or start a Linux user's group in your area. This is the best way to help yourself and to improve the situation for all of your compatriots.

VC Madness

From: frees@technologist.com
I have an application that uses its own .cshrc and .bashrc to fire up and this is done by using its own login account. Now what I would really like is for this to say select VC8 to run on and then have my normal X on VC7 as usual. Can this be done? and if so how?
--Phil

```
open -c 8 -- su - $PSEUDOUSER
```

... where PSEUDOUSER is the pseudo users whose *.shrc you want to run. Naturally you can convert the *.shrc into a normal shell script and do whatever you like with it. You have to run this as root -- (so 'su' doesn't prompt for a password) though there are ways to get around that 'runas' is available at the sunsite.unc.edu archive site and its mirrors). If launch this from another UID you'll need to ensure that this users (the launching user, not necessarily the psuedo-user) has write access to /dev/tty8 (group +w should be sufficient).

If you want to have the console visually switch to this application's VC you can just add the -s switch like so:

```
open -c 8 -s -- ....
```

... where the "--" marks the end of 'open's' arguments so that the command that follows it can unambiguously get its own arguments.

Without the -c switch the 'open' command will select the next available VC. Any subsequent 'startx' commands or other 'open' commands would then pick later ones (unless the others were freed back up).

You can have two or more copies of X running on different VC's as well. For example the command:

```
startx -- :1
```

... will create a second X session on the localhost:1 display (the first one is addressed as localhost:0 or simply :0). These X sessions can be run under different UID's and have completely different client configurations (colors, window managers, etc). There is also an 'Xnest' command that works similarly - allowing one X session to run "within" (as through a window on) one of your existing X sessions.

You can also set the terminal settings and colors using normal redirection of the form:

```
stty erase ^? > /dev/tty8
```

... and:

```
setterm -background blue -foreground yellow -bold on -store \> /dev/tty8
```

This last command would set and store a new set of default screen colors for the VC. The setterm command can also be used to control the Linux VC screen blanker's timeout (a value of 0 means "never blank").

Naturally you may want to read the man pages for all of these.

If you want to ensure that a given process will *always* be running (and will automatically be respawned when it dies) you can add it to your /etc/inittab

-- so that the init process will watch over it. This is how new 'getty' processes are spawned on your first six (or so) VC's when you boot and are respawned when you logout. Likewise if you use 'xdm' to keep a graphical (X based) login prompt on one or more of your VC's.

As you can see, its possible to do quite a bit with Linux VC's. I run 12 VC's with getty (as login consoles), have one 'xdm', one devoted to syslog, and ten more available for other purposes (such as 'startx' and 'open' commands and to to use for 'tail -f' commands when need to monitor the end of a status or log file -- from a 'make' or whatever.

The second set of 12 VC's is accessed with the *right* {Alt} key. (In case you'd never noticed, the default keyboard settings of Linux only allow you to use the *left* {Alt} key for switching VC's). I set syslog to use VC number 24 with an entry in the /etc/syslog.conf file that reads:

```
*.* /dev/tty24
```

This puts a copy of *every* syslog message on to that VC -- which is what I switch to for a quick glance and try to switch to when I leave any of my systems unattended. (That way when one does lock -- as rare as that is -- I have some idea of what the last throes of the system were).

I set that to bright red on black with the following command in my rc.local file:

```
setterm -foreground red -bold on -store > /dev/tty24
```

(I also do the same to /dev/tty12 which I customarily use only for root login's).

Hope all of that helps.
-- Jim

Linux and OSPF

From: Jose Manuel Cordova-Villanueva
jcordova@amoxcalli.leon.uia.mx

Dear Sr.

Recently I had my first contact with the Linux G. and is a big source of information, can you inform me if there are a program that can talk ospf because our ISP, is changing from RIP to OSPF and we have a linux box in one of our links, for our cisco no problem but for our Linux box?

The software you want is called 'gated' (for "gateway daemon"). This is a Unix multi-protocol router package for Linux which includes support for OSPF and other routing protocols (BGP4, IGRP, etc).

Hope that helps.
-- Jim

Security Problems with pop3 of Linux 2.1.29

From: Sam Hillman
hillman@easyway.net

Well I hope I'm posting to the right person. I have two questions, which I hope you can answer. 1. How do I setup my linux machine as a POP3 server? I can't find any FAQs or Howtos.

Usually you don't have to do anything extra to allow access to POP services. Most Linux distributions include a pop server pre-installed and appearing in the /etc/inetd.conf and /etc/services files.

A quick test is to login to the system in question and type the command:

```
telnet localhost pop-3
```

... it should respond with something like:

```
+OK your.hostname .... (some copyright info)
```

... and you can type QUIT to get out of that.

If that doesn't work you'll want to make sure that the appropriate lines appear in your /etc/services and /etc/inetd.conf files like so:

```
/etc/services:
pop-3 110/tcp # PostOffice V.3
pop 110/tcp # PostOffice V.3
```

```
/etc/inetd.conf:
pop-3 stream tcp nowait root /usr/sbin/tcpd ipop3d
```

If they appear commented out -- remove the leading hash sign(s) (or paste these samples in) and restart your inetd with a command like:

```
kill -HUP `cat /var/run/inetd.pid`
```

2. When I log on to my ISP, I download my mail and it gets dumped to the sendmail, this creates a situation where the mail is bounced back and forth until it passes the hop limit and is dumped as an error message in the postmaster box, and a nasty letter is send to the originator from MAILER-DEMON...

I think this maybe because I'm running a local area network between my two machines, the IP address of the local net is 162.blah.blah... But I also have the IP address the ISP gave me in the host file. If the ISP's IP address is the problem can I remove it from the host file, and just get a duynamic IP when I connect? Thanks in advance!

This is a bigger problem. First the 162.*.* is probably not what you want to use for you disconnected LAN. There is an RFC 1918 (originally RFC 1597) which describes and reserves a set of addresses for "non-Internet" use. These are guaranteed not to collide with any valid (routable) hosts on the 'net. Here's the list of those addresses:

192.168.*.* (255 class C address blocks)

172.16.*.* through 172.31.*.* (15 class B address blocks)

10.*.*.* (one class A address block)

... use those however you like. Be sure to keep them behind your own routers (make any hosts with those go through an IP masquerading or NAT -- network address translation -- router, or through a SOCKS or other proxy server).

The next problem is configuring sendmail for use on a disconnected system. You probably need to define your hostname (or an alias to your hostname) to match what your ISP has named you. Each ISP seems to use a different way to manage these "disconnected sendmail subdomains" -- with no standardization in site (which is why I use UUCP). I gather that some people use a scheme where they only run sendmail when they are connected. The rest of the time their MUA (mail user agents like elm, pine, mh-e, exmh, etc) just drop outgoing mail into the mqueue directory where 'sendmail' will get to it later. One problem I have with these configurations is that sendmail wants to look up these remote hosts. This seems to cause various problems for users of "disconnected" or "periodically connected" (dial-up) systems. So far the only solutions I've found are: recompile sendmail without DNS support (there used to be a sendmail.cf switch that disabled DNS and reverse DNS activity in sendmail - but that doesn't seem to work any more) use UUCP. UUCP was designed for disconnected (dial-up) and polling systems. It's what I use. The disadvantage to UUCP is that it's a bit hard to set up the first time -- and you have to find a provider that's willing to be your MX/SMTP to UUCP gateway. There are still some people out there where will do this for free or at only a nominal fee. But they are increasingly hard to find. I use a2i Communications in San Jose. You

could use a non-local provider if you want to use UUCP over TCP as the transport mechanism (UUCP is pretty flexible about the underlying transports -- you could probably use tin cans and string as far as its concerned).

There are several HOWTO's that try to cover this topic. Try browsing through some of these:

ISP Hookup HOWTO: Basic introduction to hooking up to an ISP.

Electronic Mail HOWTO: Information on Linux-based mail servers and clients.

Mail Queue mini-HOWTO: How to queue remote mail and deliver local mail.

Offline Mailing mini-HOWTO: How to set up email addresses without a dedicated Internet connection.

UUCP HOWTO: Information on UUCP software for Linux.

Sendmail+UUCP mini-HOWTO: How to use sendmail and UUCP together.

-- Jim

Thanks so much for the detailed suggestions. We have installed a newer version of pop3 on our server for now and we will look into the feasibility of implementing some of your suggestions for a final cure. Thanks again, James, we really appreciate it.

- Sam Hillman, Service Manager, Easyway Communications.

An Interesting De-Referencing Problem

From: Kevbo simitar@lvnworth.com
Here's the brain teaser I read about and promptly forgot the solution (because I thought it would never happen to me). How does one delete a FILE named "../" I have the following at the root directory.

./ a dir (which is normal)

../ a dir (which is normal)

../ a file (which is not normal)

How this happened I don't know. How to remove this not-bothering-me file has me stumped. Got an answer?

I suspect that this file is actually named something like: "../ " (note the trailing space!). In any event you can remove this with a command like

```
find / -type f -maxdepth 1 -name '..*' -print0 | xargs -0 rm -i
```

Note: you must use the GNU versions of find, xargs, and rm to ensure that these features (-print0, -0, and -i) are available. (They may be available in other implementations -- but you must check first).

The find parameters here specify files (not directories, symlinks, device nodes, sockets, or FIFO's) and force it to only search the named directory (or directories if you list more than just /). The -print0 force it to be written as a null-terminated strings (thus the receive process on the other end of the pipe must be able to properly interpret null-terminated arguments -- which is what the -0 to xargs accomplishes).

The -i on rm is just a little extra protection to prevent any other unexpected side effects. It forces rm to interactively inquire about each argument before removing it.

-- Jim

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.