Watching and Debugging Network Activity with the `tcpdump` Utility

tcpdump is a very useful program for monitoring network activity. It is installed by doing

```
sudo apt-get install tcpdump
```

and usually used with `-n` and `-i` options to show messaging activity on a particular interface. For example

```
sudo tcpdump -n -i eth0
```

will show all activity on the interface `eth0`. Without the `-n` option `tcpdump` will try to replace IP addresses with the corresponding host names. This greatly increases the network traffic and usually makes it harder to follow. `tcpdump` keeps working until stopped with a <`Ctrl-c`>.

Besides the Ethernet interface `eth0`, there is also a local interface `lo` and probably a wireless interface `wlan0`. Let's start playing with the simplest of these, the `lo` interface, which is used to communicate within a computer while using networking protocols.

Open up two separate command-line windows on your Pi desktop. In the left one, enter the command

```
sudo tcpdump -n -i lo
```

and in the right one type

```
ping -c 1 127.0.0.1
```

The right window will then show:

```
PING 127.0.0.1 (127.0.0.1) 56(84) bytes of data.
64 bytes from 127.0.0.1: icmp_req=1 ttl=64 time=0.960 ms

--- 127.0.0.1 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.960/0.960/0.960/0.000 ms
```

and the left window will show:

```
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on lo, link-type EN10MB (Ethernet), capture size 65535 bytes
12:32:12.947624 IP 127.0.0.1 > 127.0.0.1: ICMP echo request, id 2556, seq 1, length 64
12:32:12.948152 IP 127.0.0.1 > 127.0.0.1: ICMP echo reply, id 2556, seq 1, length 64
```

In the right window, you sent a single `ping` message packet to your computer and it responded in less than a millisecond. In the left window, the `tcpdump` program shows the transaction. There are three types of messages defined in the Internet Protocol:

ICMP	messages between computer systems to deal with coordination and communication problems.
UDP	messages that provide information that is not critical if lost.
TCP	messages that must be received correctly, with resending of packets as needed.

Entering the commands `man tcpdump`, `man ping`, `man ICMP`, `man UDP`, and `man TCP` provide more information.

If you have the `apache2` web server with PHP installed on your Pi as suggested in my earlier notes, the following test shows how a web page is accessed. Upon typing

```
wget 127.0.0.1
```

the right window then shows

```
--2015-06-28 13:06:51--  http://127.0.0.1/
Connecting to 127.0.0.1:80... connected.
HTTP request sent, awaiting response... 200 OK
Length: 269 [text/html]
Saving to: `index.html'

100%[======================================================================>] 269         --.-K/s   in 0s

2015-06-28 13:06:51 (3.70 MB/s) - `index.html' saved [269/269]
```

The file that was saved, `index.html`, contains

```
<html><body><h1>It works!</h1>
```

```
<p>This is the default web page for this server.</p>
<p>The web server software is running but no content has been added, yet.</p>
<br />
<p>The following date and time is produced by php code:<br /></p>
2015-06-28 13:06:51</body></html>
```

It came from the default page of the apache2 server which had the following in the file `/var/www/index.php`:

```
<html><body><h1>It works!</h1>
<p>This is the default web page for this server.</p>
<p>The web server software is running but no content has been added, yet.</p>
<br />
<p>The following date and time is produced by php code:<br /></p>
<?php
 echo date('Y-m-d H:i:s');
?>
</body></html>
```

The PHP interpreter converted the PHP code between `<?php` and `?>` into the HTML code seen in `index.html`.

The left window then shows the full transaction (the chksum errors should be ignored in this case)

```
13:06:51.582068 IP (tos 0x0, ttl 64, id 8486, offset 0, flags [DF], proto TCP (6), length 60)
    127.0.0.1.54383 > 127.0.0.1.80: Flags [S], cksum 0xfe30 (incorrect -> 0x4da0), seq 1500980538, win 43690, options [mss
65495,sackOK,TS val 197660 ecr 0,nop,wscale 6], length 0
12:49:05.512588 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto TCP (6), length 60)
    127.0.0.1.80 > 127.0.0.1.54383: Flags [S.], cksum 0xfe30 (incorrect -> 0x44ba), seq 3940751826, ack 1500980539, win
43690, options [mss 65495,sackOK,TS val 197660 ecr 197660,nop,wscale 6], length 0
13:06:51.584886 IP (tos 0x0, ttl 64, id 8487, offset 0, flags [DF], proto TCP (6), length 52)
    127.0.0.1.54383 > 127.0.0.1.80: Flags [.], cksum 0xfe28 (incorrect -> 0x15a9), seq 1, ack 1, win 683, options [nop,nop,TS
val 197660 ecr 197660], length 0
13:06:51.589335 IP (tos 0x0, ttl 64, id 8488, offset 0, flags [DF], proto TCP (6), length 167)
    127.0.0.1.54383 > 127.0.0.1.80: Flags [P.], cksum 0xfe9b (incorrect -> 0xfcef), seq 1:116, ack 1, win 683, options
[nop,nop,TS val 197660 ecr 197660], length 115
13:06:51.590018 IP (tos 0x0, ttl 64, id 7630, offset 0, flags [DF], proto TCP (6), length 52)
    127.0.0.1.80 > 127.0.0.1.54383: Flags [.], cksum 0xfe28 (incorrect -> 0x1535), seq 1, ack 116, win 683, options
[nop,nop,TS val 197661 ecr 197660], length 0
13:06:51.816343 IP (tos 0x0, ttl 64, id 7631, offset 0, flags [DF], proto TCP (6), length 569)
    127.0.0.1.80 > 127.0.0.1.54383: Flags [P.], cksum 0x002e (incorrect -> 0x514b), seq 1:518, ack 116, win 683, options
[nop,nop,TS val 197683 ecr 197660], length 517
13:06:51.816806 IP (tos 0x0, ttl 64, id 8489, offset 0, flags [DF], proto TCP (6), length 52)
    127.0.0.1.54383 > 127.0.0.1.80: Flags [.], cksum 0xfe28 (incorrect -> 0x12f3), seq 116, ack 518, win 700, options
[nop,nop,TS val 197683 ecr 197683], length 0
13:06:51.835947 IP (tos 0x0, ttl 64, id 8490, offset 0, flags [DF], proto TCP (6), length 52)
    127.0.0.1.54383 > 127.0.0.1.80: Flags [F.], cksum 0xfe28 (incorrect -> 0x12f0), seq 116, ack 518, win 700, options
[nop,nop,TS val 197685 ecr 197683], length 0
13:06:51.838289 IP (tos 0x0, ttl 64, id 7632, offset 0, flags [DF], proto TCP (6), length 52)
    127.0.0.1.80 > 127.0.0.1.54383: Flags [F.], cksum 0xfe28 (incorrect -> 0x12fe), seq 518, ack 117, win 683, options
[nop,nop,TS val 197685 ecr 197685], length 0
13:06:51.841501 IP (tos 0x0, ttl 64, id 8491, offset 0, flags [DF], proto TCP (6), length 52)
    127.0.0.1.54383 > 127.0.0.1.80: Flags [.], cksum 0xfe28 (incorrect -> 0x12ec), seq 117, ack 519, win 700, options
[nop,nop,TS val 197686 ecr 197685], length 0
```

Web pages are provided using the TCP protocol. In those transactions, a port address (usually 80) is specified by the requesting computer and placed after the IP address, `127.0.0.1.80` in this case. A port is simply a way of asking for a particular kind of service. Secure shell service (ssh) uses port 22, mail transport (smtp) uses port 25, network time protocol (ntp) uses port 123, secure web pages (https) use port 443, etc. A complete list can be seen by looking at the file `/etc/services`. Various daemon programs watch their assigned ports and handle requests sent to those ports. The requesting computer sends its message out through a port using some very large available port number, `582068` in the initial transaction shown above. Large files are sent in smaller packets, each with a packet number, which may travel different paths through the Internet, but which are reassembled correctly upon receipt using sequential packet numbers.

When networking fails to work as expected, `tcpdump` can help isolate the problem by examining traffic on each side of various interfaces. A firewall may be blocking the communication, a server may not be running to handle the request, the request may be improperly constructed, etc.

Be sure to also play with `tcpdump` on your `eth0` and `wlan0` interfaces and make it an integral part of your computer toolbox.

An extremely detailed book on communication within and between computers is *UNIX Network Programming – Networking APIs: Sockets and XTI, Volume 1, 2nd Edition* (ISBN 0-13-490012-X) by the late W. Richard Stevens. There may be newer books available now, but this author explained subtle details usually ignored by most authors. What a pity he died of an accident in his prime.