Packet Crafting

OpenFest, Sofia
November, 2013

Shteryana Shopova,
syrinx@FreeBSD.org

Disclaimer

!!! Use at your own risk !!!

Shteryana Shopova,
syrinx@FreeBSD.org


Motivation
Shteryana Shopova,
syrinx@FreeBSD.org

Motivation (2)

Shteryana Shopova,
syrinx@FreeBSD.org

Agenda

- ping(8), traceroute(8), telnet(1), nmap(1)
- ng_source(4), tcpdump(1)
- bridge(4), vlan(4), svlan(4), trunk(4)/lagg(4)
- yersinia(8)
- nemesis(1), libnet
- hyenae
- pierf, Scapy
- netmap(4)
- iperf
- PF_PACKET, etc... sockets, C code
- OS network stack & daemons

Shteryana Shopova,
syrinx@FreeBSD.org

Let's start

- ping(8)
  - options: flood, quiet, multicast, source route, etc
- traceroute(8)
  - incrementing IP TTL probes
  - UDP, but may be random IP protocol number
- telnet(1)
  - TCP port 23, unless custom specified
- nmap(1)
  - network security scanner
- and more
  - arping, nslookup/host/dig, ...

Shteryana Shopova,
syrinx@FreeBSD.org

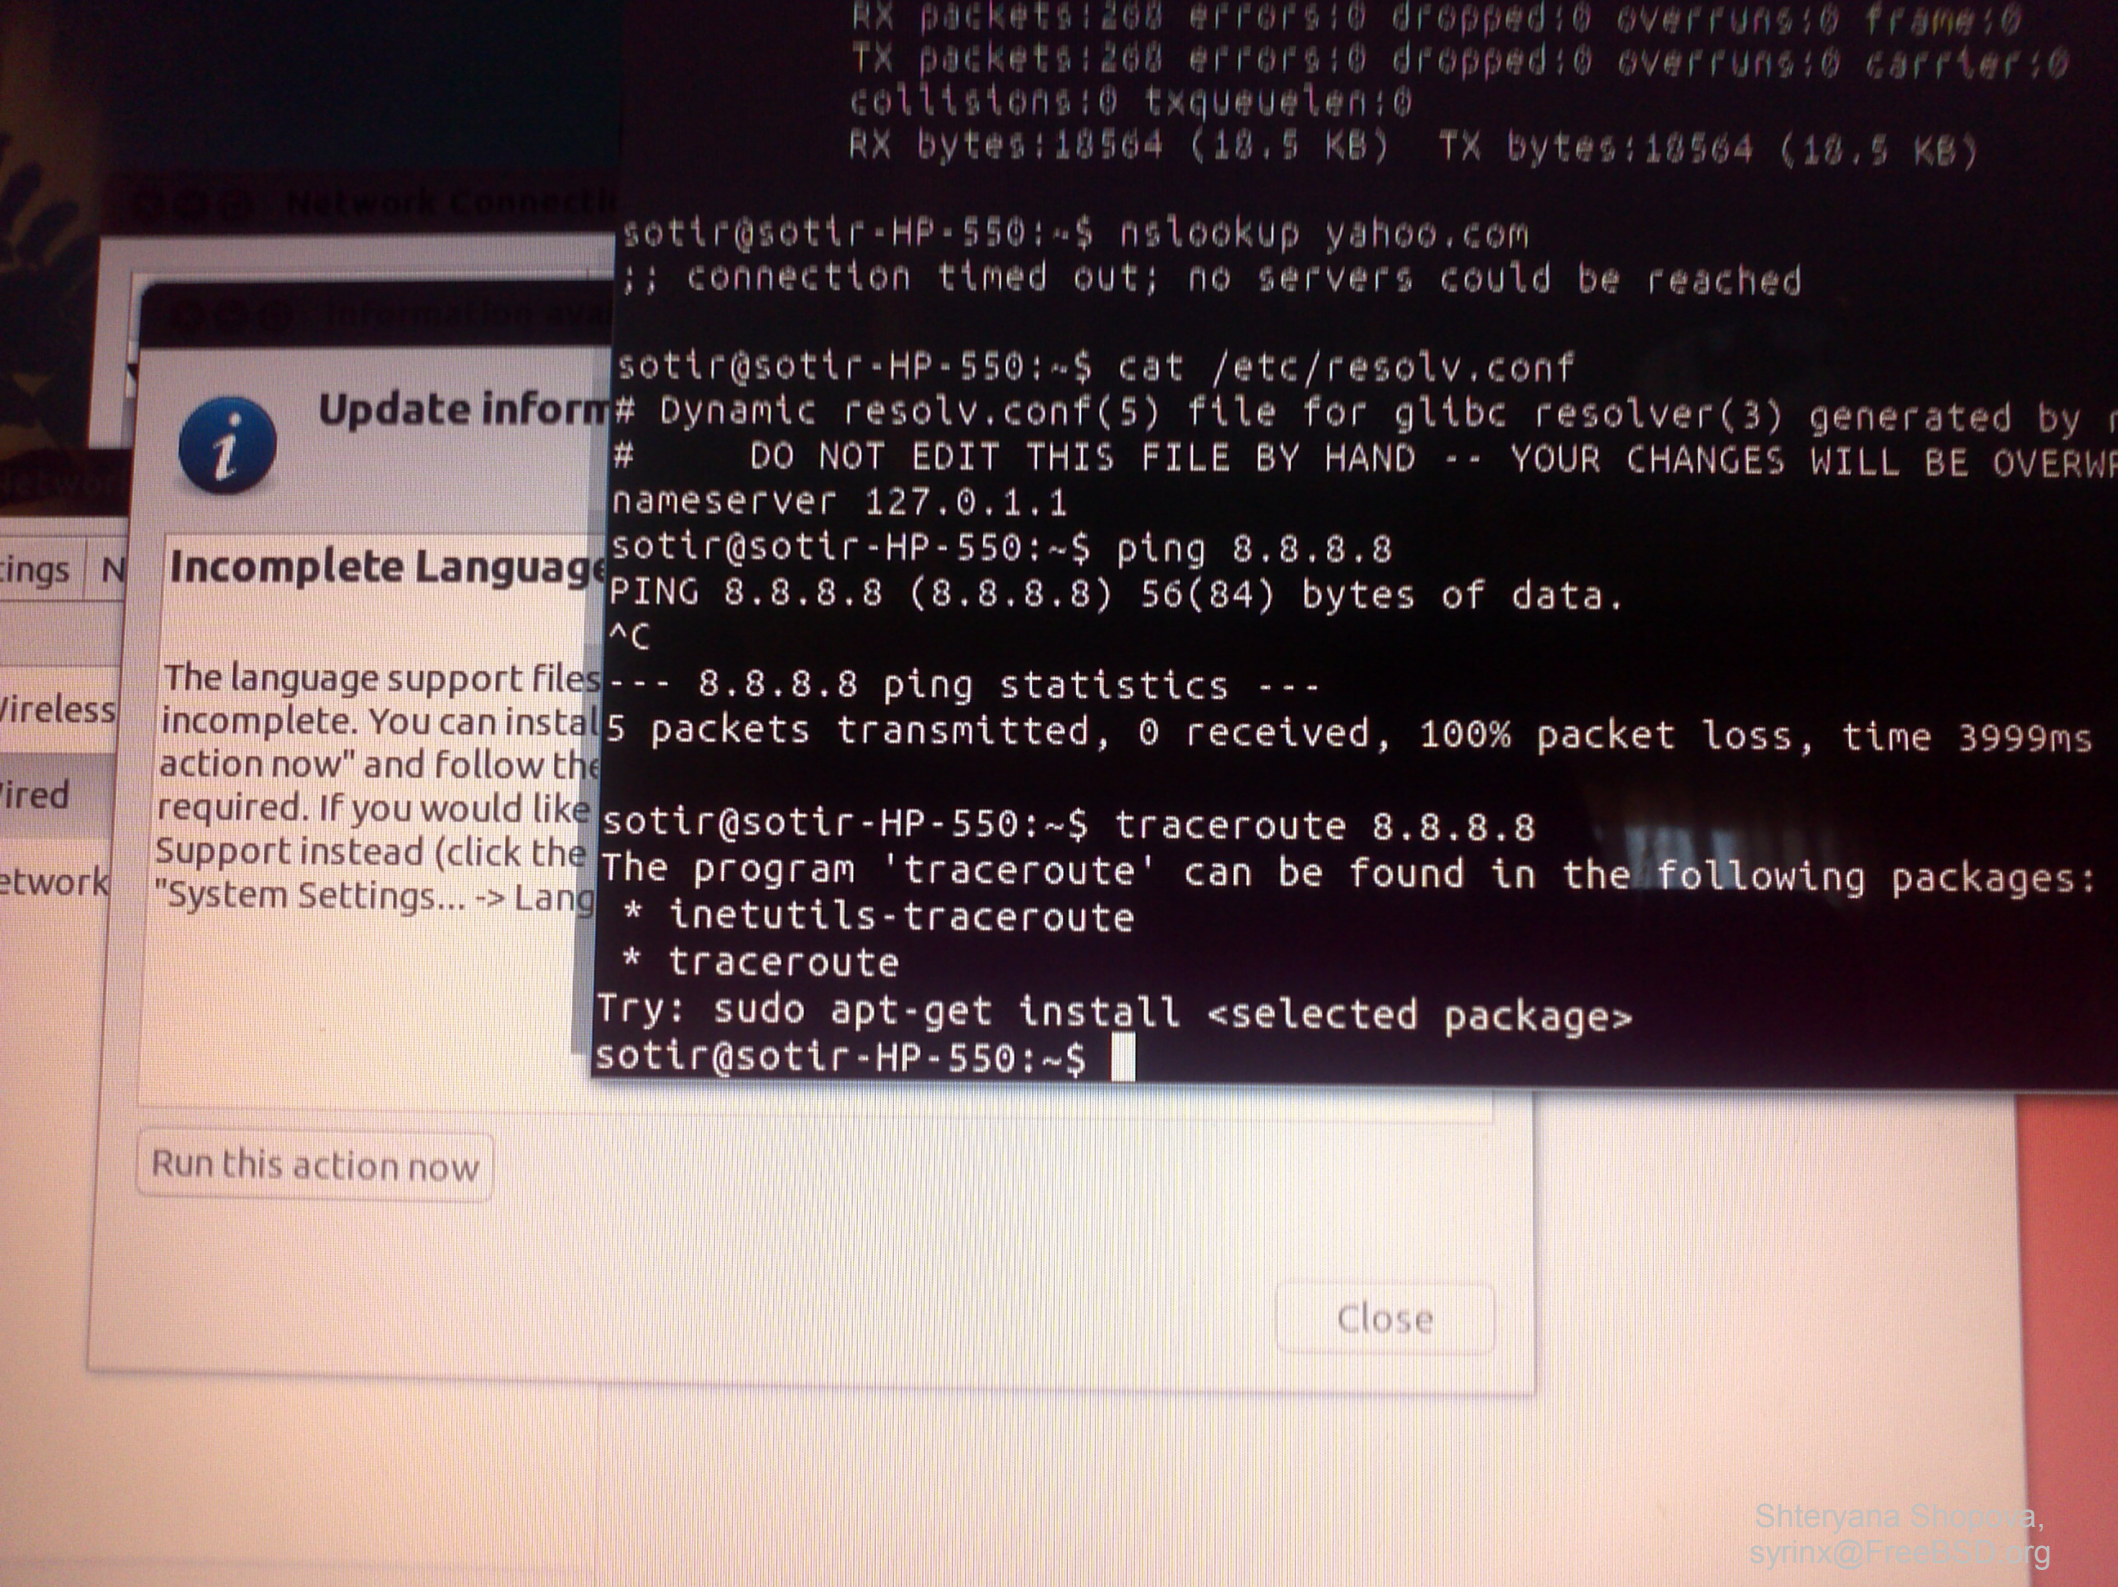

ng_source(4)

- tcpdump(8)

```
11:40:23.411341 00:e0:0c:11:95:22 > 00:00:00:77:77:78, ethertype MPLS unicast (0x8847), length 82: MPLS (label 28673,
        0x0000:  0000 0077 7778 00e0 0c11 9522 8847 0700
        0x0010:  11ff 0100 5e02 020b 0040 9530 450b 8100
        0x0020:  0001 0800 4500 002e fdc5 0000 4011 35b9
        0x0030:  6401 0132 e002 020b f28f 4908 001a 9d1b
        0x0040:  476f 7420 4d75 6c74 6963 6173 743f 213f
        0x0050:  0a00
```

  - there's wireshark too
- netgraph(3)
  - graph based kernel networking subsystem of FreeBSD
- ng_ether(4)
  - a node automatically created for every Ethernet interface in the system

Shteryana Shopova,
syrinx@FreeBSD.org

ng_source(4) contd.

```
#!/bin/sh
#########################################################
# TODO: insert some nice comment here
#########################################################
ECHO=/bin/echo
NGHOOK=/usr/sbin/nghook
HOOKNAME=input

usage () {
$ECHO     "Usage: $0 <ng-node> <count> <packet-gen> <initial-arg> <increment-gen>"
}

...

$ECHO "Injecting $LIMIT packets into $NGNODE node"

# Loop from 1 to 255
while [ "$var1" -le $LIMIT ]
do
DSTIP=${ARGS}
$ECHO "Injecting packet with DST IP ${DSTIP}"

${PACKGEN} ${DSTIP} | ${NGHOOK} ${NGNODE} ${HOOKNAME}
let "var1 = var1 + 1" >> /dev/null  # let "var1+=1"
ARGS=`${INCGEN} ${DSTIP}`
done
exit 0
```

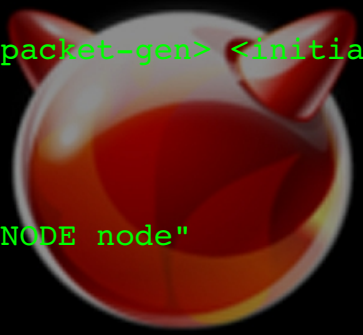

Shteryana Shopova,
syrinx@FreeBSD.org

Ether-like interfaces

- bridge(4)
- vlan(4), svlan(4)
- trunk(4)/lagg(4)
- tap(4)

Shteryana Shopova,
syrinx@FreeBSD.org

Yersinia

- http://www.yersinia.net
- Supported protocols
  - STP, CDP, DTP, DHCP, HSRP, IEEE 802.1Q, IEEE 802.1X, ISL, VTP
- UI
  - GTK, ncurses, command line, netclient
- predefined scenarios
- netclient
  - TCP port 12000
  - passwd/enable passwd : root/tomac

Shteryana Shopova,
syrinx@FreeBSD.org

nemesis(1)

```
NEMESIS -=- The NEMESIS Project Version 1.4 (Build 26)

NEMESIS Usage:
  nemesis [mode] [options]

NEMESIS modes:
  arp
  dns
  ethernet
  icmp
  igmp
  ip
  ospf (currently non-functional)
  rip
  tcp
  udp
```

hyenae

- No *BSD packages
- available on SF
- assumes /usr/include & /usr/lib
- ...

Shteryana Shopova,
syrinx@FreeBSD.org

hyenae (contd.)

- “ERROR: Root privileges required”
- Supported protocols
  - ARP, PPPoE, ICMP, TCP, UDP, DNS, DHCP, HSRP
- UI
  - command line, QT-frontend
- Example

```
$ sudo hyenae -a pppoe-discover -i em0 -s 01:00:82:00:00:c2 -d
01:00:82:00:00:c3
* Initializing
* Opening network interface (em0)
* Launching attack

  Press any key to stop
```

Shteryana Shopova,
syrinx@FreeBSD.org

Scapy

- “Powerful interactive packet manipulation program in python”
- Python intrerpreter disguised as a Domain Specific Language
- Fast Packet Designing
- Interactive packet and result manupulation
- Fast packet generator ? :dd

Shteryana Shopova,
syrinx@FreeBSD.org

Scapy - ls()

```
Welcome to Scapy (2.2.0)
>>>
>>> ls()
ARP         : ARP
ASN1_Packet : None
BOOTP       : BOOTP
CookedLinux : cooked linux
DHCP        : DHCP options
DHCP6       : DHCPv6 Generic Message)
DHCP6OptAuth : DHCP6 Option - Authentication
DHCP6OptBCMCSDomains : DHCP6 Option - BCMCS Domain Name List
DHCP6OptBCMCSServers : DHCP6 Option - BCMCS Addresses List
DHCP6OptClientFQDN : DHCP6 Option - Client FQDN
DHCP6OptClientId : DHCP6 Client Identifier Option
DHCP6OptDNSDomains : DHCP6 Option - Domain Search List option
DHCP6OptDNSServers : DHCP6 Option - DNS Recursive Name Server
DHCP6OptElapsedTime : DHCP6 Elapsed Time Option
DHCP6OptGeoConf :
DHCP6OptIAAddress : DHCP6 IA Address Option (IA_TA or IA_NA suboption)
...
<6 more screens of output>
```

Shteryana Shopova,
syrinx@FreeBSD.org

Scapy - lcs()

```
>>> lsc()
bind_layers           : Bind 2 layers on some specific fields' values
corrupt_bits          : Flip a given percentage or number of bits from a string
corrupt_bytes         : Corrupt a given percentage or number of bytes from a string
defrag                : defrag(plist) -> ([not fragmented], [defragmented],
defragment            : defrag(plist) -> plist defragmented as much as possible
etherleak             : Exploit Etherleak flaw
fragment              : Fragment a big IP datagram
fuzz                  : Transform a layer into a fuzzy layer by replacing some default values by
random objects
hexdiff               : Show differences between 2 binary strings
hexdump               : --
hexedit               : --

rdpcap                : Read a pcap file and return a packet list
send                  : Send packets at layer 3
sendp                 : Send packets at layer 2
sendpfast             : Send packets at layer 2 using tcpreplay for performance
sniff                 : Sniff packets
split_layers          : Split 2 layers previously bound
sr                    : Send and receive packets at layer 3
srbt                  : send and receive using a bluetooth socket
srflood               : Flood and receive packets at layer 3
srloop                : Send a packet at layer 3 in loop and print the answer each time
srp                   : Send and receive packets at layer 2
traceroute            : Instant TCP traceroute
wrpcap                : Write a list of packets to a pcap file

<some commands ommitted, such as wireshark>
```

Shteryana Shopova,
syrinx@FreeBSD.org

Scapy – sending a packet

```
Welcome to Scapy (2.2.0)
>>> a=Ether()/IP(dst="www.slashdot.org")/TCP()/"GET /index.html
HTTP/1.0 \n\n"
>>> hexdump(a)
0000   00 22 68 5A AA E4 74 E5  0B E3 39 68 08 00 45 00   ."hZ..t...9h..E.
0010   00 43 00 01 00 00 40 06  41 3B AC 14 00 12 D8 22   .C....@.A;.....".
0020   B5 30 00 14 00 50 00 00  00 00 00 00 00 00 50 02   .0...P........P.
0030   20 00 84 38 00 00 47 45  54 20 2F 69 6E 64 65 78    ..8..GET /index
0040   2E 68 74 6D 6C 20 48 54  54 50 2F 31 2E 30 20 0A   .html HTTP/1.0 .
0050   0A                                                 .
>>> sendp(Ether()/IP(dst="1.2.3.4",ttl=(1,4)), iface="wlan0")
....
Sent 4 packets.
```

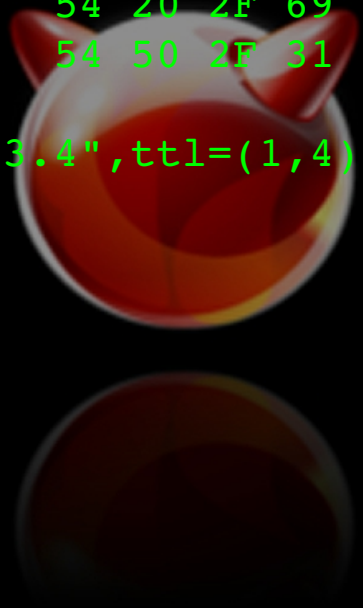

Shteryana Shopova,
syrinx@FreeBSD.org

Scapy – writing your own class

- TBD
- (homework)

Shteryana Shopova,
syrinx@FreeBSD.org

netmap(4)

- fast packet I/O framework
- reduce syscalls to improve packet processing performance
- in many cases, shared buffers between user and kernel space
- open(“/dev/netmap”)
- ioctl(.., NIOCREGIF, ..) to bind to an interface
- fill in available buffers with data
- non-blocking ioctl(.., NIOCTXSYNC) /ioctl(.., NIOCRXSYNC) to transmit/receive data
- supported drivers - em(4), igb(4), ixgbe(4), re(4)

Shteryana Shopova,
syrinx@FreeBSD.org

netmap(4)'s pkt-gen- example

```
shteryana@aphrodite:/usr/src/tools/tools/netmap %
shteryana@aphrodite:/usr/src/tools/tools/netmap % ./pkt-gen
main [1461] missing ifname
Usage:
pkt-gen arguments
        -i interface            interface name
        -f function             tx rx ping pong
        -n count                number of iterations (can be 0)
        -t pkts_to_send         also forces tx mode
        -r pkts_to_receive      also forces rx mode
        -l pkts_size            in bytes excluding CRC
        -d dst-ip               end with %n to sweep n addresses
        -s src-ip               end with %n to sweep n addresses
        -D dst-mac              end with %n to sweep n addresses
        -S src-mac              end with %n to sweep n addresses
        -a cpu_id               use setaffinity
        -b burst size           testing, mostly
        -c cores                cores to use
        -p threads              processes/threads to use
        -T report_ms            milliseconds between reports
        -P                              use libpcap instead of netmap
        -w wait_for_link_time   in seconds
shteryana@aphrodite:/usr/src/tools/tools/netmap %
```

Shteryana Shopova,
syrinx@FreeBSD.org

netmap(4) - example

```
char *buf;
int fd;
struct pollfd      fds;
struct nmreq       nmr;
struct netmap_if *nifp;

fd = open("/dev/netmap", O_RDWR);

strlcpy(nmr.nr_name, "em0", sizeof(nmr.nr_name));
nmr.nr_version = NETMAP_API;
ioctl(ng.fds.fd , NIOCGINFO, &nmr);

buf = mmap(NULL, nmr.nr_memsize, PROT_READ | PROT_WRITE,
    MAP_SHARED, fd, 0));

ioctl(fd , NIOCREGIF, &nmr);

fds.events = POLLOUT | POLLWRNORM | POLLWRBAND;
fds.fd = fd;

/* prepare packet in buf */
poll(&fds, 1, INFTIM);

ioctl(fd, NIOCTXSYNC, NULL); /* optional */
```

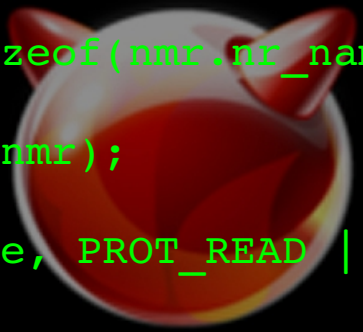

Shteryana Shopova,
syrinx@FreeBSD.org

iperf

- commonly used network testing tool
- creates UDP/TCP streams to measure throughput
- iperf (July, 2010) vs iperf3 (March, 2013)
- client and server
- uni-/bi-directional tests

Shteryana Shopova,
syrinx@FreeBSD.org

PF_PACKET socket & friends

- “The Guru of the Unix gurus”
- PF_PACKET
  - Linux-only
  - SOCK_DGRAM or SOCK_RAW modes
- BPF
- socket options
  - IP_TOS, IP_TTL,
    MCAST_JOIN/LEAVE_GROUP ,
    IP_ADD/DROP_(SOURCE_)MEMBERSHIP, etc

Slteryana Shopova,
syrinx@FreeBSD.org

References

http://www.kohala.com/start/unpv12e.html

http://www.unpbook.com/src.html

http://www.secdev.org/projects/scapy/

http://info.iet.unipi.it/~luigi/netmap

http://www.secdev.org/conf/scapy_pacsec05.pdf

http://docs.python.org/2/tutorial/introduction.html

http://www.secdev.org/projects/scapy/doc/usage.html#interactive-t

Shteryana Shopova,
syrinx@FreeBSD.org

Questions?

Shteryana Shopova,
syrinx@FreeBSD.org

Thank you!

* sources from this presentation available at http://people.freebsd.org/~syrinx/pktgen/

Shteryana Shopova,
syrinx@FreeBSD.org