

CITI Technical Report 98-6

## **Sifting the Network: Performing Packet Triage with NFR**

*M. Undy*

mundy@umich.edu

*C.J. Antonelli*

cja@umich.edu

### *ABSTRACT*

We describe a set of Network Flight Recorder scripts designed to detect network intrusions. After developing scripts that detect some known attacks, we focus on *sifting* scripts that attempt to remove "normal" traffic from a packet stream, leaving behind fewer packets requiring manual inspection. We conclude with a description of our experiences with the NFR product.

November 23, 1998

# Sifting the Network: Performing Packet Triage with NFR

*M. Undy*

mundy@umich.edu

*C.J. Antonelli*

cja@umich.edu

## INTRODUCTION

Passive protocol analysis is useful in detecting attacks on networked machines. Known attacks can be detected directly by looking for them in the network packet stream, using a network monitoring *protocol analysis engine* to look for unique bit patterns or sequences of packets unique to such attacks. Unknown attacks can be detected indirectly by separating for manual analysis suspicious traffic from traffic generated by customary network applications. For example, artificially fragmented packets, out-of-order TCP segments, overlapping segments, and packets containing incorrectly presented options generally do not occur during "normal" operations. A protocol analysis engine that sifts such "suspicious" packets from the network for manual inspection, eliding the normal traffic, is successful if the majority of network traffic is normal.

Filters for known attacks are useful in today's environment of widely-distributed attack scripts for rapidly identifying the existence of an attack and narrowing the space an administrator must examine before coping with the attack. For example, even if source addresses are forged, knowing that a "Ping of Death" attack is occurring allows an administrator to filter all ICMP echo request packets at the firewall for some period of time, or provide cumulative evidence that pings should be filtered permanently.

In this work we describe a protocol analysis engine both to detect known attacks and to sift the packet stream for suspicious traffic. We have selected the Network Flight Recorder [1] as a suitable development base, since it possesses an expressive scripting language, runs on UNIX platforms, and is available in source code form.

Ptacek and Newsham point out flaws in passive protocol analysis that can force a network monitoring machine to see packets that other hosts on the same network do not, prevent a network monitor from seeing traffic directed at other hosts, or attack a network monitor directly [2]. Nevertheless, we believe a protocol analysis engine is useful, within the limitations they outline, as their attacks are not as well-known as and are harder to construct than the attacks being monitored, particularly when faced with multiple analysis engines running on different operating system platforms.

## OVERVIEW OF NFR

NFR consists of a number of components, each responsible for a specific activity:

- *Packet suckers* capture packets from the network interface.
- A *decision engine*, written in the N-code scripting language, checks packets against a list of filters.
- *Backends* forward selected packets to storage or statistical processing.
- A *query interface* permits examination of stored data.

NFR is a multi-purpose network monitoring tool, usable for intrusion detection, usage analysis, and troubleshooting by system administrators and the hacker community alike. The interpreted N-code language allows a user to write arbitrarily complex scripts for analyzing incoming packets, limited only by the timing constraints imposed by the NFR engine to ensure that filters share system resources fairly. To ease the task of the script writer, the N-code language includes primitives for analyzing packet header and data fields, and the NFR engine supports internal operations such

as TCP reassembly. Functions called within the scripts pipe data to user specified backend databases. The graphical query interface uses HTML and Java to query these databases. Summarizing the NFR data flow, network packets are captured, N-code scripts are applied, and script outputs are sent to backends for storage and other processing.

### **NFR Testbed**

We constructed an intrusion detection testbed, consisting of a 90 Mhz Pentium with 32 MB RAM, 1.2 GB disk, and a 3Com 3C509 Ethernet card connected to the CITI production network via a switch port configured to monitor all network packets.

After some initial problems with the NFR development version running on OpenBSD 2.2, we settled on running the 1.6.2 release of NFR on FreeBSD 2.2.6-RELEASE. Along the way we identified an error in the development version, which we passed along to NFR, Inc.

### **KNOWN-ATTACK SCRIPTS**

After coordinating with NFR, Inc. staff to avoid duplication of effort, we wrote five known-attack scripts. The scripts themselves contain more detailed commentary. See Appendix A.

#### **Smurf attack**

This filter originated as a vehicle for learning how to write N-code. The filter detects the classic Smurf attack, looking for ICMP echo request packets destined for the broadcast address of a network. A known filter already detects Smurf-like attacks, but ours limits itself to ICMP echo to generate fewer false positives.

#### **Sendmail**

This filter looks for attempted and established connections on the SMTP port of machines on the local network whose IP addresses are not enumerated in the filter as legitimate mail servers. This filter thus detects both sendmail port scans and unexpected mail servers; we found such a server on our network. This filter is easily changed to monitor other ports.

#### **Buffer Overflow**

This filter looks for a specific pattern in all packet payloads going to a specified port. The current implementation detects a buffer overflow attack against named in Linux and FreeBSD. This filter can form the basis for general pattern matching within a packet.

#### **Teardrop**

This filter detects the Teardrop attack, looking for overlap in a collection of fragmented TCP packets. The obvious choice of data structure, a list of lists, is not supported by NFR, requiring us to carve up a single array manually. Our solution collects fragmented packets and uses fragment offsets and lengths to check for overlaps; a timeout discards older fragments to limit memory consumption. We collect packets as they arrive and process collected packets periodically; we have not compared this to an implementation that processes each packet as it arrives.

#### **Profiler**

We developed a profiling filter to categorize IP traffic. ICMP packets are counted by type; TCP and UDP packets are counted by port. NFR's associative arrays are useful here. This filter gives us a rough characterization of network traffic and helps focus our attention on traffic appropriate to the sifting work. As an immediate result, we observed that CITI has made a lot of progress in replacing `telnet` with `ssh`.

The above scripts were running intermittently against CITI's network traffic over a period of several weeks. Attack programs were used to verify correct filter behavior.

### **SIFTING SCRIPTS**

Here we concentrate on separating suspicious traffic from ordinary traffic. Our strategy is to observe a given stream (say, DNS traffic) for a period of time to note "normal" behavior, say a stream of "A" requests and responses, then write NFR scripts to filter out packets corresponding to this behavior. Playing network traffic through such a filter reduces its volume by suppressing normal traffic to leave "suspicious" traffic behind for manual inspection. This sifting strategy is successful to the extent that normal traffic

dominates suspicious packets.

Initially, we based our selection on protocols for which our profiling filter measured high counts. Unfortunately, some of those (`ssh` and `telnet`) defy sifting, and we chose not to inspect others (SMTP) for privacy reasons. Instead, we focused on protocols that have recognizable normal behavior: DNS, FTP, and HTTP. See Appendix B for script listings.

### Tcpdump

We wrote an NFR backend that accepts packets from the decision engine and writes them to a log file in `tcpdump` format. This permits other tools, such as `tcpdump` or `tcptrace` [3], to post-process NFR output. It also permits the output of one NFR engine to feed the input of another, which we want to explore further.

We wrote a version of this backend that accepts blocks of packets, which permits N-code scripts to accumulate packet streams before writing them out, in those cases where packets must be accumulated before the decision to log them can be made.

### DNS

Using our `tcpdump` backend, we wrote a filter that sifts DNS packets and retains suspicious packets in a log for subsequent inspection. The script inspects A, PTR, and MX requests. It logs packets that contain names that are too long, have too many labels, odd options, or an excessive number of records. DNS requests other than these are also logged.

With this simple filter we are able to remove over 90% of all DNS traffic, sifting out 61,092 of 67,048 packets seen in a testing period conducted 27-28 August 1998.

### FTP

This filter examines the FTP control stream for normal commands and begins logging the control and the associated data stream if suspicious commands are seen. Control streams are buffered so that normal commands preceding a suspicious command are logged.

A command found in the stream is first matched with a short list of normal commands, *e.g.* `CDUP` or `PWD`, and discarded if found. Then the arguments and argument lengths are checked, with some combinations flagged as suspicious, such as `RETR /etc/passwd`. Along the way, excessively long commands or those containing control characters are also flagged as suspicious.

With this filter we are able to remove about 25% of the packets from the control stream, sifting out 1,059 of 4,193 packets seen during the two-day period. Our filter does not keep enough state to allow the number of sifted data packets to be calculated. Adding `MACB`<sup>†</sup> to the list of normal commands improved the sifting rate to 80% in a small sample of 274 control packets.

These results are somewhat disappointing; it seems that the characterization of normal traffic varies from one FTP client to another, so the set of traffic common across clients is very small.

### HTTP

A coarse filter for sifting HTTP traffic sifts out all `GET` or `POST` commands and responses with length under 512 bytes. This filter sifts 98% of HTTP traffic, removing 25,301 of 25,889 HTTP packets seen during the two-day period. We did not validate our assumption that packets smaller than the triggering packet length are normal, so this filter is of limited use in its present form.

### AFS

We wrote a filter that captures all fragmented IP packets and passes them to our `tcpdump` backend. As a simple application, we used this filter to detect AFS servers sending packets larger than the MTU of the destination host's network.

## EXPERIENCES

N-code has many powerful features not found in other scripting languages, such as a list structure and a `foreach` function to walk the elements of a list. This is very helpful in looking for one of a group of items in a packet, perhaps comparing a list of IP addresses to the packet's address, or

<sup>†</sup> `MACB` is used to ask an FTP server to employ the MacBinary protocol when transferring files.

looking for a command in an FTP protocol stream. Lists are also useful to store items captured out of a packet stream. Sparse arrays make a convenient associative function for storing and retrieving items.

N-code has functions for reading arbitrary bytes of a packet, and logic functions for masking and rotating those bytes. This makes it easy to examine fields of protocol headers, even those not preconfigured into NFR. For example, fragment bits of an IP header, not available as a builtin variable in NFR, are easily recovered using `byte` and logical functions.

The NFR architecture separates data selection from backend reporting and processing functions. This makes it possible to generate new backends for novel purposes.

There are some syntactic weaknesses in the language. A ":" not separated by spaces in a timeout filter causes the engine to go into an infinite loop, interpreting this as a request to read the value of a variable from another file every 0 seconds. If the characters "\_" or "-" are used in a file name, the functions it contains cannot be referenced from within another file, as these characters collide with the N-code subtraction and hierarchy operators.

Adding a selective branch construct to the N-code language would greatly improve code complexity and readability.

The decision to support implicitly declared variables in the N-code language caused us to expend a considerable amount of debugging time, as unreferenced variables caused by simple misspelling of identifiers are difficult to find. Requiring every variable to be declared would alleviate this problem.

## CONCLUSION

We found NFR and N-code to be powerful tools for writing complex scripts for detecting known attacks, and for writing sifters to help us in detecting unknown attacks.

## ACKNOWLEDGEMENTS

We thank NFR, Inc. staff members Delores Quade, Dodge Mumford, and Mike Stolarchuk for their assistance in avoiding duplication of effort, interpreting N-code semantics, fixing bugs, and helping us to get started. We also thank Marcus Ranum for appreciating our efforts. This work was partially supported by NFR, Inc.

## References

1. Marcus J. Ranum, Kent Landfield, Mike Stolarchuk, Mark Sienkiewicz, Andrew Lambeth, and Eric Wall, "Implementing A Generalized Tool For Network Monitoring," pp. 1-8 in *Eleventh Systems Administration Conference (LISA '97)*, San Diego (October 1997.).
2. Thomas H. Ptacek and Timothy N. Newsham, *Insertion, Evasion, and Denial of Service: Eluding Network Intrusion Detection*, Secure Networks, Inc. (January 1998).
3. Shawn Ostermann, `tcptrace`, Ohio University. <http://jarok.cs.ohiou.edu/software/tcptrace>

## Appendix A

### smurf.nfr

```
1#-*-c-*-
2#This filter detects packets commonly used in a smurf attack by looking for any
3#icmp echo request packet sent to the broadcast address (any address ending in
4#255). This pattern can easily be changed to reflect the broadcast address on
5#your network, and this could also be altered to look for all traffic to the
6#broadcast address from an outside network.
7
8bcastpat=regcomp("255$");
9filter smurf icmp()
10{
11#Test to see if this packet is an ICMP echo request
12  if (icmp.type == 8)
13    {
14#We are testing to see if the last byte of the address is 255, a common
15#broadcast address
16      if (regexexec(bcastpat,cat(ip.dst)))
17        {
18          echo ("Possible Smurf attack on ",ip.src, " Using ",
19              ip.src, "->" , ip.dst, "0);
20        }
21    }
22}
```

**sendmail.nfr**

```

1#*-c-*
2#This filter detects attempts to connect to the sendmail port of any machine on your
3#network that is not a sendmail server. Could also easily be adapted to find other
4#"doorknob rattling." Put all of the mail servers on you network in this list:
5smtpservers=listadd(NULL,141.211.92.141,141.211.92.197,141.211.92.199);
6filter badsendmailtcp ip()
7{
8  if(ip.proto==6)
9    {
10     if(short(ip.blob,2)==25)
11     {
12      if(!(ip.dest inside smtpservers)&&(ip.dest inside values:my_network))
13      {
14       {
15        echo("TCP connect to port 25(sendmail) attempted",
16            "(not sendmail server) ", ip.src,"->",
17            ip.dst,"0);
18       }
19      }
20     }
21    }
22}
23filter badsendmailtcpcon tcp()
24{
25  if(ip.proto==6)
26    {
27     if(short(ip.blob,2)==25)
28     {
29      if(!(ip.dest inside smtpservers)&&(ip.dest inside values:my_network))
30      {
31       if(tcp.conn)
32       {
33        echo("TCP connect to port 25(sendmail) established",
34            "(not sendmail server) ", ip.src,"->",
35            ip.dst,"0);
36       }
37       else
38       {
39        echo("TCP connect to port 25(sendmail) attempted",
40            "(not sendmail server) ", ip.src,"->",
41            ip.dst,"0);
42       }
43      }
44     }
45    }
46}
47
48filter badsendmailudp ip()
49{
50  if(ip.proto==17)
51    {

```

```
52     if(short(ip.blob,2)==25)
53     {
54         if(!(ip.dest inside smtpservers)&&(ip.dest inside values:my_network))
55         {
56             echo("UDP send to port 25(sendmail) attempted",
57                 "(not sendmail server) ", ip.src,"->",ip.dst,"0);
58         }
59     }
60 }
61 }
```



**bufferoverflow.nfr**

```

1#-*-c-*
2#detects a specific named buffer overflow targeted at linux and freebsd platforms
3#This filter detects all TCP packets sent to a designated port, then searches them
4#for a specific string or attack signature. We must search for the string,
5#since it may be at many offsets within the packet depending on the version of the
6#program under attack. This filter is configured to detect a buffer overflow attack
7#on a named which runs under linux and freebsd. By changing the string definitions
8#and the port numbers below, this can be re-configured to detect other attacks.
9
10attack_name_linux="Linux named buffer overflow";
11attack_name_freebsd="FreeBSD named buffer overflow";
12
13attack_sig_linux0=blobbytes(0x31,0xc0,0xb0,0x3f,0x31,0xdb,0xb3);
14#the 8th byte of the attack may vary, so we must not look for it.
15attack_sig_linux1=blobbytes( 0x31,0xc9,0xcd,0x80,0x31,0xc0,
16    0xb0,0x3f,0xb1,0x01,0xcd,0x80,0x31,0xc0,0xb0,0x3f,0xb1,0x02,0xcd,0x80,
17    0xeb,0x24,0x5e,0x8d,0x1e,0x89,0x5e,0x0b,0x33,0xd2,0x89,0x56,0x07,0x89,
18    0x56,0x0f,0xb8,0x1b,0x56,0x34,0x12,0x35,0x10,0x56,0x34,0x12,0x8d,0x4e,
19    0x0b,0x8b,0xd1,0xcd,0x80,0x33,0xc0,0x40,0xcd,0x80,0xe8,0xd7,0xff,0xff,
20    0xff, "/bin/sh");
21
22attack_sig_freebsd0=blobbytes(0xeb,0x6e,0x5e,0xc6,0x06,0x9a,0x31,0xc9,0x89,0x4e,
23    0x01,0xc6,0x46,0x05,0x07,0x88,0x4e,0x06,0x51,0x31,0xdb,0xb3);
24#the 23rd byte of the attack may change, so must not look for it.
25attack_sig_freebsd1=blobbytes(0x53,
26    0x66,0xc7,0x46,0x07,0xeb,0xa7,0x31,0xc0,0xb0,0x5a,0x50,0xeb,0x50,
27    0xfe,0xc1,0x51,0x53,0xc6,0x46,0x08,0xb6,0x31,0xc0,0xb0,0x5a,0x50,0xeb,
28    0x41,0xfe,0xc1,0x51,0x53,0xc6,0x46,0x08,0xc5,0x31,0xc0,0xb0,0x5a,0x50,
29    0xeb,0x32,0xc7,0x46,0x07,0x2f,0x62,0x69,0x6e,0xc7,0x46,0x0b,0x2f,0x73,
30    0x68,0x21,0x31,0xc0,0x88,0x46,0x0e,0x8d,0x5e,0x07,0x89,0x5e,0x0f,0x89,
31    0x46,0x13,0x8d,0x5e,0x13,0x53,0x8d,0x5e,0x0f,0x53,0x8d,0x5e,0x07,0x53,
32    0xb0,0x3b,0x50,0xeb,0x05,0xe8,0x8d,0xff,0xff,0xff);
33filter bufferoverflowlinux tcp()
34{
35  if(tcp.dport==53)
36    {
37      if( (index(tcp.blob,attack_sig_linux0)!=-1)
38      &&(index(tcp.blob,attack_sig_linux1)!= -1)) {
39        echo("Found ",attack_name_linux," attack code in TCP connection",ip.src,
40          ":",tcp.sourceport,"->",ip.dest, ":",tcp.destport,"0");
41      }
42    }
43}
44
45filter bufferoverflowfreebsd tcp()
46{
47  if(tcp.dport==53)
48    {
49      if( (index(tcp.blob,attack_sig_freebsd0)!= -1)
50      &&(index(tcp.blob,attack_sig_freebsd1)!= -1)) {
51        echo("Found ",attack_name_freebsd," attack code in TCP connection",ip.src,
52          ":",tcp.sourceport,"->",ip.dest, ":",tcp.destport,"0");

```

53 }  
54 }  
55 }

**teardrop.nfr**

```

1#-*-c-*
2#This filter detects a teardrop attack by re-assembling a fragmented packet
3#stream and checking for overlap of fragments within it. We keep an array
4#containing the source, destination, protocol, id, number of fragments
5#collected for that packet and the time a fragment was last received for each
6#fragment stream, and the length and offset of each fragment of each stream is
7#recorded in another array. Since the structure of a packet's information is
8#a fixed length, we simply index into the array containing packet stream
9#information using an index multiplied by the structure's length. We take
10#advantage of nfr's support of a sparse address space to index streams of
11#fragments into the fragment array based on their index into the packet array
12#left shifted 14 bits. We multiply the fragment's index into the array of
13#fragments for that fragment stream by the length of the fragment structure
14#and OR this with the left shifted index into the packet array to obtain the
15#index into the fragment array. This gives each packet a space for as many
16#fragments as are allowed in the offset field of an IP header. This also
17#leaves the potential for an overflow attack if over 2^13 fragments of a
18#particular stream are directed at the filter.
19
20#Once these fragments are collected, a timeout trigger periodically sorts the
21#fragment array, then walks the sorted array and checks adjacent packets for
22#overlap. If overlap is found, it is reported. After a given pair of
23#packets is compared, the first one is marked as having been used in a
24#comparison, and if an overlap is found where both packets have already been
25#used in a comparison, this is not reported, because it was previously
26#reported. We currently only check linearly and do no checking within groups
27#of packets with the same offset, so some attack on the monitor is
28#possible here. Of course packets with the same offset overlap, so these are
29#reported. If the packet stream has not received a packet within a given
30#time period, it is deleted and the last packet stream in the array is moved
31#into its place.
32
33#Obviously this filter has an engine to collect and re-assemble packets, and
34#it could easily be modified to check for other anomalies between fragments
35#or assemble based on different rules.
36
37
38pkt_struct_len=6;
39frag_struct_len=3;
40pkt_count=0;
41stale_timeout=120;
42last_time=0;
43
44filter teardrop ip()
45{
46
47 $inlist=0;
48
49 $flags=byte(eth.blob,6)>>5;
50 $offset=short(eth.blob,6) & 0x1fff;
51
52

```

```

53#we have a fragmented paket
54 if(((($flags & 2)!=2)&&!(((($flags & 1)!=1)&&($offset==0))))
55 {
56     # echo("Fragmented Packet Detected0);
57     #display("verbose");
58     $tlen=(short(eth.blob,2) );
59     $hlen=4*(byte(eth.blob,0)&0x0f);
60     $ipid=short(eth.blob,4);
61     $offlag=short(eth.blob,6);
62     $len=$tlen-$hlen;
63     #echo("offlag=", $offlag, " len=", $len, "0);
64#search packlist for packet
65     $current=0;
66     $inlist=0;
67     while($current<pkt_count)
68     {
69         $pktoff=$current*pkt_struct_len;
70         if((packlist[$pktoff]==$ipid)&&(packlist[$pktoff+1]==ip.src)
71             &&(packlist[$pktoff+2]==ip.dst)&&(packlist[$pktoff+3]==ip.proto))
72         {
73#if found add frag
74             $inlist=1;
75             #fragment=fraglist[$current];
76             fraglist[(($current<<14)+
77                 (packlist[$pktoff+5]*frag_struct_len)]=$offlag;
78             fraglist[(($current<<14)+
79                 ((packlist[$pktoff+5]*frag_struct_len)+1)=$len;
80             fraglist[(($current<<14)+
81                 ((packlist[$pktoff+5]*frag_struct_len)+2)]=0;
82             packlist[$pktoff+5]=packlist[$pktoff+5]+1;
83             break;
84         }
85         $current=$current+1;
86     }
87
88     if($inlist != 1)
89     {
90#if not found, add packet
91#add new packet
92     $pktoff=pkt_count*pkt_struct_len;
93     packlist[$pktoff]=$ipid;
94     packlist[$pktoff+1]=ip.src;
95     packlist[$pktoff+2]=ip.dst;
96     packlist[$pktoff+3]=ip.proto;
97     packlist[$pktoff+4]=system.time;
98     packlist[$pktoff+5]=1;           #;//size of fragment list array
99     fraglist[(pkt_count<<14)]=$offlag;
100    fraglist[(pkt_count<<14)+1]=$len;
101    fraglist[(pkt_count<<14)+2]=0;
102    pkt_count=pkt_count+1;
103    }
104 }
105
106 }

```

```

107
108
109 filter checkoverlap timeout(sec : 20, repeat)
110 {
111   #echo("fraglist=",fraglist,"0");
112
113 #sort and re-assemble
114   $index=0;
115   while ($index<pkt_count)
116     {
117       $k=0;
118       while($k<packlist[((($index*pkt_struct_len)+5)])
119         {
120           $l=0;
121           while($l<frag_struct_len)
122             {
123               #   echo("fraglist[" ,($index<<14)+($k*frag_struct_len)+$l,"]=" ,
124                 #   fraglist[($index<<14)+($k*frag_struct_len)+$l], "0");
125                 $l=$l+1;
126             }
127             $k=$k+1;
128         }
129 #sort
130   # $frag=fraglist[$index];
131   $fragind=$index<<14;
132   $size=((packlist[((($index*pkt_struct_len)+5)]) * frag_struct_len);
133   $current=0;
134   while($current<$size)
135     {
136       $i=$current;
137       while($i<$size)
138         {
139           if((fraglist[$i+$fragind]&0x1fff)<
140             (fraglist[$current+$fragind]&0x1fff))
141             {
142               $j=0;
143               while($j<frag_struct_len)
144                 {
145                   $temp=fraglist[$i+$j+$fragind];
146                   fraglist[$i+$j+$fragind]=fraglist[$current+$j+$fragind];
147                   fraglist[$current+$j+$fragind]=$temp;
148                   $j=$j+1;
149                 }
150             }
151             $i=$i+frag_struct_len;
152         }
153
154       $current=$current+frag_struct_len;
155     }
156 #*****Do we also have to sort by length?
157 #walk and check for overlap and duplicates
158   $current=frag_struct_len;
159   #echo("fraglist=",fraglist,"0");

```

```

160     while($current<$size)
161     {
162         $curoff=(fraglist[$current+$fragind]&0x1fff)*8;
163         $prevoff=(fraglist[$current+$fragind-frag_struct_len]&0x1fff)*8;
164         $curlen=fraglist[$current+$fragind+1];
165         $prevlen=fraglist[$current+$fragind-frag_struct_len+1];
166         $curalert=fraglist[$current+$fragind+2];
167         $prevalert=fraglist[$current+$fragind-frag_struct_len+2];
168         #echo("curoff",$curoff," prevoff",$prevoff," curlen",$curlen,
169         #     " prevlen",$prevlen," curalert",$curalert," prevalert",
170         #     $prevalert,"0);
171         if(($prevoff+$prevlen)>$curoff)
172             &&(($curalert==0)||($prevalert==0))
173 #overlapping packets print src, dst, both offsets and lengths
174     {
175         echo(system.time,": Overlapping packet fragments detected ",
176             "from ",packlist[($index*pkt_struct_len)+1], " to ",
177             packlist[($index*pkt_struct_len)+2]," ID:",
178             packlist[$index*pkt_struct_len]," Protocol ",
179             packlist[($index*pkt_struct_len)+3],
180             " first fragment offset:",
181             (8*(fraglist[$current-frag_struct_len+$fragind]&0x1fff)),
182             " len:",
183             fraglist[$current+1-frag_struct_len+$fragind],
184             " second fragment offset:",
185             8*(fraglist[$current+$fragind]&0x1fff),
186             " len:",fraglist[$current+1+$fragind],"0);
187         fraglist[$current+2+$fragind-frag_struct_len]=1;
188         if(($current+frag_struct_len)==$size)
189         {
190             fraglist[$current+2+$fragind]=1;
191         }
192     }
193     $current=$current+frag_struct_len;
194 }
195
196 if((packlist[($index*pkt_struct_len)+4])<(system.time-stale_timeout)
197 &&((packlist[($index*pkt_struct_len)+4])!=NULL))
198 {
199     echo("Deleting packet index ",$index ,packlist[($index*pkt_struct_len)+4],"0);
200
201
202     $pktoff=(pkt_count-1)*pkt_struct_len;
203     $newpktoff=$index*pkt_struct_len;
204 #this will have to move the elements now
205 #delete elements to be deleted, then move and delete elements moved into
206 #that space
207     $i=0;
208     while($i<packlist[$newpktoff+5])
209     {
210         $j=0;
211         while($j<frag_struct_len)
212         {
213             fraglist[($fragind+($i*frag_struct_len))+$j]=NULL;

```

```

214         $j=$j+1;
215     }
216     $i=$i+1;
217 }
218 if(pkt_count>1)
219 {
220     $i=0;
221     while($i<packlist[$pktoff+5])
222     {
223         $j=0;
224         while($j<frag_struct_len)
225         {
226             fraglist[(($fragind+($i*frag_struct_len))+$j)
227 =fraglist[((pkt_count-1)<<14)
228             +($i*frag_struct_len))+$j];
229             fraglist[((pkt_count-1)<<14)
230                 +($i*frag_struct_len))+$j]=NULL;
231             $j=$j+1;
232         }
233         $i=$i+1;
234     }
235
236     packlist[$newpktoff]=packlist[$pktoff];
237     packlist[$newpktoff+1]=packlist[$pktoff+1];
238     packlist[$newpktoff+2]=packlist[$pktoff+2];
239     packlist[$newpktoff+3]=packlist[$pktoff+3];
240     packlist[$newpktoff+4]=packlist[$pktoff+4];
241     packlist[$newpktoff+5]=packlist[$pktoff+5];
242 }
243 packlist[$pktoff]=NULL;
244 packlist[$pktoff+1]=NULL;
245 packlist[$pktoff+2]=NULL;
246 packlist[$pktoff+3]=NULL;
247 packlist[$pktoff+4]=NULL;
248 packlist[$pktoff+5]=NULL;
249
250
251     pkt_count=pkt_count-1;
252 }
253 $index=$index+1;
254 }
255
256 }

```

**profile.nfr**

```
1 smtp=0;
2 ftp=0;
3 telnet=0;
4 ssh=0;
5 http_t=0;
6 http_u=0;
7 nntp=0;
8 ntp_t=0;
9 ntp_u=0;
10 ping=0;
11 dns_u=0;
12 dns_t=0;
13 locus_map_t=0;
14 locus_conn_t=0;
15 netbios_datagram_t=0;
16 netbios_name_t=0;
17 locus_map_u=0;
18 locus_conn_u=0;
19 netbios_datagram_u=0;
20 netbios_name_u=0;
21 rip=0;
22 xwin_t=0;
23 xwin_u=0;
24 afs_t=0;
25 afs_u=0;
26 dec_notes_t=0;
27 print_spool_t=0;
28 dec_notes_u=0;
29 print_spool_u=0;
30 p1865u=0;
31 p1046u=0;
32 kerb_iv_t=0;
33 netbios_session_t=0;
34 kerb_iv_u=0;
35 netbios_session_u=0;
36 igmp=0;
37 sql_t=0;
38 rpcbind_t=0;
39 sql_u=0;
40 rpcbind_u=0;
41 login=0;
42 whod=0;
43 pop3_t=0;
44 pop3_u=0;
45 unknown_tcp=listadd(NULL,NULL);
46 unknown_udp=listadd(NULL,NULL);
47
48 last_time=0;
49 last_report=0;
50
51 filter countdata ip()
52 {
```



```

53  if(ip.proto==6)##;tcp
54  {
55      $sport=short(ip.blob,0);
56      $dport=short(ip.blob,2);
57      if(($sport==25)||($dport==25))
58      {
59          smtp=smtp+1;
60          last_time=system.time;
61      }
62      else
63      if(($sport==21)||($dport==21))
64      {
65          ftp=ftp+1;
66          last_time=system.time;
67      }
68      else
69      if(($sport==23)||($dport==23))
70      {
71          telnet=telnet+1;
72          last_time=system.time;
73      }
74      else
75      if(($sport==22)||($dport==22))
76      {
77          ssh=ssh+1;
78          last_time=system.time;
79      }
80      else
81      if(($sport==80)||($dport==80))
82      {
83          http_t=http_t+1;
84          last_time=system.time;
85      }
86      else
87      if(($sport==119)||($dport==119))
88      {
89          nntp=nntp+1;
90          last_time=system.time;
91      }
92      else
93      if(($sport==123)||($dport==123))
94      {
95          ntp_t=ntp_t+1;
96          last_time=system.time;
97      }
98      else
99      if(($sport==53)||($dport==53))
100     {
101         dns_t=dns_t+1;
102         last_time=system.time;
103     }
104     else
105     if(($sport==127)||($dport==127))
106     {

```

```
107         locus_map_t=locus_map_t+1;
108         last_time=system.time;
109     }else
110 if (($sport==125)||($dport==125))
111     {
112         locus_conn_t=locus_conn_t+1;
113         last_time=system.time;
114     }else
115 if (($sport==138)||($dport==138))
116     {
117         netbios_datagram_t=netbios_datagram_t+1;
118         last_time=system.time;
119     }else
120 if (($sport==137)||($dport==137))
121     {
122         netbios_name_t=netbios_name_t+1;
123         last_time=system.time;
124     }
125
126     else
127 if (($sport==3333)||($dport==3333))
128     {
129         dec_notes_t=dec_notes_t+1;
130         last_time=system.time;
131     }
132     else
133 if (($sport==515)||($dport==515))
134     {
135         print_spool_t=print_spool_t+1;
136         last_time=system.time;
137     }
138     else
139 if (($sport==139)||($dport==139))
140     {
141         netbios_session_t=netbios_session_t+1;
142         last_time=system.time;
143     }
144     else
145 if (($sport==750)||($dport==750))
146     {
147         kerb_iv_t=kerb_iv_t+1;
148         last_time=system.time;
149     }
150     else
151 if (($sport==1498)||($dport==1498))
152     {
153         sql_t=sql_t+1;
154         last_time=system.time;
155     } else
156 if (($sport==111)||($dport==111))
157     {
158         rpcbind_t=rpcbind_t+1;
159         last_time=system.time;
160     } else
```

```

161         if(($sport==513)||($dport==513))
162         {
163             login=login+1;
164             last_time=system.time;
165         } else
166         if(($sport==110)||($dport==110))
167         {
168             pop3_t=pop3_t+1;
169             last_time=system.time;
170         }
171         else
172             if(((($sport>5999)&&($sport<6064))
173             ||(($dport>5999)&&($dport<6064)))
174             {
175                 xwin_t=xwin_t+1;
176                 last_time=system.time;
177             }
178         else
179             if(((($sport>6999)&&($sport<7010))
180             ||(($dport>6999)&&($dport<7010)))
181             {
182                 afs_t=afs_t+1;
183                 last_time=system.time;
184             }
185
186         else
187         {
188             if(!($sport inside unknown_tcp)&&($sport<10000))
189             {
190                 echo ("Unrecognized protocol in TCP header",
191                 $sport,"->", $dport, " nonce ",
192                 packet.nonce,"0");
193                 unknown_tcp=listadd(unknown_tcp,$sport);
194             }
195             if(!($dport inside unknown_tcp)&&($dport<10000))
196             {
197                 echo ("Unrecognized protocol in TCP header",
198                 $sport,"->", $dport, " nonce ",
199                 packet.nonce,"0");
200                 unknown_tcp=listadd(unknown_tcp,$dport);
201             }
202
203
204 #display("verbose");
205         }
206     }
207     else
208         if(ip.proto==17)#;udp
209         {
210             $sport=short(ip.blob,0);
211             $dport=short(ip.blob,2);
212             if(($sport==80)||($dport==80))
213             {
214                 http_u=http_u+1;

```

```
215     last_time=system.time;
216 }
217 else
218 if(($sport==123)||($dport==123))
219 {
220     ntp_u=ntp_u+1;
221     last_time=system.time;
222 }
223 else
224 if(($sport==53)||($dport==53))
225 {
226     dns_u=dns_u+1;
227     last_time=system.time;
228 }
229 else
230 if(($sport==127)||($dport==127))
231 {
232     locus_map_u=locus_map_u+1;
233     last_time=system.time;
234 }
235 else
236 if(($sport==125)||($dport==125))
237 {
238     locus_conn_u=locus_conn_u+1;
239     last_time=system.time;
240 }
241 else
242 if(($sport==138)||($dport==138))
243 {
244     netbios_datagram_u=netbios_datagram_u+1;
245     last_time=system.time;
246 }
247 else
248 if(($sport==137)||($dport==137))
249 {
250     netbios_name_u=netbios_name_u+1;
251     last_time=system.time;
252 }
253 else
254 if(($sport==520)||($dport==520))
255 {
256     rip=rip+1;
257     last_time=system.time;
258 }
259 else
260 if(($sport==3333)||($dport==3333))
261 {
262     dec_notes_u=dec_notes_u+1;
263     last_time=system.time;
264 }
265 else
266 if(($sport==515)||($dport==515))
267 {
268     print_spool_u=print_spool_u+1;
```

```

269         last_time=system.time;
270     }
271     else
272     if (($sport==1865)||($dport==1865))
273     {
274         p1865u=p1865u+1;
275         last_time=system.time;
276     }
277     else
278     if (($sport==1046)||($dport==1046))
279     {
280         p1046u=p1046u+1;
281         last_time=system.time;
282     }
283     else
284     if (($sport==139)||($dport==139))
285     {
286         netbios_session_u=netbios_session_u+1;
287         last_time=system.time;
288     }
289     else
290     if (($sport==750)||($dport==750))
291     {
292         kerb_iv_u=kerb_iv_u+1;
293         last_time=system.time;
294     }
295     else
296     if (($sport==1498)||($dport==1498))
297     {
298         sql_u=sql_u+1;
299         last_time=system.time;
300     } else
301     if (($sport==111)||($dport==111))
302     {
303         rpcbind_u=rpcbind_u+1;
304         last_time=system.time;
305     } else
306     if (($sport==513)||($dport==513))
307     {
308         whod=whod+1;
309         last_time=system.time;
310     } else
311     if (($sport==110)||($dport==110))
312     {
313         pop3_u=pop3_u+1;
314         last_time=system.time;
315     }
316     else
317     if ((( $sport>5999)&&($sport<6064))
318         ||(( $dport>5999)&&($dport<6064)))
319     {
320         xwin_u=xwin_u+1;
321         last_time=system.time;
322     }

```

```

323         else
324         if(($sport>6999)&&($sport<7010))
325             ||(($dport>6999)&&($dport<7010))
326             {
327             afs_u=afs_u+1;
328             last_time=system.time;
329             }
330
331
332
333
334         else
335         {
336         if(!($sport inside unknown_udp)&&($sport<10000))
337             {
338             echo ("Unrecognized protocol in UDP header ",
339                 $sport,"->", $dport, " nonce ", packet.nonce,
340                 "0");
341             unknown_udp=listadd(unknown_udp,$sport);
342             }
343         if(!($dport inside unknown_udp)&&($dport<10000))
344             {
345             echo ("Unrecognized protocol in UDP header ",
346                 $sport,"->", $dport, " nonce ",
347                 packet.nonce,"0");
348 #display("verbose");
349             unknown_udp=listadd(unknown_udp,$dport);
350             }
351
352 #we need to add these to a list, and only report once
353             }
354     }
355     else
356     if(ip.proto==1)#;icmp
357     {
358     $type=byte(ip.blob,0);
359     if(($type=0)||($type=8))
360     {
361     ping=ping+1;
362     last_time=system.time;
363     }
364     else
365     {
366     echo ("Unrecognized type in ICMP header ", $type,
367         " nonce ", packet.nonce,"0");
368     #display("verbose");
369     }
370     }
371     else
372     if(ip.proto==2)
373     {
374     igmp=igmp+1;
375     last_time=system.time;
376     }

```

```

377     else
378     {
379         echo ("Unrecognized protocol in IP header ",ip.proto,
380             "nonce ",packet.nonce,"0");
381         #display("verbose");
382     }
383
384 }
385
386 filter output timeout(sec : 120,repeat)
387 {
388     if(last_time>last_report)
389     {
390         last_report=system.time;
391         echo("smtp=",smtp," ftp=",ftp," telnet=",telnet," ssh=",ssh," http_t=",
392             http_t," http_u=",http_u," nntp=",nntp," ntp_t=",ntp_t," ntp_u=",
393             ntp_u," ping=",ping," dns_t=",dns_t," dns_u=",dns_u,
394             " locus_map_t=",locus_map_t,
395             " locus_conn_t=",locus_conn_t,
396             " netbios_datagram_t=",netbios_datagram_t,
397             " netbios_name_t=",netbios_name_t,
398             " locus_map_u=",locus_map_u,
399             " locus_conn_u=",locus_conn_u,
400             " netbios_datagram_u=",netbios_datagram_u,
401             " netbios_name_u=",netbios_name_u," rip=",rip,
402             " xwin_t=",xwin_t,
403             " xwin_u=",xwin_u,
404             " afs_t=",afs_t,
405             " afs_u=",afs_u,
406             " dec_notes_t=",dec_notes_t,
407             " print_spool_t=",print_spool_t,
408             " dec_notes_u=",dec_notes_u,
409             " print_spool_u=",print_spool_u,
410             " p1865u=",p1865u,
411             " p1046u=",p1046u,
412             " kerb_iv_t=",kerb_iv_t,
413             " netbios_session_t=",netbios_session_t,
414             " kerb_iv_u=",kerb_iv_u,
415             " netbios_session_u=",netbios_session_u,
416             " IGMP=",igmp,
417             " sql_t=",sql_t,
418             " rpcbind_t=",rpcbind_t,
419             " sql_t=",sql_t,
420             " rpcbind_t=",rpcbind_t,
421             " login=",login,
422             " whod=",whod,
423             " pop3_t=",pop3_t,
424             " pop3_u=",pop3_u,
425
426
427         "0);
428     }
429 }

```

## Appendix B

## output.c

```

1/*
2This backend logs packets sent to it one at a time to a tcpdump/pcap format file. It takes
3as argument the file name to write to, and expects input as produced by:
4
5record blobbytes(cat(123456),32,cat(packet.sec),32,cat(packet.usec),32,cat(packet.len),32,
6blobbytes(packet.blob)) to test_recorder;
7*/
8
9#include <stdio.h>
10#include <sys/types.h>
11#include <sys/time.h>
12#include <unistd.h>
13#include <fcntl.h>
14#include <signal.h>
15#include <pcap.h>
16#include <stdlib.h>
17#include <limits.h>
18#define TCPDUMP_MAGIC 0xa1b2c3d4
19int output;
20int errlog;
21char * signal_sucks;
22void cleanup()
23{
24 close(output);
25 close(errlog);
26 /*printf("exiting
27 exit(0);
28 }
29void crashed()
30{
31 close(output);
32 close(errlog);
33 printf("exiting
34 abort();
35 }
36main(int argc,char **argv)
37{
38 char * a,*b;
39 int size=1;
40 int * data_size;
41 int * type;
42 int a0;
43 char foo;
44 struct pcap_file_header hdr;
45 struct pcap_pkthdr pkt_hdr;
46 signal_sucks=argv[0];

```



```

47  signal(SIGTERM, (void*)&cleanup);
48  signal(SIGSEGV, (void*)&crashed);
49  a=(char*)malloc(1560);
50  b=a;
51  data_size=malloc(sizeof(int));
52  type=malloc(sizeof(int));
53  hdr.magic = TCPDUMP_MAGIC;
54  hdr.version_major = PCAP_VERSION_MAJOR;
55  hdr.version_minor = PCAP_VERSION_MINOR;
56  /*this should be set to the time zone*/
57  hdr.thiszone = -4;
58  hdr.snaplen = 1500;
59  hdr.sigfigs = 0;
60  /*this should be set to the real link type*/
61  hdr.linktype = DLT_EN10MB;
62
63  /*This really should do some error checking.*/
64  if(argc>2)
65      {
66          printf("Usage
67              exit (1);
68          }
69  if(argc==1)
70      output=1;
71  else
72      output=open(argv[1],O_WRONLY|O_TRUNC|O_CREAT,0x1ff);
73  if(output<1)
74      {
75          printf("Error opening file
76              exit(1);
77          }
78  write(output,&hdr,sizeof(hdr));
79  while (1)
80      {
81          a0=fread((char*)type,1,4,stdin);
82          iffeof(stdin))
83              {
84                  cleanup();
85              }
86          a0=fread((char*)data_size,1,4,stdin);
87          iffeof(stdin))
88              {
89                  cleanup();
90              }
91          if(*type==1)
92              {
93              /*  a0=*data_size;
94              if(*data_size>1500)
95                  printf("data_size=%d0,*data_size);
96              if(a0!=*data_size)
97                  {
98                  printf("data_size=%x a0=%h",*data_size,a0);
99                  crashed();
100              }*/

```

```

101     a=b;
102     size=fread(a,1,*data_size,stdin);
103     if(!feof(stdin))
104     {
105         cleanup();
106     }
107     if(strtol(a,&a,10)==123456)
108     {
109         pkt_hdr.ts.tv_sec=strtol(a,&a,10);
110         pkt_hdr.ts.tv_usec=strtol(a,&a,10);
111         pkt_hdr.len=strtol(a,&a,10);
112         a++;
113         pkt_hdr.caplen=pkt_hdr.len;
114         /*if(size<(pkt_hdr.len+20))
115         {
116             printf("Read in file size=%x[a[0-3]]=%d %d %d %d0,size,a[0],a[1],a[2],a[3]);
117             printf("pkt_hdr.ts.tv_sec=%lx pkt_hdr.ts.tv_usec=%lx pkt_hdr.len=%lx 0,pkt
_hdr.ts.tv_sec, pkt_hdr.ts.tv_usec, pkt_hdr.len);
118         }
119         */
120         write(output,&pkt_hdr,sizeof(pkt_hdr));
121         write(output,a,pkt_hdr.len);
122     }
123     else
124     {
125         errlog=open("output_err.log",O_WRONLY|O_CREAT,0x1ff);
126         write(errlog,a+7,size-7);
127         printf("Read in file size=%x0,size);
128         close (errlog);
129     }
130 }
131 }
132
133 }

```

## encap\_out.c

```

1 /*
2 This backend logs packets sent to it in blocks to a tcdump/pcap format file. It takes
3 as argument the file name to write to and expects input as produced by:
4 $data_blob=blobbytes(NULL);
5 #The next 2 statements are repeated to add packets to the block to be dumped(or discarded.)
6 This scheme is useful for holding on to data that you might want to log but won't know if
7 you do until some time after the packets arrive.
8 $dump_blob=blobbytes(cat(123456),32,cat(packet.sec),32,cat(packet.usec),32,cat(packet.len),
9 32,blobbytes(packet.blob));
10 $data_blob=blobbytes($data_blob,cat(strlen($dump_blob)),32,$dump_blob);
11 record $data_blob to test_recorder;
12 */
13 #include <stdio.h>
14 #include <sys/types.h>
15 #include <sys/time.h>
16 #include <unistd.h>
17 #include <fcntl.h>
18 #include <signal.h>
19 #include <pcap.h>
20 #include <stdlib.h>
21 #include <limits.h>
22 #define TCPDUMP_MAGIC 0xa1b2c3d4
23 int output;
24 int errlog;
25 char * signal_sucks;
26 void cleanup()
27 {
28   close(output);
29   close(errlog);
30   /*printf("exiting
31   exit(0);
32 }
33 void crashed()
34 {
35   close(output);
36   close(errlog);
37   printf("exiting
38   abort();
39 }
40 main(int argc, char **argv)
41 {
42   char * a,*b;
43   int size=1;
44   int * data_size;
45   int * type;
46   int a0,i,total_read,x,blob_size;
47   char foo;

```

```

48 struct pcap_file_header hdr;
49 struct pcap_pkthdr pkt_hdr;
50 signal_sucks=argv[0];
51 signal(SIGTERM,(void*)&cleanup);
52 signal(SIGSEGV,(void*)&crashed);
53 a=(char*)malloc(1560);
54 b=a;
55 data_size=malloc(sizeof(int));
56 type=malloc(sizeof(int));
57 hdr.magic = TCPDUMP_MAGIC;
58 hdr.version_major = PCAP_VERSION_MAJOR;
59 hdr.version_minor = PCAP_VERSION_MINOR;
60 /*this should be set to the time zone*/
61 hdr.thiszone = -4;
62 hdr.snaplen = 1500;
63 hdr.sigfigs = 0;
64 /*this should be set to the real link type*/
65 hdr.linktype = DLT_EN10MB;
66
67 /*This really should do some error checking.*/
68 if(argc>2)
69     {
70     printf("Usage
71     exit (1);
72     }
73 if(argc==1)
74     output=1;
75 else
76     output=open(argv[1],O_WRONLY|O_TRUNC|O_CREAT,0x1ff);
77 if(output<1)
78     {
79     printf("Error opening file
80     exit(1);
81     }
82 write(output,&hdr,sizeof(hdr));
83 while (1)
84     {
85     a0=fread((char*)type,1,4,stdin);
86     if(feof(stdin))
87         {
88         cleanup();
89         }
90     a0=fread((char*)data_size,1,4,stdin);
91     if(feof(stdin))
92         {
93         cleanup();
94         }
95     total_read=0;
96     if(*type==1)
97         {
98         blob_size=*data_size;
99         while(total_read<blob_size)
100         {
101         a0=0;

```

```

102     i=0;
103     a++;
104     a[-1]=0;
105     while(a[i-1]!=32)
106     {
107         x=fread(a+i,1,1,stdin);
108         a0+=x;
109         total_read+=x;
110         if(feof(stdin))
111         {
112             cleanup();
113         }
114         i++;
115     }
116     *data_size=strtol(a,NULL,10);
117
118     /*  a0=*data_size;
119         if(*data_size>1500
120         printf("data_size=%d0,*data_size);
121         if(a0!=*data_size)
122         {
123             printf("*data_size=%x a0=%h",*data_size,a0);
124             crashed();
125         }*/
126     a=b;
127     size=fread(a,1,*data_size,stdin);
128     total_read+=size;
129     if(feof(stdin))
130     {
131         cleanup();
132     }
133     if(strtol(a,&a,10)==123456)
134     {
135         pkt_hdr.ts.tv_sec=strtol(a,&a,10);
136         pkt_hdr.ts.tv_usec=strtol(a,&a,10);
137         pkt_hdr.len=strtol(a,&a,10);
138         a++;
139         pkt_hdr.caplen=pkt_hdr.len;
140         /*if(size<(pkt_hdr.len+20))
141         {
142             printf("Read in file size=%xa[0-3]=%d %d %d %d0,size,a[0],a[1],a[2
],a[3]);
143             printf("pkt_hdr.ts.tv_sec=%lx pkt_hdr.ts.tv_usec=%lx pkt_hdr.len=%lx
0,pkt_hdr.ts.tv_sec, pkt_hdr.ts.tv_usec, pkt_hdr.len);
144         }
145         */
146         write(output,&pkt_hdr,sizeof(pkt_hdr));
147         write(output,a,pkt_hdr.len);
148     }
149     else
150     {
151         errlog=open("output_err.log",O_WRONLY|O_CREAT,0x1ff);
152         write(errlog,a+7,size-7);
153         printf("Read in file size=%x0,size);

```

```
154         close (errlog);
155     }
156 }
157 }
158 }
159
160 }
```

**alldns.nfr**

1#This is a filter to reject suspected good DNS packets and log all others to the file rawbinary.bin. It does this with the help of the tcpdump format file writing backend, output. Requests and replies for A, PTR and MX records are checked for names that are too long, names with too many elements, and odd options or excessive numbers of records and if any of these conditions are found, they are logged. All other DNS packets are also logged.

```

2
3
4#currently removes about 90%(62322 removed, 5385 logged) of the packet stream. Could do
better by applying some rules to eliminate known benign zone transfer and "all" transfer traffic,
as well as SOA records.
5
6dns_types=[1,12,15];#;,2,6,12,15,16,33];#252,255);
7dns_type=[0];
8total_qcounts=0;
9max_name_char=64;
10max_name_labels=8;
11pkts_logged=0;
12pkts_sifted=0;
13tcp_nonce_list=listadd(NULL);
14
15test_schema=library_schema:new(1,["str"],scope());
16test_recorder=recorder("bin/output rawbinary.bin","test_schema");
17func tcpdump_packet
18{
19#echo("in tcpdump_packet0);
20record blobbytes(cat(123456),32,cat(packet.sec),32,cat(packet.usec),32,cat(packet.len),32,bl
obbytes(packet.blob))
21    to test_recorder;
22}
23
24func dns_name
25{
26####this dose not currenty handle compressed names
27#this function takes the offset of a dns format name from the begining of the
28#data portion of the ip packet, and returns an array with element 0 a string
29#which is the name passed to it, and elemnt 1 being the length of the total
30#name in the packet.
31$offset=$1;
32$y[0]=" ";
33$y[2]=0;
34$labels=0;
35if($2==6)
36    $x=byte(tcp.blob,$offset);
37    if($2==17)
38    $x=byte(udp.blob,$offset);
39
40# echo("$x=", $x, " $offset=", $offset, "0);
41    $offset=$offset+1;
42    $y[1]=1;
43    while($x>0)
44{
45 if(($x>max_name_char))

```

```

46  {
47    $y[2]="label too long";
48  }
49  $labels=$labels+1;
50  if($labels>max_name_labels)
51  {
52    $y[2]="too many labels";
53  }
54  if($labels>50)
55  {
56    break;
57  }
58  if($2==6)
59    $y[0]=cat($y[0],substr(tcp.blob,$offset,$x));
60
61  if($2==17)
62    $y[0]=cat($y[0],substr(udp.blob,$offset,$x));
63  $y[1]=$y[1]+$x;
64  $offset=$offset+$x;
65  if($2==6)
66    $x=byte(tcp.blob,$offset);
67  if($2==17)
68    $x=byte(udp.blob,$offset);
69  $offset=$offset+1;
70  $y[1]=$y[1]+1;
71  if($x!=0)
72    $y[0]=cat($y[0],".");
73 }
74   return $y;
75 }
76
77filter dns_tcp tcp(port: 53)
78 {
79 # display("verbose");
80 #we want to check for a query/response size larger than the given query or
81 #response. It would involve a lot of overhead to check this on multiple
82 #packet tcp streams, so we will only check if it is a single packet query
83 #or response.
84
85  $dns_size=short(tcp.blob,0);
86  if(($dns_size>(packet.len-34))&&(packet.len<1500))
87    $weird_size=1;
88  else
89    $weird_size=0;
90  $id=short(tcp.blob,0+2);
91  $qr=(byte(tcp.blob,2+2)>>7)&0x1;
92  $opcode=(byte(tcp.blob,2+2)>>3)&0xf;
93  if($opcode==0) $opcode="QUERY(0)";
94  if($opcode==1) $opcode="IQUERY(1)";
95  if($opcode==2) $opcode="STATUS(2)";
96  $aa=(byte(tcp.blob,2+2)>>2)&0x1;
97  $tc=(byte(tcp.blob,2+2)>>1)&0x1;
98  $rd=(byte(tcp.blob,2+2))&0x1;
99  $ra=(byte(tcp.blob,3+2)>>7)&0x1;

```



```

100 $z=(byte(tcp.blob,3+2)>>4)&0x7;
101 $rcode=(byte(tcp.blob,3+2)&0xf;
102 $qdcnt=short(tcp.blob,4+2);
103 $ancnt=short(tcp.blob,6+2);
104 $nscnt=short(tcp.blob,8+2);
105 $arcnt=short(tcp.blob,10+2);
106
107 if(!(tcp.connhash inside tcp_nonce_list))
108 {
109     tcp_nonce_list=listadd(tcp_nonce_list,tcp.connhash);
110     $a[0]=0;
111     $a[1]=system.time;
112     $a[2]=0;
113     tcp_nonce[tcp.connhash]=$a;
114 }
115 else
116 {
117     $a=tcp_nonce[tcp.connhash];
118     $a[0]=$a[0]+1;
119     $a[1]=system.time;
120     tcp_nonce[tcp.connhash]=$a;
121 }
122
123 #we have to make sure this is the first in a stream before trying to decode it
124 if(($a[0]==0&&$qr==0)||($a[0]==1&&$qr==1))
125 {
126     $y=dns_name(12+2,6);
127     $qtype=short(tcp.blob,12+2+$y[1]);
128     $qclass=short(tcp.blob,12+2+$y[1]+2);
129     $header=1;
130 }
131 else
132 {
133     $header=0;
134     $y[2]=$a[2];
135 }
136
137 if(((1)||($qdcnt!=1)||($y[2]!=0)||($qclass!=1)||(strcasecmp($opcode,"QUERY(0)"))||($y[1
]>256)||$weird_size||$a[2])&&$header)||$a[2])
138 {
139     #record packet.blob to test_recorder;
140     #record "test" to test_recorder;
141     #display("verbose");
142     tcpdump_packet();
143     pkts_logged=pkts_logged+1;
144     $a=tcp_nonce[tcp.connhash];
145     $a[2]=1;
146     tcp_nonce[tcp.connhash]=$a;
147     if(header==1) #This assumes the query and first response are in order
148     {
149         echo("TCP DNS packet ",ip.src,":",tcp.sport,"->",ip.dst,":",tcp.dport," ID=", $id,"
Q/R=", $qr," Opcode=", $opcode," AA=", $aa," TC=", $tc," RD=", $rd," RA=", $ra," Z=", $z," RCODE=", $rcode,
" QDCOUNT=", $qdcnt," ANCOUNT=", $ancnt," NSCOUNT=", $nscnt," ARCOUNT=", $arcnt," $y[2]=",$y[2])
;
150

```

```

151     $qcount=0;
152     $total_length=0;
153     while($qcount<$qdcnt)
154     {
155 #print query contents
156     $y=dns_name(12+$total_length+2,6);
157     $qname=$y[0];
158     $qtype=short(tcp.blob,12+2+$y[1]+$total_length);
159     $qclass=short(tcp.blob,12+2+$y[1]+$total_length+2);
160     if((!($qtype inside dns_types))||($qdcnt!=1)||($y[2]!=0)||($qclass!=1)||(str
casecmp($opcode,"QUERY(0)")||($y[1]>256))
161     echo(" Query ",$qcount," QNAME=", $qname," QTYPE=", $qtype," QCLASS=", $qclass)
;
162     if(total_qcounts[$qtype]==NULL)
163     total_qcounts[$qtype]=1;
164     else
165     total_qcounts[$qtype]=total_qcounts[$qtype]+1;
166     $total_length=$total_length+$y[1];
167     $qcount=$qcount+1;
168     }
169     if((!($qtype inside dns_types))||($qdcnt!=1)||($y[2]!=0)||($qclass!=1)||(strcase
cmp($opcode,"QUERY(0)")||($y[1]>256))
170     echo("0");
171     }
172     else
173     {
174     pkts_logged=pkts_logged+1;
175     }
176     }
177
178 else
179     {
180     pkts_sifted=pkts_sifted+1;
181     }
182 }
183
184
185
186 filter dns_udp udp(port: 53)
187 {
188 # echo(packet.sec, " ",packet.usec, " ",packet.len);
189 #record blobbytes(cat(123456),32,cat(packet.sec),32,cat(packet.usec),32,cat(packet.len),32
,blobbytes(packet.blob))
190 # to test_recorder;
191
192 #tcpdump_packet();
193 #display("verbose");
194 $id=short(udp.blob,0);
195 $qr=(byte(udp.blob,2)>>7)&0x1;
196 # $qr=(byte(udp.blob,2));
197 $opcode=(byte(udp.blob,2)>>3)&0xf;
198 if($opcode==0) $opcode="QUERY(0)";
199 if($opcode==1) $opcode="IQUERY(1)";
200 if($opcode==2) $opcode="STATUS(2)";
201 $aa=(byte(udp.blob,2)>>2)&0x1;

```

```

202 $tc=(byte(udp.blob,2)>>1)&0x1;
203 $rd=(byte(udp.blob,2)&0x1;
204 $ra=(byte(udp.blob,3)>>7)&0x1;
205 $z=(byte(udp.blob,3)>>4)&0x7;
206 $rcode=(byte(udp.blob,3)&0xf;
207 $qdcnt=short(udp.blob,4);
208 $ancnt=short(udp.blob,6);
209 $nscnt=short(udp.blob,8);
210 $arcnt=short(udp.blob,10);
211     $y=dns_name((12),17);
212 #echo("$y=", $y, "0");
213     $qname=$y[0];
214     $qtype=short(udp.blob,12+$y[1]);
215     $qclass=short(udp.blob,12+$y[1]+2);
216
217     if((!($qtype inside dns_types))||($qdcnt!=1)||($y[2]!=0)||($qclass!=1)||(strcasecmp(
$opcode,"QUERY(0)")||($y[1]>256)||($tc==1)||($z!=0)||(($rcode!=0)&&($rcode!=3))||($ancnt>8)||($nscnt>8)||($arcnt>8))
218 #if((!($qtype inside dns_types))||($qdcnt!=1)||($y[2]!=0)||($qclass!=1)||(strcasecmp($opcode,"QUERY(0)")||($y[1]>256)||($tc==1)||($z!=0)||(($rcode!=0)&&($rcode!=3))||($ancnt>1)&&($aa==0)||($qtype!=15))||($ancnt>8)||($nscnt>8)||($arcnt>8))
219     {
220         #display("verbose");
221         tcpdump_packet();
222         pkts_loged=pkts_loged+1;
223         echo("UDP DNS packet ",ip.src,":",udp.sport,"->",ip.dst,":",udp.dport," ID=", $id, "
Q/R=", $qr, " Opcode=", $opcode, " AA=", $aa, " TC=", $tc, " RD=", $rd, " RA=", $ra, " Z=", $z, " RCODE=", $rcode,
" QDCOUNT=", $qdcnt, " ANCOUNT=", $ancnt, " NSCOUNT=", $nscnt, " ARCOUNT=", $arcnt, " $y[2]=", $y[2])
;
224     }
225     else
226     {
227         pkts_sifted=pkts_sifted+1;
228     }
229     $qcount=0;
230     $total_length=0;
231     while($qcount<$qdcnt)
232     {
233 #print query contents
234         $y=dns_name((12+$total_length),17);
235 #echo("$y=", $y, "0");
236         $qname=$y[0];
237         $qtype=short(udp.blob,12+$total_length+$y[1]);
238         $qclass=short(udp.blob,12+$total_length+$y[1]+2);
239         if((!($qtype inside dns_types))||($qdcnt!=1)||($y[2]!=0)||($qclass!=1)||(strcasecmp($opcode,"QUERY(0)")||($y[1]>256)||($tc==1)||($z!=0)||(($rcode!=0)&&($rcode!=3))||($ancnt>8)||($nscnt>8)||($arcnt>8))
240         {
241             echo(" Query ", $qcount, " QNAME=", $qname, " QTYPE=", $qtype, " QCLASS=", $qclass)
s);
242         }
243         if(total_qcounts[$qtype]==NULL)
244             total_qcounts[$qtype]=1;
245         else
246             total_qcounts[$qtype]=total_qcounts[$qtype]+1;

```

```

247         $total_length=$total_length+$y[1];
248         $qcount=$qcount+1;
249     }
250     if((!($qtype inside dns_types))||($qdcnt!=1)||($y[2]!=0)||($qclass!=1)||(strcase
cmp($opcode,"QUERY(0)")||($y[1]>256)||($tc==1)||($z!=0)||(($rcode!=0)&&($rcode!=3))||($ancount>8)||
($nscount>8)||($ancount>8)||($arcount>8))
251     {
252         echo("0");
253     }
254
255 }
256
257
258 filter printit timeout(sec : 240,repeat)
259 {
260 #echo("total_qcounts=",total_qcounts,"0");
261 echo("DNS: Packets Sifted=",pkts_sifted," Packets Logged=", pkts_logged,"(",(100*pkts_sifted)/
(pkts_sifted)+(pkts_logged)),"% Removed)0);
262
263 }
264
265 filter garbcoll timeout(sec : 600,repeat)
266 {
267 new_list=listadd(NULL);
268 foreach $x inside (tcp_nonce_list)
269     {
270     $a=tcp_nonce[$x];
271     if(($a[1]+300)<system.time)
272     {
273         tcp_nonce[$x]=NULL;
274
275     }
276     else
277     {
278         new_list=listadd(new_list,$x);
279     }
280 }
281 tcp_nonce_list=new_list;
282 new_list=NULL;
283 }

```

**ftp.nfr**

1#This filter discards suspected good control traffic on the FTP port, port 21, and logs all other control packets to the file "ftp.tcpdump" in tcpdump format. All packets, including all previous control packets on that connection, are logged to this file once a suspicious packet has been detected. Additionally, the data on the port(s) of the data connection(s) for that control stream are logged for all subsequent packets. This data is logged through the file "alldns.nfr" to the log file "rawbinary.bin".

2

3#The command is parsed out and checked to see if it is a recognized command. The length of the command is also checked to detect buffer overflow attacks. The arguments to some commands are checked and the arguments to other commands are assumed to be harmless as long as they aren't too long. For example, any password is ok, but a user name of root, or a get of passwd is suspicious. When checking arguments, we check for cursor control characters, since these can be used to mask a malicious command.

4

5#NOTE: This filter assumes that the packet variable tcp.nonce is unique. This is not guaranteed but "likely" according to the documentation. If it is a counter, we should be ok. If it is some sort of hash, we could be counting on uniqueness that does not exist in some cases.

6

7#Currently, we are removing about 40% of the packets from the control stream (690 packets filtered 1052 packets logged 16778 data packets logged. We are not keeping enough state to calculate the data packets removed). If we add some rules to sift out pasv commands from netscape ftp streams, we can probably do better.

8

9#We have added a filter to remove the pasv traffic created by netscape, but it does allow a potential attack by a client posing as a netscape client.

10

11#Also removing "MACB" command, now removes 80%(221 packets sifted, 53 packets logged, 0 associated data packets logged). A bug caused data packets to not be logged).

12

13#we are also removing "quit".

14

15good\_commands=listadd(NULL,"LIST","NLST", "QUIT","SYST","CWD","SIZE","MDTM","AUTH","REST","C  
DUP","PWD","MACB","quit","MKD");#,"TYPE","PORT","USER","RETR","STOR");

16bad\_files=listadd(NULL,"passwd","shadow","host");

17test\_schema=library\_schema:new(1,["str"],scope());

18test\_recorder=recorder("bin/encap\_out ftp.tcpdump","test\_schema");

19#test\_recorder=recorder("dd of=ftp.raw","test\_schema");

20packets\_sifted=0;

21packets\_loged=0;

22data\_packets\_loged=0;

23

24func tcpdump\_packet

25{

26 record blobbytes(cat(123456),32,cat(packet.sec),32,cat(packet.usec),32,cat(packet.len),32,  
blobbytes(packet.blob))

27 to test\_recorder;

28}

29filter ftpclient tcp(client, dport : 21)

30{

31 declare \$suspect\_stream inside tcp.connSym;

32 declare \$data\_blob inside tcp.connSym;

33 declare \$con\_pkts inside tcp.connSym;

34 declare \$pasv inside tcp.connSym;

```

35 declare $mozilla inside tcp.connSym;
36 $pass=0;
37 if($con_pkts!=NULL)
38   {
39     $con_pkts=$con_pkts+1;
40   }
41 else
42   {
43     $con_pkts=1;
44   }
45 $dump_blob=blobbytes(cat(123456),32,cat(packet.sec),32,cat(packet.usec),32,cat(packet.len)
,32,blobbytes(packet.blob));
46 $data_blob=blobbytes($data_blob,cat(strlen($dump_blob),32,$dump_blob);
47 $cmd_start=index(tcp.blob,cat(0xff));
48 if($cmd_start!=-1)
49   {
50 #check to see if this is an abort(242) or status?, otherwise flag
51 #we don't check status
52   if(
53     strcasecmp(substr(tcp.blob,$cmd_start+1,1),blobbytes(242))||
54     strcasecmp(substr(tcp.blob,$cmd_start+2,1),blobbytes(255)))
55     $suspect_stream=1;
56   }
57 else
58   {
59     $good=0;
60 #we are checking for cursor control charechters that would let malicious
61 #arguments otherwise pass through the detector
62     foreach $a inside(listadd(NULL,blobbytes(8),blobbytes(3)))#we should also add DEL
63     {
64       if(index(tcp.blob,$a)!=-1)
65         $suspect_stream=1;
66     }
67     foreach $a inside(good_commands)
68     {
69       if(index(tcp.blob,$a)!=-1)
70         {
71           $good=1;
72 #           if(index($a,"PASS")!=-1)
73 #             $pass=1;
74         }
75     }
76
77     if(($good!=1)||($suspect_stream==1))
78     {
79 #test other commands that require argument testing or processing here
80 #use substr to find begining of comand, add to get arg, compare this to bad values or add po
rt
81     #echo(substr(tcp.blob,0,10),"0");
82     if(index(tcp.blob,"TYPE")!=-1)
83     {
84       $a=index(tcp.blob,"TYPE");
85 #     echo ("offset after type",substr(tcp.blob,$a+5,10),"0");
86     if((index(substr(tcp.blob,$a+5,1),"I")==0)||((index(substr(tcp.blob,$a+5,1),"A"
)==0))

```

```

87     {
88     #   echo ($a,"0");
89     #   echo (substr(tcp.blob,$a+5,10),"0");
90     $good=1;
91     }
92   }
93   else
94     if($a=index(tcp.blob,"PORT")!=-1)
95     {
96     $arg=substr(tcp.blob,$a+4,tcp.length-$a-4);
97     $j=0;
98     $n="";
99     $i=0;
100    while($j<4)
101    {
102
103        $b=NULL;
104        while((strcasecmp($b,"")!=0)&&($i<15))
105        {
106            $k=0;
107            $b=substr($arg,$i,1);
108            if($b!="")
109                $n=cat($n,$b);
110            $i=$i+1;
111            $k=$k+1;
112        }
113        if($k>3)
114        {
115            $suspect_stream=1;
116            $bad_address=1;
117            break;
118        }
119        if($j<3)
120            $n=cat($n,".");
121        $j=$j+1;
122    }
123    $address=host($n);
124    $j=0;
125    # display ("verbose");
126    while($j<3)
127    {
128        $k=0;
129        $n=0;
130        $b=NULL;
131        $c=100;
132        while((strcasecmp($b,"")!=0)&&($k<3)&&($c>47))
133        {
134            $b=substr($arg,$i,1);
135            $c=short(blobbytes(0,$b,0,0),0);
136            if((strcasecmp($b,"")!=0)&&($c>47))
137            {
138                $n=(10*$n)+($c-48);
139            }
140            $i=$i+1;

```

```

141         $k=$k+1;
142     }
143     if($k>3)
144     {
145         $suspect_stream=1;
146         $bad_address=1;
147         break;
148     }
149     if($j<1)
150         $port=$n*256;
151     else
152         $port=$n+$port;
153     $j=$j+1;
154 }
155 if(($suspect_stream==1)&&($bad_address!=1))
156 {
157 #add to list of connections to dump
158     bad_ports[cat($address,$port)]=tcp.connhash;
159     bad_hash[tcp.connhash]=listadd(bad_hash[tcp.connhash],
160                                     cat($address,$port));
161 }
162 else
163     $good=1;
164 }
165 else
166     if($a=index(tcp.blob,"USER")!=-1)
167     {
168         if((index(substr(tcp.blob,$a+4,tcp.length-$a-4),"root")==-1))
169         {
170             $good=1;
171         }
172     }
173 else
174     if($a=index(tcp.blob,"PASV")!=-1)
175     {
176         $good=1;
177         if($pasv!=0)
178         {
179             $suspect_stream=1;
180         }
181         $pasv=1;
182     }
183 else
184     if($a=index(tcp.blob,"PASS")!=-1)
185     {
186         $good=1;
187
188         if(index(tcp.blob,"mozilla")!=-1)
189         {
190             $mozilla=1;
191             $pasv=0;
192         }
193     }
194 else

```



```

195         if($a=index(tcp.blob,"RETR")!=-1)
196         {
197             foreach $b inside(bad_files)
198             {
199                 if((index(substr(tcp.blob,$a+4,tcp.length-$a-4),$b)!=-1))
200                     $suspect_file=1;
201             }
202             if($suspect_file!=1)
203                 $good=1;
204         }
205     else
206     {
207         if($a=index(tcp.blob,"STOR")!=-1)
208         {
209             foreach $b inside(bad_files)
210             {
211                 if((index(substr(tcp.blob,$a+4,tcp.length-$a-4),$b)!=-1))
212                     $suspect_file=1;
213             }
214             if($suspect_file!=1)
215                 $good=1;
216         }
217     }
218 }
219 }
220 if($good!=1)
221 {
222     $suspect_stream=1;
223 }
224
225 #parse stream and compare caommands,
226 #flag bad commands, and find associated data stream from port command
227
228
229     if(tcp.length>500)
230     {
231         $suspect_stream=1;
232     }
233 }
234
235 if($suspect_stream==1)
236 {
237     #if($pass!=1)
238     #display("verbose");
239
240 }
241
242 }
243 filter ftpserver tcp(server, dport : 21)
244 {
245     declare $suspect_stream inside tcp.connSym;
246     declare $data_blob inside tcp.connSym;
247     declare $con_pkts inside tcp.connSym;
248     declare $pssv inside tcp.connSym;

```

```

249 declare $mozilla inside tcp.connSym;
250 if($con_pkts!=NULL)
251 {
252     $con_pkts=$con_pkts+1;
253 }
254 else
255 {
256     $con_pkts=1;
257 }
258 $dump_blob=blobbytes(cat(123456),32,cat(packet.sec),32,cat(packet.usec),32,cat(packet.len)
,32,blobbytes(packet.blob));
259 $data_blob=blobbytes($data_blob,cat(strlen($dump_blob),32,$dump_blob);
260 $resp_start=index(tcp.blob,cat(0xff));
261 # if($resp_start!=-1)
262 # {
263 #     if(strcasecmp(substr(tcp.blob,$cmd_start+1,1),blobbytes(242))
264#         ||strcasecmp(substr(tcp.blob,$cmd_start+2,1),blobbytes(255)))
265#         $suspect_stream=1;
266#         #check to see if this is part of a break, or status, otherwise flag
267
268#     }
269# else
270
271 if(tcp.length>1000)
272 {
273     $suspect_stream=1;
274
275 }
276
277
278 }
279filter embeded_open tcp(noticesession)
280 {
281     declare $data_blob inside tcp.connSym;
282     declare $pasv inside tcp.connSym;
283     $pasv=5;
284     $data_blob=blobbytes(NULL);
285 }
286filter embeded_close tcp(discardsession)
287 {
288     declare $suspect_stream inside tcp.connSym;
289     declare $data_blob inside tcp.connSym;
290     declare $con_pkts inside tcp.connSym;
291     declare $pasv inside tcp.connSym;
292     declare $mozilla inside tcp.connSym;
293 #note that this test opens the server to attacks on the anonymous account(or any other
account that has been set to take mozilla as a password using the passive command. This may be
an ok compromise, but should be removed for the best security.
294     if(($pasv==1)&&($mozilla!=1))
295     {
296         $suspect_stream=1;
297     }
298     if($suspect_stream==1)
299     {

```

```

300     packets_logged=packets_logged+$con_pkts;
301     record $data_blob to test_recorder;
302     #echo("$data_blob=", $data_blob, "0");#/*
303     #*/ /*#write the commands and responses to a file through a tcpdump output recorder
304     #record data_blob*/
305
306     }
307 else
308     if($con_pkts!=NULL)
309         packets_sifted=packets_sifted+$con_pkts;
310 foreach $a inside (bad_hash[tcp.connhash])
311     {
312         bad_ports[$a]=NULL;
313     }
314
315 bad_hash[tcp.connhash]=NULL;
316 }
317 filter dump_bad_data tcp()
318 {
319
320
321     if ((bad_ports[cat(ip.src,tcp.sport)]!=NULL)||
322         (bad_ports[cat(ip.dst,tcp.dport)]!=NULL))
323     {
324         #echo("suspect data packet detected0);
325         data_packets_logged=data_packets_logged+1;
326         id_alldns:tcpdump_packet();
327     }
328 }
329
330 filter printit timeout(sec : 240,repeat)
331 {
332     echo("FTP: Packets Sifted=", packets_sifted, " Packets Logged=", packets_logged, "((", (100*packets
333     _sifted)/((packets_sifted)+(packets_logged)), "% Removed) Data Packets Logged=", data_packets_logged, "0
334 }

```

**http.nfr**

1#This filter discards all HTTP traffic of type "get" or "post" that is under a given length (512 bytes in this case). All other http packets are logged through alldns.nfr. This is a very coarse filter, in need of refinement. Perhaps some parsing of the arguments is possible.

```

2
3good_types=listadd(NULL,GET,POST);
4packets_loged=0;
5packets_sifted=0;
6max_req_len=512;
7filter http tcp(client, dport : 80)
8{
9  good=0;
10 bad=0;
11  if (tcp.length > max_req_len)
12    {
13      #echo ("File too large");
14      bad=1;
15    }
16  foreach $type inside (good_types)
17    {
18      if(index($type,substr(tcp.blob,0,25))!=-1)
19        {
20
21          good=1;
22        }
23
24    }
25  if((good==0)||(bad==1))
26    {
27#      echo("*****HTTP Traffic*****0);
28#      display("verbose");
29      id_alldns:tcpdump_packet();
30      packets_loged=packets_loged+1;
31    }
32  else
33    {
34      packets_sifted=packets_sifted+1;
35    }
36}
37filter printit timeout(sec : 240,repeat)
38{
39echo("HTTP: Packets Sifted=",packets_sifted," Packets Loged=", packets_loged,"(",100*packet
s_sifted)/((packets_sifted)+(packets_loged)),"% Removed)0);
40}

```

**allfrag.nfr**

1#This filter captures all fragmented IP packets and writes them to the file fragedPkts.pcap using the special backend that writes to a pcap/tcpdump format file, output. I used this filter to isolate some AFS servers that were sending packets larger than the MTU of the network of the destination host.

```
2
3test_schema=library_schema:new(1,["str"],scope());
4test_recorder=recorder("bin/output fragedPkts.pcap","test_schema");
5
6func tcpdump_packet
7{
8record blobbytes(cat(123456),32,cat(packet.sec),32,cat(packet.usec),32,cat(packet.len),32,bl
obbytes(packet.blob))
9    to test_recorder;
10}
11filter allfrags ip()
12{
13    if(short(eth.blob,6)&0xbfff)
14        {
15            tcpdump_packet();
16#            echo("Begining of fragmented packet0);
17#            display("verbose");
18#            echo("End of fragmented packet0);
19
20        }
21
22}
```