

Networking - Assignment 3

Packet Filtering and Encryption

Johan IJsveld & Rick van der Zwet
<jijsveld@xs4all.nl> & <hvdzwet@liacs.nl>

LIACS
Leiden University
Niels Bohrweg 1
2333 CA Leiden
The Netherlands

29 april 2009

1 Implementation

We have chosen for a single emphC program without any arguments to keep it simple, meaning for the encryption key *0x12345678* will be hard-coded into the program and queue number are hard-coded *0* for input and *1* for output.

2 Global working

At first a package arrives at *netfilter* which is a Linux packet filter implementation. This will push the packet using callback function implemented at *inlibnetfilter_queue* to our implementation. let's call our application *xorfilter* for a moment. The *xorfilter* performs it's 'magic' which will be in this case a mangle of the data payload using a XOR computation with the key above. Next it recalculates the TCP checksum to ensure the checksum matches the altered payload, leading to proper validation of the packet on the other end.

3 Usage xorfilter

At first make sure to enable the *xorfilter*, this needed to be running under the root user, due to the enhanced privileges it needs at start by 'hooking' into *iptables*. Next alter *iptables* using the CLI script, for a reference sample, check *setqueue.sh*.

4 Usage echod

Once started the program *echod* will listen at port *12345*. It's only purpose it serves (hences the 'd' of daemon in echod) is to reply the data it receives. One could connect using the program telnet to the program using *telnet localhost 12345* and everything you type will be returned as well.

5 xorfilter encryption

The key consist of 4 bytes, so for every 4 bytes, the data is XORed with the key, which is first translated using *htonl* to make sure we are using big endian instead of small endian as all network traffic is big endian by definition. Next in case of a remainder -when the length of the data is not a power of 4 bytes- the remainder is XORed with the relevant byte of the key. Example:

```
DATA: AB CD EF 34 23 21 34 (7 bytes)
Step 1 (key repeated each 4 bytes, done using for loop on words):
KEY : 12 34 56 78
IMM : B9 F9 B9 4C 23 21 34
Step 2 (remainder key, done using for loop on bytes):
KEY :           12 34 56
IMM : B9 F9 B9 4C 31 15 62
```

Appendix

echod (1 file), xorfilter (9 files), netfilter CLI example (2 files)

```
01: /* Netlab 2009 - group 1 - assignment 3
02: * Johan IJsveld
03: * Rick van der Zwet
04: */
05: #include <sys/types.h>
06: #include <netdb.h>
07: #include <sys/types.h>
08: #include <stdio.h>
09: #include <string.h>
10: #include <unistd.h>
11:
12: #include <sys/socket.h>
13: #include <netinet/ip.h>
14: #include <netinet/tcp.h>
15:
16: #define LISTEN_BACKLOG 0
17: #define BUF_SIZE 1000
18:
19:
20: #include <signal.h>
21:
22:
23: int main() {
24:     int conn_fd;
25:     char buf[BUF_SIZE];
26:     ssize_t recv_size;
27:     int retno;
28:
29:     /* Create a tcp/ip socket */
30:     int mysocket = socket(PF_INET, SOCK_STREAM, 0);
31:
32:     struct sockaddr_in my_addr;
33:     struct sockaddr_in their_addr;
34:     int sin_size;
35:
36:
37:     // ip(7) Create listen port, NOTE addresses in byte-order BYTEORDER(3)
38:     my_addr.sin_family = AF_INET; // By default
39:     my_addr.sin_port = htons(12345);
40:     my_addr.sin_addr.s_addr = htonl(INADDR_ANY);
41:     memset(my_addr.sin_zero, '\0', sizeof my_addr.sin_zero);
42:
43:     // get connected
44:     retno = bind(mysocket, (struct sockaddr *) &my_addr, sizeof(struct sockaddr_in));
45:     retno = listen(mysocket, LISTEN_BACKLOG);
46:     sin_size = sizeof(struct sockaddr_in);
47:
48:     // Make it a server, e.g. never ending
49:     while(1) {
50:         conn_fd = accept(mysocket, (struct sockaddr *)&their_addr, &sin_size);
51:         if (fork() == 0) {
52:             // child will help to echo the stuff
53:             while(1) {
54:                 retno = recv(conn_fd, &buf, BUF_SIZE, 0);
55:                 printf("%s\n", buf);
56:                 if (retno <= 0)
57:                     break;
58:                 retno = send(conn_fd, &buf, retno, 0);
59:                 if (retno <= 0)
60:                     break;
61:             }
62:             close(conn_fd);
63:         }
64:     }
65:     close(mysocket);
66:
67:
68: return EX_OK;
69: }
```

```
01: #include "checksum.h"
02:
03: long checksum( unsigned short *addr, int count )
04: {
05:     /* Compute Internet Checksum for "count" bytes
06:      * beginning at location "addr".
07:      */
```

```
08:     long checksum;
09:     register long sum = 0;
10:     while( count > 1 ) {
11:         /* This is the inner loop */
12:         sum += * (unsigned short *) addr++;
13:         count -= 2;
14:     }
15:
16:     /* Add left-over byte, if any */
17:     if( count > 0 )
18:         sum += * (unsigned char *) addr;
19:
20:     /* Fold 32-bit sum to 16 bits */
21:     while (sum >> 16)
22:         sum = (sum & 0xffff) + (sum >> 16);
23:     checksum = ~sum;
24:     return checksum;
25: }
26:
01: #ifndef CHECKSUM_H
02: #define CHECKSUM_H
03:
04: /* Compute Internet Checksum for "count" bytes
05:  * beginning at location "addr".
06:  */
07: long checksum( unsigned short *addr, int count );
08:
09: #endif
10:
001: #include "filter.h"
002: #include "xorencrypt.h"
003: #include "checksum.h"
004:
005: #include <stdio.h>
006: #include <stdlib.h>
007: #include <string.h>
008: #include <unistd.h>
009:
010: #include <netinet/in.h>
011: #include <linux/types.h>
012: #include <linux/ip.h>
013: #include <linux/tcp.h>
014: #include <linux/netfilter.h> /* for NF_ACCEPT */
015: #include <arpa/inet.h>
016: #include <inttypes.h>
017:
018: const unsigned int key = 0x12345678; /* The key used for encrypting */
019:
020: /* Returns the IP header from the datagram
021:  */
022: struct iphdr *get_ip_hdr(char *datagram)
023: {
024:     return (struct iphdr *)datagram;
025: }
026:
027: /* Returns the TCP header from the datagram
028:  */
029: struct tcphdr *get_tcp_hdr(char *datagram)
030: {
031:     char *tcpdata = datagram + get_ip_len(datagram);
032:     return (struct tcphdr *)tcpdata;
033: }
034:
035: /* Returns pointer to TCP payload data from the datagram
036:  */
037: char *get_tcp_payload(char *datagram)
038: {
039:     struct tcphdr *tcphdr = get_tcp_hdr(datagram);
040:     int iphdr_len = get_ip_len(datagram);
041:     int tcphdr_len = get_tcp_len(tcphdr);
042:     return datagram + iphdr_len + tcphdr_len;
043: }
044:
045: /* Returns the total length of the datagram
046:  */
047: _u16 get_ip_tot_len(struct iphdr *ip)
048: {
049:     return ntohs(ip->tot_len);
050: }
051:
052: /* Returns the total length of the IP header
```

```

053: */
054: __u16 get_ip_len(char *datagram)
055: {
056:     struct iphdr *iphdr = get_ip_hdr(datagram);
057:     return iphdr->ihl * 4; /* ihl is in words */
058: }
059:
060: /* Returns the length of the IP options field
061: */
062: __u16 get_ip_opt_len(char *datagram)
063: {
064:     return get_ip_len(datagram) - sizeof(struct iphdr);
065: }
066:
067: /* Returns the total length of the TCP header
068: */
069: __u16 get_tcp_len(struct tcphdr *tcphdr)
070: {
071:     return tcphdr->doff * 4; /* data offset is in words */
072: }
073:
074: /* Returns the length of the TCP options field
075: */
076: __u16 get_tcp_opt_len(struct tcphdr *hdr)
077: {
078:     return get_tcp_len(hdr) - sizeof(struct tcphdr);
079: }
080:
081: /* Returns the length in bytes of the tcp payload data
082: */
083: __u16 get_tcp_data_len(char *datagram)
084: {
085:     struct iphdr *iphdr = get_ip_hdr(datagram);
086:     struct tcphdr *tcphdr = get_tcp_hdr(datagram);
087:     int total_len = get_ip_tot_len(iphdr);
088:     int iphdr_len = get_ip_len(datagram);
089:     int tcphdr_len = get_tcp_len(tcphdr);
090:     return total_len - iphdr_len - tcphdr_len;
091: }
092:
093: /* Returns the source address */
094: __u16 get_tcp_source(struct tcphdr *hdr)
095: {
096:     return ntohs(hdr->source);
097: }
098:
099: /* Returns the destination address */
100: __u16 get_tcp_dest(struct tcphdr *hdr)
101: {
102:     return ntohs(hdr->dest);
103: }
104:
105: /* the pseudo header used in calculating checksum */
106: struct pseudo_hdr {
107:     __u32 source;
108:     __u32 dest;
109:     __u8 zero;
110:     __u8 ptcl;
111:     __u16 tcplen;
112: };
113:
114: void print_datagram(char *datagram)
115: {
116:     struct iphdr *iphdr = get_ip_hdr(datagram);
117:     struct tcphdr *tcphdr = get_tcp_hdr(datagram);
118:     unsigned int iphdr_len = get_ip_len(datagram);
119:     unsigned int tcphdr_len = get_tcp_len(tcphdr);
120:     unsigned int total_len = get_ip_tot_len(iphdr);
121:     unsigned int data_len = get_tcp_data_len(datagram);
122:     char *data = get_tcp_payload(datagram);
123:     printf("datagram length: %d, iphdr length: %d, tcphdr length: %d, data length: %d, data
124:     pointer: %p, sizeof(struct pseudo_hdr): %d\n",
125:           total_len, iphdr_len, tcphdr_len, data_len, data, sizeof(struct pseudo_hdr));
126: }
127:
128: /* Encrypts TCP payload data using XOR in the datagram pointed to by datagram
129: */
130: int encrypt_payload(char *datagram)
131: {
132:     struct iphdr *iphdr = get_ip_hdr(datagram);
133:     struct tcphdr *tcphdr = get_tcp_hdr(datagram);
134:     unsigned int tcphdr_len = get_tcp_len(tcphdr);
135:     unsigned int total_len = get_ip_tot_len(iphdr);
136:     unsigned int data_len = get_tcp_data_len(datagram);
137:     char *data = get_tcp_payload(datagram);
138:     struct pseudo_hdr *phdr;
139:     long crc;
140:
141:     print_datagram(datagram);
142:     if (iphdr->protocol != 6) {
143:         printf("not a tcp datagram\n");
144:         return total_len;
145:     }
146:
147:     /* memset(data, 0, data_len); */
148:     /* xor payload data */
149:     printf("Before crypt: '%s'\n", data);
150:     xorencrypt(data, data_len, htonl(key));
151:     printf("After crypt: '%s'\n", data);
152:
153:     printf("crc before: %d\n", tcphdr->check);
154:     /* calculate new crc value by making new datagram with pseudo header */
155:     tcphdr->check = 0;
156:     char *copy = malloc(sizeof(struct pseudo_hdr) + tcphdr_len + data_len);
157:
158:     /* set pseudo header */
159:     phdr = (struct pseudo_hdr *)copy;
160:     phdr->source = iphdr->saddr;
161:     phdr->dest = iphdr->daddr;
162:     phdr->zero = 0;
163:     phdr->ptcl = iphdr->protocol;
164:     phdr->tcplen = tcphdr_len + data_len;
165:     phdr->tcplen = htons(phdr->tcplen);
166:
167:     /* copy tcp header minus options */
168:     memcpy(copy + sizeof(struct pseudo_hdr), (char *)tcphdr, tcphdr_len);
169:     /* copy payload data */
170:     memcpy(copy + sizeof(struct pseudo_hdr) + tcphdr_len, data, data_len);
171:
172:     /* calculate and set new crc */
173:     crc = checksum((unsigned short *)copy, sizeof(struct pseudo_hdr) + tcphdr
174:     data_len);
175:     free(copy);
176:
177:     tcphdr->check = (__u16)crc;
178:     printf("crc after: %d\n", tcphdr->check);
179:     return total_len;
180: }
181:
182: #ifndef FILTER_H
183: #define FILTER_H
184: #include <linux/types.h>
185:
186: /* Returns the IP header from the datagram */
187: struct iphdr *get_ip_hdr(char *datagram);
188:
189: /* Returns the TCP header from the datagram */
190: struct tcphdr *get_tcp_hdr(char *datagram);
191:
192: /* Returns pointer to TCP payload data */
193: char *get_tcp_payload(char *datagram);
194:
195: /* Returns the total length of the datagram */
196: __u16 get_ip_tot_len(struct iphdr *ip);
197:
198: /* Returns the total length of the IP header */
199: __u16 get_ip_len(char *datagram);
200:
201: /* Returns the length of the IP options field */
202: __u16 get_ip_opt_len(char *datagram);
203:
204: /* Returns the total length of the TCP header */
205: __u16 get_tcp_len(struct tcphdr *tcphdr);
206:
207: /* Returns the length of the TCP options field */
208: __u16 get_tcp_opt_len(struct tcphdr *hdr);
209:
210: /* Returns the length in bytes of the tcp payload data */
211: __u16 get_tcp_data_len(char *datagram);
212:
213: /* Returns the source address */
214: __u16 get_tcp_source(struct tcphdr *hdr);
215:
216: /* Returns the destination address */
217: __u16 get_tcp_dest(struct tcphdr *hdr);
218:
219: /* Encrypts TCP payload data using XOR in the datagram pointed to by data */
220: int encrypt_payload(char *datagram);
221:
222: #endif
223:
224: #include "queue.h"
225:
226: #include <stdio.h>
227: #include <stdlib.h>
228: #include <unistd.h>
229: #include <netinet/in.h>
230: #include <sys/types.h>
231: #include <sys/socket.h>
232: #include <linux/types.h>
233: #include <linux/netfilter.h> /* for NF_ACCEPT */
234:
235: #include <libnetfilter_queue/libnetfilter_queue.h>
236:
237: #define MAX_QUEUES 3
238:
239: /* data for queue's */
240: struct queue_t {
241:     struct nfq_q_handle *qh; /* handle to the queue */
242:     callback cbp; /* callback function to handle packets from this queue */
243:     int n;
244: };
245:
246: /* the queue's that are setup */
247: static struct queue_t queues[MAX_QUEUES];
248:
249:

```

```

026: /* The main netlib handle and file descriptor */
027: static struct nfq_handle *h = NULL;
028: static int fd;
029:
030: /*
031:      struct nfnl_handle *nh;
032: */
033:
034: void init()
035: {
036:     printf("opening library handle\n");
037:     h = nfq_open();
038:     if (!h) {
039:         fprintf(stderr, "error during nfq_open()\n");
040:         exit(1);
041:     }
042:
043:     printf("unbinding existing nf_queue handler for AF_INET (if any)\n");
044:     if (nfq_unbind_pf(h, AF_INET) < 0) {
045:         fprintf(stderr, "error during nfq_unbind_pf()\n");
046:         exit(1);
047:     }
048:
049:     printf("binding nfnlqueue as nf_queue handler for AF_INET\n");
050:     if (nfq_bind_pf(h, AF_INET) < 0) {
051:         fprintf(stderr, "error during nfq_bind_pf()\n");
052:         exit(1);
053:     }
054:
055:     fd = nfq_fd(h);
056: }
057:
058: void deinit()
059: {
060:     printf("closing library handle\n");
061:     nfq_close(h);
062: }
063:
064: /* callback called by netfilter_queue when packet is available */
065: static int cb(struct nfq_q_handle *qh, struct nfgenmsg *nfmsg, struct nfq_data *nfa,
066: void *data)
067: {
068:     char *datagram;
069:     int n = *(int *)data;
070:     u_int32_t id = print_pkt(nfa);
071:     u_int32_t verdict;
072:     int len;
073:
074:     printf("entering nfq callback\n");
075:
076:     if (nfq_get_payload(nfa, &datagram) < 0) {
077:         fprintf(stderr, "error could not get payload data\n");
078:         exit(1);
079:     }
080:
081:     verdict = queues[n].cbp(datagram, &len);
082: }
083:
084: void setup_queue(int n, callback cbp)
085: {
086:     if (n >= MAX_QUEUES) {
087:         fprintf(stderr, "error not enough queues\n");
088:         exit(1);
089:     }
090:
091:     printf("binding this socket to queue %d\n", n);
092:     queues[n].n = n;
093:     queues[n].qh = nfq_create_queue(h, n, &cb, &(queues[n].n));
094:     if (!queues[n].qh) {
095:         fprintf(stderr, "error during nfq_create_queue()\n");
096:         exit(1);
097:     }
098:
099:     printf("setting copy_packet mode\n");
100:     if (nfq_set_mode(queues[n].qh, NFQNL_COPY_PACKET, 0xffff) < 0) {
101:         fprintf(stderr, "can't set packet_copy mode\n");
102:         exit(1);
103:     }
104:     queues[n].cbp = cbp;
105: }
106:
107: void close_queue(int n)
108: {
109:     if (n >= MAX_QUEUES) {
110:         fprintf(stderr, "error not enough queues\n");
111:         exit(1);
112:     }
113:
114:     printf("unbinding from queue %d\n", n);
115:     nfq_destroy_queue(queues[n].qh);
116: }
117: #ifdef INSANE
118:     /* normally, applications SHOULD NOT issue this command, since
119:     * it detaches other programs/sockets from AF_INET, too ! */
120:     printf("unbinding from AF_INET\n");
121:     nfq_unbind_pf(h, AF_INET);
122: #endif
123: }
124:
125: void receive_data()
126: {
127:     int rv;
128:     char buf[66000] __attribute__((aligned));
129:
130:     while ((rv = recv(fd, buf, sizeof(buf), 0)) && rv >= 0) {
131:         printf("\n--- pkt received ---\n");
132:         nfq_handle_packet(h, buf, rv);
133:     }
134: }
135:
136: /* Print packet and return packet id */
137: u_int32_t print_pkt(struct nfq_data *tb)
138: {
139:     int id = 0;
140:     struct nfqnl_msg_packet_hdr *ph;
141:     struct nfqnl_msg_packet_hw *hwph;
142:     u_int32_t mark, ifi;
143:     int ret;
144:     char *data;
145:
146:     ph = nfq_get_msg_packet_hdr(tb);
147:     if (ph) {
148:         id = ntohl(ph->packet_id);
149:         printf("hw_protocol=0x%04x hook=%u id=%u ",
150:             ntohs(ph->hw_protocol), ph->hook, id);
151:     }
152:
153:     hwph = nfq_get_packet_hw(tb);
154:     if (hwph) {
155:         int i, hlen = ntohs(hwph->hw_addrlen);
156:
157:         printf("hw_src_addr=");
158:         for (i = 0; i < hlen-1; i++)
159:             printf("%02x:", hwph->hw_addr[i]);
160:         printf("%02x ", hwph->hw_addr[hlen-1]);
161:     }
162:
163:     mark = nfq_get_nfmark(tb);
164:     if (mark)
165:         printf("mark=%u ", mark);
166:
167:     ifi = nfq_get_indev(tb);
168:     if (ifi)
169:         printf("indev=%u ", ifi);
170:
171:     ifi = nfq_get_outdev(tb);
172:     if (ifi)
173:         printf("outdev=%u ", ifi);
174:     ifi = nfq_get_physindev(tb);
175:     if (ifi)
176:         printf("physindev=%u ", ifi);
177:
178:     ifi = nfq_get_physoutdev(tb);
179:     if (ifi)
180:         printf("physoutdev=%u ", ifi);
181:
182:     ret = nfq_get_payload(tb, &data);
183:     if (ret >= 0)
184:         printf("payload_len=%d ", ret);
185:
186:     fputc('\n', stdout);
187:
188:     return id;
189: }
190:
191: #ifndef QUEUE_H
192: #define QUEUE_H
193:
194: #include <netinet/in.h>
195: #include <linux/netfilter.h> /* for NF_ACCEPT */
196:
197: #include <libnetfilter_queue/libnetfilter_queue.h>
198:
199: /* callback receives a datagram and returns a verdict (either NF_ACCEPT or NF_DROP)
200: typedef u_int32_t (*callback)(char *datagram, int *len);
201:
202: /* initialize netfilter */
203: void init();
204:
205: /* deinitialize netfilter */
206: void deinit();
207:
208: /* bind to queue n */
209: void setup_queue(int n, callback cbp);
210:
211: /* unbind from queue n */
212: void close_queue(int n);
213:
214: /* start receiving packets */
215: void receive_data();
216:
217: /* Print packet and return packet id */
218: u_int32_t print_pkt (struct nfq_data *tb);
219:
220: #endif
221: #include "xorencrypt.h"

```

```

02:
03: typedef unsigned int key_t;
04:
05: /* Encrypt data of len bytes with XOR and key */
06: void xorencrypt(void *data, int len, key_t key)
07: {
08:     int i;
09:     key_t *word_data = (key_t *)data;
10:     int word_len = len / sizeof(key_t);
11:     int remainder = len % sizeof(key_t);
12:
13:     for (i = 0; i < word_len; ++i) {
14:         printf("%i\n", i);
15:         word_data[i] ^= key;
16:     }
17:
18:     for (i = len - remainder; i < len; ++i) {
19:         ((unsigned char *)data)[i] ^= ((unsigned char*)&key)[i];
20:     }
21: }
22:
1: #ifndef XORENCRYPT_H
2: #define XORENCRYPT_H
3:
4: /* Encrypt data of len bytes with XOR and key */
5: void xorencrypt(void *data, int len, unsigned int key);
6:
7: #endif
8:
01: #include <stdio.h>
02: #include <string.h>
03: #include "queue.h"
04: #include "filter.h"
05:
06: #include "xorencrypt.h"
07:
08: static u_int32_t handle_input(char *datagram, int *len)
09: {
10:     printf("encrypting input payload\n");
11:     *len = encrypt_payload(datagram);
12:     return NF_ACCEPT;
13: }
14:
15: static u_int32_t handle_output(char *datagram, int *len)
16: {
17:     printf("encrypting output payload\n");
18:     *len = encrypt_payload(datagram);
19:     return NF_ACCEPT;
20: }
21:
22: int main(void)
23: {
24:     init();
25:
26:     setup_queue(1, handle_input);
27:     setup_queue(2, handle_output);
28:
29:     receive_data();
30:
31:     close_queue(1);
32:     close_queue(2);
33:
34:     deinit();
35:
36:     return 0;
37: }
38:
1: #!/bin/sh
2: /sbin/iptables -t mangle -A INPUT -i lo -j NFQUEUE --queue-num 0
3: /sbin/iptables -t mangle -A OUTPUT -o lo -j NFQUEUE --queue-num 1
4:
5: /sbin/iptables -t mangle -A INPUT -i eth0 -j NFQUEUE --queue-num 0
6: /sbin/iptables -t mangle -A OUTPUT -o eth0 -j NFQUEUE --queue-num 1
7:
8: #/sbin/iptables -t mangle -A INPUT -i eth0 -d 10.211.55.4 -j NFQUEUE --queue-num 1
9: #/sbin/iptables -t mangle -A OUTPUT -o eth0 -s 10.211.55.4 -j NFQUEUE --queue-num 2
1: #!/bin/sh
2: /sbin/iptables -F
3: /sbin/iptables -t mangle -F
4:

```