

Session-based packet marking and auditing for network forensics

Omer Demir, The City University of New York, NY, USA

Ping Ji, The City University of New York, NY, USA

Jinwoo Kim, The City University of New York, NY, USA

Article originally published in "International Journal of Digital Evidence", V. 6 (2007), n. 1, <http://utica.edu/academic/institutes/ecii/ijde/articles.cfm?current=1>
Reprinted with permission.

ABSTRACT. The widely acknowledged problem of reliably identifying the origin of network data has been the subject of many research works. Due to the nature of Internet Protocol, a source IP can be easily falsified, which results in numerous problems, including the infamous Denial-Of-Service (DOS) attacks. In this paper, two light-weight novel approaches are proposed to solve this problem by providing simple and effective logging and IP-Traceback mechanism: Session-Based packet Logging (SBL) and SYN-based Packet Marking (SYNPM). The contribution of these schemes lies in the fact that they are easy to be implemented with little overhead, and are practical under sensitive privacy regulations, since they do not need to access detailed contents of each individual communication session. Currently, SBL and SYNPM approaches support only TCP sessions.

KEYWORDS: *Connection information, Individual connection, Informatic crime, Network forensics, Network packet*

Introduction

The nature of the Internet Protocol has been known to make it difficult to reliably identify the actual source of information in cyberspace. Information over the Internet travels in small units, often called packets, and the source of the information - the IP address of the source host - is contained inside the packet with other data contents. Since the IP address of the *original* source computer is often one of the most critical pieces of evidence in online investigation, law enforcement agencies and system administrators encounter enormous challenges, due to this loophole in network forensics.

IP-Traceback is a method for reliably determining the origin of a packet on the Internet (Wikipedia, 2007). The major problem of determining the real source address of an incoming packet is that the Internet routers may replace the source IP address with their own addresses in network packets, and there is no provision in TCP/IP to discover the true origin of a packet (Morris, 1985).

The consequence, therefore, is an easy alteration or a masquerade of source address from an ill-minded sender to conduct criminal activities in cyberspace. One of the worst cases that this loophole can lead to is Denial-Of-Service (DOS) attacks. Under DOS attacks, targeted machines suffer significantly from exhausting valuable resources by flooded spoofing messages from senders, who often hide their identities through incorrect or altered IP addresses. In most situations, this kind of criminal activity is extremely hard to prevent and trace, once it is committed. The current techniques to trace the IP address of a remote machine are far from those used in telecommunication infrastructure, which is efficient and light-weight in terms of the amount of information to be logged. Specifically, telephone logging systems usually include the time of the call, the number of the caller and receiver, the duration of the call, and the place of the call. Good topography of relations in the telephone system can be built based on the logged information. However, in the Internet, no IP-Traceback technology has adopted an approach that is similar to the one used in telephone auditing systems. In addition, current IP-Traceback logging techniques do not provide a long term logging mechanism, and the duration of logging is usually on the order of hours at most. If there were call logs for computer communications, it would be very helpful to law enforcement agencies for later forensics investigation.

In this paper, two novel schemes are proposed to alleviate the challenging problem of IP-Traceback: Session-Based packet Logging (SBL) technique, and SYN-based Packet Marking (SYNPM) Technique. Not only are the proposed techniques novel, but they are quite practical solutions for the following reasons. First, both of the proposed techniques effectively provide log information for longer time periods (potentially on the order of months) compared to current IP-Traceback techniques, which provide only up to hours of log information. Second, these techniques respect the privacy of communications, given that they do not log detailed contents of each communication session. This makes these schemes ideal to provide effective addressable evidence to the court under sensitive privacy regulations.

The SBL technique, the first proposed scheme in this paper, addresses and solves the problems residing in IP-Traceback from different perspectives, compared to the previous approaches. The problem is examined by applying conventional schemes

for telephone communication system in forensics networks. To this end, existing schemes being used in telecommunication investigation are adapted to the Internet Protocol to realize simple and efficient packet logging. In the SBL approach, only critical information in network forensic investigation, the IP addresses and the duration of communication sessions are recorded by the SYN and FIN packets over the logging period. This technique can be easily deployed at network entities for monitoring TCP traffic, and requires much less storage space over conventional schemes. The second proposed technique, SYNPM, is another effective tool for the purpose, but differs from the first approach in the following manner. In the SYNPM approach, routers insert distinguishable *identifiers*, a special signature of a router, in the SYN packet whenever it routes packets. The uniqueness of this scheme from previously proposed packet marking scheme is its efficiency, as it marks only the first SYN packet for a session, as it is the packet that initializes the session and contains sufficient information for IP-Traceback in forensics networks.

The proposed scheme in this paper currently supports only TCP sessions, but these approaches could be extended to UDP connections, which have many inherent network security problems. The rest of this paper is organized as follows: in Section II, several related IP-Traceback approaches, that tried to remedy discussed problems, are presented; the Session-Based packet Logging (SBL) scheme is described in detail in Section III; Section IV details the evaluation results that are derived from experiments of SBL; in Section V, SYN-based Packet Marking Protocol is presented; section VI concludes this paper and discusses potential directions of future work.

Related work

There have been several IP-Traceback techniques. Snoeren and colleagues present a hash-based IP-Traceback approach called Single packet IP-Traceback. In this technique, a hash-based scheme is offered for single packet IP-Traceback and it is claimed to require a space of 0.5% of link capacity for recording (Snoeren et al., 2002). This technique enables single packet IP-Traceback, but it has several limitations. First, the original attacking packet has to be found, in order to search the audit. The original packet's digest is acquired, and then the query is done using that digest. It is not common in forensic analysis to have an original copy of an attacking

packet. Most of the time only the source IP address and the time of the packet are known. Another problem is false positives. In the case of having the original packet for IP-Traceback, the possibility of false positives will make it difficult to see the log in the court. In order to use them, they have to be certain, with no chance of false positives. Memory limitation is another drawback. The amount of time, during which queries can be supported, is directly dependent on the amount of memory dedicated to its implementation.

Savage and colleagues describe a general purpose IP-Traceback mechanism based on a probabilistic packet marking in the network. Their approach allows a victim to identify the network path(s) traversed by attack traffic without requiring interactive operational support from Internet Service Providers (ISPs). Moreover, they claim it is possible to trace back the packet after the incident has finished with its scheme (Savage et al., 2001). They proposed several different marking algorithms, but there are some problems with their approach. The first is a backward compatibility issue, since the IP header encoding has several practical limitations, and it might negatively impact users that require fragmented IP datagram. The second problem is its effectiveness under distributed attack. The implementation of this technique inherently has serious limitations, due to the difficulty in correctly grouping fragments together. Consequently, the probability of misattributing an edge, as well as the amount of state needed to evaluate this decision, increases very quickly with the fan out of an attack. The third problem is with path validation. Some number of the packets sent by the attacker are unmarked by intervening routers. The victim cannot differentiate between these packets and genuine marked packets. Therefore, an attacker could insert "fake" edges by carefully manipulating the identification fields in the packets it sends. The fourth problem is detection of the attack's origin. While this IP-level traceback algorithm could be an important part of the solution for stopping Denial-Of-Service attacks, it is by no means a complete solution. This algorithm attempts to determine the approximate origin of attack traffic. There are a number of reasons why this may differ from the true source of the attack. Attackers can hide their true identities by "laundering" attacks through third parties, either indirectly - for example, "smurf attacks" (CERT Advisory, 1998) or "DNS reflectors" (CERT Incident Note, 2000) or directly via compromised "stepping stone" machines or IP-in-IP tunnels.

Ingress filtering is another approach towards eliminating the ability to forge source addresses. As the name implies, the incoming packets are filtered at the routers. Routers filter the packets depending on the source address and the real network segment that the router is connected to. One of the possible problems of this method is that it requires routers with knowledge of legitimate and illegitimate traffic. Ingress filtering must be deployed as much as possible for best and effective filtering. Another problem is that even if ingress filtering were universally deployed at the customer-to-ISP level, attackers could still forge addresses from the hundreds or thousands of hosts within a valid customer network (CERT Advisory, 2002). Generally, IP-Traceback techniques start from the router closest to the victim, and independently test its upward links until they determine all that are used to transmit the attackers' traffic. This procedure is repeated recursively on the upward routers until the source is founded. This scheme is generally called link testing. There are two varieties of link testing schemes, input debugging and controlled flooding. In input debugging, the victim is supposed to recognize that it is being attacked, and develop an attack signature that describes common features of all the attack packets. The victim sends this signature to a network operator. The network operator then installs a corresponding input debugging filter on the victim's upstream output port. This way the router from which this packet has arrived can be determined. The process is then repeated recursively on the upstream router, until the originating site is reached or the trace leaves the ISP's. In case of leaving ISP border, the upstream ISP must be contacted and the same procedure must be repeated. Unfortunately, management is a considerable problem for this method. Controlled flooding is a link-testing IP-Traceback technique developed by Burch and Cheswick which does not require any support from network operators (Burch, Cheswick, 2000). This technique is called controlled flooding, since it tests links by flooding them with large bursts of traffic, and observing how this affects traffic from the attacker. The inherent problem with this technique is that controlled flooding itself happens to be a Denial-Of-Service attack.

Another technique for IP-Traceback is to log packets at routers and then use data mining techniques to determine the path that the packets traversed. This technique could be quite effective, as it allows tracing an attack even long after the attack has completed

(Sager, 1998; Stone, 2000). But the problem is the storage of these log files. Another potential question which might be raised by privacy advocate groups is that the log file might contain sensitive data of people's private conversations.

Session-based IP-Traceback

In this section, a light-weighted IP-Traceback scheme - Session-Based packet Logging (SBL) - that can be easily deployed at network entities for monitoring TCP traffic is proposed. Since storage is one of the major concerns in logging systems of network forensics, SBL is designed to record only the necessary contents of a communication session. Specifically, in network forensics, the data transmitted and the session numbers or the fragments of a transaction are not usually needed. All that is necessary are the source and destination addresses and the time stamps of the communication duration. Once the traces of illegal actions in the victim system (for example, IP address) are known, it is only necessary to check the IP addresses against the log to investigate the attack. Since the purpose here is to log the IP addresses and the duration of *communication sessions* between hosts, it is enough to log the TCP sessions by recording the first session establishment packet (SYN packet) and the session termination packet (FIN packet) for each connection.

In the SBL scheme, the source IP, Destination IP, incoming port of the router, outgoing port of the router and the time stamps are logged for each logged packet. When keeping the logs, the packet encapsulation which happens at each network entity (for example, a router) must be considered. When performing encapsulation, instead of replacing a previous IP with its own IP address, each router envelopes the IP packet sent by the neighbouring network entity (for example, end system, adjacent routers) as data payload into a new packet, and attaches its IP header to the new packet. In this situation the original source IP (for example, the IP address from the previous hop) is located in the first four bytes of the payload, in the encapsulated packet. Therefore, in the SBL logging approach, the header information and the first four bytes of the data payload of each logged packet are recorded.

The storage space (in the unit of bytes) that each captured packet may need must be considered next. The source and destination IP addresses for each SYN packet must be recorded, which

consumes 8 bytes space in total. In addition, as discussed, the first four bytes of the data payload should be recorded. It is safe to assume that there are 256 incoming ports and 256 outgoing ports for each network router. Thus, two more bytes are needed to record the input and output port numbers. Therefore, there are 14 bytes of information to log. It is also important to keep track of the time stamp of each packet. Since one day has 86,400 seconds, which can be represented by 3 bytes with a 1/128 sec precision, combined with the 14 bytes information, for each packet, 17 bytes are required to record the information of interest. Provided that a SYN packet size is 62 bytes, and a FIN packet is 54 bytes, the ratio of the logged data size to the total size of a logged packet is “17/62” or “17/54”, which saves the storage space significantly. In the experiments, the log files were kept with each data entry as shown in Table 1. Ethereal was used to monitor the traffic, where a SYN packet can be observed, when the corresponding “Syn” field is set to be “1”; and a FIN packet can be observed, if the “Fin” field is set to be “1” (Ethereal Packet Sniffer, 2007). Figure 1 shows the partial contents of packet information obtained by Ethereal, where Figure 1(a) corresponds to a SYN packet, and Figure 1(b) corresponds to a FIN packet.

Table 1. Sample table for logs

Time	Incoming port	Source IP	First 4 bytes of IP payload	Destination IP	Destination port
00000,001	2	X.Y.Z.D	X.X.X.X	T.B.G.V	6
00000,012	3	A.B.C.D	Y.Y.Y.Y	E.F.G.H	7

Figure 1. Sample portion of packet logging file from Ethereal Labs

- a Observation of a SYN packet
- b Observation of a FIN packet

```

...0...= Push: Not set
...0...= Reset: Not set
...1...= Syn: Set
...0...= Fin: Not set
Window size: 16384
Checksum: 0xdad2 (correct)
Options: (8 bytes)
Maximum segment size: 1460 bytes
NOP

```

a

```

Header length: 20 bytes
Flags: 0x0011 (FIN,ACK)
0...0...= Congestion Window Reduced (CWR): Not set
..0...0...= ECN-Echo: Not set
0...0...= Urgent: Not set
...1...= Acknowledgment: Set
...0...= Push: Not set
...0...= Syn: Not set
...0...= Fin: Set
Window size: 17328
Checksum: 0xd399 (correct)
SEQ/ACK analysis

```

b

Given the packet format of SYN and FIN packets, we define the filter format for SBL session logs as the following:

1. Set SYN bit to "1" and the Ack bit to "0". This is to get rid of the SYN ACK packets sent by the server back to the client, as it would not provide additional information other than the server's acceptance of the connection;
2. FIN bit must be set to "1".

A sample capture filter for session logging for Ethereal is shown below:

```
"tcp[tcpflags] & (tcp-syn) = 1 and tcp[tcpflags] & (tcp-ack) = 0  
OR tcp[tcpflags] & (tcp-ack) = 0"
```

Experimental results for session-based IP-Traceback

In the experiments, Ethereal was utilized to sniff network communications between one static residence host and other network end hosts. For each round of the experiment, the SYN and FIN packets were captured using the SBL scheme, as mentioned in the previous section. The experiment was run 33 times, with varied duration for each run. Ideally, the SBL session logging scheme should be implemented at not only the end hosts, but also intermediate routers. However, tabbing fast machines to the network to sniff the aforementioned information is more efficient than putting additional computation burdens on the routers. Thus, the experiments using Ethereal based on end-to-end systems are sufficient to investigate the performance of SBL logging. Table 2 shows the experiment results of the 33 experiments conducted. For each individual experiment, the number of packets captured, the number of logged packets, the size of all captured packets, the size of logged packets, and the duration of the experiment was recorded. Figure 2(a) shows the number of total packets vs. the number of logged packets for each run of the experiment. Log-scale is used in the y-axis for Figure 2(a) to display the data clearly. Figure 2(b) shows the ratio of the number of logged packets and the number of total packets in different experiments. Furthermore, Figure 3(a) shows the packet size of total packets vs. the packet size of logged packets in log-scale for each run of experiment, and Figure 3(b) shows the ratio of the packet size of logged packets to the total packets for each experiment.

From the figures, we observe that the average logged packet number and the average logged packet size are very low (0,0097 and 0,0027 respectively). Even for a large amount of total number of packet (or total packet size), the size of logged file remains limited. The maximum ratio between logged number of packets and the total number of packets is only around 0,043, and the maximum ratio between logged packet size and the total packet size is only 0,025.

Table 2. Experiment data

Experiment ID	# of all packets	# of logged packets	Size of all packets (KB)	Size of logged packets (KB)	Duration of the experiment (hh:mm:ss)
1	263370	1242	102402	95	24:10:21
2	113036	183	102402	14	1:12:07
3	111939	84	102007	7	2:46:58
4	135588	678	102401	52	15:14:14
5	12016	509	5988	39	1:03:57
6	119	4	40	1	0:00:51
7	146143	1221	102401	94	16:28:27
8	88193	1671	33850	128	3:59:37
9	173895	1044	99423	80	6:52:46
10	2190	35	1025	3	0:39:57
11	1529	15	1025	2	0:10:58
12	2373	40	1025	4	22:47
13	5967	11	1025	1	18:23
14	11204	1	1025	1	6:37
15	8554	8	1025	1	5:03
16	9060	11	1025	1	5:20
17	7178	8	1025	1	3:51
18	2863	10	1025	1	1:02
19	5607	8	1025	1	3:12
20	3898	5	1025	1	1:54
21	7232	9	1025	1	3:42
22	9934	6	1025	1	5:34
23	10490	4	1025	1	6:11
24	11091	19	1025	2	13:05
25	1936	30	1025	3	4:17
26	2031	50	1025	4	8:18

27	2318	70	1025	6	3:57
28	2808	58	1025	5	0:26:34
29	1101	15	473	2	0:29:37
30	151003	439	102401	34	3:42:28
31	29551	115	23946	9	2:17:33
32	105662	2018	54827	154	23:59:54
33	50264	428	11971	33	

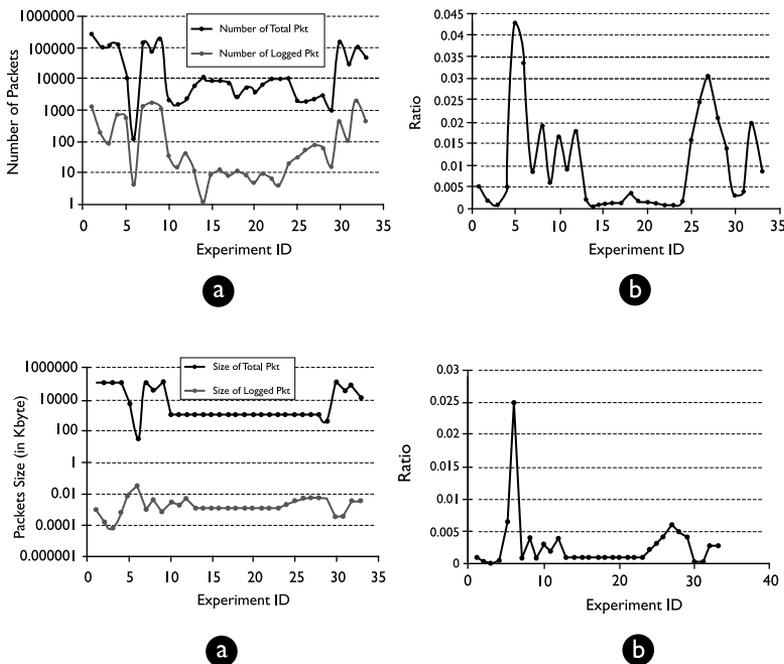


Figure 2.

- a Number of packets observed over different experiments
- b The ratio between number of logged packets and total packets over different experiments (average ratio = 0,0097)

Figure 3.

- a The size of packets (in KB) observed over different experiments
- b The ratio between the size of logged packets and the size of total packets over different experiments (average ratio = 0,0027)

This experiment method is suitable for a regular user. However, for large ISPs, the log files may be huge, and stored at multiple network entities (e.g., intermediate routers). Nonetheless, from a simple analysis, using SBL for a large amount of data logging, the storage space consumed by the log files is also reasonable. For example, for the network communication with a total data of 782KB, an experiment with about 110 hours in duration may be needed. Accordingly, an average log size for the experiment should be about $782/110 \times 24 = 171$ KB per day. Since the fraction between the size of logged data and the size of a SYN packet is "17/62",

the average size of the corresponding log file can be estimated as $171KB \times 17/62 \approx 50KB$. For a 1000 user environment the log file size would be approximately $50KB \times 1000 = 50MB$ per day, and about 1.5GB for a 30-day month. For a 1-million user environment, it would take approximately 50 GB per day and 1500 GB per month. These values are very reasonable for the log files recording session-based traffic for large ISPs. When compression techniques are used, another 2/3 of the storage space can be saved.

A SYN-based Packet Marking scheme

The session-based packet logging scheme requires all the routers along the communication path to keep a log for the session traffic, which requires a significant amount of storage resource at each router. When IP-Traceback is performed for a particular session, every single router along the route must be visited, and specific data must be retrieved. However, by using the SYN packet-based logging scheme, the storage space required for keeping the logs has been significantly reduced, and the searching algorithm for identifying specific sessions in the database is simplified, the overall task of log-based IP-Traceback for the entire network is still complicated. Therefore, in this section, another Protocol, the SYN-based Packet Marking (SYNPM) Protocol is proposed to enhance the performance of the session based IP-Traceback approach. In the SYN-based Packet Marking Protocol, instead of logging packets, the routers insert distinguishable *identifiers* in the SYN packet hop-by-hop. The identifiers are special signatures of routers, and are appended to the packets to record the fact that the corresponding router is along the communication path of the TCP session. However, in the meantime, no router keeps any log of the passing-by packets. The data storage burden is pushed all the way to the end systems. By doing this, the storage problem existing in a packet logging scheme is resolved easily. Other packet marking techniques have also been proposed along this line of research. However, in the existing packet marking schemes, either all of the packets or statistically chosen packets are marked. In this work, it is proposed that session-based packet marking scheme by marking only the first SYN packet for a session be conducted, as it is the packet that initializes the session and contains sufficient information for the IP-Traceback scheme in forensics networks. Using SYNPM, routers attach their *1-Byte Identifiers* to the SYN

packet of each session; thus, every connection is logged with the reverse path to the source, which makes some malicious attacks such as Denial-Of-Service attacks easy to be detected. The SYNPM scheme is presented in detail here.

- At first, the communication initiating peer of a session (e.g., the TCP client) generates the SYN packet for establishing a TCP connection. It is assumed that the initiating TTL value set by the end system is T , and the variable t is used to denote the changing TTL value of the SYN packet. Thus the SYNPM Protocol originally reserves T bytes in the data payload for packet marking to be conducted along the communication path. The T bytes of data are initialized as all zeros. For instance, if the TTL value is set to be 255 by the system, then the initiating peer would reserve 255 bytes of space in the data payload of the first SYN packet, and initialize these data into zeros.
- At each router, a piece of software is designed so that the router may perform a very simple modification on the data payload, and attach its identifier to the data field of the SYN segment; thus, when the SYN packet arrives at a router, the router attaches a 1-byte unique identifier at the last *available* byte field (i.e., the t -th byte field) of the reserved bytes (e.g., the 255-th byte for the first hop router when the initial TTL value is set to be 255) in the data. In the meantime, the router decrements the TTL value by one (i.e., $t=t-1$).
- The SYN packet is thereafter marked by routers attaching their special 1-byte identifiers one hop after another onto the t -th byte of the reserved packet marking field in data, where t is in fact used to keep track of the updated TTL value. This marked packet will eventually be stored in the target computer (for example, the data sender) as the log entry of the session.

By conducting the above steps, a trace of routers along the communication path of a session has been logged in a reverse order at the end of the reserved data payload in the SYN segment, and a log entry has been created by the destination system of the SYN packet. In order to conduct IP-traceback for a session, an end system takes one log entry as input and checks its recorded final TTL value (stored in variable t) which can be used to calculate the number of hops (i.e., $T-t$) that the packet traversed easily. The

information that should be used for IP-traceback thereafter is stored from the t -th byte to the last recorded byte of the data payload in the logged SYN packet. For example, if for a log entry the variable t is recorded as 242, it indicates that from the 242nd byte until the very last recorded byte in the data payload of the corresponding SYN packet, *unique identifiers* are attached by the routers along the communication path of that particular session. After obtaining the useful data portion from a SYN log, the end system may initiate an IP-traceback process by sending out a *trace command* with a unique query ID (qID) to the routers attached directly to the end system. Each router that receives this command extracts the *first* byte of the input data, which is supposed to be the identifier of the first-hop router of the end system, and the router checks whether the identifier matches one of its own. If a router has identified that itself is indeed a router logged by the SYN packet for the corresponding session, the router conducts the following three steps:

1. The router erases its ID (the first byte of the data portion in the query) from the query.
2. The router broadcasts the same trace command with the same query ID (qID) to the neighbouring routers attached to it. In this new trace command, modified query data is attached, where the router has erased its own record from the logged trace data.
3. The router sends its own IP address directly to the query-originating end host with the query number (qID).

The above process is conducted hop-by-hop iteratively from the “first-hop” router(s) attached to the query-initiating host until the *original source* recorded by the SYN packet. By having every identified router send back its real IP address, the query originator may recover the original routing trace of a logged session and discover the authentic source of the TCP connection.

Figure 4 shows a sample topology to demonstrate the Packet Marketing mechanism. In Figure 4, R_t represents Router, C1-C4 represents computers 1 to 4, and the numbers 1, 2 and 3 beside the arrows represent the connection initialization requests from different computers. The TTL field is originally set as 255.

of SBL by effectively simplifying the overall task of log-based IP-Traceback through inserting unique signatures of each router during the SYN segment transmission. Since a router does not need to keep any log information of the passing-by packets, the storage problem existing in traditional packet logging scheme is resolved easily. Furthermore, in this scheme, only the first SYN packet for a session is marked and recorded, which is different from other pattern marking schemes that usually mark either all of packets or statistically chosen packets. Here, only the first SYN packet is marked, as it is the packet that initializes the session and it contains a sufficient amount of information for IP-Traceback in forensics networks.

The two approaches proposed in this paper focus on logging packets, particularly in TCP sessions. However, many network security problems reside in UDP connections, and having session-based logs for UDP traffic would be very useful for investigating UDP-based attacks. In our future work, we would like to extend our research on session-based logging scheme, in the direction of recording UDP traffic efficiently.

References

BurchHal, Cheswick Bill (2000), *Tracing anonymous packets to their approximate source*. Proceedings of the 14th USENIX Conference on Systems Administration (LISA), New Orleans, LA, USA, 3-8 December, pp. 319-327

CERT Advisory CA-2000-01(2002), *Denial-Of-Service Developments*
<http://www.cert.org/advisories/CA-2000-01.html>

CERT Advisory CA-98-01(1998), *Smurf IP Denial-Of-Service Attacks*
<http://www.cert.org/advisories/CA-1998-01.html>

CERT Incident Note IN-2000-04 (2000), *Denial-Of-Service attacks using name servers*, http://www.cert.org/incident_notes/IN-2000-04.html

EtherealPacket Sniffer (2007), <http://www.ethreal.com>

Morris Robert T. (1985), *A weakness in the 4.2BSD Unix TCP/IP Software*, AT&T Bell Laboratories, Murray Hill, NJ, USA [Bell Labs Computer Science Technical Report 117]

All URLs checked
December 2007

Sager Glenn (1998), *Security fun with OCxmon and cflowd*, Internet 2 Working Group Meeting, 20 November

Savage Stefan, Wetherall David, Karlin Anna, Anderson Tom (2001), *Network support for IP traceback*, "IEEE/ACM Transactions on Networking", V. 9, n. 3, pp. 226–237

Snoeren Alex C., Partridge Craig, Sanchez Lui A., Jones Christine E., Tchakountio Fabrice, Schwartz Beverly, Kent Stephen T., Strayer Timothy W. (2002), *Single-Packet IP Traceback*, "IEEE/ACM Transactions on Networking", V. 10, n. 6, pp. 721-734

Stone Robert (2000), *CenterTrack: an IP overlay network for tracking DoS floods*. Proceedings of the 9th Conference on USENIX Security Symposium, Denver, CO, USA, 14-17 August, V. 9, p. 15

Sintesi

Uno dei maggiori problemi relativi all'utilizzo di tecnologie informatiche è la protezione delle connessioni internet (IP – Internet Protocol) da intrusioni nel proprio sistema operativo, dal blocco di un servizio, dai furti di risorse informative e altro ancora. Il problema riguarda in particolare la possibilità di determinare esattamente l'origine di un certo protocollo internet (network packet), soprattutto ai fini delle investigazioni riguardo i crimini informatici. È infatti appurato che ad ogni "IP address" (cioè ad ogni identificabile "centro di smistamento") sono collegabili illimitati indirizzi personali. Rintracciare, tra questi ultimi, l'autore di un determinato crimine informatico è piuttosto difficile, anche perché possono verificarsi problemi di privacy. Vengono quindi presentati due nuovi sistemi di tracciabilità informatica: la Session-Based packet Logging (SBL) e la SYN-based Packet Marking (SYNPM). Le loro caratteristiche, generali e innovative, consistono nel fatto che sono facilmente implementabili, non pongono problemi di privacy, in quanto non necessitano di dettagli per ogni connessione individuale e, a differenza degli usuali sistemi di tracciabilità, mantengono le informazioni di connessione per un lungo periodo di tempo. Il primo, SBL, si basa sull'adattamento delle tecniche utilizzate per le indagini sulle comunicazioni telefoniche nell'informatica forense: si basa quindi sulla memorizzazione dell'indirizzo IP e della durata della connessione ben oltre il momento del login. Il sistema SYNPM affida a ciascun router (periferiche che smistano le informazioni ricevute da più computer connessi ad una stessa rete) una distinta firma identificativa. Il sistema è cioè diretto a "marcare" in modo distinguibile la prima connessione di ciascun utente. Gli autori, oltre a descrivere dettagliatamente i due sistemi, comprensivi dei

risultati ottenuti e dei possibili sviluppi futuri, offrono anche un'interessante disamina sui sistemi maggiormente utilizzati fino ad oggi, basati sull'IP-Traceback, evidenziandone le criticità. Nella maggior parte dei casi, si tratta infatti di tecniche che determinano dei falsi positivi, che non riescono ad identificare l'esatta origine dell'attacco informatico o che si limitano a bloccare l'attacco su segnalazione dell'utente. In genere, sono sistemi che difficilmente riescono a funzionare in poco tempo, concedendo un grande vantaggio temporale per il criminale informatico che può cancellare o alterare le tracce del proprio passaggio.

In definitiva, i due sistemi presentano molti vantaggi applicativi, sia dal punto di vista dell'efficienza sia dal punto di vista del risparmio dello spazio di memoria. Il limite, attualmente, è rintracciabile nel loro utilizzo solo per i protocolli TCP. Per i protocolli UPD sono attualmente in corso sperimentazioni e proposte di adattamento che fanno ben sperare per la completa efficienza delle investigazioni informatiche.