A SIGNALING ARCHITECTURE AGAINST DOS ATTACKS

Ahmad Fadlallah, Ahmed Serhrouchni
Ecole Nationale Supérieure des Télécommunications, GET - CNRS UMR 5141 – LTCI
Departement Informatique et Reseaux (InfRes)
46, rue Barrault - 75 634 Paris Cedex 13 – France

Keywords: Network Security, Denial of service, signaling protocols.

Abstract: Denial of service (DoS) attacks figure highly among the dangers that face the Internet. Many research studies deal with DoS, proposing models and/or architectures to stop this threat. The proposed solutions vary between prevention, detection, filtering and traceback of the attack. The latter (attack traceback) constitutes an important part of the DoS defense. The most complex issue it has to face is related to the fact that attackers often used spoofed or incorrect IP addresses, thus disguising the true origin. In this work, we propose a signaling architecture and a security-oriented signaling protocol named 3SP (Simple Security Signaling Protocol). This solution makes it easier to trace both the DoS and other types of attack back to their sources; it is simple, robust and efficient against IP spoofing, and thus constitutes a novel and efficient approach to deal with the attack traceback problem.

1 INTRODUCTION

The Internet is on track to becoming the backbone network for all telecommunication. Internet security is of critical importance in today’s computing environment, given that our society, government, and economy are increasingly relying on the Internet. Denial of service attacks are considered among the hardest security problems that threaten the “Digital society”; according to the CST/FBI 2004 computer crime and security survey, denial of service attacks occupied the second position (behind virus attacks) on the attack list in terms of the caused losses.

According to the CERT (CERT, 1997), “a denial of service attack is characterized by an explicit attempt by attackers to prevent legitimate users of a service from using that service”. Its distributed form: the distributed denial of service (DDoS) attack is one in which a multitude of compromised systems attack a single target thereby causing denial of service.

Many solutions have been proposed to stop such type of attacks, which can be classified into three categories: attack prevention, detection, and response. The latter can be mainly split into filtering attack packets and tracing back to the attack source(s).

In this paper, we propose a signaling architecture, which facilitates the attack traceback and filtering processes.

Our contribution is three-fold. First, our scheme is efficient in defending against denial of service attacks (even when they use concealment techniques such as IP spoofing). Second, this mechanism does not need the global cooperation of the whole Internet community but an increase in deployment increases its benefit. Third, our scheme is applicable to all traffic types, and thus can be used to traceback other attack types.

The remainder of this paper is organized as follows. Section 2 presents the denial of service attacks. Section 3 reviews the related work in this field with special regard to the main techniques proposed for attack traceback. Section 4 explains the reasoning behind our proposal, describes the architecture and the 3SP protocol. Section 5 discusses the advantages of our solution and some open issues. Finally, section 6 concludes the paper.

Fadlallah A. and Serhrouchni A. (2005).
A SIGNALING ARCHITECTURE AGAINST DOS ATTACKS.
In *Proceedings of the Second International Conference on e-Business and Telecommunication Networks*, pages 216-221
DOI: 10.5220/0001415702160221

2 DENIAL OF SERVICE ATTACKS

Denial of Service (DoS) is an attack designed to render a computer or network incapable of providing normal services (Stein, 2003). The attack may exhaust a key resource by exploiting software vulnerability (vulnerability attacks) or by simply sending – by an attacker – a higher volume of traffic than the victim is provisioned to handle (flooding attacks). This kind of attack is possible to defend by installing patches provided by operating system or software manufacturers, or by simple configuration of network equipment (routers, firewalls, IDS…).

In the summer of 1999, a new breed of attack has been developed called distributed denial of service (DDoS) attack. Simply put, a DDoS attack saturates a network. It simply overwhelms the target server with an immense volume of traffic that prevents normal users from accessing the server. These distributed attacks rely on recruiting a fleet of compromised "zombie" computers (also named agents) that unwittingly join forces to flood the victim server.

Moreover, attackers can render distributed denial-of service attacks more difficult to defend against by bouncing their flooding traffic off reflectors; that is, by leading zombies to spoof requests from the victim to a large set of Internet servers that will in turn send their combined replies to the victim. This type of attacks is known as distributed reflector denial of service (DRDoS) (Paxon, 2001).

3 DENIAL OF SERVICE DEFENSE

Because of the seriousness of the problem many defense mechanisms have been proposed to combat these attacks. There are three lines of defense against the attack: attack prevention and preemption (before the attack), attack detection (during the attack), and attack response (during and after the attack).

The first line of defense is obviously to prevent DDoS attacks from taking place. The detection of an attack is responsible for identifying DDoS attacks or attack packets. Once an attack has been detected, an ideal response is to filter attack traffic and identify attack source.

Traceback and filtering can be complementary tasks; since traceback allows us to go nearer to the attack source(s), where filtering is easier.

Filtering attack traffic is implicit in most detection solutions like (Mirkovic & al., 2002) and (Gil & al., 2001); still there is the problem of eliminating false positives that usually lead to collateral damage. Unfortunately, there is no easy way to track IP traffic to its source due to the statelessness of the IP protocol, and to IP spoofing. In order to address this limitation, several approaches have been proposed.

The IP marking approaches (Savage & al., 2000) (Song & al., 2001) (Park & al.-1, 2001) enable routers to mark packets with partial path information and try to reconstruct the complete path from the packets that contain the marking.

ICMP traceback (iTrace) (Bellovin, 2001) proposes to introduce a new message "ICMP trace back" (or an iTrace message) so that routers can generate iTrace messages to help the victim or its upstream ISP to identify the source of spoofed IP packets. An intention-driven iTrace is also introduced to reduce unnecessary iTrace messages and thus improve the performance of iTrace systems (Mankin & al., 2001).

(Dean & al., 2001) proposes an algebraic approach to transform the IP traceback problem into a polynomial reconstruction problem, and uses techniques from algebraic coding theory to recover the true origin of spoofed IP packets.

CenterTrack (Stone, 2000) is an IP overlay network that selectively reroutes suspect IP packets directly from edge routers to special tracing routers, which can easily determine the ingress edge router by observing from which tunnel the packets arrive.

(Sanchez & al., 2001) develop another traceback solution named Source Path Isolation Engine (SPIE); it stores the message digests of recently received IP packets and can reconstruct the attack paths of given spoofed IP packets. There are also many other techniques and issues related to IP traceback (e.g., approximate traceback (Burch & al., 2000), legal and societal issues (Lee & al., 2001), and vendors' solutions (Cisco, 1999).

There are – in general – two main limitations of existing traceback solutions; first, the need for more complex algorithms for path reconstruction, and second the need for a global deployment – of some mechanisms – in order to be efficient.

4 SIGNALING ARCHITECTURE FOR DOS ATTACKS

In this section, we present a signaling architecture, which can be situated as a response mechanism that makes it easier to traceback and filter a DDoS attack. While designing this solution, we tried to keep in mind all the lessons learned from existing proposals, trying to re-use their strong points as much as possible and avoid their weaknesses.

The first question to answer is why using a signaling approach? or to put it in another way: how can signaling facilitate the difficult tasks of traceback and filtering?

In fact, signaling is the best solution to identify the source of a given traffic; a reliable signaling mechanism is also very efficient against concealment techniques – such as IP spoofing – used by attackers. A typical example is in telephone networks where one of the main reasons for the absence of DoS is that each call can be traced back – through signaling – and the caller identified. Moreover – as previously stated – the ability to identify the source of the attack (or at least to go the nearest possible) will make it easy to filter attack packets.

4.1 Architecture

In the signaling architecture, signaling messages are only received, processed and sent by traceback signaling entities (TSE). These entities can be placed in the network devices. Depending on their functionalities, we differentiate between two types of TSE:

- Light traceback signaling entities (LTSE): It processes and forwards signaling messages. It is stateless in a sense that it does not maintain any information about a signaled flow.
- Full traceback signaling entities (FTSE): In addition to the LTSE functionalities, an FTSE monitors internet traffic, and may initiate a signaling session (generate signaling messages) for a given flow. It also maintains per-flow state information (statefull).

Filtering functions can also be added to both signaling entities, in order to filter/rate-limit attack packets when detected.

4.2 Simple Security Signaling Protocol (3SP)

The signaling entities communicates through the simple security signaling protocol (3SP), which – as its name indicates – is a simple protocol that provides a traceback-oriented signaling. To this purpose, it provides a set of messages enabling signaling between signaling entities before and during the transmission of a given flow.

We first present the general behavior of the protocol in a traceback scenario, and then describe the messages used in the different protocol exchanges.

4.2.1 Traceback scenario

In the following, the source TSE designates the TSE that initiates the signaling session. The destination TSE is the signaling entity that ends the signaling session, while the forwarder TSE represents an intermediate TSE (between source and destination).

The behavior of the 3SP protocol is presented under simplifying assumptions; we consider an attack with one attacker and one target host. The same mechanism can work in the case of multiple attackers.

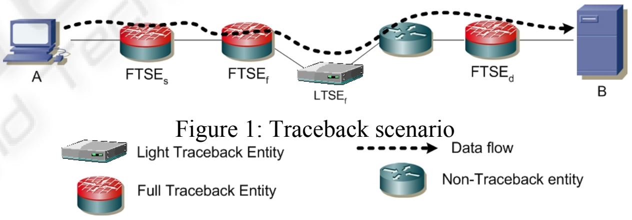

Figure 1: Traceback scenario

The figure above illustrates a traceback scenario in which appear the different TSEs.

A user (A) (that might be an attacker) sends traffic with destination to (B). The source FTSE (the first one along the data path) (FTSE$_s$) sends a signaling request towards (B). This message contains the description of the signaled flow.

When a forwarder FTSE (FTSE$_f$) receives the signaling request, an authentication process is performed with the previous TSE, and if successful, the forwarder adds it address in the router address list (contained in the request), installs a signaling state and forwards the request towards (B).

An LTSE acts similarly to a forwarder FTSE but without maintaining a signaling state. The non-signaling routers simply forward messages as normal packets based on the destination IP address.

When the signaling request reaches the destination FTSE (FTSE$_d$), an authentication process is performed with the last TSE, and if successful, a signaling state is installed.

A signaling state contains information about the signaled flow (sender/receiver IP address/port, higher level protocols); it also contains a session identifier that uniquely identify a signaled flow (used also for refresh mechanism), the router list and some authentication information.

When receiving a signaling request, an FTSE acts as a forwarder or a destination depending on the destination of the request; if the request is intended to a host which is out of its "protection scope", the FTSE forwards the message (forwarder), otherwise it acts as a destination FTSE.

It is important to note that the "state" term is somehow different from the same term used in reservation signaling protocols; a state – here – is used to avoid redundancies in signaling information and especially for long-period flows. In this context, 3SP takes a "soft state" approach (control states in hosts and routers will expire if not refreshed within a specified amount of time) to managing the state in routers; all the information is – no doubt – logged before their deletion. Note that a 3SP state is refreshed by repeating the same first signaling process. Figure 3 illustrates the different 3SP exchanges between the signaling entities.

The traceback mechanism is simple; once attack traffic is detected, a simple search based on the traffic description allows the detection of the signaling entities (router list) along the attack path. This search can be carried out either in the existing signaling states in the case of an ongoing attack or in the logged signaling states in the case of a traceback after an attack.

It is essential to have a fast lookup process. There are many solutions; for example one solution may be to use an appropriate hash function to build a single hash table. Another solution is to use a Bloom filter (Bloom, 1970), etc.

4.2.2 3SP messages

The table below describes the different messages used for router communication.

Table 1: 3SP protocol messages

N	Message Name	Description
1	SIGNAL_REQUEST	sent by source FTSE to "signal" a given flow
2	SIGNAL_AUTHENTICATION	contains the information needed to perform the authentication process
3	SIGNAL_ERROR	used to signal an occurred error

Each 3SP message is composed of a "common" header and a payload consisting of variable length, typed objects. In the following, we define the format of the common header and each of the 3SP message types.

The 3SP common header includes the following fields:

Vers: 4 bits
Protocol version

Flags: 8 bits
Reserved for future enhancements

Msg Type: 4 bits
Indicates the message type.

1 = SIGNAL_REQUEST
2 = SIGNAL_AUTHENTICATION
3 = SIGNAL_ERROR

3SP checksum: 16 bits
Contains the checksum, computed on the entire message.

3SP length: 16 bits
The total length of the 3SP message in bytes, including the common header and the variable-length objects that follow.

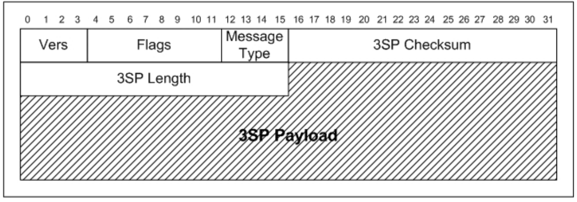

Figure 2: 3SP packet

Each 3SP message has to satisfy a number of

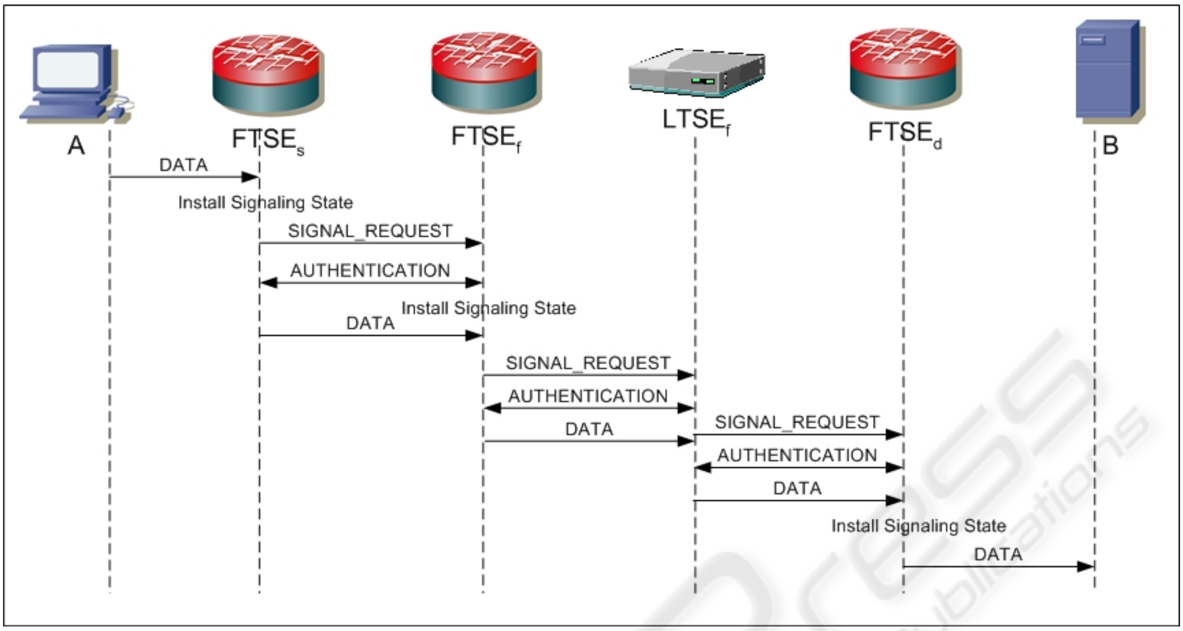

Figure 3: 3SP message flow

syntax rules. These rules are specified using Backus-Naur Form (BNF). The BNF implies an order for the objects in a message. However, in our case, object order makes no logical difference.

<SIGNAL_REQUEST Message> ::=
<3SP header> <ID_SESSION> <Traffic Descriptor>

<SIGNAL_AUTHENTICATION Message>::=
<3SP header><ID_SESSION> <AUTH_INFO>

<SIGNAL_ERROR> ::= <3SP header>
<ID_SESSION><Error Descriptor><AUTH_INFO>

5 DISCUSSION

5.1 Advantages

Compared with other traceback solutions, the proposed architecture presents several advantages:

- As its name indicates, 3SP is simple in terms of design and implementation. Moreover, this simplicity is also reflected in the attack path computation process which is a very important criterion for traceback mechanisms. 3SP makes this easy to achieve; there is no need for complex path reconstitution algorithms such those used in packet marking based traceback solutions.
- The traceback mechanism is efficient, even when attackers use concealment techniques such as IP spoofing. It is also able to trace a single packet back to its source.
- The solution is robust in a sense that it does not generate any false positive path.
- 3SP can be easily extended and coupled with other defense mechanisms such as detection and filtering mechanisms, within the framework of a global defense solution. 3SP facilitates – in this way – the interoperation between heterogeneous defense systems.
- The traceback mechanism is able to work properly even in the case where some of the routers does not support it. In the presence of an edge-backbone-edge network infrastructure (which is the most common nowadays), just two 3SP-capable routers are enough for the protocol to function. The 3SP benefits grow incrementally as number of deployment points increases.
- 3SP can be useful for other purposes than the DoS attack traceback; in fact, 3SP signaling can be associated with specific flows and thus used to traceback them. In this context 3SP can be useful for other attack types, such as spam.
- 3SP is a network signaling protocol, so there is no need for any modifications from the client side.

5.2 Open issues

While this scheme appears to be a promising direction, we recognize that there are some issues to be further addressed in the future. First, the 3SP is not able to efficiently handle transformed attack packets, second the authentication scheme must be carefully chosen in order to maintain a good performance while preserving a high level of security for the overall signaling mechanism.

6 CONCLUSION

Denial of service attacks are a great threat that faces the Internet today. Many solutions were proposed to combat these attacks, including prevention, detection, filtering and traceback solutions.

In this paper, we presented a signaling architecture, which facilitates the traceback of Internet attacks and in particular DDoS attacks, even with concealment techniques such as IP spoofing.

The traceback mechanism can be easily coupled with other defense mechanisms such as attack detection and filtering, and thus constitutes a fundamental part of global defense architecture.

The signaling protocol 3SP has the capability of functioning properly even in the presence of non 3SP-capable routers, thus enabling incremental deployment of the protocol itself.

An experimental implementation of the 3SP protocol is currently in progress, in order to evaluate its performances and test the different possible authentication mechanisms.

REFERENCES

Bellovin, S.M., 2001. ICMP traceback messages, Internet draft.

Bloom, B. H., 1970. Space/time tradeoffs in hash coding with allowable errors. *Communications of the ACM*, 13(7):422– 426.

Burch, H., Cheswick, B., 2000. Tracing Anonymous Packets to Their Approximate Source, *in Proceedings of the 14th USENIX Systems Administration Conference*.

CERT® Coordination Center, 1997. “Denial of Service Attacks”. Available from:

http://www.cert.org/tech_tips/denial_of_service.html

Cisco Systems, 1999. Characterizing and Tracing Packet Floods Using Cisco Routers.

Dean, D., Franklin, M., Stubblefield, A., 2001. An algebraic approach to IP Traceback. *In Proceedings of the 2001 Network and Distributed System Security Symposium*.

Gil, T.M., Poleto, M., 2001. MULTOPS: a data-structure for bandwidth attack detection, *in Proceedings of 10th Usenix Security Symposium*.

Lee, S. C., Shields, C., 2001. Tracing the Source of Network Attack: A Technical, Legal and Societal Problem, *in proceedings of the 2001 IEEE Workshop on Information Assurance and Security*.

Mankin, A., Massey, D., Wu, C., Wu, S. F., Zhang, L., 2001. On Design and Evaluation of Intention-Driven ICMP Traceback, *In Proceedings of IEEE International Conference on Computer Communications and Networks*.

Mirkovic, J., Prier, G., Reiher, P., 2002. Attacking DDoS at the source. In *Proceedings of ICNP 2002*, pp. 312–321

Park, K., Lee, H., 2001. On the Effectiveness of Probabilistic Packet Marking for IP Traceback, *In Proceedings of IEEE INFOCOM* 2001.

Paxon, V., 2001. An analysis of using reflectors for distributed denial-of-service attacks. *Computer Communication Review.*

Sanchez, L.A., Milliken, W.C., Snoeren, A.C., Tchakountio, F., Jones, C.E,. Kent, S.T., Partridge, C., Strayer, W.T., 2001. Hardware Support for a Hash-Based IP Traceback, *in Proceedings of DARPA Information Survivability Conference & Exposition*.

Savage, S., Wetherall, D., Karlin, A., Anderson, T., 2000. Practical network support for IP Traceback, *In Proceedings of 2000 ACM SIGCOMM Conference*.

Song, D. X., Perrig, A., 2001. Advanced and authenticated marking schemes for IP Traceback, *IEEE INFOCOM 2001*.

Stein, L.D., Stewart, J.N., 2003. *The World Wide Web Security FAQ, v 1.7*.

Stone, R., 2000. CenterTrack: An IP Overlay Network for Tracking DoS Floods, *In Proceedings of 9th Usenix Security Symposium*.