ANON: An IP-Layer Anonymizing Infrastructure

Chen-Mou Cheng, H.T. Kung, Koan-Sin Tan, and Scott Bradner
Harvard University
doug@eecs.harvard.edu, kung@harvard.edu, freedom@acm.org, sob@harvard.edu

Abstract

This exhibition demonstrates an IP-layer anonymizing infrastructure, called ANON, which allows server addresses to be hidden from clients and vice versa. In providing address anonymity, ANON uses a network resident set of IP-layer anonymizing forwarders that can forward IP packets with nested encryption and decryption applied to their source and destination addresses. To prevent adversaries from compromising the anonymity by learning the forwarding path, ANON incorporates a suite of countermeasures, including link padding and non-malleable, semantically secure link encryption. To prevent denial of service (DoS) attacks through the anonymizing infrastructure itself, ANON uses rate limiting. Finally, to increase the resilience against attacks and infrastructure failures, ANON uses redundant forwarders with anycast addresses and a fault-tolerant overlay network to connect forwarders.

1. Introduction

Over the current Internet, when a client acquires services from an application server, called a *target server* subsequently, packets sent and received by the client reveal the server's IP address in the packet headers. There are a number of situations where it would be useful for an application to be able to communicate with a destination without revealing the IP address of the destination to the source, the IP address of the source to the destination, or both. The IP address of a destination may also want to be hidden from the public, beyond just clients. For example, a Web site may want to hide its IP addresses to reduce the risk of denial of service (DoS) attacks aimed at these addresses, or an organization may want to ensure its anonymity by not revealing its IP addresses.

One way to achieve this anonymity, as described in this exhibition, is to use an overlay network connecting a network resident set of IP-layer servers that can forward IP packets, with nested encryption and decryption applied to their source and destination addresses. We will call these network resident IP-layer servers *anonymizing forwarders*, or simply *forwarders*, and an IP anonymizing infrastructure based on these anonymizing forwarders an *forwarding infrastructure,* or simply *ANON*.

Using ANON, a client can send and receive packets to and from target servers by using their *handles* rather than IP addresses, where handles are information strings from which forwarders in the ANON infrastructure can translate into their corresponding IP addresses. This is analogous to an organization sending and receiving U.S. mail using its P.O. Box number rather than street address, where the P.O. Box number is the handle to the organization and postal offices correspond to ANON forwarders.

The design of ANON assumes that it will be used mainly for low- to medium-bandwidth signaling and data applications, not data transfer that may require high bandwidth. This assumption makes ANON countermeasures against various security threats effective. There are many applications that fit the model defined here, including signaling protocols such as connection setup and termination, user authentication and authorization, service discovery and registration, and instant messaging.

2. Threat Model

The ANON forwarding infrastructure provides countermeasures to the following three types of threats:

- Type 1 threat (unauthorized address discovery). The forwarding infrastructure may leak address information that it is supposed to hide.
- Type 2 threat (inband DoS attacks). The forwarding infrastructure may be used as a conduit to launch DoS attacks on forwarders or target servers.
- Type 3 threat (outband DoS attacks). Forwarders in the infrastructure may themselves be subject to DoS attacks through external network paths.

We assume that forwarders are managed by trustworthy third parties so that they cannot be compromised. In addition, we assume that the location and addresses of

Proceedings of the DARPA Information Survivability Conference and Exposition (DISCEX'03)
0-7695-1897-4/03 $17.00 © 2003 IEEE

forwarders are not publicized. This means that adversaries will not know the addresses of forwarders beyond the first hop, without tracing the forwarding infrastructure.

However, after having located a forwarder, we assume that adversaries can monitor traffic in and out of the forwarder. We feel that this strong adversary model is justified in view of the fact that it is possible to monitor a particular set of links even without physical wiretapping. For example, by tricking routers to think that there is a shorter path to or from a forwarder, an adversary will be able to direct traffic to or from the forwarder to his own networks for monitoring and recording purposes.

An example attack related to type 1 threat works as follows. Acting as a legitimate client, an adversary sends probe packets to a target server whose address he intends to discover. By specially marking his packets or transmitting them according to certain timing patterns, referred to as *packet tagging* and *traffic tagging*, respectively, and by using link monitoring, the adversary can try to identify these packets on links of the forwarding infrastructure and trace through these links to discover the address of the target server. ANON uses techniques such as link encryption and link padding to defend against these attacks.

In type 2 threat, an adversary, again acting as a legitimate client, sends a large number of packets to a target server with the intention of swamping the server or forwarders on the forwarding path. Because the attack uses the infrastructure itself, we call it an *inband DoS attack*. ANON uses rate limiting to curtail these attacks by having upstream forwarders drop excessive traffic.

In type 3 threat, an adversary, after having discovered the address of a forwarder, sends DoS attack packets to the address using network paths external to the forwarding infrastructure. Since the attack does not use the forwarding infrastructure, we call it an *outband DoS attack*. To address type 3 threat, ANON forwarders need to be resilient to DoS attacks, which can be achieved through the use of anycast-style addresses, or the use of a fault-tolerant transport network such as Chord.

Our assumption that forwarders can not be compromised significantly simplifies our threat model. This assumption has allowed us not to be concerned with attacks originating from compromised forwarders and other forwarder-facilitated attacks. We are looking into a revised threat model where some of the forwarders might have been compromised.

3. ANON Design Objectives

There are two main objectives for the design of the ANON forwarding infrastructure:

- Stateless, real-time decrypting and forwarding of IP packets with nested encryption and decryption applied to their source and destination addresses.
- Providing defense mechanisms against the three types of threats described above.

4. The ANON Infrastructure

The ANON infrastructure consists of a set of anonymizing forwarders and some number of initialization servers, as depicted in Figure 1. Forwarders will encrypt and decrypt IP addresses, whereas initialization servers will provide clients with handles to target servers, which consist of addresses of entry forwarders and nestedly encrypted addresses of target servers. An overlay network connects the forwarders. That is, a pair of forwarders may be connected using a path involving multiple IP routers. Request packets from a client to a target server will be forwarded over a forwarding path consisting of a subset of these forwarders. Reply packets from the target server to the client will use the same path in the reverse direction. For different reply-request sessions, different forwarding paths may be used.

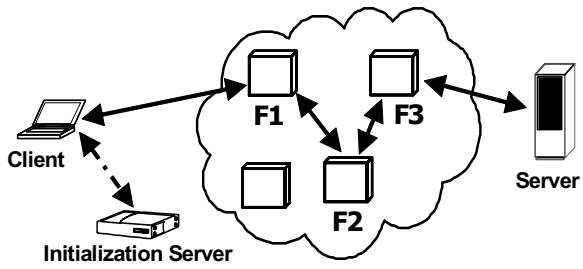

Figure 1. The ANON infrastructure. F1, F2, and F3 are anonymizing forwarders, and the solid arrows indicate an instance of a packet forwarding path.

To increase availability, forwarders may use anycast-style addressing so that any of a number of forwarders using the same anycast address may forward a packet sent to it.

The role of initialization servers is to provide clients with handles to target servers. Thus, initialization servers and the handles provided by them need to be properly authenticated, possibly using digital certificates, to ensure that these encrypted addresses will be trustworthy. In addi-

tion, initialization servers may need to be replicated in various locations to ensure their high availability.

Consider, for example, the case of hiding the IP address of a target server from clients. In this case, the use of ANON will involve three usage steps:

- Server registration. A target server whose IP address needs to be hidden will invoke a process that selects a sequence of forwarders, builds a handle for the target server, and registers the results to initialization servers. The sequence of forwarders can be selected manually or automatically, as well as statically or dynamically. Note that initialization servers do not know the real IP addresses of target servers, so compromising an initialization server does not lead to compromise of target servers' address anonymity.
- Client initialization. Given a target server to which a client wishes to access, the client obtains the handle to the target server from an initialization server.
- Packet forwarding. Based on the information obtained from the client initialization, ANON forwards packets to and from the target server over the selected sequence of forwarders.

5. Testbed Implementation

We have implemented a laboratory testbed for ANON at Harvard supporting basic functions such as link encryption, link padding, protocol camouflaging and rate limiting. The testbed consists of seven forwarders. Nodes in the testbed are implemented on top of FreeBSD operating system. The FreeBSD divert socket is used to implement various header processing operations at the user-level. For the symmetric key algorithm, the testbed uses the AES reference implementation from NIST.

The testbed includes NAT gateways, GWc for clients and GWs for servers. These gateways allow existing clients and servers to use the testbed without modifications. Clients and servers may connect to their respective gateways directly or via a VPN connection.

The current testbed implementation can achieve a throughput of 5 Mbps. This performance, adequate for signaling protocols and low-bandwidth data applications, is made possible mainly because we have managed to avoid using public key encryption and decryption in packet forwarding.

6. Summary and Concluding Remarks

We have developed an anonymizing infrastructure at the IP layer. The infrastructure is specially designed for low- to medium-bandwidth applications such as authentication, authorization and instant messaging. By employing a suite of countermeasures, such as non-malleable, semantically secure packet encryption, link padding and rate limiting, we have shown that even if an adversary is capable of monitoring links, it would be difficult for him to compromise the anonymity provided by the infrastructure. Likely, the only way an adversary can succeed is to take on the direct attack of compromising forwarders one by one. By using trustworthy third parties to manage the forwarding infrastructure, in a way similar to how current backbone routers are managed, we can make sure that compromising forwarders would be very difficult.

To lower the cost of link padding, we have designed two novel algorithms that maintain the same level of anonymity but generates padding traffic in an economical manner. We have also developed an approach of using fault-tolerant overlay networks to enhance the resilience of the anonymizing infrastructure against attacks and failures. To demonstrate the implementation feasibility of ANON, we have developed a laboratory testbed.

An important next step that we plan to carry out is application trials of this anonymizing infrastructure. When sufficient experiences have been learned from these application experiments, we will consider the possibility of incorporating some of the anonymizing features into routers.

Acknowledgment

This work was supported in part by DARPA through AFRL/IFKD under contract F33615-01-C-1983