FONet: A Federated Overlay Network for DoS Defense in the Internet

Jinu Kurian and Kamil Sarac
Dept of Computer Science, University of Texas at Dallas
jinuk@student.utdallas.edu, ksarac@utdallas.edu

1 Introduction

Despite years of research and industrial interest towards preventing them, Denial of Service (DoS) attacks continue to pose a significant threat to the health and utility of the Internet. Over the years, several DoS defense approaches have been proposed by the research community. Broadly classified, these methods can be either *reactive* methods or *proactive* methods. Reactive methods are usually based on anomaly detection, pattern detection or statistical methods and aim to alleviate the effect of an attack on the victim site by filtering out or pushing back attack traffic. Proactive methods on the other hand aim to prevent the possibility of a DoS attack happening or enabling a victim site to continue to provide service to legitimate users even when it is under attack.

In recent years, proactive methods have found increasing interest in the research community over reactive methods. This is mainly due to the difficulty in reactive methods of distinguishing legitimate traffic from attack traffic and the poor longevity of many distinguishing mechanisms. Receiver-controlled communication methods [6, 7] have been proposed recently as a proactive defense mechanism. This method aims to provide an alternative to the push-based, sender-controlled nature of unicast. The advantage of this method is that the use of receiver-control effectively prevents the possibility of a sender (attacker) launching attacks at a receiver (victim) site. However, the previously proposed approaches require extensive modifications to the Internet for their deployment.

Another example of a proactive defensive mechanism is the overlay-based approach [4, 2]. This method uses overlay routing combined with lightweight filtering to reduce the probability of a successful attack on a protected target. The advantage of overlay-based mechanisms is that they provide an effective defense mechanism while requiring minimal changes in the network for their deployment. However in existing proposals, the overlay nodes themselves are vulnerable to attack. So, the architecture has to rely on resource redundancy and expensive overlay routing to protect itself from attack.

In this paper we present a snapshot of our proposal for an architecture that can provide DoS resistant communication services in the Internet. We aim to build this architecture by combining the advantages of overlay-based and receiver-control based approaches while avoiding the disadvantages of their previous implementations. Also pivotal to the design of our architecture is the observation that many of the DoS solutions that have been proposed over the years (including basic mechanisms like ingress/egress filtering) have not motivated ISPs to deploy them in their networks. To this end we additionally require our architecture to be cost-effective in that it imposes minimal changes in the network. It should also be able to provide different levels of DoS protection services to suit the security requirements of a wide spectrum of customers; *e.g.* Web Servers, Content servers, FTP servers etc. These services will be *value-added* in nature and can be used to cater to the needs of a growing population of Internet users who are sensitive to cyber-attacks and are willing to pay for better protection. Deployment of these services can provide customers with a cost-effective method to alleviate their losses due to DoS attacks. This gives them an incentive to utilize the FONet services and thus an added financial incentive to ISPs to deploy FONet nodes in their networks.

2 Architecture

Our architecture (Fig 1) consists of overlay nodes (called FONet nodes) individually deployed by ISPs in their domains. FONet nodes communicate with each other using a new network layer protocol, p2cast, which provides DoS resistant channels between them. Management of individual overlay nodes is federated to the ISPs deploying them while the network as a whole is shared between all deploying domains.

Within a domain there are two types of FONet nodes: *(1)* Access FONet nodes which authenticate users of protected servers and forward user data to a local Routing FONet, and *(2)* Routing FONet nodes which receive data from its local Access FONet nodes or neighboring Routing FONet nodes and forward the data towards the destination.

P2Cast which is used to communicate between FONet nodes uses an address range (232.254/16) that is different from the unicast address range. This makes it extremely easy to distinguish unicast traffic from p2cast traffic. FONet nodes receive data from outside their domains only from their FONet

neighbors via p2cast. Unwanted unicast traffic from outside the domain towards these nodes can be filtered out completely. This effectively prevents an attacker from launching remote flooding attacks towards FONet nodes. This highlights a unique feature of the architecture in that it uses DoS protected overlay nodes which communicate with each other using DoS resistant tunnels. Thus, unlike previous overlay-based approaches, FONet aims to individually protect overlay nodes from DoS attacks. The advantages of this approach are that resource redundancy can be reduced compared to previous methods; expensive overlay routing becomes unnecessary; and the architecture by itself becomes much more secure.

FONet provides DoS resistance through receiver-control to protected servers. It achieves this by restricting access through the overlay to the server only to users with the proper credentials. The credentials used depends on the type of service (see section 2.2) being provided by FONet. It can be either a simple proof that the user is human (using Graphical Turing Tests (GTT)); or an authentication cookie installed in an authorized user's machine by the server; or a high-level username/passsword combination. These authentication mechanisms are stateless at the FONet nodes making the architecture easily scalable to a large number of users. Finally, the use of overlays makes the changes required in the network for its deployment minimal.

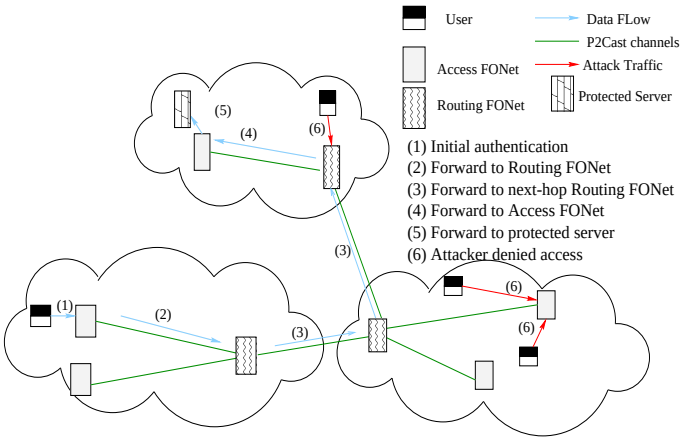

Figure 1. High level architecture and operation.

2.1 P2Cast

In FONet, overlay nodes securely communicate with each other using a novel network layer primitive, p2cast [5] which has been designed to provide a receiver-controlled communication model. Architecturally, p2cast is similar to PIM-SSM [3] in that *(i)* the receiver uses an (S,G) channel to communicate with the sender, *(ii)* channel creation is initiated by the receiver and the sender cannot send until the channel has been created, and *(iii)* reverse path forwarding (RPF) is used at routers in the channel path to prevent sender spoofing. Unlike PIM-SSM, we have designed p2cast to provide a one-to-one communication model. The receiver joins the (S,G) channel in a manner similar to PIM-SSM but modifications are made in the join procedure to ensure that there is only one receiver per (S,G) channel. MPLS or IPSec based VPNs' could possibly be used as an alternative to p2cast; but their cost and implications to the FONet architecture require further investigation.

2.2 Secure Services Provided

Recognizing that different customers have different security needs, we aim to utilize the FONet architecture to provide a number of different protection services to suit customer requirements: *(i)* **Strict Protection Service (SPS)** offers the highest level of protection in which protected servers communicate with users via FONet only. Undesired unicast traffic can be completely dropped without affecting legitimate users' traffic. *(ii)* **Partial Protection Service (PPS)** is aimed towards more open servers (like E*TRADE.com) and allows for communication via FONet for privileged (paying) users and unicast for normal (non-paying) users. Privileged users receive secure, DoS resistant, higher bandwidth communication via FONet. Normal users have to use unicast (which is rate-limited in preference to FONet traffic) and hence receive none of the FONet benefits. *(iii)* **Basic Protection Service (BPS)** is aimed towards open-access servers (like Google.com) and employs GTTs' at the Access FONet nodes. Attackers (zombies) will be filtered out at the access points while legitimate (human) users can continue to receive service.

3 Final Discussion

We have discussed an architecture that can provide proactive, cost-effective protection services to interested sites against DoS attacks. Currently, we are working towards evaluating FONet to further understand its performance and characteristics. For further information, please visit [1].

References

[1] Fonet homepage. http://www.utdallas.edu/∼jinuk/fonet.htm.

[2] D. Andersen. Mayday: Distributed Filtering for Internet Services. In *4th Usenix Symposium on Internet Technologies and Systems*, Seattle, WA, USA, Mar 2003.

[3] S. Deering, D. Estrin, D. Farinacci, V. Jacobson, G. Liu, and L. Wei. PIM architecture for wide-area multicast routing. *IEEE/ACM Transactions on Networking*, Apr 1996.

[4] A. Keromytis, V. Misra, and D. Rubenstein. SOS: Secure Overlay Services. In *Proceedings of ACM SIGCOMM '02*, Pittsburgh, PA, USA, Aug 2002.

[5] K. Sarac. SSM-based receiver-controlled communication in the Internet. In *South Central Information Security Symposium*, Denton, TX, USA, Apr 2003.

[6] A. Yaar, A. Perrig, and D. Song. SIFF: A Stateless Internet Flow Filter to Mitigate DDoS Attacks. In *IEEE Symposium on Security and Privacy*, Oakland, CA, USA, May 2004.

[7] X. Yang, D. Wetherall, and T. Anderson. A DoS-limiting Network Architecture. In *Proceedings of ACM SIGCOMM '05*, Philadelphia, PA, USA, Sept 2005.