

FONet : A Federated Overlay Network for DoS Defense in the Internet (A Position Paper)

Jinu Kurian
Dept. of Computer Science
University of Texas at Dallas
Richardson, Texas 75083
Email: jinuk@student.utdallas.edu

Kamil Sarac
Dept. of Computer Science
University of Texas at Dallas
Richardson, Texas 75083
Email: ksarac@utdallas.edu

Abstract—We propose a novel service architecture to provide DoS resistant communication services in the Internet. The architecture consists of a large scale federated overlay network with DoS protected tunnels established between overlay nodes. Individual overlay nodes are deployed and maintained by the domains hosting them. The overlay network as a whole is shared by all participating domains. This architecture is designed to be secure against DoS attacks and can provide different levels of DoS protection as value-added communication services on a large scale.

I. 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. Recently there has been a marked growth in the frequency and the ferocity of DoS attacks. Attacks have also become increasingly professional in nature, with attackers using botnets consisting of tens of thousands of compromised machines (bots) to launch DoS attacks as a method of extortion [1] or in some cases to cripple business competitors [2].

Much of the current growth and success of the Internet can be attributed to unicast's inherently open model. Unicast provides a *sender-initiated* communication model where a sender can send messages without obtaining prior permission from a receiver. The underlying network infrastructure is designed to forward these messages to the receiver, which has to reactively decide how to handle them. Until the packet is received, the receiver has no way of determining if the packet is legitimate or malicious. This model thus makes the receiver inherently vulnerable to flooding based DoS attacks.

Over the years, several DoS defense approaches [3] have been proposed by the research community. Earlier approaches fall into three categories: 1) detection, 2) traceback, and 3) pushback. Detection aims to identify an ongoing attack through anomaly, pattern, or statistical analysis [4] [5] [6]. With attackers increasingly modeling attack traffic to appear as legitimate, detection methods are becoming largely ineffective and suffer from lack of longevity of the identifying mechanism. Traceback [7] [8] aims to locate the origin of attack packets in a DoS attack. Traceback requires modifications in the routers which makes it expensive to deploy. In detection and traceback, once an attack or attacker has been identified, a response mechanism can be initiated to alleviate the effect

of the ongoing attack. Pushback [9] [10] has been proposed as a possible response mechanism. In pushback, once the attack flow has been identified, a local rate limit or a filter can be setup to attenuate the attack. When used with traceback, these filters can be deployed close to the attacker to 'push back' the attack traffic away from the victim. Pushback may suffer from false positives in the presence of spoofing and reflector based attacks [11].

More recently, several researchers have proposed alternative approaches with a goal of eliminating the vulnerabilities that make DoS attacks possible in the Internet. Some of these methods redesign the Internet to remove the vulnerabilities entirely [12] [13]. Others mask the vulnerabilities and thereby provide DoS resistance to compliant end systems [14] [15] [16]. Although disparate in their methods, both approaches are similar in that they seek to eliminate the vulnerabilities by providing *receiver-controlled* communication support to end systems. In this approach, end systems are provided with the ability to 1) explicitly choose who to communicate with by pre-authorizing remote senders and 2) to thwart an attack from an authorized sender with minimal damage.

Existing receiver-controlled communication approaches can be classified into three categories:

- 1) **Overlay based approach:** These combine overlay routing with lightweight filtering to provide DoS resistance for pre-approved clients of a single (or a small number of) protected server(s). Examples of this approach include SOS [14], Mayday [15], and the i^3 -based approach [16]. These approaches demonstrate the feasibility of receiver-control for DoS defense. An advantage of the overlay based approach is that the use of overlays makes changes required in the network for their deployment minimal.
- 2) **Capabilities based approach:** These introduce a new network layer service to enforce receiver-control. Examples of this approach include TVA [13] and SIFF [17]. This approach is a long term solution that incorporates security into the design of the underlying network for the next generation Internet.
- 3) **Connection based approach:** These are "off by default" approaches in which a receiver can choose to establish a network connection with a sender on demand. P2Cast [18] and Off by Default [19] are examples of

this approach.

The approaches above consider several types of DoS resistant communication services. These are:

- 1) **Strict Protection Service (SPS):** In SPS, the protected node is open only to traffic from well known or previously authorized users. SOS [14] and Mayday [15] are example methods which provide SPS. In these methods, users are authenticated at an access point and their traffic is routed through a set of overlay nodes to the target. Unauthenticated traffic is dropped at the access point or at filters deployed around the protected node.
- 2) **Partial Protection Service (PPS):** In PPS, the protected node is open to both authorized and unauthorized users. Authenticated user traffic is given preferential treatment in the network to the unauthenticated traffic. TVA [13] and SIFF [17] are example method which provide PPS. In these methods, authorized traffic carries a token to distinguish it from unauthorized traffic in the network.
- 3) **Basic Protection Service (BPS):** In BPS, the protected server has no prior knowledge of the users. The objective here is to distinguish between human users and automated tools (bots) which may be used to launch DoS attacks. WebSOS [20] and the i^3 -based approach [16] are example methods which provide BPS. In these methods, lightweight authenticators like GTTs [21] or client puzzles [22] are administered to ensure that only human users are allowed access to the protected server.

In this paper, we present a new service architecture that can provide DoS resistance in the Internet via receiver-control. Our architecture consists of a large scale *federated* and *secure* overlay network (FONet) that provides DoS resistant communication as a *value-added* service in the Internet. We create this architecture by building on the advantages of the previously proposed overlay based approaches and extending them to provide different levels of DoS resistant services in a large scale. 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 [23]) have not motivated ISPs to deploy them in their networks. The absence of suitable financial incentives has often been cited [15] as a primary reason for this lack of motivation on the part of the ISPs. To this end, the design of our architecture is leveraged towards providing DoS resistance as a *value-added* service in the Internet. Additionally, we design our architecture to be *cost effective* in that it imposes minimal changes in the network; supports partial deployment; and can provide different levels of DoS protection services to suit the security requirements of a wide spectrum of customers in a large scale. These properties provide a natural incentive for ISPs to deploy our solution in their domains.

The rest of the paper is organized as follows. Section II introduces the approach and the value-added services it can provide. Section III gives a high level architectural overview of the system. Section IV describes its operation. Section V presents possible deployment models. Section VI concludes the paper.

II. FONET OVERVIEW

The FONet architecture has three unique characteristics:

- **Federated:** Individual overlay nodes are deployed in their networks by ISPs. Management and operation of these overlay nodes are leased out to higher layer service providers (FONet Service Providers or FSPs). By interconnecting the FONet nodes in different ISP domains, FSPs create large scale overlay service networks on top of the existing ISP networks. This helps FSPs solve the coordination problem among ISPs in providing end-to-end value-added services.
- **Secure:** The overlay nodes are made resistant to DoS attacks by closing them to unwanted traffic from outside their domains. The only traffic allowed to a FONet node from outside its domain is from neighboring FONet nodes that it has established overlay level connections with. Communication between overlay nodes is through DoS resistant tunnels established between them. Thus, the overlay nodes by themselves and the overlay traffic are both completely protected against DoS attacks. The advantage of this approach is that resource redundancy can be minimized, expensive circuitous overlay routing becomes unnecessary, and the overlay architecture by itself becomes more scalable and secure.
- **Value-added services provided:** The value-added services approach fits into today's Internet realities in that not every Internet user is security aware at the same level, although, there is a growing population of users who are sensitive to cyber-attacks and are willing to pay for better protection. These features provide a natural financial incentive for ISPs to deploy these services in their domains. In our solution, we combine several DoS resistant services into a single architecture and provide them as value-added services on a larger scale. This reduces resource redundancy, makes the architecture cost effective and provides users with a richer spectrum of choices compared to previous approaches.

Before continuing with our discussion, the terminology we will use throughout the rest of the paper needs to be fixed. There are four entities to be considered in this context : 1) *ISPs*, 2) *FSPs*, 3) *customers*, and 4) *users*. *ISPs* deploy FONet nodes to participate in the overlay. *FSPs* are the entities that provide FONet services to the *customers*. *Customers* are Internet sites that provide online services to their remote *users*. Depending on their business model, *customers* may charge their remote *users* for the use of the FONet-based access service. Alternatively, *customers* may provide FONet-based access service free of charge to their privileged remote *users*.

Recognizing that not everyone in the Internet has the same security needs, we aim to utilize the FONet architecture to provide support to different types of DoS protection services including SPS, PPS, and BPS. SPS is ideal for *customer* sites with a known set of *users*. SPS can also be used to create a highly secure network environment for the protection of distributed components of critical national infrastructures. An SPS *customer* site is made open to remote access via FONet only. Undesired unicast traffic to a SPS *customer* site is

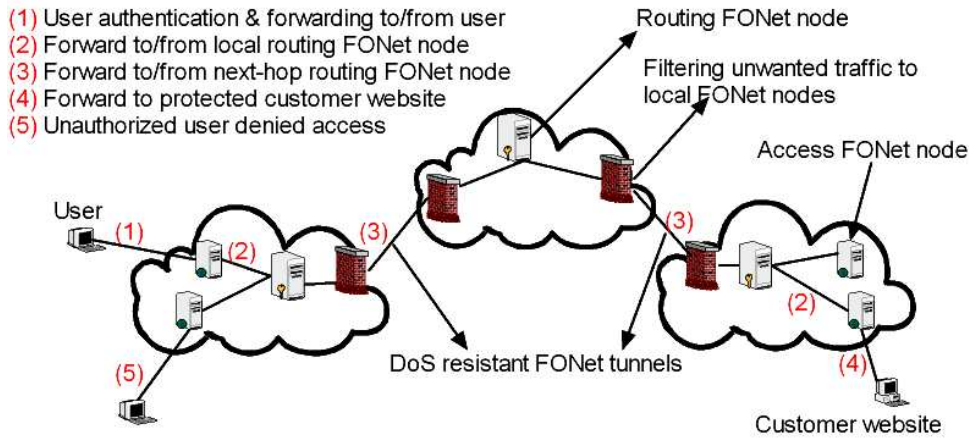


Fig. 1. High-level FONet architecture and operation

completely dropped without affecting legitimate remote *users* traffic via FONet. PPS is geared towards more open *customer* sites, like E*TRADE.com. The authorized (privileged) *users* are provided DoS resistant access via the FONet overlay. Unauthorized (non-privileged) users can access the PPS *customer* site via rate-limited unicast connection only. BPS is geared towards open access *customer* sites, like Google.com. The edge nodes of the overlay (or Access FONet nodes as we will describe in the next section) can employ traffic modeling mechanisms [5] to enforce expected legitimate *user* behavior on incoming traffic into the overlay. Additionally, these nodes can be configured to administer lightweight authentication mechanisms, like GTTs [24] or client puzzles [22], to weed out attackers (bots) while allowing legitimate (human) *users* access. The *customer* site under attack can block or severely rate limit unicast traffic destined to it without affecting legitimate *user* traffic via FONet.

III. ARCHITECTURAL OVERVIEW

The FONet architecture consists of three main components:

FONet nodes: FONet nodes are high end server machines capable of handling large amounts of Internet traffic. Two types of FONet nodes may be present in a domain (1) Access FONet nodes and (2) Routing FONet nodes. Access FONet nodes serve as the entry/exit point to the FONet overlay. They are made accessible only to local *users* and the local Routing FONet nodes. Routing FONet nodes forward *user* data over the overlay network towards the *customer*. They are made accessible only to the local Access FONet nodes in their domain and Routing FONet nodes in the neighboring domains.

DoS resistant tunnels: These are used to connect neighboring FONet nodes to each other. There are many possible tunneling candidates including MPLS based tunnels, P2Cast [18] tunnels, IPSec tunnels, leased lines, etc. Each of these techniques has its own unique advantages and disadvantages which need to be considered by ISPs before deployment.

Filtering support: Filtering support is provided at the domain boundaries to prevent unwanted unicast traffic from outside the domains to the FONet nodes. In the case of a Routing FONet, Access Control Lists (ACLs) [25] deployed at the edges of

the domain can be used. These ACLs can be configured to drop all unicast traffic to the Routing FONet node while allowing traffic from neighboring Routing FONets to enter. For example, in the case of MPLS based tunnels, traffic can be filtered out based on the appropriate MPLS label (which cannot be spoofed [26]) of the neighboring Routing FONet nodes.

For Access FONet nodes, all traffic from outside the domain can be dropped at the edges. Null routing can be used to statically route all unicast traffic destined to the Access FONet nodes to a null interface. Alternatively, Access FONet nodes can be made unicast unreachable outside their domains by using domain-local private addresses or by not advertising them via BGP, thereby stopping attacks at their origin.

So far we have discussed how overlay nodes are protected from DoS attacks from outside their domains. Attacks from within the domain also need to be considered. Routing FONet nodes are closed to all unicast traffic from within the domain. This makes unicast based attacks against the routing infrastructure impossible. Access FONet nodes are closed to unicast from outside the domain but are open to traffic from within the domain, and can hence be attacked. However, such local attacks will have limited effect and can be easily detected and provisioned against. Even if an attack is successful, it will have only a local effect and will not affect the rest of the overlay infrastructure. An attacker might try to disable inter-FONet traffic by attacking a router in the overlay path. However, bandwidth provisioning mechanisms [27] [28], depending on the technology used, can be employed to ensure that such an attack would not affect inter-FONet traffic. Even if the attack is successful, routing at the FONet level could choose an alternate overlay path to continue packet forwarding.

IV. HIGH-LEVEL FUNCTIONAL OVERVIEW

Figure 1 shows the overlay architecture and its functional overview. From a *user* perspective, if she is accessing a PPS or SPS protected *customer*, she is required to pre-register with the *customer* to obtain a set of authentication cookies. Once registered, the *user* can then contact the nearest Access FONet in her domain where she is authenticated before being allowed

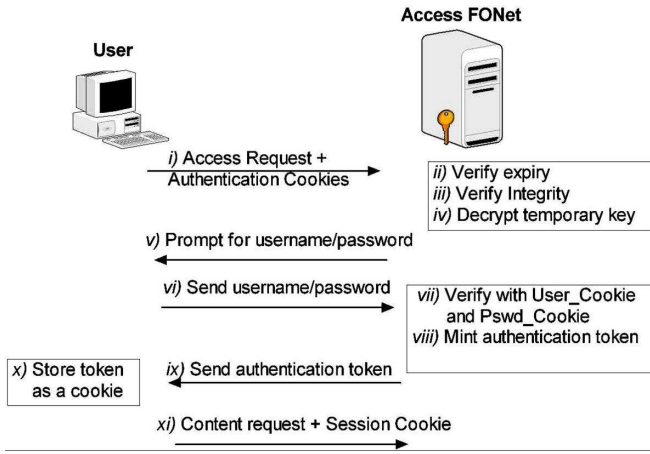


Fig. 2. Initial user authentication.

access into the FONet overlay. Once authenticated, *user* traffic is routed through the overlay network hop-by-hop until it reaches the remote *customer's* Access FONet and finally to the *customer* site.

A. User Authentication

The objective of *user* authentication is to securely authenticate remote and distributed *users* of a *customer* while maintaining minimal *user* information at the Access FONet nodes and without requiring that all authentication requests be forwarded to the *customer* site for approval. *User* authentication has three steps 1) pre-registration (for SPS and PPS) to obtain a set of authentication cookies, 2) initial authentication at the *user-site* Access FONet node with the authentication cookies, and 3) session maintenance after the initial authentication. **Pre-registration:** Pre-registration is required for SPS and PPS *users* to obtain a set of authentication cookies for a remote *customer* site. To pre-register, a *user* contacts the *user-site* Access FONet which mediates the *user's* registration request with the remote customer. This pre-registration process allows unauthenticated and possibly malicious *users* access to the overlay and the *customer* site. So, possible vulnerabilities need to be carefully examined. An attacker can spoof its IP address or use a large number of bots and flood the *user-site* Access FONet with registration requests. The attacker can thus force the *user-site* Access FONet to generate multiple bogus pre-registration requests to overload the remote *customer* and its Access FONet. The *user-site* Access FONet uses SYN cookies [29] before establishing connections with the *user* to prevent spoofing. Once the connection is established, it checks if the *user* is known to be malicious by verifying with a locally maintained *blacklist*. This *blacklist* is maintained as a Bloom filter [24] [30] and is used to temporarily deny access to *users* who have been confirmed as malicious during the authentication process, or during a session with the *customer*. Next it performs a lightweight authentication (using GTTs [24], client puzzles [22], traffic shaping, etc.) of the *user* to weed out bots. If the *user* fails the lightweight authentication, it is temporarily added (for e.g. *user's* IP address, requested

customer site) to the local *blacklist*. These measures ensure that the *users* allowed access for pre-registration are not discernibly malicious.

Once the *user* has completed the lightweight authentication, the *user-site* Access FONet forwards the request via the FONet overlay to the *customer* server along with the *user-site* Access FONet's public key. The *user's* credentials are verified by the *customer* server and allowed to register. Once the *user* has successfully registered, the *customer* server generates a set of authentication cookies encrypted with the *user-site* Access FONet's public key. The generation of cookies (the most expensive step in the pre-registration process since it uses public keys) is performed only after the *user* has been confirmed as legitimate. Thus, computational overload attacks on this step become implausible. The cookies are returned to the *user* and installed in her machine to complete pre-registration.

Initial user authentication: When a pre-registered *user* wants to access a *customer* server, it contacts the *user-site* Access FONet for authentication. In the case SPS and PPS, authentication is performed with the authentication cookies installed in the *user's* machine during pre-registration (see Figure 2). The cookie set used in FONet is similar to the cookie set suggested by Park et al [31]. Additionally, we include a FONet Cookie which contains the AS number of the domain of the Access FONet that installed them (see Figure 3). In the BPS case, lightweight authentication mechanisms like GTTs [24], client puzzles [22], or source end traffic validation mechanisms [5] can be used. Once initial authentication is completed, the Access FONet mints a unique authenticator and installs it as a temporary session cookie in the *user's* machine.

Session Maintenance: Session maintenance allows the Access FONet to statelessly verify the validity of the *user's* requests. The session cookies installed during initial *user* authentication are retrieved and verified for this purpose. The cookies used are standard session cookies [32] with an explicit expiration time and a unique data value. This data value is made specific to the Access FONet which installed them, thereby restricting the cookie's validity to a single Access FONet. During a session, if the authorized *user* turns malicious, the *customer* server can propagate a *deny-access* request through its local Access FONet to the *user-site* Access FONet. The *user* is *blacklisted* at both ends to prevent further abuse of the *user's* access privileges.

B. Protecting FONet Traffic

In Section III, we described how the FONet infrastructure by itself is protected against attacks by closing FONet nodes to unwanted traffic. To protect the traffic between FONet nodes from attack, DoS resistant tunnels are established between them. Several candidate mechanisms for creating these tunnels were mentioned in Section III. We consider some of these methods below.

MPLS VPNs are standardized in RFC2457bis [33] and allow service providers to deploy scalable VPNs in their domains. It uses a 16-bit Route Distinguisher (RD) which extends IPv4 addresses into VPN-IPv4 addresses. When used

Domain Flag Path Cookie Name Value Secure Date

Domain Name	Flag	Path	Cookie Name	Value	Secure	Date
Domain Name	True	/	Pswd_Cookie	Hashed Password	False	31-12-2005
Domain Name	True	/	User_Cookie	Hashed Username	False	31-12-2005
Domain Name	True	/	Key_Cookie	Encrypted Key	False	31-12-2005
Domain Name	True	/	FONet_Cookie	FONet number	False	31-12-2005
Domain Name	True	/	Seal_Cookie	Keyed Hash of Cookies	False	31-12-2005

Fig. 3. Authentication cookies.

in FONet, this allows a FONet node to be a part of several distinct address and traffic separated VPNs with its neighboring FONet nodes. MPLS VPNs are DoS resistant when deployed correctly and are resistant to address spoofing [26]. Furthermore, MPLS VPNs can choose to provide integrated Class of Service (CoS) support [27]. This allows ISPs to provide bandwidth guarantees for FONet tunnels to protect FONet traffic from attacks.

In addition to MPLS VPNs, several other technologies can be used to implement FONet tunnels. IPsec or SSL based VPNs can be used as alternatives with different security and overhead levels that we plan to investigate further. In addition, existing IP multicast service can potentially be used to establish spoofing and DoS resistant tunnels between neighboring FONet nodes [18]. Finally, private leased communication lines can be used as another alternative with much better security and bandwidth guarantees. The tunneling mechanism to be used for interconnecting FONet nodes depends on availability of the appropriate technologies, preferences of local ISPs, and the inter-ISP policies and SLAs. We believe that the above mentioned tunneling mechanisms provide a rich spectrum of choices to use for the implementation of these tunnels.

C. Routing

Unlike many of the previous overlay systems that use path probing [34], Chord based routing [14] etc, FONet performs overlay routing with the use of a simple path-vector protocol to disseminate reachability information. Path probing techniques are not very scalable [35] and can create conflicts with underlying network layer routing process [36]. Chord routing can improve overlay security via indirection and obscurity, but it introduces resource redundancy and end-to-end latency. The objective of routing in FONet is not to optimize routing performance but to minimize routing cost and overhead. This makes the routing scalable to a large number of nodes and allows *ISPs* to easily deploy and maintain FONet nodes.

In FONet, we use a constrained version of BGP to disseminate routing information between FONet nodes. BGP UPDATE messages are created by FONet nodes and exchanged through BGP sessions established between FONet peers. Currently, BGP is used for inter-domain routing in the Internet where it is associated with policy based decisions and input and output filtering of routes advertised between peers to conform to individual ISP policies. As we mentioned previously, a FONet overlay is operated by a single *FSP* and therefore forms a single intra-domain routing environment. Therefore, contrary to the use of BGP in inter-domain routing

in the Internet, FONet avoids such filtering and policy based routing decisions and utilizes the BGP UPDATE messages to choose optimal overlay paths based on a uniform decision criterion. One possible criterion for the optimal overlay path could be the shortest AS_PATH between overlay nodes. Once the optimal route is decided, routing and forwarding tables are created and maintained at the FONet nodes.

As FONet deployment increases, the topology will eventually evolve to match the underlying AS level network topology. In this full deployment scenario, Routing FONet nodes can use local network layer routing information to build and maintain their forwarding tables. The best AS level paths chosen by the local routing process can be used for data forwarding in FONet. Therefore, in a full deployment scenario, FONet does not need to run a routing protocol at the FONet level and instead uses the existing unicast routing information for traffic forwarding. This further reduces the overhead introduced by a FONet overlay on the underlying networks.

V. ECONOMIC AND DEPLOYMENT MODEL

As with any novel technology proposal, a solid economic and deployment model is required to make FONet a feasible approach for DoS defense in the Internet. In the design of FONet we have ensured that the architecture design is *ISP friendly* in many aspects. These include the value-added services approach which provides financial incentives for deployment; and the federated overlay approach which simplifies deployment and management of the overlay network. Additionally, the use of overlays requires minimal changes in the network for deployment. In the rest of the section, we will consider two possible deployment models for FONet in the Internet. Both models have many similarities, so the actual deployment model adopted could also be a combination of both in the global scope.

The first deployment model is similar to the Akamai [37] service model. One or more *FSPs* deploy FONet nodes in various *ISP* networks. *ISPs* provide tunneling, routing and filtering support within their domains, and charges the *FSP* for this support. *FSPs* in turn charge their *customers* who purchase value-added services to better serve their *users*. Depending on *customer* policy and type of service provided, *users* may in turn be charged for these value-added services.

In the second deployment model, *ISPs* can choose to deploy FONet nodes within their domains, establish tunnels with their neighboring *ISPs*, and provide routing and filtering support. *FSPs* lease the use of FONet nodes from various *ISPs* and manage and coordinate their usage at the overlay level. *Customers* purchase value-added services as before from the *FSP* and provide them to their *users*.

During the initial deployment phase, not all ISPs will participate in FONet services. In this environment, non-participating ISPs will provide tunneling services to connect FONet nodes in remote domains and will charge for this service. The method of the tunneling used depends on the locally available technologies and in the worst case can be leased lines. As more and more FONet tunnels are created over a non-participating ISP network, the ISP will potentially realize the lucrative

business opportunity and will likely join into the FONet overlay.

Note that our architecture design allows an ISP to participate in a FONet overlay by running a Routing FONet node only. Therefore, transit ISPs having no stub networks as their customers can participate in a FONet overlay by running Routing FONet nodes only. However, ISPs providing connectivity service to stub networks would likely support both Access and Routing FONet nodes.

VI. CONCLUSION

In this paper we have presented an architecture that provides a cost effective solution to interested *customer* sites against DoS attacks. The unique contributions of our work include:

- 1) the federated overlay approach which simplifies management, reduces resource redundancy and cost of deployment,
- 2) the ability to provide different types of services to suit the needs of different *customers*,
- 3) the value-added nature of our proposal that provides financial incentive to *ISPs* to deploy it, and
- 4) the use of DoS resistant tunnels between DoS protected nodes which makes the architecture by itself secure.

Currently we are working on a prototype implementation of our FONet overlay system to deploy it on a test network environment. After the initial testing, we plan to move it into PlanetLab environment to perform more practical real life experiments.

REFERENCES

- [1] D. Pappalardo and E. Messmer, "Extortion via DDoS on the rise," May 2005. <http://www.networkworld.com/news/2005/051605ddos-extortion.html>.
- [2] "PCWorld : Web of Crime Series," August 2005. <http://www.pcworld.com/news/article/0,aid,122258,00.asp>.
- [3] J. Mirkovic and P. Reiher, "A Taxonomy of DDoS Attacks and DDoS Defense Mechanisms," *ACM SIGCOMM Computer Communications Review*, vol. 34, pp. 39–54, April 2004.
- [4] T. Peng, C. Leckie, and R. Kotagiri, "Protection from Distributed Denial of Service Attacks Using History-based IP Filtering," in *IEEE International Conference on Communications (ICC)*, (Anchorage, AK, USA), May 2003.
- [5] J. Mirkovic, G. Prier, and P. Reiher, "Attacking DDoS at the source," in *Proceedings of the IEEE International Conference on Network Protocols*, (Paris, FRANCE), November 2002.
- [6] P. Barford, J. Kline, D. Plonka, and A. Ron, "A Signal Analysis of Network Traffic Anomalies," in *Proc. ACM SIGCOMM Internet Measurement Workshop*, (Marseille, FRANCE), pp. 71–82, November 2002.
- [7] A. Snoeren, C. Partridge, L. Sanchez, C. Jones, F. Tchakountio, B. Schwartz, S. Kent, and W. Strayer, "Single-packet IP traceback," *IEEE/ACM Transactions on Networking*, vol. 10, December 2002.
- [8] A. Yaar, A. Perrig, and D. Song, "FIT: Fast Internet Traceback," in *Proceedings of IEEE INFOCOM*, (Miami, FL, USA), March 2005.
- [9] K. Argyraki and D. R. Cheriton, "Active Internet Traffic Filtering: Real-time Response to Denial of Service Attacks," in *USENIX Annual Technical Conference*, (Anaheim, CA, USA), April 2005.
- [10] R. Manajan, S. Bellovin, S. Floyd, J. Ioannidis, V. Paxson, and S. Shenker, "Controlling High Bandwidth Aggregates in the Network," *Computer Communications Review*, vol. 32, July 2002.
- [11] M. Handley and A. Greenhalgh, "Steps towards a DoS-resistant Internet Architecture," in *Proc. ACM SIGCOMM workshop on Future directions in Network Architecture*, (Portland, OR, USA), August 2004.
- [12] T. Anderson, T. Roscoe, and D. Wetherall, "Preventing Internet Denial-of-Service with Capabilities," in *Proceedings of Second ACM SIGCOMM HotNets Workshop*, (Cambridge, MA, USA), November 2003.
- [13] X. Yang, D. Wetherall, and T. Anderson, "A DoS-limiting Network Architecture," in *Proc. ACM SIGCOMM*, (Philadelphia, PA, USA), August 2005.
- [14] A. Keromytis, V. Misra, and D. Rubenstein, "SOS: An Architecture for Mitigating DDoS Attacks," *IEEE Journal on Selected Areas in Communications (JSAC), Special Issue on Service Overlay Networks*, vol. 22, January 2004.
- [15] D. Andersen, "Mayday: Distributed filtering for Internet services," in *4th Usenix Symposium on Internet Technologies and Systems*, (Seattle, WA, USA), March 2003.
- [16] D. Adkins, K. Lakshminarayanan, A. Perrig, and I. Stoica, "Taming IP Packet Flooding Attacks," in *Proceedings of Second ACM SIGCOMM HotNets Workshop*, (Cambridge, MA, USA), November 2003.
- [17] A. Yaar, A. Perrig, and D. Song, "SIFF: A Stateless Internet Flow Filter to Mitigate DDoS Flooding Attacks," in *IEEE Symposium on Security and Privacy*, (Oakland, CA, USA), May 2004.
- [18] K. Sarac, "SSM-Based Receiver-Controlled Communication in the Internet," in *Proceedings of South Central Information Security Symposium*, (Denton, TX, USA), April 2003.
- [19] H. Ballani, Y. Chawathe, S. Ratnasamy, T. Roscoe, and S. Shenker, "Off by default!!!," in *Proceedings of ACM HotNets 2005*, (College Park, MD, USA), November 2005.
- [20] W. Morein, A. Stavrou, D. Cook, A. Keromytis, V. Misra, and D. Rubenstein, "Using Graphic Turing Tests to Counter Automated DDoS Attacks Against Web Servers," in *Proceedings of the 10th ACM International Conference on Computer and Communications Security (CCS)*, (Washington, DC, USA), October 2003.
- [21] L. Von Ahn, M. Blum, and J. Langford, "Telling Humans and Computers Apart Automatically," *Communications of ACM*, vol. 47, pp. 56–60, February 2004.
- [22] T. Aura, P. Nikander, and J. Leiwo, "DoS-resistant Authentication with Client Puzzles," in *8th International Workshop on Security Protocols*, (Cambridge, UK), April 2000.
- [23] P. Ferguson and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks Which Employ IP Source Address Spoofing," Internet Engineering Task Force (IETF), RFC 2267, January 1998.
- [24] S. Kandula, D. Katabi, M. Jacob, and A. W. Berger, "Botz-4-Sale: Surviving Organized DDoS Attacks That Mimic Flash Crowds," in *2nd Symposium on Networked Systems Design and Implementation (NSDI)*, (Boston, MA, USA), May 2005.
- [25] Cisco systems, "Infrastructure Access Control Lists: Protecting your core." Available at <http://www.cisco.com>.
- [26] J. Pultz and N. Richard, "Analysis of MPLS-based VPN technology." Available at <http://www.cisco.com>.
- [27] Cisco Systems, "MPLS Bandwidth Protection White Paper." Available at <http://www.cisco.com>.
- [28] Cisco Systems, "Committed Access Rate." Available at [Cisco.com](http://www.cisco.com).
- [29] D. Bernstein, "SYN cookies." <http://cr.yt.to/syncookies.html>.
- [30] A. Broder and M. Mitzenmacher, "Network Applications of Bloom Filters: A Survey," in *Proceedings of 40th Annual Conference on Communication, Control, and Computing*, (Allerton, IL, USA), October 2002.
- [31] J. S. Park and R. S. Sandhu, "Secure Cookies on the Web," *IEEE Internet Computing*, vol. 4, pp. 36–44, July 2000.
- [32] K. Fu, E. Sit, K. Smith, and N. Feamster, "Do's and Don'ts of Client Authentication on the Web," in *Proceedings of the 10th USENIX Security Symposium*, (Washington, DC, USA), August 2001.
- [33] E. Rosen and Y. Rekhter, "BGP/MPLS VPNs." Internet Engineering Task Force (IETF), RFC 2547, March 1999.
- [34] D. Andersen, H. Balakrishnan, M. Kaashoek, and R. Morris, "Resilient Overlay Networks," in *Proceedings of 18th ACM SOSP*, (Banff, CANADA), October 2001.
- [35] S. Rewaskar and J. Kaur, "Testing the Scalability Limits of Overlay Routing Infrastructures," in *Proceedings of the Fifth Passive and Active Measurements Workshop (PAM'04)*, (Juan-les-Pins, FRANCE), April 2004.
- [36] R. Keralapura, N. Taft, C. Chuah, and G. Iannaconne, "Can ISP's Take the Heat from Overlay Networks?," in *ACM HotNets Workshop*, (San Diego, CA, USA), November 2004.
- [37] "Akamai Technologies, Inc." <http://www.akamai.com>.