Security Mechanisms for a Cooperative Firewall

Hammad Kabir, Raimo Kantola, Jesús Llorente Santos
Department of Communications and Networking
Aalto University
Helsinki, Finland
{hammad.kabir, raimo.kantola, jesus.llorente.santos}@aalto.fi

***Abstract*— Customer Edge Switching (CES) is a proposed replacement of Network Address Translators (NAT) that overcomes the drawbacks of traditional NAT traversal schemes. CES enabled networks assure policy based reachability of hosts in private realms, without requiring keep-alive signaling. CES aims at improving security in the Internet by balancing the interests of the receiver with the interests of the sender, unlike the traditional best effort Internet that solely attends to the interests of the sender. The architecture substantially helps with the scalability limitations of IPv4 due to the generalization of private addressing of the hosts. This paper relates to the specifics of security in Customer Edge Switches and presents security models that protect hosts in private realms against attacks. The presented work is a part of a larger project that addresses many issues of the current Internet and proposes the use of CES as collaborative firewalls to reduce volume of unwanted traffic and mitigate Denial of Service (DoS) attacks in the Internet.**

Keywords*— *Security; Trust; Policy; NAT traversal; DoS, DDoS; Cooperative firewall.

I. INTRODUCTION

Recent surveys by ITU-T reveal an impressive growth of mobile users and mobile broadband subscriptions, which are effectively replacing fixed phone and broadband subscriptions [1]. This growing number of mobile users calls for support of mobility within the Internet and demands more addresses from the already depleted IPv4 address space. The deployment of NAT at the network edges alleviated the address exhaustion issue at the cost of introducing the reachability problem. The reachability problem prevents a host in the public realm from unilaterally initiating a connection with a host in the private network. Over the years, various NAT traversal proposals have attempted to solve the reachability problem. In [2], we analyse these NAT traversal schemes and propose our own solution.

By default, the reachability problem of NAT contributes to the network security by dropping inbound packets from the Internet that do not belong to a known connection. From security perspective, it makes difficult for an attacker to intrude the host. However, irrespective of all network and host-based security methods, a host today receives an increasing number of flows that it classifies as unwanted or malicious. A firewall on the host to block unwanted traffic after it has reached the device itself is not an efficient solution, as it results in battery drain and clutters the air interface with unwanted traffic.

Malicious hosts are abusing the current best effort paradigm of the Internet communications to launch attacks on their victims. The marginal interest towards security in the Internet and absence of authentication mechanisms in the traditional TCP/IP stack has hurt the Internet in many ways, including long periods of dis-connectivity due to DoS attacks. Huge volume of spam, man-in-the-middle (MITM) attacks, Internet fraud and a wide range of malicious activities owe themselves to feeble security implementations in the Internet. The rising volume and increased sophistication of the latest attacks [4] demands better security methods and heuristics to protect the network connections that have become valuable due to convergence of data, networks and people [3].

We propose Customer Edge Switching (CES) [5] as an architecture that aims to improve security in the Internet, besides addressing the challenges like IPv4 address space exhaustion, reachability issue and scalability of the core routing system. CES aims at improving security in the Internet due to the cooperative nature of CES firewalls.

CES employs routing locators (RLOC) for edge-to-edge routing and host identifiers for flow admission to end-hosts. In opposition to the claims that such a split architecture would weaken the security by compromising the host identifiers, we present a model that leverages well-understood security mechanisms with a tolerable performance penalty. The Host Identity Protocol (HIP) [6] also provides a secure way of separating identities and addresses. Unlike HIP, which requires changes in many entities and is focused on hosts, CES limits all the changes solely to the edge nodes and therefore, it can be deployed one customer network at a time.

This paper will discuss the specifics of security related to Customer Edge Switching. The paper proposes a set of security models that aim to secure CES and the hosts served by CES against attacks that are inherent in the Internet, in particular source address spoofing, Distributed Denial of Service (DDoS) attacks and man-in-the-middle (MITM) attacks.

The rest of the paper is structured as follows. Section II briefly presents Customer Edge Switching. Section III takes a look at related work. Section IV presents the security vulnerabilities. Section V presents the security models employed to secure CES and its users against inherent Internet attacks. Section VI evaluates the performance of the security models and Section VII concludes the paper.

II. CUSTOMER EDGE SWITCHING

CES deployment at network edges splits the Internet into customer network (CN) and service provider network (SPN), as depicted in Figure 1. The split architecture enforces the separation of edge-to-edge routing locator and host identifiers.

Hosts are no longer identified by IP addresses, but with their Fully Qualified Domain Name (FQDN). As a consequence, communication is triggered by name resolution queries.

Fig. 1. CES Architecture

CES functionality at network edges can be considered as an extension of a stateful firewall into a cooperative firewall. Unlike traditional firewalls that either admit or drop a received packet, CES may issue additional queries prior to taking the final accept/drop decision. These queries are issued in accordance with the policy-based reachability defined for the hosts. A CES node acts as a connection broker for hosts located in its network, and it negotiates with remote edge node via Customer Edge Traversal Protocol (CETP) [8] to ensure that the interests of the receiver are met by only admitting the flows that fulfil the policy requirements of the destination. CES also support PRGW [2] for interaction with legacy hosts.

A CETP packet flow between two edge nodes is uniquely identified by Source Session Tag (SST) and Destination Session Tag (DST) carried in packets of the flow. The CETP header carries several control type-length-value (TLV) [8] elements that allow the edge nodes to make an informed decision on flow admission/rejection.

Figure 2 presents a CETP communication, where the sender reaches a host in a private realm by resolving the FQDN of the destination. Upon receiving the NAPTR record in the DNS response, the outbound CES (oCES) learns RLOCs of the CES node that hosts the destination. Next, the oCES node initiates policy negotiations by forwarding the sender's policy elements towards the inbound CES (iCES). The oCES reserves a state within its connection table to process the subsequent response packets from the iCES, and hence assumes stateful operation.

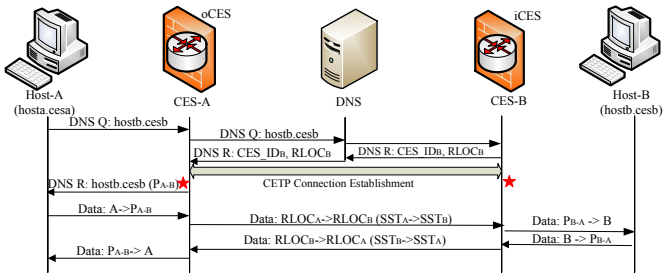

A: Private IP of Host-A	CES_IDA: CES Id of Host-A
B: Private IP of Host-B	CES_IDB: CES Id of Host-B
RLOCA: Routing Locator of CES-A	PA-B: Proxy-address representing Host-B to Host-A
RLOCB: Routing Locator of CES-B	PB-A: Proxy-address representing Host-A to Host-B
hosta.cesa: Domain Name of Host- A	SSTA: Session Tag for session initiated in CES-A
hostb.cesb: Domain Name of Host- B	SSTB: Session Tag for session initiated in CES-B

Fig. 2. CES-to-CES connection establishment

The remote CES processes the received packet as iCES, upon not having a prior connection state. The connection is successfully established when both the oCES and the iCES nodes successfully fulfil each other's policy requirements. After the connection is established, the remote host is locally represented with a proxy-address in the connection state. The DNS response bearing the proxy-address is returned to the sender host, and both hosts exchange the data packets via the states created within the respective CES nodes.

A CETP negotiation may complete in 1 Round Trip Time (RTT), if the inbound CETP packet successfully fulfil the policy requirements of the destination. However, if the first inbound packet does not meet the policy requirements of the destination, the iCES node informs the oCES about all the required policy elements that the next inbound packet from the sender should satisfy to establish the connection in 2RTT.

III. Related Work

CES proposes a future Internet architecture with explicit mechanisms for tackling source address spoofing and DDoS. Therefore, in this section, we limit the discussion of related work to other proposals and methods that are focused on overcoming spoofing and DDoS.

The classical solution for tackling source address spoofing is ingress filtering [9]. Unfortunately, it has not been adopted universally, possibly because it is the receiver or its ISP that suffer from the spoofing while other entities are supposed to invest in configuring and processing the ingress filtering. Therefore, benefits and investments are not well aligned.

SIFF [10] proposes a stateless solution to mitigate DoS attacks. It divides the network traffic into privileged and non-privileged flows. The privileged flows are tagged and the SIFF enabled routers prioritize them over unprivileged flows. PBS [12] provides a signalling architecture that requires the sender to acquire authorization from the receiver prior to sending a packet flow. Next, the permission states are installed within the PBS nodes (end-hosts and routers), for subsequent data flow. The proposal requires changes in both end hosts and network nodes, which we believe could be detrimental to its adoption.

StopIt [11] presents a DoS resistant system, which enables a receiver to upload filters for blocking the unwanted inbound traffic. The StopIt servers communicate with each other to block the reported traffic at the source, and punish the misbehaving hosts. However, the lack of an Internet wide trust reporting system results in difficult design choices.

In operational practise, BGP withdrawal updates with an agreed community attribute are used to sink-hole or drop DDoS packets closer to the target network. The weakness of this method is that the blocking applies to all protocols and all hosts, including legitimate sources [14].

IV. Network Vulnerabilities

Address spoofing and Denial of Service (DoS) attacks are inherent to the current Internet. There have been many attempts to tackle the issues. However, the problem still remains unsolved. The idea of Customer Edge Switching is to put the responsibility and the tools for solving these inherent problems at the hands of the receiver. Given an appropriate policy, a host served by an iCES node will never receive spoofed packets from outside the CN. A DoS attack from a host served by an oCES can be traced back to the customer network of the host.

Since the CES RLOCs are routable Internet addresses, this raises a security concern where a legacy element connected to the Internet can send forged CETP traffic towards the CES. Figure 3 presents the case where a legacy host with a CETP attack module impersonates CES-A, by sending forged packets towards CES-B. Upon receiving a CETP packet that fulfils the policy requirements of the destination, the iCES node establishes a connection with the sender.

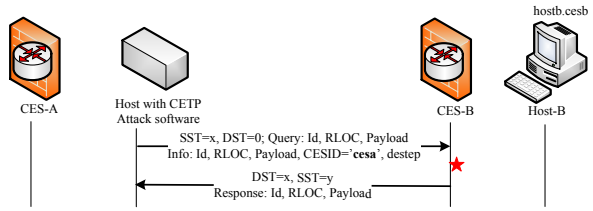

Fig. 3. Attack from legacy elements in the Internet

If the sender is a spoofing source, this leaves CES vulnerable to DoS attacks. For a non-spoofing source, presented in Figure 3, the attacker can successfully portray itself as a legitimate CES and can access the victim host behind the attacked CES.

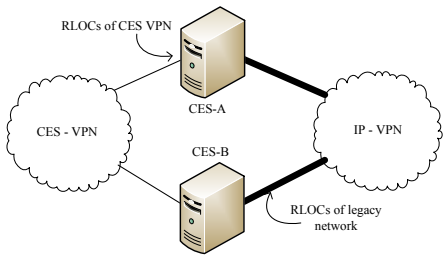

Fig. 4. CES deployment model at network edges

A CES node can step up the security by advertising several RLOCs associated with different technologies at different layers: L4 (SSL/TLS), L3 (VPN) etc. As a consequence, packet filtering becomes more effective at network edges thus minimizing the risk of accepting non-legitimate flows.

A CES node can benefit from the model presented in Figure 4, where CETP traffic from CES nodes and IP traffic from legacy hosts is expected over a separate set of interfaces. The IP traffic from legacy hosts is served by PRGW, whereas the CES processes the CETP traffic received from other CES networks. By ingress filtering the traffic received over the legacy interface, CES can drop the CETP attack traffic originated from a legacy host. However, when the CES node is not large or it doesn't have multiple interfaces e.g. in an ADSL modem, it is vulnerable to the attacks presented in Figure 3.

V. Security of Customer Edge Switches

The key to improving security in the Internet comes from authentication and non-repudiation of communication. A fine-grained packet admission policy can effectively be applied after the receiver has clearly determined the sender's identity. Therefore, CES deployment at network edges should enable the receiver to collect and attribute the evidence of good or bad behaviour against the sender's network. Hence, a malicious sender can be immediately traced back, when trust is violated.

Adhering to these requirements, we define a light-weight cookie mechanism. Whereas, CETP header signature and Home Subscriber Server (HSS) based verification are defined to ascertain the legitimacy of the CETP packet source.

A. CETP Cookie

We define CETP cookie as light-weight mechanism that prevents an inbound CES from opening a connection upon receiving a spoofed CETP packet. The mechanism draws its inspiration from the SYN cookie algorithm [14], developed to mitigate SYN flooding attacks. Hence, upon receiving a CETP packet, the iCES node responds with a cookie computed using

$$64\text{-bit DES (Last 4-bytes of hash (SST, DST, Host-ID, Destination-ID, CES-ID, SECRET)} + T_0) \quad (1)$$

The next packet from the sender must bear the same cookie, verified by equation (1), for the sender to be determined as a non-spoofing source. This provides iCES sufficient protection against DoS attacks that employ source address spoofing as their primary launch point.

The cookie mechanism bears three security layers against forgery attempts: 1) a local SECRET, 2) symmetric-key encryption and 3) timeout. Besides eliminating spoofing, the timeout value in the cookie mechanism prevents a replayed packet from establishing a connection with the inbound CES.

B. CES Registration/Verification

CES nodes deployed at network edges must be univocally identified in order to differentiate them from legacy elements present in the Internet infrastructure i.e. routers, NAT, servers etc. This prevents a legacy component from sending forged CETP flows towards a CES node. This paper presents two CES registration schemes that enable the receiver to determine the legitimacy of the CETP packet source. A verification failure in these mechanisms, after spoofing is eliminated, enables the CES to attribute the attack to the packet source.

1) Centralized Registration

The mechanism involves maintaining a list of all the CES nodes and their respective RLOCs in a centralized database e.g. HSS. After spoofing is eliminated, the CES node can verify the CETP flow against this database to determine if the sender is indeed a legitimate CES. A remote CES only need to be validated once and the subsequent CETP packets from the oCES are accepted without performing the CES verification.

2) Decentralized Registration

Because the centralized CES registration mechanism described above is considered an additional infrastructure, we propose a relatively decentralized CES registration method that utilizes the existing Internet infrastructure i.e. Certificate Authority, to determine the legitimacy of the packet source.

The current version of X.509 certificate defines various extension fields that provide additional information about the certificate and put constraints to the certificate usage. These extensions in X.509 certificate among others include: Basic

Constraints, Key Usage and Extended Key Usage fields. When the Basic Constraint identifies the certificate as an end-entity certificate e.g. client, server etc. the Extended Key Usage field according to [15] "*indicates one or more purposes for which the certified public key may be used, in addition to the basic purposes indicated in the Key Usage extension*".

We propose that the certificate issued to a CES node must carry "CES Verification" as object identifier in Extended Key Usage field in order to differentiate a CES certificate from rest of the certificates. This prevents a certificate bearing legacy host from imitating as a CES, by sending forged CETP packets.

The CES certificates are used in conjunction with the CETP header signature to ascertain the legitimacy of the packet source. Upon receiving the first CETP packet from the sender, the iCES node requests the CETP header signature and the public-key certificate from the sender. If the Extended Key Usage field of the received certificate bears "CES Verification" and the received CETP header signature can be verified with the sender's public-key, the sender is trusted as a legitimate CES. With spoofing eliminated, a failure in header signature verification attributes the attack evidence against the packet source. We will here omit further details, but in [6] we also tested the use of the CES certificate to thwart MITM attacks.

C. *CES Security model*

As mentioned before, because of deployment constraints it will not always be possible to process the CETP traffic and the legacy IP traffic on different interfaces. Rather, in some deployments they can share the same interface.

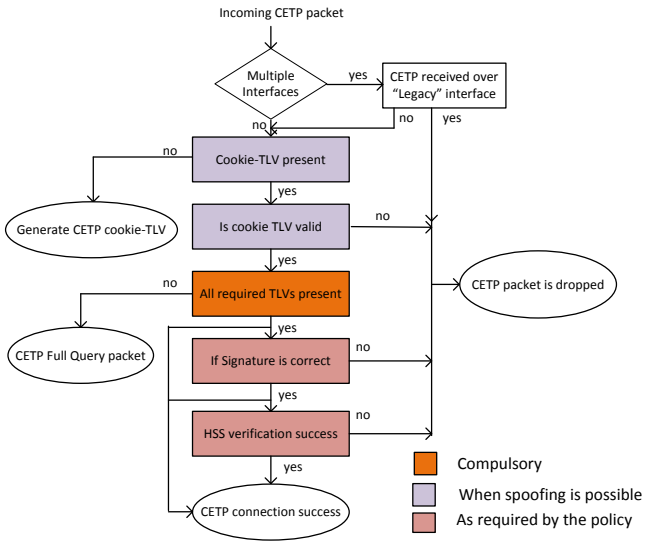

Fig. 5. Inbound CES security model

In case of a shared interface, the received CETP flow is processed as per the security model depicted in Figure 5. The received packet is first checked for presence of a CETP cookie. The received cookie is validated as per the CETP cookie mechanism. The successful cookie verification guarantees that the sender is a non-spoofing entity. Whereas, a failure in cookie verification leads to the CETP packet drop as the packet may have come from an attack source. For a packet that has not arrived with a cookie, the iCES node sends a CETP cookie encoded within a CETP packet towards the sender, along with the policy requirements of the destination host.

After spoofing is eliminated, the iCES node determines the reception of the required policy elements within the inbound packet. For a missing required policy element, the iCES encodes a CETP packet listing all the policy requirements of the destination and sends it towards the sender. However, if all the required policy elements have been received, the iCES node determines the legitimacy of the packet source via either HSS based verification or CETP header signature.

With spoofing eliminated, a failure in signature verification or HSS based verification identifies the source of the attack. Next, the CES attributes the attack evidence to the sender's identity and blacklists the sender for time 'T_I'.

The outbound CES also employs a similar but relatively simpler security model, due to its stateful nature. An oCES keeps record of each (SST, DST) pair that CES has sent towards a destination. A received packet is admitted in oCES only if the DST of the received packet matches with SST of the connection state (SST, DST=0). The absence of a query TLV projects the admitted packet as the last packet of the connection establishment. This would trigger the oCES to execute either of the verification mechanisms to ascertain authenticity of the remote CES node, similar to the iCES security model. A successful verification would lead to connection establishment with the remote CES.

VI. PERFORMANCE ANALYSIS

In this section, we evaluate the performance of our CES prototype after introducing the security models. The cost of the security is evaluated in terms of processing delay introduced by new processing modules in the connection establishment process. Figure 6 presents a comparison of delay for 80 CES-to-CES connection establishments before and after the security.

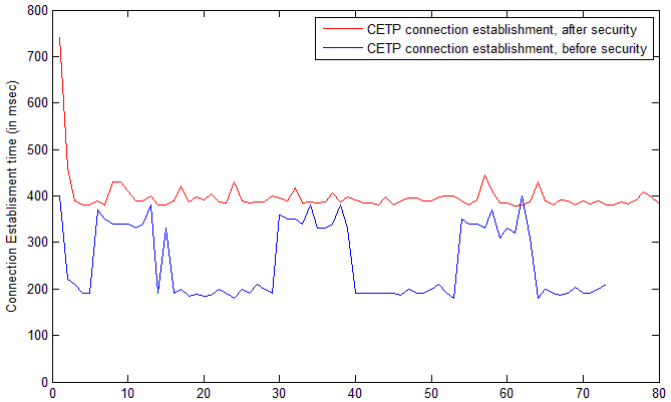

Fig. 6. Comparison of CETP connection establishment delay

The connection setup delay in the figure accounts for CETP part of the connection establishment, i.e. leaving out the DNS query and response duration. In a real network, edge-to-edge latency for 2 to 4 CETP messages would have to be added in 1RTT and 2RTT flows to make the total duration for connection setup. The connection setup delay after the security is slightly more than the setup time prior to the security. The

rise and fall of the connection setup delay before the security is reflective of a connection establishment in either 1RTT or 2RTT, respectively. However, after security, the delay remains almost constant as all the connections are established in 2RTT. This is because to its first CETP packet the oCES receives a cookie in response, which must be relayed back by the oCES in the next inbound packet along with the requested policy elements for the connection to establish.

The difference of nearly 35 milliseconds mainly comes from the HSS based ID-verification of the remote host. Whereas, the rest of security modules introduce less than 2 milliseconds of processing delay to the connection setup.

Table I presents the time duration that a received CETP packet is processed in the security model before a decision is reached. The CES reacts to a spoofed/forged packet with CETP cookie mechanism in less than 5 microseconds, which makes it an ideal mechanism to counter the spoofing attacks and DoS attempts. The first CETP packet from an oCES requires that the sender's claim of having a legitimate CES-ID is verified by either of CES registration mechanisms.

TABLE I. **TIME IN SECURITY MODEL FOR A RECEIVED PACKET**

	Response duration	***Outcome***
CETP packet with forged cookie	0.00433 msec	Packet drop
CETP packet without cookie	0.00373 msec	Respond with cookie
HSS-CES verification (1st packet)	16.00 msec	Accept/Deny
Signature- CES verification (1st packet)	~ 2 msec	Accept/Deny
CES-ID verification (subsequent packets)	< 0.001 msec	Accept/Deny

A centralized CES registration mechanism modelled via HSS takes nearly 16 milliseconds. However, the CES-ID verification reduces drastically to 2 milliseconds if we employ the certificate based CES verification. Once the legitimacy of the sender is ascertained, the subsequent packets from the sender are accepted/denied in a fraction of a millisecond. It is pertinent to mention that these delay values are computed within CES at algorithmic level, i.e. they do not account for the time spent in acquiring, packetizing and sending a packet.

We note that our prototype is implemented in Python. We use the Scapy [16] for access to the packets. Because the CES prototype employs control plane and data plane separation to handle the CETP packets, most of the processing time is spent in processing packets up and down in the stack. We expect that the processing time could be significantly reduced by for example a C-based implementation. We further note that the delay of about 400ms in the CETP session setup usually fits in with the DNS response waiting time at the host. Windows would normally make the first DNS re-attempt in 1s. If the host re-attempts the DNS query before receiving the response within the first timeout, the oCES state machine can smoothly absorb the query while the CETP process is converging.

VII. CONCLUSION

We have presented Customer Edge Switching as an architecture that aims at improving security in the Internet, by putting the receiver in charge of the communication. The architecture addresses the inherent vulnerabilities of the Internet, and it safeguards the network against attacks by blocking packets with spoofed source addresses and DoS from reaching the hosts in the private realm. The security models pave a way for a proactive approach, which enables a CES node to attribute the attack evidence against the packet source.

The verification presented here shows that we are able to secure the ID/address split architecture by reusing well-known security mechanism on the level of interaction between senders and receivers, and that the performance penalty is tolerable. The method of blocking DDoS using CETP is more fine-grained than BGP withdrawal updates: the blocking can be applied to the attack protocols only and to the oCES nodes that serve the attacking hosts. The network administrators on the sender side can further improve the accuracy of the blocking by applying the ingress filtering for their hosts. Moreover, we have secured CES against MITM attacks using the CES certificates.

REFERENCES

[1] ITU-T ICT STATISTICS. Free statistics. [Retrieved on Aug.2013] Available: http://www.itu.int/en/ITU-D/Statistics/Pages/stat/default.aspx

[2] J. Llorente, R. A. Kantola, N. Beijar, and P. Leppäaho, "Implementing NAT Traversal with Private Realm Gateway", Communications (ICC), 2013 IEEE International Conference, 2013, pp. 3581-3586.

[3] "2013 Cisco Annual Security Report," CISCO, 2013.

[4] "2014 Cisco Annual Security Report," CISCO, 2014.

[5] R. Kantola, "Implementing Trust-to-Trust with Customer Edge Switching,", AMCA in connection with AINA 2010, Perth, Australia, April 2010.

[6] R. Moskowitz and P. Nikander, "Host Identity Protocol (HIP) Architecture," IETF RFC 4423, May 2006.

[7] H. Kabir, "Security Mechanisms for a Cooperative Firewall," MSc. Thesis, Aalto University, Espoo, Finland, 2014.

[8] M. Pahlevan," Signaling and Policy Enforcement for Co-operative Firewalls," M.Sc. Thesis, Aalto University, Department of Communications and Networking, Apr. 2013.

[9] P. Ferguson and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing," RFC 2827, May 2000.

[10] A. Yaar, A. Perrig, and D. Song, "SIFF: a Stateless Internet Flow Filter to Mitigate DDoS Flooding Attacks," –in Proc. IEEE Symposium on Security and Privacy (SP 2004), pp. 130–143, May 2004.

[11] X. Liu, X. Yang, and Y. Lu, "To Filter or to Authorize: Network-Layer DoS Defense Against Multimillion-Node Botnets," –in SIGCOMM Comput. Commun. Rev., vol. 38 (4) pp. 195–206, Aug. 2008.

[12] S. G. Hong and H. Schulzrinne, "PBS: Signaling architecture for network traffic authorization," –in IEEE Communications Magazine, vol. 51 (7), Jul. 2013.

[13] D. Turk, "Configuring BGP to Block Denial-of-Service Attacks," RFC 3882, September 2004.

[14] W. Eddy, "TCP SYN Flooding Attacks and Common Mitigations," RFC 4987, August 2007.

[15] D. Cooper, S. Santesson, S. Farrell, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile," RFC 5280, May 2008

[16] (2014, Mar.) SCAPY. http://www.secdev.org/projects/scapy/