PREPRINT submitted to Security and Communication Networks (Wiley)

Policy Based Communications for 5G Mobile with Customer Edge Switching

Raimo Kantola, Jesús Llorente Santos, Nicklas Beijar
Department of Communications and Networking
Aalto University
{firstname.lastname}@aalto.fi, Tel +358 40 7501636, Fax: +358 9 47122474

ABSTRACT

Besides more capacity and faster connections, 5G is expected to provide ultra-reliable services for example for machine to machine communications. In this paper we advocate that 5G must do its best to eliminate malicious traffic as a cause of failure of legitimate services. This paper proposes that all communications in 5G should be controlled by policy. The policies facilitate cooperation of customer networks against misbehaving actors and collecting evidence of malicious activity. Dynamic policies can react to hosts that are used in attacks. We propose a system controlled by policy that overcomes the classical weaknesses in the Internet, namely source address spoofing and denial of service attacks. We propose to improve the mobile device experience by new methods of Network Address Translator traversal suitable for battery powered mobile devices. We believe that 5G will be the major driver for the Future Internet, which is why we relate our approach to other proposals for Future Internet architecture. Our approach can be deployed one network at a time as it limits the changes to edge nodes, no compulsory changes are proposed to hosts. The paper reports the experience from experimentation, evaluates scalability and security including initial results on performance.

Keywords: 5G, Future Internet, locator/identifier split, customer edge switching, NAT traversal, policy management.

1. Introduction

5G is the next major step in the development of mobile networks. Besides more capacity and faster connections, 5G is expected to smoothly support the Industrial Internet and the Internet of Things [1, 2]. In addition, new services in critical areas such as e-health, e-banking or machine-to-machine communications will require ultra-high network reliability and availability [2]. 5G networks will be built on top of Software Defined Networks (SDN), Network Function Virtualization (NFV), Cloud technologies combined with an evolved radio access making use of new spectrum and energy efficiency [3, 4]. We argue that 5G must do its best to eliminate malicious traffic as a cause of failure of legitimate services. The reason is that in many scenarios human life or the operation of industrial machinery depend on network communications. 5G networking will be about automating many activities that now depend on humans, not just about making phone calls or accessing web content. It is also about digital commercial transactions between the communicating parties with the need to minimize risks, to limit liabilities and lower the transaction costs. We believe that for these reasons, 5G must make significant progress compared to the present day Internet in terms of providing a predictable service.

Failures occur due to hardware failure, software bugs, human error and malicious human activity. For the sake of providing a more predictable service, 5G should overcome the classical

Internet weaknesses of source address spoofing and the ease of launching denial of service (DoS) attacks. The goal should be diminishing the role of malicious human activity as a cause of failure of legitimate network use by attributing malicious network activity to the responsible actors and automatically containing the harm they are capable of causing. We also believe that it is essential to blend the boundary of closed and open networking; who can communicate with whom should be uniformly managed by a policy. This is because (a) under favorable network conditions we may be open to accept any communication; however, under network duress we would want to use additional network protection prior to accepting a flow. Also, (b) in a service-oriented society where many things are sold "as-a-service", it should be easy to outsource responsibility for assets at homes or in the Industrial Internet to a service provider. This could be handled by highly automated policy management in a scalable and efficient manner.

Based on the high level goals, we advocate a particular approach to network security: by design, network protocols and algorithms should do their best to identify attacks, what resources and which methods the attacker is using. We propose to aggregate and share this information, so that the scope and impact of the attacks can be minimized.

This paper describes how we can meet the goals described above without modifying the hosts, using Customer Edge Switching (CES) [5]. CES enables cooperation between edge nodes and enforces user-defined policies for controlled flow admission. We argue that the deployment of the proposed solution is feasible one network at a time, especially for networks that serve mobile hosts and devices connected to the Internet of Things. The paper also shows how the solution enables communication across address realms and consequently solves the problem of traversing NATs. In addition, CES helps deal with the IPv4 address exhaustion, scalability of the core routing system and mutual trust of the communicating parties. The paper also presents an initial evaluation of scalability, security and performance.

Our overall solution for policy based communication for 5G has two tiers: (1) the interaction of customer networks and (2) evidence collection, aggregation, and reputation of all entities. The latter could be implemented by an Internet wide or alliance wide trust management system [6, 7]. Here we concentrate on the first tier. The tier that brings the required level of control to the interaction of hosts and customer networks is CES. The goal is to replace NATs with CES nodes. CES aggregates two main components: (1) Customer Edge Traversal Protocol (CETP) [8] and (2) a Realm Gateway (RGW) [9, 10]. CETP is used for session signaling and tunneling packets between two networks that both have deployed CES. The RGW allows communication between legacy Internet hosts and hosts behind CES nodes. The solution is supported by identity and policy functions in the management plane. As proof of concept, we have prototyped the main aspects of CES, verified the correctness of its algorithms, and studied the limitations of our approach. The instructions to access the demonstrator of CES are publicly available [11].

To summarize, the contribution presented here is a new system controlled by user-defined policies that overcomes the classical weaknesses in the Internet, namely source address spoofing and DDoS, and allows to ubiquitously collect evidence on network misbehavior thus helping to contain the harm caused by infected Internet hosts. Particular advantages of our proposal are that it can be deployed one network at a time, the costs of deployment are well aligned with the benefits and the system is well suited to the needs of mobile devices.

2. Motivation and Objectives

The motivation and objectives stem from the needs of the mobile users, in particular in 5G. We relate them also the Internet principles, because 5G will be a major step in the development of the Internet.

The Internet is based on best-effort service, provided by the IP layer. Best-effort implies that the network will do its best to deliver the packet of the senders to their intended recipients, i.e. the network serves the interests of the senders. The interests of the receivers differ from the interests of the senders by the amount of unwanted traffic that the receivers will discard. We argue that networks should do better in order to best serve the interests of both senders and receivers, thus delivering only the desired traffic. The idea that the network will deliver only the desired traffic is particularly important to a wireless, battery-powered receiver. The interests of both parties can be expressed in the form of policies. We call Best-Effort Communications (BEC) the principle by which a network equally meets the needs of both sender and receiver. To implement the principle, cooperation between networks is needed. We set the objective that all communication is managed by policies defined by the users. Because the objective is to serve mobile devices and users, the policies must be executed by edge nodes serving those users.

This gives raise to the architecture of Fig. 1 that is applied to Internet communications for enforcing trust and security.

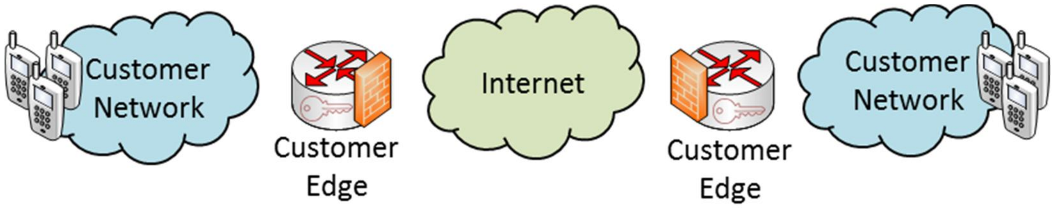

Figure 1 - Architecture for Trust and Security for the Internet

In the architecture of Fig. 1 we call each customer network a trust domain. The Internet connecting the customer networks forms a federated trust domain. We call the edge node serving the initiator of communication the outbound edge node. The destination of the communication is served by an inbound edge node.

The model imposes cooperative behavior between customer networks in order to establish the level of trust required by the receiver. A flow is admitted only when the sender meets all the requirements set by the receiver in session negotiation. We argue that due to a web of business contracts between customers and network providers and between network providers, it will be more feasible to enforce cooperative behavior among the customer networks instead of directly between the connected hosts. Fully analyzing this argument is out of scope of this paper. We intend to study this further using game theory as in [12, 13].

If a sender uses source address spoofing, the receiver is unable to identify the real source of the packet. As a result, attributing evidence of behavior to either hosts or networks is a challenging task. Internet allows any host to send packets to any destination address. Under this premise, any user can launch a DoS attack against any other user. If the malicious user can spoof its address, or hide behind another host, the attack is hard to trace back and the likelihood of getting away without any punishment is high. The objective is to give tools to the receiver to block spoofing with ease. A related objective is to be able blame the source network when a malicious source is detected. A receiver's policy should be able to state conditions like: prior to creating state in the inbound edge, eliminate spoofing, type of identity the sender must present prior to admission, what checks are executed on the presented identity, which exact source identities are admitted.

The content and importance of communication differs from application to application. Therefore, the level of required assurance prior to admission should be decided by the receiver's policy. The administration of these policies also needs to be considered. An improvement to the current state-of-the-art would be managing the admission policies just like the quality of service and charging policies are managed by the 3GPP policy management architecture. A mobile operator now uses the same firewall rules for all mobile users or a large group of users. Under policy-based communications, each user is treated individually.

The use of unique and robust identifiers is essential for attributing blame of malicious activity. For many applications (e-mail, web, etc.) a locally significant ID, whose semantics is known to the allocating network, would be sufficient. We believe that for privacy reasons such a locally significant ID is suitable for the purpose of attributing evidence of behavior. This is because in most cases the owner of the host is also a victim. End hosts may have one or more globally unique e.g. Fully Qualified Domain Names (FQDN) or Mobile Subscriber ISDN telephony numbers (MSISDN). These are used for reachability.

To put all the above together, the edge node will have flow state similar to NAT binding, where the host's private address, locally significant and global ID, the routing locators of the edge nodes and communication port numbers come together. The binding state allows checking many conditions before flow admission.

For mobile communications we consider three major challenges: (a) control of the communications; (b) reachability of the nodes; and (c) network-enabled power saving.

We propose uniform policy control of communications to actively block unwanted traffic. The enforcement of the policy is a firewall function. However, it is a function of the policy management to allow easy modification of the admission rules when a user installs a new application.

Network-based firewalls are preferred over host-based because they block traffic before it reaches the mobile device; avoid cluttering the radio interface; increase the network capacity available for legitimate traffic; and avoid disrupting idle cycles in the mobile devices. However, network based middle-boxes have been traditionally seen to hinder end user reachability by blocking inbound connections originated in the Internet. To overcome the limitation we propose to manage reachability by user defined policies.

For the purpose of reachability, it is assumed that mobile nodes will roam across different stub networks, typically behind firewalls or Network Address Translator (NAT). To achieve reachability, a clear distinction has to be made between host-addressing and host-identification. For the latter we can use e.g. Fully Qualified Domain Names (FQDN) and MS-ISDN numbers. In addition, the network must provide mechanisms that guarantee the unilateral establishment of inbound connections in order to allow hosting services in the private network.Mäenpää showed experimentally that a round of Interactive Connectivity Establishment (ICE) uses 100 STUN/TURN messages for NAT traversal and the two rounds of ICE for Peer-to-Peer SIP can introduce a session setup delay of more than 20s. Mäenpää also showed how ICE can be optimized to cut the delay to about 10s [14]. At the same time, the classical ITU-T delay requirement for a national call setup is 2s [15]. The question of power saving in mobile nodes is related to computational requirements, usage of radio interface and cycles in idle state. The network must provide mechanisms that enable host communications minimizing the battery consumption of the mobiles. Our objective is to remove the need for ICE and STUN/TURN. At the same time, applications that want to use them should be able to do so for compatibility.

We argue, that for most new networking innovations, a prerequisite for adoption is that it can be deployed one network at a time. Each network that makes the deployment decision should be able to justify the investment independent of other investors. We strive to reach this goal with Customer Edge Switching and Realm Gateway. Our solution introduces policy-based communication using trusted identities for minimizing the risk of communication as well as end-to-end connectivity across private stub networks without any cumbersome mechanism for NAT Traversal. This means hosting servers in private networks, which is most useful for applications where reachability is a key feature.

3. Related work

Ubiquitous policy based communication is a new Internet Architecture. It favors the receiver by letting it check several conditions prior to flow admission to the air interface and the device. More precisely, both parties of the communications can check several conditions before the edge-to-edge session is established and data can start to flow. Many solutions for the Future Internet have been proposed. Few of the well-known proposals are particularly motivated by mobile use. It is more popular to start from the issue of scalability of the core Internet. Most of the proposed solutions require changes in several parts of the network or even in the end devices, thus raising the question of incentives for deployment. As an example, IPv6 has demonstrated how the process of adoption is slowed down when synchronized changes are required from different stakeholders.

Opposed to the traditional best-effort approach, the Pub/Sub architecture [16, 17] tries to prioritize the wishes of the receiver, i.e. nothing is delivered to the receiver unless it has previously subscribed to the content. The main challenge lies in how to accurately define the host's wishes of blocking undesired traffic and allowing the intended traffic.

Several proposals are based on the premise of decoupling location and identification functions of the IP address and/or introducing an additional protocol layer to all nodes or at least to the hosts. TRIAD [18] proposes a combination of name and source routing across address realms. IPNL [19] presents a NAT-extended architecture that combines globally routable IP addresses with domain names for end-host identification by adding an additional IP layer at the cost of revealing the private IP addresses of the hosts. Similarly, Shim6 [20] includes an additional locator below the transport protocol to provide multi-homing and load sharing. HIP [21] also requires adding a new layer to the stack in hosts; introduces identity tags and cryptography to authenticate peers and protect them from DoS attacks. I3 [22] presents an overlay network based on indirection with a rendezvous server that enables decoupling senders from receivers and provides support for mobility, anycast and multicast. In addition to the ID/Locator split, MILSA [23] proposes a new architecture to distinguish the functional roles between trust domains (organizations) and connectivity domains (service providers). In [24] trust is already recognized as the most important new requirement compared to the original Internet design.

Work on Locator/Identifier Separation Protocol (LISP) [25] has also gained momentum in the IETF. LISP calls host addresses Endpoint Identifiers (EID) while the addresses of network nodes called tunnel routers are routing locators (RLOC). EIDs are not decoupled from routing since they are used for addressing in edge networks. LISP relies on a Domain Name System (DNS) request originating from the sender host to trigger the signaling operations edge-to-edge in order to locate the destination and enable subsequent data forwarding. The protocol implements mapping and encapsulation for IPv4/IPv6 and address rewriting for IPv6. The solution is transparent to end hosts but is not exempt from security concerns [26].

With a stronger focus on security and tackling DoS attacks we find SIFF [27], StopIt [28] and PBS [29]. In SIFF the authors propose a proactive stateless solution for minimizing the effects of DoS attacks by tagging and prioritizing privileged packets, however it requires changes in both hosts and routers and is vulnerable to brute force attacks. StopIt enforces filtering at the edges and enables a receiver to dynamically upload network filters thus blocking unwanted incoming traffic; StopIt servers communicate with each other in order to block the reported traffic at the source before leaving the network and to establish punishments to misbehaving hosts. StopIt lacks an Internet wide trust processing system and makes heavy design choices to counter system attacks. PBS aggregates both proactive and filtering mechanisms and relies on the NSIS [30] protocol suite for carrying signaling messages. Unfortunately, PBS also requires synchronized changes both in hosts and network edge nodes, which we believe is detrimental to its adoption. The paper also reveals how a heavy cryptographic solution for preserving confidentiality edge-to-edge easily leads to high delays thus making it unusable for delay sensitive applications.

In 5G, flagship 5G-PPP project in Europe, Metis [2], advocates that network control should be based on SDN [4]. For the subject of this paper, we fully embrace the proposal that the new network functions are aligned with the idea of SDN.

Since none of the proposed solutions for the Internet architecture have been widely adopted, we argue that we should look at the problem from a different angle, namely trust. To help adoption, we differ from all the mentioned proposals in that we limit the changes to the edge; we provide the possibility of deploying one edge network at a time while there are no compulsory changes in hosts.

4. The Architecture of Customer Edge Switching

Our preliminary work on Customer Edge Switching (CES) [5] describes an implementation of the trust-to-trust principle advocated by David Clark [31]. In this paper, we present an enhanced solution for the case of IPv4 based core transport and for use in 5G.

4.1. Overview

A CES function can be seen as an extension of a stateful firewall. It is a co-operative firewall that leverages policy negotiation allowing secure communications between hosts. In addition to the two decision alternatives of current firewalls, i.e., allowing or dropping a given packet, the CES function can issue additional queries to satisfy a policy prior to making its final decision. For this dialogue we propose a protocol called Customer Edge Traversal Protocol (CETP). Furthermore, one of the key features of CES is precisely that it does not require changes to either end hosts or protocols and assures reachability to applications without polling mechanisms.

Additional motivation for CES originates from the lack of available public IPv4 addresses and the slow penetration of IPv6. The address exhaustion problem can be temporarily tackled with the deployment of Network Address Translator (NAT), which implies connecting the hosts to private networks. However, NATs can become especially burdensome for those hosts that require to be reachable at all times, which often means reverting to NAT-Traversal mechanisms. The officially recommended methods of NAT traversal are based on the so call Unilateral Network Address Fixing (UNSAF) [32]. For this purpose the STUN [33] and TURN [34] protocols are used in different ways including the Interactive Connectivity Establishment (ICE) [35]. These methods are cumbersome because (a) they require application specific code; (b) they add significant delay to flow or session establishment; (c) they add messages to the air interface; and (d) they interfere with the power saving methods used on mobile devices.

Since 5G is driven by mobile networking and already today most of the Internet users are mobile, the network should be able to serve these devices better: i.e. no application specific code for the sake of reachability, immediate session setup, no additional messaging over the air interface for the sake of NAT traversal and full alignment of methods of communication with the way power saving works on mobile devices.

CES enables global communication between hosts located in customer networks using private IP addresses, acting as an on-demand routing protocol for traversing multi-homed stub network boundaries. Communicating sessions between hosts are established in a co-operative fashion, thus enhancing trust between the edge nodes and enforcing the policies of both the sender and the receiver. This negotiation is performed via the CETP protocol, which defines a number of control Type/Length/Value (TLV) elements for representing the policies in play. The negotiation process is required only once, and it guarantees that no data will be exchanged between sender and receiver unless both policies are satisfied.

A CES node is reachable in the public domain using a set of Routing Locators (RLOCs) that follow the same semantics as in LISP [25]. For robust connectivity of the customer networks to the core, CETP supports on-demand routing through the customer edge. In the case of connecting it to an IP core, a CES node behaves like a LISP tunnel router with the exception that it supports several types of identities for the purpose of communication rather than EIDs (which actually are addresses). A CES node is similar to a realm boundary node in TRIAD [18] but instead of source routing over the realm boundaries CES uses ID to address mapping to cross the boundary.

CES provides several mechanisms for establishing flow legitimacy, including the negotiation of a variety of IDs, return routability checks, and the use of secure routing locators. The conjunction of these operations allows CES to identify both the sender and its network on a required level of assurance, prior to delivering data packets to the destination.

When the two communicating hosts are connected via their respective CES nodes, the negotiation of CETP policies, that must be satisfied, precedes the relaying of user data between the hosts. However, for legacy communications i.e. there is only one CES node involved, Customer Edge Switching offers a better-than-NAT service provided by Realm Gateway (RGW). RGW allows incoming connections towards the private network based on standard domain resolutions allowing smooth connectivity to the communicating parties.

Our experiments reveal that CES and CETP communications are de facto compatible with NAT-friendly applications and do not require NAT Traversal techniques. However, there are a number of applications that are broken when the communication is handled by a middle-box. These issues are commonly related to the use of IP address literals as well as split connections for signaling and data. In [36] we studied the impact of NAT-unfriendly protocols on CES communication and provide design guidelines to build Application Layer Gateways (ALGs).

4.2. DNS as a Communication Trigger

The CES architecture is tightly coupled with DNS, such that communication in CES is granted via domain resolution. Hosts attempting to communicate with remote parties are required to issue DNS queries in order to decouple identifiers from routing locators. However, using DNS allows the system to decouple endpoint identifiers from routing locators. Consequently, CES maintains a delegated DNS zone of authority with the records of the served hosts.

DNS resolutions are used by RGW for inbound connection establishment, achieving NAT traversal and allowing inbound connections to the private network. For CES to CES communications, the phase of service discovery precedes the CETP connection establishment.

Service discovery is triggered by an outbound CES (oCES) node and determines whether the destination is hosted by another CES and the appropriate course of action to establish the connection. The procedure relies on standard Naming Authority Pointer (NAPTR) [37] DNS queries to ascertain the availability of the CETP service on the remote CES. An example of a valid NAPTR response offered by an inbound CES (iCES) node is as follows:

```
b.ces. 30 IN NAPTR 10 6 "U" "CETP+cesid"
"!^(.*)$!cesid:1=cesb.ces.?ip=192.0.2.10?alias=IXP!" .
```

The value "`CETP+cesid`" indicates that the CETP service is available, however, different services might also become available in the future i.e. TCP/TLS, HIP. The routing locator "`ip=192.0.2.10`" and "`alias=IXP`" identifies the connected network and the location of the service within that network. The alias field enables the use of private transit links instead of less preferred links or networks. The NAPTR response may also convey a valid CES identifier "`cesid:1=cesb.ces.`" indicating ID type and value to perform routability checks prior to initiating the CETP signaling.

From the architecture perspective, NAPTR records could be cached in intermediate proxies or delegated to third-party DNS servers; however, if the CES node retains the authority over the DNS leaf, it would improve the control for the CES functionality and administration. In addition, DNS provides mechanisms for redirecting other domains towards the ISP's controlled zone by means of Name Server (NS) and Canonical Name (CNAME) records.

At host level, decoupling endpoint identifiers from routing locators means that an IP address can no longer be used for identifying a node. To overcome this addressing challenge we leverage the concept of proxy-address for the representation of remote hosts. These proxy-addresses are allocated from the IPv4 private address space [38] or using the IPv6 Unique Local Address (ULA) [39] regarding the nature of the host.

4.3. Customer Edge Traversal Protocol

Customer Edge Traversal Protocol (CETP) has been devised to convey signaling and data across CES nodes. The signaling is exclusively CES related information required to establish a data connection for two communicating hosts across their respective CES nodes. The information exchanged corresponds to the policies defined for the users. On the other hand, the data payload strictly carries user-generated data tunneled within a CETP session. Instead of native CETP tunneling, other tunneling methods for data could be used in some environments.

Ingress filtering at the edge nodes contributes to reducing the amount of spoofed traffic in the network. Bogus traffic generated by misbehaving hosts can be blocked closer to the source thus reducing the negative impact on the network.

CETP includes session identifiers to distinguish between different user connections. Therefore, it is possible to use different links to carry separately the signaling and the data for a higher level of assurance and improved heuristics of the algorithms. However, we are aware that this might not be feasible in all cases, and that is why CES is equipped with a number of extended checking mechanisms.

A CETP packet contains a mandatory fixed header of 4 bytes, source and destination session tag fields of variable length and optional signaling or payload elements. Both of these elements have been defined following a flexible TLV encoding schema. Furthermore, CETP packets are agnostic to the underlying technology and can be transported over a number of network and transport protocols, i.e. Ethernet, IPv4, IPv6, TCP or UDP.

TLV elements are classified attending to the Type value, which can be further dissected into operation, group and code values. The Value field of Length bytes is always padded to a 32-bit boundary for faster processing with a minimal impact on overhead.

The operation field defines three values: (a) `info` indicates the value of an element; (b) `query` requests the value of an element; and (c) `response` is compulsory to a previous query and might be empty if such element is not available.

The TLV elements currently defined for CETP are presented in Table 1.

Table 1: Type/Length/Value elements in CETP

Group	Code	Description
control	cesid	The CES ID
	dstep	The destination endpoint ID
	caces	The CA address for validating a CES
	caep	The CA address for validating a host
	terminate	Contains the terminating information
	warning	Contains the warning information
	ack	The acknowledgement sequence number
	ttl	The time-to-live for the session
	ratelimit	The rate limitation for the session (bps or pps)
	headersignature	The signature of the CETP packet
id	fqdn	The FQDN
	maid	The Mobile Assured ID
	moc	The Mobile Operator Certificate
	msisdn	The MSISDN number of the host
rloc	ipv4	An IPv4 address of the CES
	ipv6	An IPv6 address of the CES
	eth	A MAC address of the CES
payload	ipv4	The payload contained in an IPv4 packet
	ipv6	The payload contained in an IPv6 packet
	eth	The payload contained in an Ethernet frame

A policy is made of three different vectors: offer, requirement and available. These vectors store the TLV elements used to define each specific policy.

The scenario depicted in Fig. 2 illustrates a connection setup of two end hosts behind their respective CES nodes. The policies defined are the following:

```
Outbound Policy Host-A:
        Offer        = {control.cesid, id.fqdn, rloc.ipv4}
        Requirements = {control.cesid, rloc.ipv4}
        Available    = {control.cesid, id.fqdn, rloc.ipv4}
```

```
Inbound Policy Host-B:
       Offer        = {} #There is no offer for inbound
       Requirements = {control.cesid, rloc.ipv4}
       Available    = {control.cesid, id.fqdn, id.maid, rloc.ipv4, rloc.eth}
```

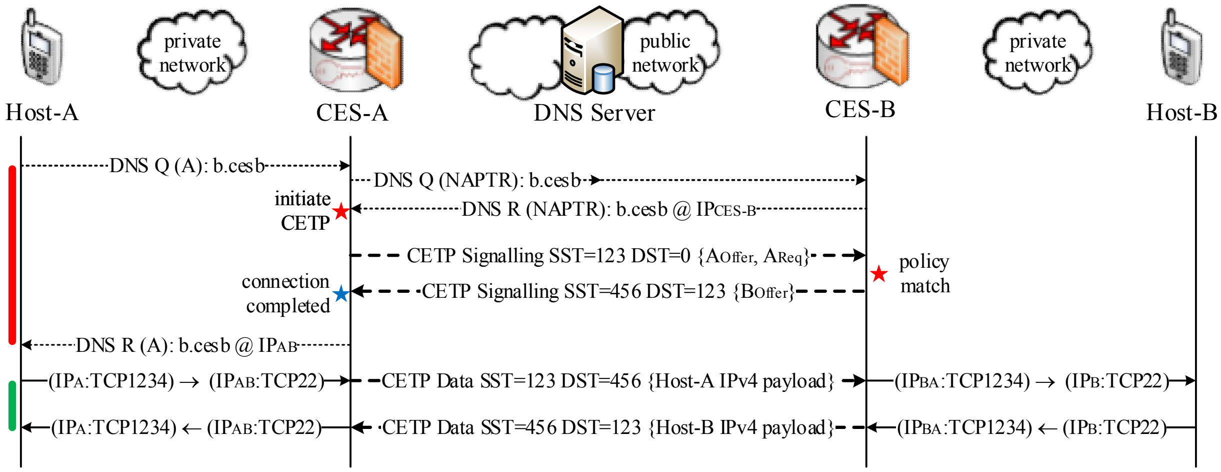
private network
public network
private network
Host-A
CES-A
DNS Server
CES-B
Host-B
DNS Q (A): b.cesb
DNS Q (NAPTR): b.cesb
initiate CETP
DNS R (NAPTR): b.cesb @ IPCES-B
CETP Signalling SST=123 DST=0 {AOffer, AReq}
policy match
connection completed
CETP Signalling SST=456 DST=123 {BOffer}
DNS R (A): b.cesb @ IPAB
(IPA:TCP1234) → (IPAB:TCP22)
CETP Data SST=123 DST=456 {Host-A IPv4 payload}
(IPBA:TCP1234) → (IPB:TCP22)
(IPA:TCP1234) ← (IPAB:TCP22)
CETP Data SST=456 DST=123 {Host-B IPv4 payload}
(IPBA:TCP1234) ← (IPB:TCP22)

Figure 2 - CES to CES Communications with CETP

We can identify three different phases during the connection setup. The following actions take place when Host-A attempts to connect to the SSH service of Host-B.

- *DNS*: Host-A resolves the FQDN b.cesb via CES-A to communicate with Host-B. CES-A initiates CETP discovery with a NAPTR resolution. CES-B confirms the availability of a CETP endpoint answering the NAPTR query.
- *CETP Signaling*: CES-A initiates a CETP connection with session tags (SST=123,DST=0). The request also includes the policy vectors for Host-A $\{A_{Offer}, A_{Req}\}$ and the destination FQDN – b.cesb. CES-B checks for policy match and answers with (SST=456,DST=123) and computes an answer to the A_{Req} with the B_{Ava}. CES-A validates the answer and completes the connection. Both CES nodes allocate proxy addresses for their respective hosts, IP_{AB} and IP_{BA}.
- *DNS*: CES-A answers Host-A with the allocated address IP_{AB}.
- *CETP Data*: Host-A and Host-B are able to communicate through their respective proxy IP addresses. The CES devices perform CETP tunneling of data packets with the respective session tags, across core networks, on the negotiated RLOCs.

The iCES node executes its admission policy and decides whether to accept the offer as is or else to respond to the oCES node with its own requirements. iCES can postpone a connection establishment until any possibility of source RLOC spoofing or source ID spoofing has been eliminated. This is possible by executing return routability checks over addresses or names, using signed RLOCs to prevent spoofing or demanding certified IDs prior to accepting the new connection. The evaluation of two policies always results in either success or failure; this result is typically achieved within 1-3 round trips depending on the policies in use.

We conducted additional testing and analyzed the setup delays for new flows using CETP in our development environment with a proprietary Python data plane implementation. These tests follow the architecture represented in Fig. 2 and aim at identifying possible bottlenecks and perform further optimizations to the architecture.

In the experiment, we established a hundred new CETP flows between the CES devices and measured the delays perceived on the originating host. The results in Fig. 3 represent the delay perceived by the originator in two stages: DNS resolution of the destination domain name and forwarding of a data packet. The figure also distinguishes between *connection establishment* (*e*) and *connection reuse* (*r*).

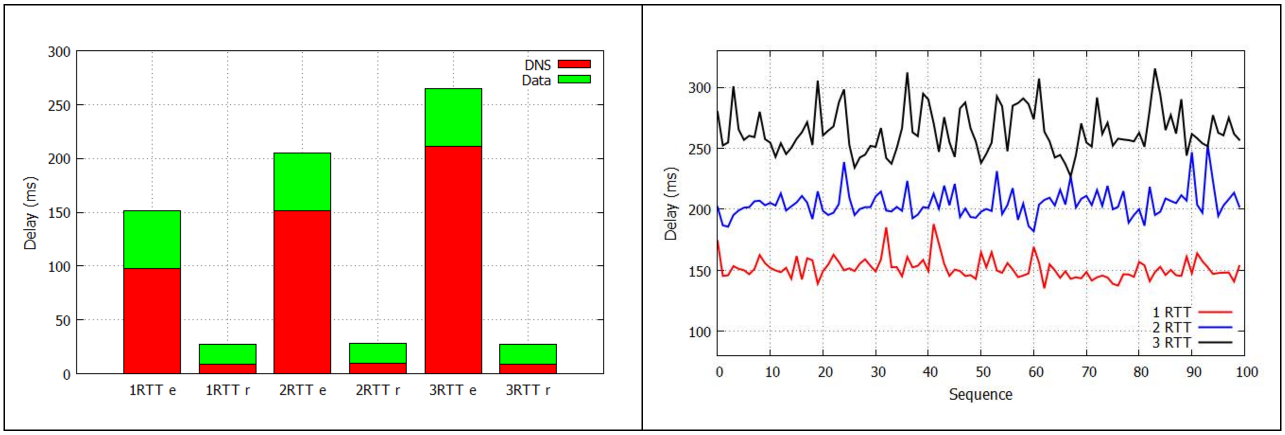

Figure 3 - Session setup delays using CETP

For new connections, the results reveal that most of the delay is attributed to the signaling phase for the establishment of such connections. The reason is due to the separation of control and data planes, which introduces additional packet processing and due to the python implementation. Moreover, the signaling delay is also increased with the complexity of the policies, as additional round trips are required for a successful match. On the other hand, the delay for data packets remains constants regardless of these policies. When a host reuses an existing connection and the connection state already exists in control and data planes, the delay is reduced to its minimum.

The scenario is open to optimization using dynamic policies that maximize connection reuse and automatically react to attacks or lack thereof. Another way to speed up the session setup is optimizing the code itself.

4.4. Realm Gateway

Realm Gateway (RGW) [9, 10] is the solution we created for providing connectivity between hosts in IP legacy networks and hosts in trust domains or private networks. It can be integrated with CES or used as a standalone solution that can be deployed independently. The RGW resembles the behavior of traditional NAT devices allowing the hosts located in the private network to connect to public networks sharing a single public IP address. However, opposite to NATs, the RGW allows unilaterally initiated inbound connections from public networks towards private hosts via the Circular Pool of Public Addresses (CPPA). The CPPA algorithm is triggered by incoming DNS queries and solves the reachability problem introduced by NATs.

The RGW provides three different methods for establishing inbound connections: (a) a general purpose connection is triggered by an incoming DNS query and managed by the CPPA, the query may convey a service name, e.g. `b.rgw` (FQDN), `ssh.b.rgw` (Service FQDN); (b) incoming HTTP(S) traffic is handled by a reverse HTTP proxy, triggered by a particularly formed DNS query concerning the destination host, e.g. `www.b.rgw`, and (c) specific inbound mapping on public IP addresses, similar to port forwarding or demilitarized zone typical of NATs for high load traffic premium subscribers.

The CPPA operates transparently to the communicating hosts and relies on standard DNS resolutions. Therefore it does not require any changes to the current hosts, protocols or applications. Upon receiving a DNS query, the CPPA dynamically allocates an available address from the pool, locks it for a maximum time T_{Tout} (typically 2s), and replies to the DNS query with the allocated address and TTL zero to avoid caching. Afterwards, temporary state information is created. The state ($H:iP_H$, $R_X:oP_H$, $P_{protocol}$, T_{Tout}) is unique and includes: the IP address and port of the private host ($H:iP_H$), the IP address and port on the public side of the RGW ($R_X:oP_H$), the protocol ($P_{protocol}$) and the lifetime of the entry (T_{Tout}). Upon receiving a new flow matching the state, a new connection is created. The address is released and returned to the pool for future use. The inbound packet is forwarded to the private host acknowledging subsequent data packets as an on-going flow.

Fig. 4 represents how several Internet hosts can initiate connections with different services hosted by hosts connected behind a Realm Gateway.

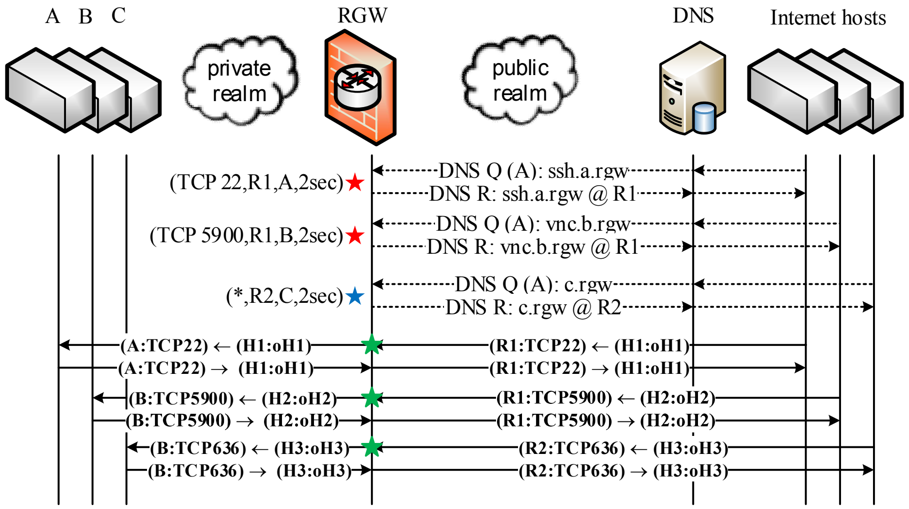

Figure 4 - Establishing Inbound Connections with RGW

Through extensive testing we isolated a scenario, however unlikely, susceptible to errors. Due to the circulating nature of the IP addresses given by the RGW, a host that reuses a socket client to connect to several servers might enter an error state. As best practice, we recommend client applications to refrain from using the SO_REUSE and let the OS choose an ephemeral port for new connections.

The scalability limitations of a RGW in terms of new inbound flows are determined by a number of factors. The pool size (N_{IP}) indicates the number of IP addresses available for allocation; the offered traffic (ρ); and the service time (T_S) refers to the time elapsed since the allocation of the address until the reception of the first data packet. Eq. 1 gives the number of connections established per time unit based on the size of the pool and the service time. Retransmissions and packet loss are intentionally excluded from the calculation; as a result, the equation yields the upper bound of the number of new connections per second for the system operating under ideal circumstances.

$$\eta = \frac{Nip * \rho}{Ts} \quad \text{(Eq. 1)}$$

The ρ variable is a generalization of the offered traffic and relates to the inbound inter-arrival times and traffic patterns with FQDN or SFQDN. When SFQDN requests are present, they can overload the same public IP address resulting in higher address reuse and efficiency. Analytically we find that the theoretical maximum for ρ corresponds to uniformly distributed traffic following the SFQDN name schema, also uniformly distributed and spread throughout all the combinations possible of ports (2^{16}) and transport protocols (UDP and TCP). In practice, real traffic would be exponentially distributed and the efficiency of the SFQDN is determined by the most popular service. For FQDN initiated traffic, due to the normally repeated 4 DNS queries, experimentally we established that the number of new flows for even a small number of addresses (3..7) is close to the upper bound in Eq. 1.

Further discussion on scalability and performance of the Realm Gateway is available [10]. In the study we analyze the impact of different traffic distributions and pool sizes, with variations of FQDN and SFQDN.

4.5. Security Mechanisms

CES is our proposal for replacing NATs in customer networks and allows policy controlled connections to end devices. The policies are explicitly defined to satisfy the communication desires of the hosts. Policy based communications aims to nullify classical weaknesses in the Internet such as source address spoofing or DoS attacks.

CES includes a number of security techniques for different communicating scenarios. We distinguish between CES-to-CES and legacy communications.

A) CES-to-CES Communications

The scenario grants end-to-end communication between hosts located in remote private networks connected to the Internet by their respective CES node. The communication between the end hosts is brokered by the CES nodes and the policies defined are enforced via the CETP protocol. The following mechanisms are proposed to minimize the risks against network abuse [8].

- *CETP Cookie TLV*: It follows a similar approach to TCP Cookie in order to mitigate TCP SYN flood attacks. CES computes a cookie value by hashing a number of fields of the CETP header, CES identity values, local secrets and timestamp. The cookie mechanism bears three security layers against forgery attempts such as a local secret, a symmetrickey encryption and a timeout. CETP Cookie contributes to eliminating address spoofing and the detection of replay attacks.

- *CES Registration*: CES nodes can use IDs or registered RLOCs for identification. Registration of CES nodes can be performed in a centralized or decentralized fashion. After address spoofing is eliminated, a centralized server can verify the authenticity of a given CES RLOC before the CETP packets are processed. In some other cases, the verification process can be leveraged with the X.509 certificate and a Certificate Authority. In conjunction with the CETP Header Signature, it is possible to determine the legitimacy of a source.
- Ultimately, the edge-to-edge protocol could facilitate mutual authentication of the hosts prior to flow admission.

We consider a communication to be securely established if we can assure the communicating endpoint is not a spoofed source, the identification is validated with a reputable source and the CETP policy negotiation completes including the use of trustworthy identifiers.

B) *Legacy Communications*

This scenario grants end-to-end communication between hosts located in private networks and the legacy Internet. CES allows inbound connections towards the private network without requiring additional NAT traversal mechanisms via the CPPA component of RGW.

The RGW could be susceptible to inherent Internet attacks involving the abuse of DNS and typical address spoofing. If not neutralized, DoS attacks utilizing DNS as a reflector and amplifier could eventually lead to a blocking state of the CPPA resulting in the depletion of the address pool. As a result, the RGW would be unable to accept any new inbound connection, however the ongoing connections would not be affected. In other cases, unwanted traffic sent by spoofed sources could hijack a connection state reserved by a legitimate source, thus leaking malicious traffic in the private network.

The RGW adheres to the following principles in order to effectively tackle inherent Internet vulnerabilities: (a) flow acceptance must be limited to verifiable sources to prevent address spoofing and resource exhaustion; (b) UDP flow initiations are admitted after the connection was signaled through a secure channel e.g. SIP(S) [40]; (c) under network stress, resource access should be granted based on the source reputation.

The RGW implements the following security techniques[1] in order to minimize the communication risks based on the principles defined:

- *Rate limitation and DNS Server Classification*: DNS queries served simultaneously to a host and from a DNS server are rate limited. DNS servers are classified into White, Grey and Black lists. Servers on each list get a different treatment from the RGW. Whitelisting is based on contracts between ISPs and enforcing of Service Level Agreements (SLA) guarantees of better serviceability. Untrusted peers enter contention for shared resources

[1] H. Kabir, J. Llorente and R. Kantola, "Securing the Private Realm Gateway," unpublished.

where weighted queues may apply. Misbehaving peers are penalized with a degraded service.

- *TCP DNS Relay*: A remote DNS server can connect to the RGW via a TCP connection, which guarantees a spoofing free communication channel. In conjunction with DNS extensions EDNS0 [41] and DNS ingress filtering on the remote site, it is possible to collect evidence of misbehavior tracing back to the originating host.
- *Monitoring*: The arrival of DNS queries and their success or failure rates are monitored continuously allowing to dynamically re-classify servers that are conveying attack traffic. The arrival of the first packets to a flow (TCP SYN) that does not exist is monitored to detect connection hijacking attempts.
- *TCP Splice*: For Internet originated incoming connections, the RGW may challenge the sender of the TCP SYN with a cookie embedded in the Initial Sequence Number (ISN) [42, 43]. If the TCP handshake completes successfully, we can ascertain the authenticity of the sender, assure no spoofing and accept the connection. However, if the sender has history of misbehavior, the connection is terminated. Data packets between source and destination are forwarded following TCP splicing techniques [44]. Furthermore, RGW solves the problem related to content-based switching and server selection with the Circular Pool algorithm, as the destination is known based on a previous DNS query.
- In addition, in a large RGW, policies can differ per inbound interface. A large RGW node has more options on reacting to attacks and is able to modify its attack surface.

The depicted mechanisms contribute to establishing safer connections with legacy networks. The interworking with legacy networks with additional security features is a fundamental capability that facilitates adoption.

5. The bigger picture and deployment

We expect that 5G network is controlled by a set of network applications [46]. One of them is the Mobility Management Application that provides Mobility-as-a-Service (MaaS). CES functionality is the heart of the Access Application that links MaaS to the Internet and to other Service Delivery networks. This Access Application takes on the functions that are now handled by the Packet Data Gateway (P-GW) in 3GPP networks. We created a demonstrator of 5G network control using a set of SDN applications [46] where the Access application is based on our CES prototype. The demonstrator carries traffic between the Internet and LTE user devices through an eNodeB that was supplied by Nokia.

The benefits of our approach to mobile access operators are that CES will make it viable to efficiently implement web servers, SIP User Agents or any other server (e.g., for interactive multimedia communication on mobile devices), giving the mobile operators new attractive services to offer to their customers. In particular, CES makes VoIP more viable for mobile hosts than with the present solutions.

The deployment of CES in corporate and access networks can take place one stub network at a time because CES behaves similarly to NAT. The CES technology benefits corporations by providing them with better means of protecting their networks from unwanted traffic and extending the lifetime of IPv4 in their networks. At the same time it makes it technically easier to deploy IPv6.

CES strengthens the mobile operator as a trust broker for its customers allowing full mutual authentication of the edge nodes prior to session setup via CETP. For mobile operator assured IDs and mobility support CES requires a Diameter Client that can query the HSS. Similarly, policy management can be delegated to the current 3GPP architecture. These policies are retrieved by the communicating CES node, cached and reused. Policy update mechanisms can also take place if any modification is required. Policies can be static or dynamic. By using dynamic policies, we can optimize the operations and reduce the delay of session setup.

CES achieves isolation of customer and core network technologies. Each trust domain can select its own forwarding technology for its network, irrespective of others. Options for forwarding technologies may include IPv4, IPv6, 802.1 (and its variants such as 802.1ah), combinations of optics with carrier grade Ethernet, MPLS-TP [47], or MPLS/IP [48]. Renumbering at the core has no impact on the corporate network. Renumbering of a customer network has no impact on the core network.

If end users are asked to manage their policies directly, usability will be poor. We suggest that instead, when a user makes decisions like, install a new application on his device, or delegate some function to a service provider at his home, the creation and installation of suitable policies is handled at the background between e.g. the mobile application store or a contracting service and the 5G operator.

Deploying a global system of trust based on stakeholder reputation will be challenging. Initially, ISPs can form trust alliances for exchanging evidence of the behavior of Internet entities within the alliance. Exchanging information on detecting intrusion patters is already taking place. When customer networks accept responsibility for the hosts they serve, this should not imply a legal liability. Accepting the responsibility should mean just making a promise of reporting evidence of misbehavior and making reasonable efforts of acting on evidence that the network has a bot that is causing harm to other users of the Internet. Making such a promise implies staking your reputation on your good behavior and performance.

To summarize, the proposed concept of CES is most suitable for deployment as a part of the 5G or other 3GPP mobile network, putting all mobile devices behind CES nodes. The role of CES in 5G would be to reduce or eliminate malicious traffic as a cause of failure of a legitimate service letting 5G to provide ultra-reliable service. The concept is equally well suited for use as a basis for providing access to objects in the IoT. The concept is well aligned with the principles of SDN. For deployment into corporate network gateways, CES functionality needs to be integrated into firewall products. Finally, compared to many previous tunneling proposals focusing on solving the core scalability issues, CES is focused on trust and providing added value to end customers, and improves the scalability of the core as a side effect of its tunneling based edge. Elaboration of the use cases is for further study.

6. Conclusions

We presented an overview of relevant solutions developed for the Future Internet and propose our own solution to 5G that we see as a major driver of the Future Internet. On the premise that cooperation [12, 13] is an effective mechanism to curb antisocial behavior in society, we start from the idea that a new flow should be admitted to a receiver only after the receiver's conditions are met. To facilitate trust establishment, we suggest that the collection of evidence of misbehavior should be ubiquitous and attributing this evidence to an entity should be easy. Once the evidence is aggregated to a reputation of an entity, receivers can react to traffic sent by entities with a poor security reputation appropriately: either deny communication or apply additional checks before admission.

Contrary to the current Internet model, where any host can send a packet to any address, in our solution all communication is granted based on policy; this policy allows the receiver to decide what traffic it wants to receive, thus minimizing its risk to an acceptable level. The policy effectively puts the receiver in charge of flow admission thus balancing the needs of the receiver with those of the sender.

CES is an extension to NAT and it separates host identifiers from routing locators. The solution supports many types of identifiers for providing different levels of assurance according to application and user needs. By confining hosts to using private addresses and by keeping them private we boost the scalability of the core network to more users and hosts. By tunneling all traffic edge-to-edge we isolate technology choices in core networks from those in customer networks, decoupling the development of each network from the others. In our solution, complex routing functionality runs in edge nodes while large wide area networks can be built solely with packet transport technology even without power hungry core IP routers.

CES imposes no changes to existing hosts or protocols and it can be deployed one network at a time. CES seeks to be maximally independent of applications but where this cannot be achieved or has not yet been achieved, we continue to support traditional methods of NAT traversal.

The proposed approach can be seen as a synthesis between the traditional networks and Pub/Sub. Best Effort Communication replaces the concept of subscription by a policy that is a set of rules for the allowing and denying of communication. In BEC, the interests of the sender and the receiver meet via a connection broker that executes the policies of the communicating parties. We refer to this connection broker as the CES function.

Our solution differs from LISP, IPNL, and TRIAD because we introduce the IDs of the communicating entities and map them on an address realm and the co-located trust boundaries to addresses. In addition, unlike TRIAD and IPNL, we preserve the semantics of private addresses, facilitating network hiding. Unlike HIP, the host itself does not need to be aware of the ID. Unlike HIP and PBS we do not require any changes in hosts.

What about end-to-end and the proposed BEC principle? Having ubiquitously adopted CES, the Internet would still be an end-to-end network; the difference with the current Internet would be that the Internet service would be carried over a slightly more intelligent bearer and that the new end-to-end protocols would be NAT-friendly by design. CES takes note of the practicality that if some function cannot be efficiently implemented in end nodes, network support is justified.

We believe that our solution is what is needed to adapt the Internet to the dominance of mobile access and make it feasible to fully merge the Internet and the mobile worlds into a single communication framework. We argue that this should be the goal of the 5G network architecture.

We are pursuing further research into CES in areas like optimization for Software Defined Networking, further development of the CETP protocol particularly for SDN, generalizing the concept of realm gateways, improving security, and incremental adoption of trust concepts. In addition, we continue our efforts towards developing the interfaces with the 3GPP architecture and deeper integration with the mobility management and support of roaming of users.

Acknowledgements: We want to express our gratitude to Marko Luoma, Jukka Manner for their insights in particular in early phase of our development of CES. Our thanks also go to our partners in the ETNA, 100GET, Finnish ICT SHOK Future Internet Programme, and in CELTIC MEVICO and CELTIC Plus SIGMONA projects and the EIT ICTLabs SDN activity.

References

[1] Ericsson, "5G radio access," white paper, June, 2013.

[2] A. Osseiran, F. Boccardi, et al., "Scenarios for 5G mobile and wireless communications: the vision of the METIS project," *in Communications Magazine, IEEE*, vol. 52 (5) pp. 26-352014.

[3] Nokia Solutions and Networks, "Looking ahead to 5G," white paper, Dec. 2013.

[4] Huawei, "5G: A technology vision," white paper, April 2013.

[5] R. Kantola, "Implementing Trust-to-Trust with Customer Edge Switching," in *Proc. 24th IEEE International Conference on Advanced Information Networking and Applications Workshops (AINA 2010),* pp. 1092-1099, Perth, Australia, 20-23 Apr. 2010.

[6] Zheng Yan, R. Kantola and Yue Shen, "Unwanted Traffic Control via Global Trust Management," in *IEEE 10th Int. Conf. on Trust, Security and Privacy in Computing and Communications (TrustCom 2011),* pp. 647-654, Changsha, China, 16-18 Nov. 2011.

[7] Zheng Yan, R. Kantola and Yue Shen, "Unwanted Traffic Control via Hybrid Trust Management," in *IEEE 11th Int. Conf. on Trust, Security and Privacy in Computing and Communications (TrustCom 2012),* pp. 666-673, Liverpool, UK, 25-27 June 2012.

[8] H. Kabir, R. Kantola and J. Llorente, "Security Mechanisms for a Cooperative Firewall ," in *International Symposium on Cyberspace Safety and Security (CSS),* Paris, France, 2014.

[9] J. Llorente, R. Kantola, N. Beijar and P. Leppäaho, "Implementing NAT Traversal with Private Realm Gateway," in *Proc. IEEE International Conference on Communications (ICC 2013),* pp. 2174-2179, Budapest, Hungary, 9-13 June 2013.

[10] J. Llorente and R. Kantola, "Transition to IPv6 with Realm Gateway 64," in *accepted for publication in Proc. IEEE International Conference on Communications (ICC 2015),* London, UK, 8-12 June 2015. Preprint available at: http://www.re2ee.org/

[11] R. Kantola. "Routing Edge-to-Edge and Through Ethernets." Internet: http://www.re2ee.org/, [Accessed January 24, 2015].

[12] R. Axelrod. *The Evolution of Cooperation (revised edition).* Basic Books. Dec. 2006.

[13] M. Nowak, "Why We Help: The Evolution of Cooperation," *in Scientific American* pp. 34-39July 2012.

[14] J. Mäenpää. "Framework Architecture for Decentralized Communications." Doctoral dissertation, Aalto University, Department of Communications and Networking, 2013. Available at http://urn.fi/URN:ISBN:978-952-60-5121-5.

[15] ITU-T, "Network grade of service parameters and target values for circuit-switched services in the evolving ISDN. recommendation E.721," International Telecommunication Union, 1999.

[16] P. T. Eugster, P. A. Felber and R. a. K. Guerraoui Anne-Marie, "The Many Faces of Publish/Subscribe," *in ACM Comput.Surv.*, vol. 35 (2) pp. 114-131, June 2003.

[17] M. Sarela, T. Rinta-aho and S. Tarkoma, "RTFM: Publish/subscribe internetworking architecture," in *Proc. ICT Mobile Summit,* Stockholm, Sweden, 10-12 June 2008.

[18] M. Gritter and D. R. Cheriton, "An Architecture for Content Routing Support in the Internet," in *Proc. 3rd conference on USENIX Symposium on Internet Technologies and Systems - Volume 3,* pp. 4-4, San Francisco, CA, USA, 26-28 Mar. 2001.

[19] P. Francis and R. Gummadi, "IPNL: A NAT-extended internet architecture," in *Proc. ACM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM 2001),* pp. 69-80, San Diego, CA, USA, 27-31 Aug. 2001.

[20] E. Nordmark and M. Bagnulo, *Shim6: Level 3 Multihoming Shim Protocol for IPv6,* IETF RFC 5533, June 2009.

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

[22] I. Stoica, D. Adkins, S. Zhuang and S. a. S. Shenker Sonesh, "Internet Indirection Infrastructure," *in IEEE/ACM Transactions on Networking*, vol. 12 (2) pp. 205-218, Apr. 2004.

[23] J. Pan, S. Paul, R. Jain and M. Bowman, "MILSA: A Mobility and Multihoming Supporting Identifier Locator Split Architecture for Naming in the Next Generation Internet," in *Proc. IEEE Global*

Communications Conference (GLOBECOM 2008), pp. 1-6, New Orleans, LA, USA, 30 Nov. 4 Dec. 2008.

[24] J. Kempf, R. Austein and IAB, *The Rise of the Middle and the Future of End-to-End: Reflections on the Evolution of the Internet Architecture,* IETF RFC 3724, Mar. 2004.

[25] D. Farinacci, V. Fuller, D. Meyer and D. Lewis, *The Locator/ID Separation Protocol (LISP),* IETF RFC 6830, Jan. 2013.

[26] D. Saucez, L. Iannone and O. Bonaventure, "LISP threats analysis," IETF Internet draft, work in progress, July 2014.

[27] 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, Oakland, CA, USA, 9-12 May 2004.

[28] 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.

[29] Se Gi Hong and H. Schulzrinne, "PBS: Signalling architecture for network traffic authorization," *in IEEE Communications Magazine*, vol. 51 (7), July 2013.

[30] R. Hancock, G. Karagiannis, J. Loughney and S. V. d. Bosch, *Next Steps in Signaling (NSIS): Framework,* IETF RFC 4080, June 2005.

[31] D. Clark, "Application design and the end-to-end arguments," MIT Communications Futures Program, Bi-annual meeting, Philadelphia, PA, USA, 30-31 May 2007.

[32] L. Daigle and IAB, *IAB Considerations for UNilateral Self-Address Fixing (UNSAF) Across Network Address Translation,* IETF RFC 3424, Nov. 2002.

[33] J. Rosenberg, R. Mahy, P. Matthews and D. Wing, *Session Traversal Utilities for NAT (STUN),* IETF RFC 5389, Oct. 2008.

[34] R. Mahy, P. Matthews and J. Rosenberg, *Traversal Using Relays around NAT (TURN): Relay Extensions to Session Traversal Utilities for NAT (STUN),* IETF RFC 5766, Apr. 2010.

[35] J. Rosenberg, *Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols,* IETF RFC 5245, Apr. 2010.

[36] P. Leppäaho, N. Beijar, R. Kantola and J. Llorente, "Traversal of the Customer Edge with NAT-Unfriendly Protocols," in *Proc. IEEE International Conference on Communications (ICC 2013),* pp. 1526-1531, Budapest, Hungary, 9-13 June 2013.

[37] M. Mealling and R. Daniel, *The Naming Authority Pointer (NAPTR) DNS Resource Record,* IETF RFC 2915, Sept. 2000.

[38] Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. d. Groot and E. Lear, *Address Allocation for Private Internets,* IETF RFC 1918, Feb. 1996.

[39] R. Hinden and B. Haberman, *Unique Local IPv6 Unicast Addresses,* IETF RFC 4193, Oct. 2005.

[40] J. Rosenberg, H. Schulzrinne, G. Camarillo, A. Johnston, J. Peterson, R. Sparks, M. Handley and E. Schooler, *SIP: Session Initiation Protocol,* IETF RFC 3261, June 2002.

[41] J. Damas, M. Graff and P. Vixie, *Extension Mechanisms for DNS (EDNS(0)),* IETF RFC 6891, Apr. 2013.

[42] W. Eddy, *TCP SYN Flooding Attacks and Common Mitigations,* IETF RFC 4987, Aug. 2007.

[43] W. Simpson, *TCP Cookie Transactions (TCPCT),* IETF RFC 6013, Jan. 2011.

[44] M. Kobayashi and T. Murase, "Asymmetric TCP splicing for content-based switches," in *Communications, 2002. ICC 2002. IEEE International Conference on,* pp. 1321-1326 vol.2, 2002.

[45] J. Costa-Requena, R. Kantola, J. Llorente, V. Ferrer, J. Manner, A. Yi-Ding, Y. Liu and S. Tarkoma, "Software Defined 5G Mobile Backhaul," in *1st International Conference on 5G for Ubiquitous Connectivity,* Levi, Finland, 26-27 Nov 2014.

[46] B. Niven-Jenkins, D. Brungard, M. Betts, N. Sprecher and S. Ueno, *Requirements of an MPLS Transport Profile,* IETF RFC 5654, Sept. 2009.

[47] E. Rosen and Y. Rekhter, *BGP/MPLS IP Virtual Private Networks (VPNs),* IETF RFC 4364, Feb. 2006.