Future Internet Is by Ethernet

Marko Luoma, Raimo Kantola, and Jukka Manner

Aalto University, Department of Communications and Networking,
Otakaari 5A, 00076 Aalto, Finland
`{firstname.lastname}@tkk.fi`
<http://www.comnet.tkk.fi>

Abstract. This is a position paper describing an approach to the creation of the Future Internet. We argue that the new architecture must respond to two key challenges: (1) increase trust among Internet stakeholders and (2) provide cost efficient scaling of the network to new levels of capacity, number of users and applications. We argue that the solution is to redesign the Internet by gradually replacing IP with a carrier grade transport system. In practice such a packet transport system can be created based on Carrier Grade Ethernet. We call the resulting network Internet by Ethernet. We make some fundamental arguments and outline the research agenda that will open based on the premises that we describe.

Keywords: Future Internet, Packet Switching, Trust-to-Trust, Carrier Grade Transport, Ethernet.

1 Introduction

We argue that (1) the Internet should be redesigned not just from the technological perspective but also from the perspective of business models and contracts that are done based on the business model. It seems pointless to keep scaling network speeds up if the network is filled with traffic that the users do not want to receive and the capacity that is available is hogged by a few percent of the users. We argue that the way forward is to build a federated trust model into business models in order to change the operational principle from "assist the sender" to a more trustworthy "make the selfish and malicious actors pay".

To achieve this goal, from the technological perspective we argue that (2) Internet needs to be redesigned from the bottom-up in terms of the protocol stack. We have seen where the continuous adaptation to the application requirements has led the Internet infrastructure. Our starting point is the concept of Carrier Grade transport which we believe to be packet based and more specifically Ethernet like. Carrier grade transport provides a predictable, reliable and trustworthy service to user's packets. Based on the Carrier Grade transport system, the operators can enhance the level of trust in their business relationship. The goal is that the operators form a federated trust domain where the cost of transport to an operator will depend on the rating of the operator in the community. This federated trust domain is surrounded by a trust boundary towards the customers and users.

The consequence of packet based carrier grade transport is profound. With the growing role of the transport system, the role of IP as the network integrator will be gradually decreased. This continues the trend started by MPLS that has become a general purpose network connectivity services platform.

We argue that (3) addressing in the Future Internet is recursive and addresses and customer (device) identifiers are separated. In recursive addressing, a chain of addresses points to a unique end-point. Recursive addressing lets users with a private address in one address space to communicate with users that are in a different private address space. This makes it possible to keep reusing IPv4 addresses. Thus the huge address space defined for IPv6 will not be needed. Address space boundaries co-inside with trust boundaries.

(4) The Future Internet by Ethernet has connection state on address space boundaries. Connection state stores everything that is needed to manipulate address and identity information in the packets. Packets are switched over the address space/trust boundaries and the switching state is managed by implicit signaling that is embedded in the normal traffic patterns that applications use [2], [5].

Inside trust domains forwarding entries can be created by the background routing process like in an IP network. Domains that provide traffic engineered protected connections may use centralized routing while domains providing best effort service may use distributed routing.

Routing is recursive and public network domains are completely isolated from consumer and corporate networks. This means that (5) the public services core does not reveal its addresses or methods of forwarding to its clients. Equally, consumer or corporate networks do not reveal their addressing information to other client networks nor to the core network [2].

Variants of Ethernet such as 802.1ah carry virtual LAN identifiers and service instance identifiers in packets. (6) This can directly be used to virtualize the public network.

(7) A Carrier Grade transport network has full OAM. It follows that connections across links and public network domains are supervised by OAM packet flows. This makes it possible to recover from link or node failures quickly and trace failures to services.

(8) Mobility should be a service on top of the transport network, i.e. mobility should be implemented on Ethernet level. We should also focus first on efficient local mobility, and add the design of global mobility on top.

Based on the arguments, a few initial comments should be made. Connection state and OAM make it possible to protect the users from malicious traffic and the public network from greedy traffic flows. Based on OAM, the traffic entering a public network can be smoothed to a fair share for the flow within network capacity. It will be possible to make use of the information gathered by the OAM system in the host's protocol stack as well as in the boundary elements leading to better end to end transport protocol algorithms and a fairer service. Exactly how far we can take this is a subject for research.

The concept of connection state on trust boundaries leverages the experience of Network Address Translators and stateful Firewalls. Instead of seeing such devices as a nuisance Internet by Ethernet sees them as key elements making Trust-to-Trust advocated by Dave Clark a networking principle rather than an afterthought. Instead of single sided switching usually implemented by a NAT, the new boundary nodes will use implicit signaling to create switching state both for the outbound and the inbound flows. For the customers to brave the creation of switching state for an incoming flow, new protocols leveraging the connection state will be needed. The goal is creating reasonable means for the target customer network to ensure that the incoming flow is legitimate. More details are in [2], [10].

The concepts of carrier grade and connection state will bring more control into the network. The result is that end users will have more freedom of creating end user services because the network is more trustworthy. The end user services must accept the fact that it is no business of the users to see how the public network is implementing its packet transport service edge to edge except in terms of the SLA parameters if such have been put into the contract.

The rest of the paper is organized as follows. Section 2 explains the principles we suggested above. Section 3 discusses migration. Section 4 comments on some obvious objections to the suggested approach. Section 5 outlines research directions that emerge based on the principles. In Section 6 gives a few conclusions and a summary.

2 Principles Explained

In this section we justify and explain the principles (1)(8) suggested. In the last subsection we discuss the scalability of what we propose.

2.1 Federated Trust-to-Trust

The problem of the current Internet is the fact that the technology supports only stakeholder relationships of a low level of trust. First the relationship of customers and network providers need to be addressed. Next the relationships between providers of the different legs on the end to end path need to be supported. We argue that we should see the end to end path as a chain of three trust domains. First is the originating user's trust domain. Next is the federated multi-operator trust domain that provides a public packet delivery service. The third leg is the trust domain of the target user. We show this concept in Fig. 1.

In an ideal world a customer knows that each network provider is tightly bound to provide reliable and secure (trusted) service. We may argue that current Internet aims to do this by the business model where transit and peering agreements are made with interconnecting provider domains. In an ideal world these contracts would form a chain of trust relationships between customers. However, this is not the reality as there are minimal incentives to secure the

service - as the service which is agreed in contract is merely a reachability service and there is no control loop for assessing the performance of individual providers in the chain. Current interconnection infrastructure lacks the functionalities allowing the building of more sophisticated services with inbuilt trust.

What shall be changed from the current infrastructure and business model? For one, there needs to be a mechanism to rate providers based on their operational processes and their reputation in the market. As an incentive this rating should be reflected onto the interconnection charges. This operation resembles the way companies are rated by rating institutions for their liquidity. This rating affects the terms and margins which companies pay for their loans.

Fig. 1. Communication over Trust Domains

With a proper certification from a rating system, a provider can convince interconnecting providers on its ability to control the influx of traffic so that there is a minimal amount of malicious flows. This certification should be dynamic and based on a continuous rating process. Reputation of the provider should be based on the trouble ticketing of other providers and their customers. These tickets should be escalated onto a reputation system. The tickets are accumulated over a period of time by the reputation system. When the result is over a certain threshold relative to ratings over all providers, the rating of the provider should decrease and automatically increase the interconnection charges of that provider. Ideally, this should be reflected on the access price of individual customers. This will change the current network principle "help the sender" to a more trustworthy principle "help the receiver" by putting variable price for the sending of a packet.

2.2 Bottom-Up

A lot of work is currently being carried out with the idea that IP is the inviolate holy network integrator. Thus, a lot of work concentrates on functions above IP and in practice above TCP/IP. The official truth is that IPv6 will be the successor to IPv4. At the same time, the community is painfully aware of the problems inherent to IP networks. We argue that IPv6 does not meet current networking requirements. Stakeholders have legitimate needs of protecting their networks and their business that is carried over and with the network. A huge address space introduced by IPv6 does not address this need [1].

The users wish to control who can communicate with whom. Let us take two examples. (a) Mobile users are even more adverse to DDOS, SPAM and other unwanted traffic than fixed users. Depending on the pricing model, mobile users may have to pay for such traffic. The mobile operators do not want to see significant amounts of unwanted traffic being sent over the expensive air interface. The idea of a host based Firewall does not fit together with battery powered devices. If unwanted traffic is received for Firewall processing, it depletes the battery. One should note that the batteries in mobile phones last days because most of the time the mobile sleeps. (b) The idea of the Real World Internet tries to put sensors and actuators e.g. in people's homes on the Internet. We do not want other people to meddle with our homes unless we have contracted them to do so. Consequently, a huge address space giving every possible device a globally reachable address as in IPv6 is not a blessing, it is a problem.

Transport for the Internet has changed several times. First IP was carried over TDM lines borrowed from PSTN. The service scaled from modem speeds to 2Mbit/s. Frame relay allowed scaling the service to a few tens of Mbit/s. ATM increased the link speeds to several hundreds of Mbit/s. The current transport method - packets of synchronous digital hierarchy - scales up to 40Gbit/s. Scaling up from this level is done by adopting wavelength division with either 10Gbit/s or 40Gbit/s lambda rates. So, the transport method has changed each time when the maximum link speeds have increased by a factor of 15 to 65. Now the question is: what transport will be used for the next upgrade? We believe that the next electronic transport system will be packet based. In practice the protocol is called Ethernet. We are participating in the work towards 100Gigabit/s Ethernet and see that there is little interest of scaling synchronous transport further than what is currently available.

The reasons for this are the flexibility in division of resources, cost and power saving. Ethernet packets carry addresses and VLAN tags as identifiers for the packet owner. Neither of these is present in SDH frames. The presence of addresses in the basic frame has fundamental implications. One is that transport is asynchronous and the speed can be flexibly adjusted according to need. This is a good starting point for saving energy and network capacity. There is ongoing work towards energy efficient Ethernet in IEEE. Finally, we point to the price difference of POS and Ethernet interface boards in high speed core routers.

In the process of becoming the transport medium for public networks, Ethernet must become carrier grade. Backbone Ethernet connections need to become similar to SDH circuits, recovery from failures should be as fast as or better than in SDH and the service should be accurately managed by each operator. Broadcasting of unknown packets, MAC learning and spanning tree protocols need to be replaced by new solutions in Carrier networks. We have created an example solution in this direction in the ETNA project [3],[9], [10]. ETNA control and data planes support both Carrier Grade intra and inter-carrier native Ethernet networks. Another approach that seems more suitable for corporate networks based on distributed operation and DHTs is being designed and implemented in the Celtic 100GET project [4].

A packet based transport network delivers addressed packets from edge to edge. Most customer and corporate networks are based on Ethernet. So, what we have is a chain of Ethernets end to end. The result is that the role of IP routing is diminished step by step. The remaining role of IP is acting as an identity protocol and for a time, a domain wide routed protocol. In the face of such dramatic changes in the underlying protocol layers, it seems that we must look at the Future Internet starting from the bottom of the protocol stack, not from the applications.

2.3 Addressing, Identification and Routing

Unallocated IPv4 addresses will be exhausted in 2012. IPv6 has long been advocated as the remedy. However, this solution makes existing users pay for the new ones yet to be connected. The dual semantics of IP addresses has led to many difficulties in networking. IPv6 does not change this.

We propose an alternative way forward: let us keep reusing IPv4 addresses as many times it takes and let us re-invent the transport infrastructure as needed. Recursive addressing is possible with protocol stacks such as IP over IP, MAC in MAC, IP over MAC-in-MAC, etc.

We must decide how to identify users, their devices and services at the same time as we choose the addressing system. We argue that we have three options for the identities: (1) identities are random values or (2) we keep using IPv4 addresses as local user identifiers or (3) we reuse identities assured by the mobile network infrastructure in a new way. Introducing a new deterministic global identification scheme for users, hosts and services is costly. Mobile access to Internet is gaining momentum: the number of broadband mobile subscriptions has overtaken the number of fixed broadband subscriptions and grows faster than fixed access. So, leveraging the mobile infrastructure for brokering trust between users is starting to make sense.

In [2], [10] we suggest using random identifiers for users, hosts and services. A quick look at the birthday paradox shows that IP over IP does not scale for random IDs [2]. The ongoing work on Locator/ID Separation Protocol or LISP uses deterministic 32 bit values as End-point Identifiers [6]. Unfortunately, LISP requires that EIDs are globally unique leading to limited scalability. Let us note that it is sufficient to be able to create edge nodes that serve up-to a few million users communicating with several million targets. Only places where the IDs are processed are these edge nodes. Therefore, the random id does not need to be globally unique. It suffices that IDs are unique with high probability in the boundary nodes that will process the IDs. In [2], based on the Birthday paradox, we show that random identities in the order of 60 to 80 bits will be sufficient.

This argument leads to the protocol stack: IP over MAC-in-MAC. In parallel to IP other identity protocols can be used. For stepwise deployment, other stacks are easily accommodated as well.

This addressing structure leads to end to end forwarding over routed or bridged originating customer network (leg 1), followed by (leg 2) a public

network that may be a routed Ethernet network like the one we created in ETNA and the terminating customer network (leg 3) that can also be routed or bridged. Leg 2 can be also for example an IP/MPLS network. This is shown in Fig. 1.

2.4 Connection State on Boundaries

The most important trust boundary is between the user (consumer or corporate) network and the public network. On this boundary we move from a private addressing scheme to a core addressing scheme. In our solution outbound packets are encapsulated to backbone frames and inbound packets are de-capsulated. This is similar to what LISP does without specifying any trust architecture.

We can break the boundary to customer edge (CE) and provider edge (PE) functions as shown in Fig. 1. In outbound packets, the customer edge replaces the source address with source id. For a given name of a target host or user or service, the ingress provider edge returns a destination id to a requester in a customer network. An outbound packet leaves the customer edge with source id and destination id information replacing address fields in the payload headers. Using connection state in the PE, the destination id will be appended by the destination address of the egress PE node. Upon reception of a packet to a destination id, the egress PE node finds the locally significant address of the destination and sends the packet to the customer edge. For an incoming packet, the CE will apply any reasonable means, we call packet access control, to find out whether the packet is legitimate or not. We have discussed such methods in [2]. The CE creates its own connection state and executes the policy created by the local network administrator.

The connection state on the boundaries can be leveraged to many uses. For example, multi-homing can be hidden from the core routing system, traffic engineering and packet access control can be stateful and thus more intelligent than simple filtering using rules.

CE processing outlined above renders NAT traversal based on UNilateral Self Address Fixing [7], [8] unnecessary. Any remote target is seen by a host as if residing in the local network. The messy application code related to NAT traversal can be dropped from applications. The downside is that new service description methods will need to be introduced. Instead of using IP addresses and port numbers (like in SDP), service description should use e.g. names and ports. For ease of deployment, the changes in hosts can be postponed by intelligent customer edge processing of packets as we show in [2].

2.5 Network Isolation

Name queries are routed through the CE and PE nodes. Part of the information returned by a name query is stored in connection state in PE and CE nodes. The principle is: *A customer network never publishes any address information. A public network does not show its addressing to customer networks.* As a result, the network structures are hidden.

Because each leg of the end to end connection uses its own forwarding /routing / switching system, network events such as routing configuration errors impacting one leg will not impact the other legs.

2.6 Network Virtualization

ISPs make their money from corporate connectivity while earning money from the consumer services is challenging. Therefore, virtualizing the infrastructure cost efficiently is important to ISPs. Nowadays, MPLS/VPNs and virtual routing are used for the purpose. Often customer Ethernet circuits are run over Pseudo-wires created using IP/MPLS. Depending on the level of traffic isolation, this can lead to a complex and thus a less than cost efficient network structure. Due to IEEE Ethernet lack of support for global networks, MPLS shim is necessary. However, our work in ETNA shows that it is rather straightforward to create a native Ethernet-based control plane that provides carrier grade control and lets the operators cost efficiently create virtual networks that span even several carriers. In another approach we are pursuing in the 100GET project, we can create virtual networks but also turn the Ethernet broadcast traffic into unicast based on Routing Bridges (rbridges) and DHTs, thus enabling the very basic Ethernet technology to scale.

2.7 OAM

In the carrier grade transport network we can run link level OAM, a less frequent OAM flow edge to edge over a domain and even less frequent OAM flow end-to-end. OAM packet frequencies can vary from milliseconds to seconds. Limiting factors are the number of parallel OAM sessions and functionalities that are bound to the sessions. It is rather straightforward to create simple fault notification on link level but to combine more intelligent aspects like resource control and quality reporting in each intermediate point to sessions is not simple.

This problem was touched by the work done with ATM available bit rate (ABR) traffic management. The optimization goal was to maximize network utilization with fair resource division by using OAM messages to convey indications of a fair amount capacity. This was, however, found to be complex to implement and impossible to manage in large networks. Optimizing the use of OAM flows and leveraging them for different needs is an ongoing research challenge.

2.8 Mobility

Most of the next 2 Billion new Internet users will be using a wireless connection. Most of them will do so through a mobile network. According to ITU, the number of mobile broadband subscriptions has overtaken the number of fixed broadband subscriptions. Any networking protocol and architecture that does not support mobile use, will be less than satisfactory and will lead to costly add-on solutions. We should design the Future Internet for mobile use in mind.

It may be argued that mobility is not the function of packet transport, but in large the result of mobility is the change of serving edge node and thus change in the transport connection. It is also true that in order to maintain maximal flexibility the base protocol should be as simple as possible and should therefore not bear the burden of individual functionalities - like mobility management. This leads to a layered structure where the protocol is not the dividing method rather the functionality that is pursued. This means that the base protocol can have extensions for performing value added tasks - like mobility, identity resolution or even charging.

In this respect, the overwhelming use for mobility is in the local access network. Thus, we need to make sure that the base packet transport supports natively the movement of the end-points, and even sub-networks - a concept that is available in IP networks as a complex add-on that barely works. Since Ethernet addresses are not tied to topology, mobility is easier than in IP. In our forthcoming design, mobility is implemented by the network without explicit support from end hosts. We are building such functionality over the rbridges technology. Support for global mobility can be added later, as the add-on when needed.

2.9 Scalability of the Approach

It is a legitimate concern whether the suggested approach with its heavy reliance on connection state on the network edge will scale. For cost efficiency reasons, we must be able to build large CE and PE nodes that can serve up-to several millions of users with their numerous applications. This is a topic for experimentation and research and a good challenge for the vendors.

An alternative is the model that operators give CE devices an address that is globally unique and the CE has a DHCP server that assigns dynamic IP addresses to users, hosts and services. These are used to identify users locally under one CES but never for global routing. The allocation of these identifiers can be e.g. per DNS request. The identifiers are stored in DDNS that can be part of the CE. The protocol stack stays the same as before, the difference is in the source and destination addresses carried in the inner Ethernet header. CE to CE forwarding is based on two NSAP addresses carried in the inner and the backbone Ethernet layers. One NSAP address points to a CES and the backbone has its own NSAP addressing for PE devices. Both NSAP addresses are carried in the DA fields of the respective Ethernet frame. The benefit is that PE processing per packet will be simplified.

In both the alternative model and the suggested one, the CE scalability has similar challenges as NATs do. In the alternative model, an identity for a user or host is allocated by the visited network. In the model using random identifiers, the id is allocated and managed by the home network. This leads to differences in the mobility architecture.

3 Migration

How to move from the present Internet to Internet by Ethernet? The suggested architecture has three tiers on which deployment and migration can progress rather independently of each other. Each of the tiers pursue a common goal of enhancing trust and providing cost efficient scaling of the Internet to higher capacity, cost efficient connection of wireless users and meeting the needs of new applications.

On the transport tier, the first step is that one operator decides to use an Ethernet core network once the technology matures. From the point of view of the carried traffic, Ethernet circuits, E-LANs and E-trees are as good as their IP/MPLS counter parts. These legacy systems can continue to be used in parallel.

On the access tier, the first step is a directory service that provides the information for the Ethernet CE and PE nodes. In another paper, we show how to use DNS for the purpose. Step 2 is that one or several Ethernet core operators start providing customer access at Ethernet layer to CE devices. In step 3, a customer network connected to such an operator can deploy a native Ethernet Customer Edge device and connect it to a PE device owned by an operator. The customer may optionally discontinue its IP connection to the Internet at the same time and access the legacy Internet through NATs hosted by an operator.

Deployment of solutions on the transport and access tiers requires no or minimal changes to hosts. Double-NAT traversal solutions may be reused. It would be clean that the applications that use IP addresses on the application level were revised to use names and other more suitable means of identification. For applications like FTP or SIP that use IP addresses as identifiers and transmit them to remote parties on a control channel, an application specific state machine in the CE function can be used eliminating the need for immediate changes in host software.

The highest in the suggested architecture is the tier of federated trust among the stakeholders. It fits nicely into the overall picture but can progress rather independently of the two lower tiers once cost efficient IP trace back becomes feasible. A way of supporting IP trace back is shown in [10].

Alternative and further strategies are also possible. For example, in the future we could design an end-host stack without IP at all. We could pursue removing IP headers from the packets traversing the core. All communication is done over Ethernet, i.e., transport protocols and their payloads are run directly over Ethernet. When communication with a legacy IP network is needed, we could use remote APIs on CE or PE devices to encapsulate the data into IP packets.

4 Objections and Counter Arguments

Networks with MAC addresses do not scale. Instead of a single MAC address, our approach uses the minimum of two independent Ethernet domains (customer and backbone) each with their own addressing. Scaling to large backbone networks

is likely to be based on other than vendor assigned addresses. Nothing stops carrying NSAP like hierarchical addresses in the backbone frames. The address allocation can be managed by the operators. Also, by removing MAC-learning and spanning trees, the basic Ethernet will be scaled up.

Why would Ethernet be any better than IP? The question of forwarding protocol is secondary to control plane structure and allocation of functions in the network. IP has so much legacy that it is better to go for a clean break than try to keep patching IP. The recursive variant of IP over IP does not scale to mobile access [2]. LISP that also uses IP over UDP/IP has scalability limitations. The lack of support for carrier grade concepts in IP is fundamental. We can instead re-invent Ethernet, make it carrier grade and build a cost efficient technology that can be economically upgraded to the future link speeds. We can not communicate without layer 2. Ethernet is becoming ubiquitous. It has addresses in each frame. It is possible to leverage those for global communication. So, why not simplify the stack and get rid of some of the cost?

We have all this functionality in IP - why redesign it all for Ethernet? This is untrue. IP is fundamentally not carrier grade. Further development of IP networking technology has become very challenging because of the erosion of its fundamental principles of end to end and IP over everything. IPv4 addresses will be exhausted soon. Move to IPv6 is not well motivated. IPv6 does not help to hide networks from each other, it does not help to create trustworthy services and it does not help to virtualize the infrastructure. An accepted solution to multi-homing that is essential for mission critical connectivity of businesses to the core Internet is still missing. People want NATs even in IPv6 networks and we may end up with the same or worse routing scalability problems in IPv6 than what we have in IPv4.

It is true that we end up redesigning many functions that are present in current IP networks. Let's not worry about this. We have been doing this always in the history of networking for each new generation of technology. This is not the first and not the last time. Instead of complaining about it let's concentrate on improving the designs.

Why would be need this federated trust system - it can become a target for attack and fail. At the moment we fight unwanted traffic largely only by defense. We use NATs and Firewalls. However, the senders of unwanted traffic are professionals and the sending has become a business. We argue that we "the nice users" are at war and we are being buried in the infoglut. Let us recall that no war has been won by defense only. The Federated trust layer is a proactive response with the goal of making sending unwanted traffic unprofitable.

5 Research Directions

The outlined architecture is a subject for technical research, prototyping and experimentation. Results of the efforts in ETNA [9] and 100GET will be made and partially have been made available for others to build upon.

Fundamental research questions are: how far can we take the suggested trust architecture in term of technology and how can we introduce it in terms of

politics and tussle among the stakeholders. The challenge is that probably a new global association of Internet Exchange is needed to govern the federated trust schema. Operation can be given to a private company on a competitive basis every few years.

Instead of starting from global politics, we may also start from better mechanisms for the customer/operator interface. There the challenge lies in finding a deployment path on which each step has an incentive for the critical stakeholders. Our proposal of Customer Edge Switching is easier to deploy than LISP because no "alternative topology" is needed. For CES that connects to an IP core, each access network and each corporate network can make an independent deployment decision [10].

Exact definition of the addressing used in Global Ethernet-like networks needs to be studied and agreed upon. We have made a proposal in ETNA.

Technical methods for protecting customer networks from unwanted incoming traffic are an important area of research although numerous methods known from other contexts can be reused here. New trust related protocols will be needed as well.

A fundamental issue is whether we trust the trust model or we concentrate on protecting the user by all means feasible while the user is engaged in communication with another user who she fundamentally does not trust irrespective of any network supported trust model. Probably, we must assume that no trust model can be absolutely relied upon.

Scalability of nodes that will store connection state will be an important topic.

New means for describing identity in applications need to be studied and defined. These include for example new guidelines of use or a new version of Service Description Protocol. Interactions of the suggested addressing and identification schemas with applications and their impact on the mobility architecture need to be studied.

Mobility management as such in the context of the Ethernet protocol family is an important topic.

Finally, a carrier grade packet transport system that is energy efficient is needed. Probably making full use of the energy efficient Ethernet that introduces variable link speeds for the sake of energy efficiency will require traffic engineering on Ethernet layer.

6 Summary

We have implemented a rather comprehensive prototype of an Ethernet based carrier grade transport architecture that this paper relies upon [1], [3], [9]. We have also implemented a corporate network Ethernet routing solution that seeks to eliminate broadcasting due to ARP and unknown target addresses from bridged networks [4].

In both contexts, we have made our first attempts to integrate mobility management into the Ethernet protocol family. More remains to be done.

This paper seeks to generalize the ideas that lay behind our prototypes.

We propose to make the trust architecture the starting point of designing the Future Internet. We propose a new path of technology development that will seek to satisfy the increasing needs in trusted and scalable communications. Based on the architecture design, we discuss migration challenges and approaches and outline a number of research topics for the networking community.

Acknowledgments. Our thanks go to our partners in ETNA and 100GET projects.

References

1. Kantola, R., Luoma, M., Lamminen, O.-P.: Transport for Carrier Grade Internet. In: 1st IEEE Below IP Networking WS at GlobeCom, Honolulu (2009)
2. Kantola, R.: Implementing Trust-to-Trust with Customer Edge Switching. In: AMCA in Connection with AINA, Perth (April 2010)
3. Lamminen, O.-P., Luoma, M., Nousiainen, J., Taira, T.: Control Plane for Carrier Grade Ethernet Network. In: 1st IEEE Below IP Networking WS at GlobeCom, Honolulu (2009)
4. Varis, N., Manner, J.: Minimizing ARP Broadcasting in Trill. in: 1st IEEE Below IP Networking WS at GlobeCom, Honolulu (2009)
5. Virtanen, L.: Communicating Globally Using Private IP Addresses, M.Sc thesis, Comnet/TKK, Espoo, Finland (2009)
6. Farinacci, D., Fuller, V., et al.: Locator/ID Separation Protocol (LISP) (January 2010) draft-ietf-lisp-06.txt
7. Behaviour Engineering for Hindrance Avoidance (Behave), <http://www.ietf.org/dyn/wg/charter/behave-charter.html> (referred January 31, 2010)
8. Daigle, L. (ed.): Informational RFC 3424, IAB Considerations for UNilateral Self-Address Fixing Across Network Address Translation (2002)
9. ETNA web site, <http://www.ict-etna.eu/index.html>
10. Routed End-to-End Ethernet, Protocols, <http://www.re2ee.org>