Investigating the Deployment and Adoption of re-ECN

Alexandros Kostopoulos
Athens University of Economics and Business
Department of Informatics
76 Patission, Athens 10434, Greece
alexkosto@aueb.gr

Ken Richardson
Roke Manor Research Limited
Old Salisbury Lane, Romsey,
Hampshire, S051 0ZN, UK
+44 (0) 1794 833491
ken.richardson@roke.co.uk

Michalis Kanakakis
Athens University of Economics and Business
Department of Informatics
76 Patission, Athens 10434, Greece
kanakakis@aueb.gr

ABSTRACT
Networking research aims to design protocols for future Internet architectures that are able to hold end hosts accountable for the congestion they cause. Re-ECN is a protocol that provides valuable information to ISPs about network congestion, and which could be used as a useful input to support the allocation of network resources more equitably. In this paper, we propose and apply an adoption framework for the re-ECN protocol. As well as the technical design aspects, we also focus on the high level challenges and opportunities for the key stakeholders and present potential deployment scenarios that might lead to the widespread adoption of re-ECN.

Categories and Subject Descriptors
C.2.1 [**Computer Communication Networks**]: Network Architecture and Design, J.4 [**Social And Behavioral Sciences**]: Economics

General Terms
Design, Economics

Keywords
Future Internet; Congestion Exposure; re-ECN; protocol adoption

1. INTRODUCTION
The Internet is rapidly evolving. As a result, current business models must evolve too in order to capture new stakeholder needs and expectations. Although the Internet was originally designed for end hosts, rather than networks, to be informed about congestion, an important goal for future Internet architectures is for networks to be able to hold end hosts accountable for the congestion they cause. In response, re-ECN has been proposed as a mechanism that builds upon ECN to provide additional information about the level of network congestion to Internet Service Providers (ISPs) [1]. This information may then be used by them to introduce congestion-based traffic contracts for appropriate end users. Such schemes provide incentives for those end users to implement rate adaptation mechanisms to regulate traffic flows in order to remain within their congestion 'allowance'. Thus, re-ECN could be an important component in future Internet architectures for congestion exposure [2].

However, there are key challenges facing the successful deployment of re-ECN, primarily concerning the misalignment of costs and benefits across the key stakeholders. This challenge is discussed in this paper. The key problem is that the full benefit of re-ECN can only be realized through universal adoption but the incentives for unilateral adoption by individual stakeholders, particularly ISPs, are weak. The major challenge for supporters of re-ECN is to identify a deployment roadmap that improves the alignment of the costs and benefits to each stakeholder so that, step by step, the goal of widespread deployment can be achieved.

After setting the scene, this paper introduces an adoption framework for re-ECN and uses this to explore the issues and incentives concerning the deployment and adoption of re-ECN by the key stakeholders. It contributes to the ongoing debate by identifying these stakeholder incentives and proposing a potential adoption scenario that may lead to the widespread deployment and diffusion of re-ECN technology.

ACM ReArch 2010, November 30, 2010, Philadelphia, USA.

2. BACKGROUND
2.1 Explicit Congestion Notification (ECN)
Traditionally, TCP/IP networks signal congestion by dropping packets. The ECN protocol [3] allows routers to mark packets that would otherwise have been dropped as having experienced congestion. It does this by using two bits in the *DiffServ* field in the IP Header. This information is then used by end hosts to reduce traffic flows without the inconvenience of suffering actual packet losses. The proportion of marked packets can also be used by the network as a measure of upstream congestion, i.e. the level of congestion already experienced by a traffic flow up to that point along the path.

When ECN is successfully negotiated between sender and receiver, an ECN-aware router may signal congestion by marking, rather than dropping, a packet. The receiver echoes the congestion indication back to the sender before appreciable queue growth has occurred. Thus, the ECN mechanism gives hosts the opportunity to automatically reduce their transfer rate to prevent packet losses and hence reduce the number of retransmissions.

2.2 Re-ECN
Re-ECN, which is currently being standardized at the IETF, is designed to provide the network with information about the

expected level of congestion along the entire path. This is done by using one of the re-ECN codepoints to mark a packet as 'red' when it experiences congestion (like ECN). The destination then reports the number of 'red' packets that it has received back to the source. In addition to responding to this congestion indication, the sender sets a packet to 'black' for each 'red' packet that has been reported by the receiver. Otherwise, the packet remains 'grey'. This gives re-ECN its name as the level of congestion as determined by ECN is 're-echoed' back to the receiver. Packets may also be marked 'green' and this is done by the sender at the start of a flow before the feedback loop has been established with the receiver [4].

In order to ensure that end users accurately declare the amount of congestion they are causing the re-ECN scheme requires a *dropper* at network egress, if users are not trusted [4]. The purpose of the *dropper* is to ensure that the sender honestly declares its expected whole-path congestion. If the sender under states the congestion it is causing, then the proportion of 'black' packets that it sees will be less than the proportion of 'red' packets. In this case, the dropper will penalize the sender by dropping packets.

A *policer* at network ingress could also be used to limit those trying to cause more congestion than they are allowed under their contract. It is not always required, since there are several possible ways of using re-ECN information to provide incentives (i.e. congestion charging, DiffServ, etc). However, a policer makes possible flat rate, congestion allowance contracts between end users and ISPs. A modified token bucket, i.e. with a maximum congestion allowance, can be used as a bulk policer to decrease a user's rate.

Within this re-ECN framework, information is revealed so that end users and networks can be held accountable for the congestion they cause. Moreover, it creates the motivation for end users to perform appropriate congestion control for non-time-critical bulk data transfers [5].

3. ADOPTION FRAMEWORK

In this paper, we apply the adoption framework proposed in [6] to the re-ECN protocol. This can be used as a roadmap in order to consider the different deployment cases of a new protocol and investigate the most critical factors that could boost its adoption.

The adoption framework consists of three steps (Figure 1). The first step is *protocol design*. More specifically, the protocol must provide a perceived benefit, be capable of incremental deployment, and embody good technical design (as reflected in the Internet's design principles). Of course, standardization (i.e. by the IETF) is also a key factor in promoting the subsequent deployment and adoption of the protocol.

The second step is the *deployment process* which aims to achieve widespread deployment of the protocol. This step identifies and investigates the pre-requisites needed in its deployment and surveys the different deployment cases. The third step is that of ensuring that the incentives for its *adoption* by all the key stakeholders are satisfied. This involves consideration of different scenarios along with different types of business model or use case that may lower the barriers of adopting re-ECN. The feedback loops shown in the diagram represent the continuous reconsideration of protocol design and deployment.

The distinction between deployment and adoption is a subtle one. However, it can be explained with reference to the number of protocols that have been deployed through inclusion in protocol stacks, but have not actually been enabled or widely used (i.e. IPv6) [7]. The factors within each of these three steps are now discussed in detail in the following sections.

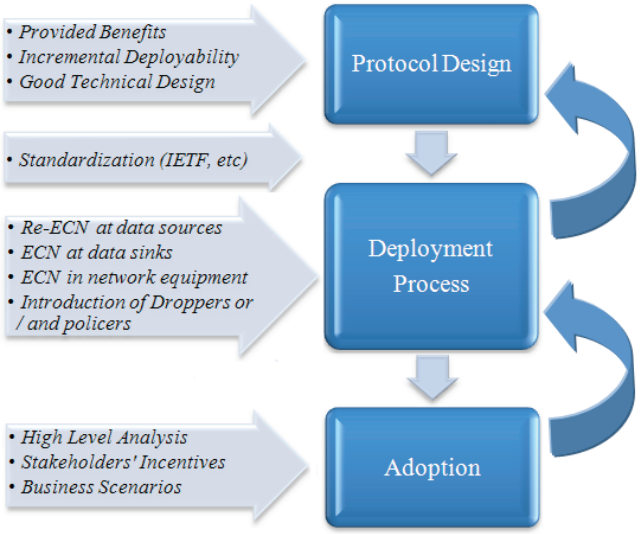

Figure 1. Re-ECN Adoption Framework.

4. PROTOCOL DESIGN

4.1 Provided Benefits

The main technical benefit of re-ECN is that it provides additional information to ISPs about network congestion. Whereas ECN provides information about upstream congestion only (i.e. the congestion a traffic flow has already experienced), re-ECN provides additional information that can be used in conjunction with ECN information to infer the level of downstream congestion (or, to be more precise, the downstream congestion experienced one round trip earlier). Thus, each network node can compute an estimate of the congestion that each flow is likely to experience between it and the final destination.

However, it is important to understand how this technical benefit translates into business benefits for ISPs through the use of alternative charging models offered to end users. The key point is that ECN by itself can only be used to support dynamic congestion pricing, i.e. pricing based on the number of ECN marks. However, this is not a very popular charging model with end users, who prefer flat rate pricing. Re-ECN allows flat-rate congestion-based contracts to be introduced, with enforceable limits (e.g. a pre-defined number of 'black' packets). This provides the incentive for end users to remain within their congestion allowance by reducing their traffic flows in response to congestion. It should be noted here that the '*receiver pays*' model is the only one possible in plain ECN, exposing the receiver to denial of funds attacks, because the congestion rate of a flow is only known at the egress of the network.

What is more, re-ECN provides sufficient information for network operators to implement more effective interconnection agreements at the trust boundaries, as both participants have full information about downstream congestion and can agree on the charges.

4.2 Incremental Deployability

A key aspect of incremental deployability is backwards compatibility with existing protocols. Re-ECN builds on plain ECN, which is already standardized as an addition to the Internet Protocol with two bits in the IP Header (v4 or v6) assigned to the ECN field. Incremental deployability is achieved by simply adopting an extra unused bit in the IP Header to signal the level of congestion expected along the whole path (as reported by the receiver based on the congestion level experienced one round trip earlier). Although the RE flag is a separate, single bit it can be read as an extension to the two-bit ECN field [3]. Re-ECN is designed so that re-ECN packets are distinguishable from non-re-ECN packets and an end host can choose which packets it sends as re-ECN enabled. Therefore, it is up to network operators to provide an incentive for them not to turn off re-ECN, by enforcing a bit-rate limit on non-re-ECN traffic.

Since re-ECN builds upon ECN, current issues related to ECN deployment are also relevant to the future deployment and adoption of re-ECN. These deployment issues fall into two areas.

Firstly, there have been problems concerning the enabling of ECN in standard Operating Systems (OSs) such as Windows and Linux. ECN was disabled by default in Windows Vista and earlier versions, but in new OSs (i.e. the Linux mainline distribution, Windows 7 and Windows Server 2008) it is only disabled at the client end while being enabled at the server end of a connection. Therefore, in the future any client will be able to ensure that ECN is enabled merely by unilaterally enabling it at the client end. To manually enable ECN, end users should use the 'netsh' command on Windows or the 'sysctl' interface on Linux.

However, mass market deployment should not be dependent on end users having to manually configure their OSs. Instead, a management application with administrator rights is the most likely deployment route for enabling ECN on behalf of the users. A patch already exists for Linux (although not in the default distribution), which tests for any offending middleboxes before enabling ECN on the behalf of the user. Such 'ECN black-hole detection' is also being considered for the next release of Windows.

Secondly, a key step in the deployment of ECN (and therefore re-ECN) will be the field testing and fixing or replacement of certain middleboxes. For instance, a non-zero ECN field triggers a bug in some home gateways, which makes them crash. Fortunately, problems such as these eventually fade into insignificance as the number of remaining problematic boxes dwindles. For instance, ECN deployment was originally hampered by firewall blockages, but most of them were either upgraded or reconfigured between 2001 and circa 2006.

4.3 Good Technical Design

Clark *et al* [8] and Ford *et al* [9] have proposed guidelines and principles to help researchers and engineers in designing successful protocols. In this section, we assess re-ECN against these design principles.

4.3.1 Design for Tussle

Clark proposes that the Internet should be designed so that the outcomes are determined at run-time [8]. Re-ECN is designed with this in mind for it ensures that the fine-grained information necessary for traffic control is shared between end hosts and the network, whereas previously it was only available to end hosts. Based on this information, network operators can choose to make users accountable for the congestion they cause through the type of contracts they offer and their congestion allowance policy. Indeed, re-ECN is designed so that end users, network operators, application developers and OS developers can choose whether to adopt it or not. Crucially, network operators can provide users with the incentive to adopt re-ECN if they themselves choose to. Therefore re-ECN is designed such that its deployment and adoption is a tussle in its own right.

4.3.2 Information Exposure

Re-ECN is an example of the application of the Information Exposure Design Principle [9]. The congestion marking of packets by network nodes ensures that information on resource scarcity is made available to end hosts and carried in the actual packet header. As a result, this aggregated information along the data path gives a full picture of the state of congestion of the path, where each node in the network knows about congestion on the path ahead. Also, the end host exposes information that it knows to the network (and the network can make it in the end host's interest to do this).

4.3.3 Fuzzy Ends

With Congestion Exposure [2], the network has sufficient information to perform per-flow congestion control on behalf of the end host. Another case where the fuzzy ends principle is applied is the use of re-ECN proxies. If users are unwilling to upgrade their network stacks to adopt re-ECN protocol (or before they have done so), a re-ECN proxy in their home router could provide congestion transparency, but would have only indirect control over the user. This might be provided by the router manufacturer, their ISP or by a third party. Proxy mechanisms can assist with the deployment of re-ECN, even though such proxies are difficult to deploy.

4.3.4 Separation of Policy from Mechanism

Re-ECN is a protocol that provides a mechanism for congestion information exposure. Such information could be a useful input for ISPs when setting policy. However, even when congestion accountability mechanisms are deployed (i.e. policers), operators can still choose to overlook the information they provide in their traffic management policies. Although this seems to be an extreme situation, there might be cases where such information could be ignored (i.e. the network provider trusts the end-device). Thus, the re-ECN ensures that policy and mechanism can be separated.

5. DEPLOYMENT PROCESS

In this section, we present a sequence of steps needed in the deployment of re-ECN which facilitate its adoption.

There are four basic steps required in the deployment of re-ECN:

i) Re-ECN implementation at data sources, along with a new rate adaptation mechanism;

ii) ECN implementation at data sinks;

iii) Enabling of ECN in network equipment by ISPs;

iv) Introduction of policers and droppers by ISPs.

The first step is the implementation of re-ECN in sender's OSs. Alongside the protocol there must be a rate adaptation mechanism that is able to respond to the excessive congestion marking of packets by reducing the level of traffic offered to the network over time and all flows (in contrast to TCP, which provides instantaneous per-flow rate adaptation).

In the second step, the receiver has to implement at least ECN (and preferably re-ECN) in its TCP/IP stack. Although re-ECN is not essential, the sender's congestion markings will be more precise if the receiver has been upgraded to incorporate re-ECN. Currently, the receiver needs to be at least ECN-enabled for re-ECN to work, but a re-design of re-ECN is in progress so that it can work independently of ECN. This involves defining a new set of re-ECN codepoints that are orthogonal to ECN[1].

Even if the receiver is only ECN compatible, the sender can still infer enough from the congestion information echoed back from the receiver to benefit from the re-ECN protocol reasonably well. This is because in ECN congestion control, the sender only needs to know if at least one of the transmitted packets during a RTT was congestion marked to reduce the sending rate. However, the re-ECN protocol prefers to receive more information, as ideally the sender has to be informed of the exact number of marked packets.

The third step, of key importance, involves the enabling of ECN in routers and its support in other middleboxes in order to improve their ability to police re-ECN. Forwarding elements (the data plane) ideally need the existing ECN standard [3] to be implemented and deployed so that packets are marked 'red' by routers as they experience congestion.

The final step is the deployment of the re-ECN network equipment needed to police the accuracy of the information provided by end users (and other networks) and to penalize those users who provide false information. Although the specifics of this are dependent upon the traffic management strategy to be pursued by ISPs, it is likely that new network equipment will only be needed at trust boundaries. However, it is to be noted that no function in the network needs to alter the re-ECN markings (unless it is a proxy for the sender). Note, also, that once a network operator decides to deploy re-ECN-based policing functions around its network, it will want to deploy ECN on as much network equipment as possible to improve the ability to verify the correctness of the re-ECN markings supplied by end users, and to encourage other network operators to do so as well.

6. ADOPTION

6.1 High Level Analysis

Before considering potential adoption scenarios, we first perform a high level analysis in order to identify who the key stakeholders are in the re-ECN adoption process, and their main motivations and drivers.

6.1.1 End Users and Content Providers

Categorizing end users is no longer a straightforward task since the nature of Internet traffic has evolved to include peer-to-peer flows and the delivery of content from distributed Content Delivery Networks (CDNs) and Data Centers. In the context of re-ECN adoption, end users should be categorized as 'light' or 'heavy' users, not in terms of traffic volume but in terms of the network congestion they cause. Re-ECN exposes congestion information to provide incentives for all end users to adapt their rate in response to the level of congestion experienced. As a result, network resources are allocated in such a way that all end users become 'light' users in congestion terms.

Content providers may also have an incentive to adopt or promote re-ECN to avoid the 'bad publicity' they may receive due to the network congestion they cause. On the other, they may also have disincentives to support re-ECN, because the rate adaptation enforcement may discourage users from downloading content from their servers.

However, the incentives for unilateral adoption of re-ECN by end users are weak but they can be increased by network operators enabling the ECN marking of packets in their routers.

6.1.2 Internet Service Providers

ISPs are particularly interested in allocating network resources between users fairly and, more specifically, avoiding 'free riders'. They aim to maximize the number of users, as well as their revenues, and fair allocation of resources is a means of achieving this without needing to add more capacity or deploy costly and controversial traffic management equipment. Re-ECN adoption results in the more efficient allocation of the available capacity and provides improved QoS to all users.

To incentivize end users to adopt re-ECN and rate adaptation, ISPs will want to offer contracts to their customers based on congestion allowances rather than volume allowances whilst maintaining a flat-rate tariff. This type of contract is compatible with economic principles, as it takes into account the externalities (congestion) created in the network.

The main drawback for the unilateral adoption of re-ECN by ISPs is the high level of investment needed in developing and deploying policers, droppers, border gateways, and end user proxies. This level of investment may only be justifiable following the widespread adoption of re-ECN and rate adaptation by end users.

6.1.3 Application and OS Developers

Specific application developers, such as BitTorrent, have expressed interest in deploying re-ECN along with their micro transport protocols (µTP). However, it will eventually need to be deployed directly in the kernel of OSs as a modification to the code of the main transport protocols (e.g. TCP).

OS vendors' primary motivation is to get users off old versions of their OS, because supporting them results in increased costs. If the network operator will be throttling non-re-ECN traffic, this degrades legacy versions versus the new ones. Consequently, OS

[1] There is an important case where the network can compare re-ECN markings to drops: the case where there is only a single bottleneck on a path (e.g. at a remote access server, such as a *BRAS*, in the downstream direction). The BRAS could compare the re-ECN markings and how much traffic it drops against each flow. If re-ECN markings were understated, it could apply additional policy-based dropping of packets in the traffic flow.

vendors could encourage a continual upgrade process, which will really help their business, without being open to blame for this degradation [11].

However, as already identified in section 6.1.1, the incentives for those in the end user community to unilaterally invest in re-ECN are weak unless ISPs also adopt ECN and/or begin to favour re-ECN traffic.

6.1.4 Infrastructure Vendors

Re-ECN will create new opportunities for infrastructure vendors because new network equipment will be needed to support its deployment, e.g. policers, droppers and border gateways. New markets for the supply of this equipment to ISPs will be created giving infrastructure vendors a strong incentive to invest in development of this equipment if they are confident of sufficient future sales to satisfy their business cases.

6.2 Potential Scenarios

This section considers possible adoption scenarios for re-ECN. Adoption requires a correct alignment of commercial and/or social interests across the key stakeholders identified in the previous section. In particular, the financial costs and benefits of re-ECN adoption must be distributed such that each stakeholder is appropriately inentivized (if not at the outset then at some later point as the adoption scenario plays out). Costs include both investment costs (for example, in new OSs and network equipment) and operating costs; benefits include improved QoS as less severe congestion is experienced by end users, and lower network costs since less network capacity is needed to provide a particular level of QoS.

However, as we have seen there are some key interdependencies between stakeholders creating something of a 'chicken and egg' situation. This is represented in Figure 2 in a simplified form but which nonetheless captures the main points.

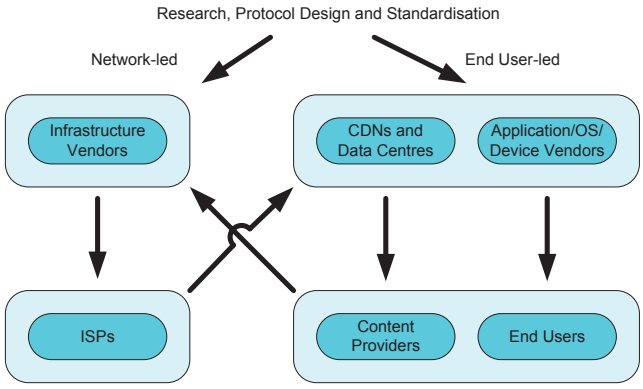

Figure 2. Re-ECN Adoption Routemap.

The diagram illustrates that there are two main routes to achieving re-ECN adoption. The first (shown on the left hand side) represents the network community making the first move and the second (shown on the right hand side) represents the end user community making the first move. Each move ought to cause the other side to respond and thus complete the deployment and adoption of re-ECN. However, we have a classic first mover problem since the incentives for each side to move first are weak.

The main problem concerning the 'network' route is that a large unilateral upfront investment in network equipment is required by ISPs (the policers, droppers and also end user proxies) whilst the initial benefits are limited. To avoid the need for this major investment by ISPs prior to the widespread adoption of re-ECN by end users, we turn our attention to the 'end user' route where a smaller upfront investment by OS vendors and others has the potential to overcome the main blocking point for ISP adoption of re-ECN.

In practice, for the most part, end users are not themselves directly responsible for the deployment and adoption of re-ECN and the associated rate adaptation protocols. Although, in principle, individual end users could choose to adopt re-ECN, it is likely that only a few power-users would manually deploy re-ECN themselves. Rather, it is CDNs and Data Centers who will do so on behalf of Content Providers, whilst OS or Device Vendors will do so on behalf of other end users, as it is they who determine what is actually incorporated and activated within OS protocol stacks.

In the end user-led scenario, the premise is that it is Content Providers (in conjunction with the CDNs and Data Centers who host and deliver their content) who first adopt the re-ECN protocol. Consequently, the packets they send are marked to provide visibility of the congestion they are causing (or rather not causing as indicated by the low number of marked packets). These users thus demonstrate themselves to be 'network friendly', and simultaneously exert peer pressure on other end users to also adopt re-ECN.

However, to provide a greater incentive for Content Providers to adopt re-ECN, ISPs should also enable ECN in their routers (and favour ECN traffic). Assuming both sender and receiver are ECN-capable then end users will also benefit from plain ECN, receiving indication of the onset of congestion prior to the dropping of packets. For ISPs, this is a lower cost first move but one that may kick-start the process of re-ECN adoption that will eventually lead to widespread realization of the benefits that re-ECN offers.

Application developers might also choose to deploy re-ECN as a strategic move to highlight how little congestion their application causes (e.g. LEDBAT-like protocols used by BitTorrent, Windows Update, Virus Update, Video Delivery Software, etc). Similarly, an OS Vendor might deploy re-ECN as the default for all applications in order to demonstrate how effective it is at providing services to applications that minimizes congestion whilst maintaining performance.

Once re-ECN congestion information is visible in a significant proportion of packets, then some networks (perhaps especially mobile operators) may start to use re-ECN information for traffic management purposes. As a first step they may choose to give an advantage to re-ECN traffic relative to non-re-ECN traffic by, for example, enforcing a bit-rate limit on non-re-ECN traffic. More significantly, they may also choose to replace volume allowances for their customers with congestion allowances whilst maintaining a flat-rate tariff. Instead of monitoring the volume of traffic sent by an end user an ISP must monitor (and police) the number of 'black' packets sent by an end user as a direct measure of the congestion they are causing. Although this now requires the development and deployment of policers, a key aspect of this scenario is that this investment is only needed after the

widespread adoption of re-ECN by end users. Where the number of 'black' packets exceeds its congestion allowance then the end user is penalized by dropping packets at the policer. A *token-bucket* approach for such policing mechanisms is described in [11]. Such congestion-based contracts thus provide an incentive for an end user to implement a rate adaptation mechanism that regulates traffic flows (including sending different traffic streams at differentiated rates if necessary) so that the congestion caused remains within their allowance.

Only at this point will CDNs, Data Centers and Heavy Users who believe they are causing significantly more than average congestion have a real incentive to turn *off* re-ECN. To counter this, networks will need to ensure that re-ECN markings cannot be understated relative to actual congestion. This requires a *dropper* to be inserted at the egress to the network, which provides the incentive for the sender to tell the truth about the level of congestion it is causing by marking the correct proportion of packets 'black'. In this scenario, this works fine for flows that are contained entirely within the ISPs own network (e.g. Flow 1 in Figure 3) but, from the ISPs perspective, is less satisfactory for flows that originate or terminate in other networks (e.g. Flows 2 or 3 in Figure 3) since packets in those flows may only be marked 'red' by congested routers in the ISPs own network rather than along the whole path if neighboring networks have not yet enabled ECN in their network routers.

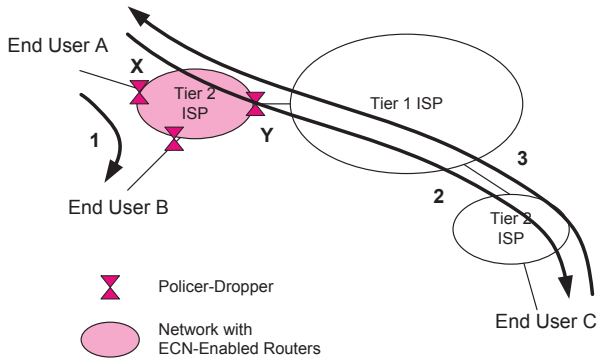

Figure 3. Traffic Flows.

As re-ECN adoption becomes more widespread, the overall benefit will increase as more senders adapt their traffic flows in response to the onset of congestion, and more ISPs enable ECN in their routers and deploy policers and droppers. The QoS experienced by end users increases and network operators do not need to invest more and more in network capacity.

7. CONCLUSIONS

In this paper, we propose an adoption framework for the re-ECN protocol which can be used as a roadmap for analyzing its potential deployment and adoption by stakeholders. We explain the benefits of re-ECN and its backward compatibility with plain ECN, which is achieved by only using one more bit for the core code-point of the protocol. We consider how known problems with middleboxes and OSs fade into insignificance, and we find that its design scores highly against the Internet's design principles.

Furthermore, we present the steps needed for the successful deployment of re-ECN, but note that the incentives for unilateral adoption by either the end user or network communities are weak. We thus outline an adoption scenario where both communities proceed step by step towards the universal or near-universal adoption of re-ECN so that its benefits can be fully realized.

We plan to extend our work by introducing new adoption scenarios that will boost re-ECN adoption. Furthermore, the stakeholder incentives will be revisited and some further conclusions drawn about the most likely route that will lead towards the widespread adoption of the re-ECN protocol.

8. ACKNOWLEDGMENTS

The authors would like to thank Bob Briscoe, Phil Eardley, Pekka Nikander, and Costas Courcoubetis for their useful comments. This research was supported by Trilogy (http://www.trilogy-project.org), a research project (ICT-216372) partially funded by the EC under its 7th Framework Programme.

9. REFERENCES

[1] Briscoe, B., Jacquet, A., Di Cairano-Gilfedder, C., Salvatori, A., Soppera, A., and Koyabe, M. Policing Congestion Response in an Internetwork using Re-feedback. Proc. of ACM SIGCOMM, Philadelphia, PA, USA, September 2005.

[2] Briscoe B., Woundy, R., Moncaster, T., and Leslie, J. Congestion Exposure Mechanism (Conex), IETF, June 2010.

[3] Ramakrishnan, K., Floyd, S. and D. Black. The Addition of Explicit Congestion Notification to IP, RFC 3168, available at 'http://www.ietf.org/rfc/rfc3168.txt', September 2001.

[4] Briscoe, B., Jacquet, A., Moncaster, T., Smith, A. Re-ECN: Adding Accountability for Causing Congestion to TCP/IP, draft-briscoe-tsvwg-re-ecn-tcp-07, March 2009.

[5] Briscoe, B., Jacquet, A., Moncaster, T., Smith, A. Re-ECN: The Motivation for Adding Accountability for Causing Congestion to TCP/IP, draft-briscoe-tsvwg-re-ecn-tcp-motivation-01, September 2009.

[6] Kostopoulos, A., Warma, H., Levä, T., Heinrich, B., Ford, A. and Eggert, L. Towards Multipath TCP Adoption: Challenges and Perspectives. NGI 2010, Paris, June 2010.

[7] Kalogiros, C., Kostopoulos, A. and Ford, A. On Designing for Tussle: Future Internet in Retrospect. EUNICE 2009, LNCS 5733, pp. 98–107, Barcelona, September 2009.

[8] Clark, D., Sollins, K., Wroclawski, J., and Braden, R. Tussle in Cyberspace: Defining Tomorrow's Internet. ACM SIGCOMM CCR 32(4)347--356, October 2002.

[9] Ford, A., Eardley, P., and van Schewick, B. New Design Principles for the Internet. IEEE ICC Future Networks, Dresden, Germany, June 2009.

[10] Briscoe, B. Using Self-interest to Prevent Malice. The Workshop on the Economics of Securing the Information Infrastructure, October 2006.

[11] Jacquet, A., Briscoe, B., and Moncaster, T. Policing Freedom to Use the Internet Resource Pool. In Proceedings of CoNext ReArch'08, Madrid, Spain, December 2008.