Towards the Future Internet
G. Tselentis et al. (Eds.)
IOS Press, 2009

doi:10.3233/978-1-60750-007-0-79

The Trilogy Architecture for the Future Internet

Louise BURNESS[a], Philip EARDLEY[a], Robert HANCOCK[b]
[a]*BT Innovate, Martlesham Heath, Ipswich, UK*
[b]*Roke Manor Research Limited, Romsey, UK*

Abstract. Socio-economic aspects are not intrinsic to the current Internet architecture and so they are handled extrinsically. This has led to increasing distortions and stresses; two examples are inter-domain scaling problems (a symptom of the way multihoming and traffic engineering are handled) and deep packet inspection (a symptom of the lack of resource accountability). The Trilogy architecture jointly integrates both the technical and socio-economic aspects into a single solution: it is thus *designed for tussle*. A Future Internet that follows the Trilogy vision should automatically be able to adapt to the changes in society's demands on the Internet as they occur without requiring permanent redesign.

Keywords. future Internet, architecture, resource accountability, resource pooling, multipath TCP, Trilogy project

1. Introduction

The current Internet is a remarkable phenomenon, not only technically, but from an economic and social perspective as well. It has grown to become the network of choice for a huge variety of distributed applications, starting from email and file transfer, through a whole range of business to consumer and business to business e-commerce systems, all the way to online news and entertainment - for many people, displacing traditional print and broadcast media. And it has achieved all this almost entirely by accident: the happy decision by the original Internet designers to implement a minimal, best-efforts data transfer capability has allowed the Internet to incorporate technological advances and adapt to new application requirements with extraordinary speed and economy.

Where the Internet has been less successful - and this has its roots in the restricted priorities of the designers [Clark88] rather than fundamental flaws in the design - is in accommodating the conflicting economic interests of its participants, both between different operators of the network infrastructure and between the users and providers of network services. This problem was identified in [Clark05] as the "tussle in cyberspace", and arises increasingly often as the demands on the Internet increase.

In the routing area, inter-domain architectures using BGP are tractable while the domain level topology away from the edge is small in scale and organisationally largely hierarchical. However, as the scale increases and topologies become more strongly meshed (for example, because of both direct peering between edge providers

and end-site multihoming), the lack of economic control over provider-provider interactions becomes more apparent. This lack shows itself technically as growth in router tables sizes and churn rates, and the inability of individual operators to manage this load without harming end to end reachability.

The Internet copes well with the demands of best efforts applications, where the network capacity can be engineered to match reasonable expectations on long-term average data transfer requirements. In recent years, these assumptions have been less and less valid. Dependency on the Internet for more demanding applications (voice, physical infrastructure control, critical business operations) has mushroomed, at the same time as peer to peer traffic has demonstrated its ability to absorb all available bandwidth and more. While integrated services approaches can solve these problems in special scenarios, they cannot scale to general purpose Internet deployments because of the technical, and even more importantly administrative, complexity of marshalling the resources of all the application user and network providers involved.

The Trilogy project in the EU 7th Framework programme has been established to address these issues for the Future Internet. The project is based on the view that the major architectural concept behind the Internet - that of a simple, ubiquitous, transparent data delivery infrastructure, which can accommodate innovation at the link and application levels - remains sound today and for the future. The focus is therefore on the network and transport layers of the standard protocol stack, but taking an economic and social perspective as well as pure engineering research. In fact, one goal of the project is to develop an architecture that embeds both technical functions and the socio-economic forces that influence them; this paper represents the first stage in the development of this architecture. Our vision is that the Future Internet will thereby automatically adapt to the changes in society's demands on the Internet as they occur without requiring permanent redesign.

In this paper we make three contributions:

- We hypothesise a baseline Trilogy architecture, which is comparable in scope to the current Internet network and transport layers, but with a subtly different internal structure. In particular, we divide the functions involving the networking infrastructure into two planes, for reachability and forwarding, and distinguish them from the transport services to which the network is totally transparent. (Section 2)
- We propose an accountability framework that is based on congestion volume. Accountability enables a rational basis for sharing resources amongst users on an Internet that is a playground for competing users, applications and businesses. (Section 3)
- We propose the concept of resource pooling, which enables separate network resources to behave like a single large pooled resource, and we propose a technique to achieve this: an end-to-end multipath transport protocol. Resource pooling enables better resilience and efficiency. (Section 4)

This paper includes some selected key pieces of our architecture (more extensive details are elsewhere). We stress that this is our initial architecture; it will change as we validate it through our own activities and through feedback from other researchers and network designers – hence we highly welcome comments on the paper.

2. Baseline Trilogy Architecture

A fundamental assumption of the architecture is that it is based on a minimal packet delivery service. The ideal case is that packets are entirely self-describing, meaning that other concepts such as connections, flows or sessions are higher level constructs, invisible to the delivery service, and the network delivers each packet independently of every other.

We decompose the packet delivery functionality into two parts:

- the reachability plane: responsible for hop-by-hop outgoing link selection and hence enabling network-wide reachability
- the forwarding plane: responsible for deciding how the transmission resource on each link is apportioned between packets.

Distinct from the packet delivery service, we identify the functions that are implemented in a pure end-to-end fashion:

- transport services: functions, such as reliability, flow control or message framing, that are totally invisible to the packet delivery service.

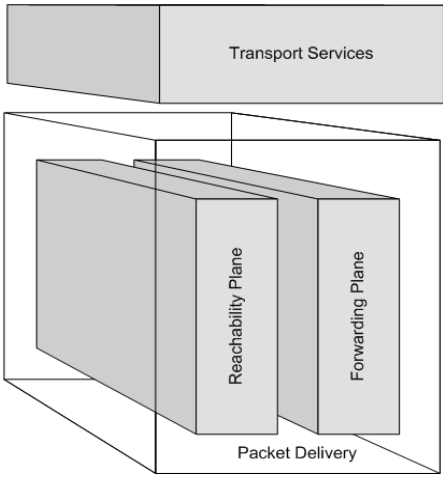

Figure 1: Trilogy baseline architecture

The reachability and forwarding planes are separate; together they achieve the packet delivery service. The key identifier space is the destination locator, which is handled in detail only by the reachability plane. The forwarding plane treats locators as opaque tags which only have to be tested for equality, in order to test for path consistency. The transport services use endpoint identifiers that label the communicating parties, but these are totally independent of the locators used in the reachability plane.

In the rest of this section we briefly describe these three elements. The accountability framework and end-to-end multipath (Sections 3 & 4) are refinements of this baseline architecture. [Del3] contains further details, including other refinements of the baseline architecture and potential extensions

2.1. The Reachability Plane

The reachability plane is responsible for hop-by-hop outgoing link selection. We initially take the information needed to route each packet to be the destination address or locator. Every allocated locator is reachable by default, and functions such as (D)DoS protection must be implemented at the receiver end system. We assume that the reachability service is implemented by a set of autonomous, interconnected administrations or domains. The internal structure of each domain may be non-trivial, but is externally invisible: all that is exposed is information about which locator spaces are reachable and under what circumstances.

Open questions include:

- Are locators from a single global namespace? This approach has the greatest engineering simplicity, and indeed is the only approach totally compatible with the packets being fully self-describing. However, it is also the root cause of many of the Internet stresses, in that it immediately places all network participants (at least, packet sources and destinations if not forwarding infrastructure) into a single "tussle space". The Loc/ID work in the IRTF'S Routing Research Group [RRG] is essentially exploring trade-offs of possible relaxations.
- What additional identifier spaces are used to manage the global topology, and how? We seek to minimise the use of locators in describing topology, to avoid extending that tussle space into inter-network operations. Inter-domain topology should ideally be described in terms of different identifiers, such as the Autonomous System (AS) number of BGP. Note that in interior routing protocols, router identifiers are typically IP addresses, but there is no actual need to couple these identifiers to the locators in the traffic being routed.
- What reachability information or control is shared, beyond the locator spaces themselves? It is notable that much of BGP operation is concerned with traffic engineering (for load balancing), which is an area where the project is exploring solutions in a different part of the architecture (see Section 4).
- What other packet information influences the path? Note that the path itself is not visible in the packet format, because we believe that the end nodes should be isolated from the network infrastructure. As well as the destination locator, we allow that other identifiers visible at the packet level may influence routing explicitly (path selector bits and service class identifiers). The reason is to enhance path consistency between packets, which is an important property for endpoint-managed resource control algorithms.

2.2. The Forwarding Plane

The forwarding plane decides how transmission resource on each link are apportioned between packets. Information about resource allocation along the path can either be implicit (delay or loss) or explicit ('marking' in the packet header); end systems either measure the implicit path properties or read the explicit information, and modify their sending rates in response. The allowed set of responses is left very open at this stage, and specific issues of accountability are discussed further in Section 3.

The fundamental packet delivery service is best-efforts: the network delivers packets to their destination, without guarantees on ordering or throughput. However, there is an implicit requirement that an end system generating a sequence of packets must be able to depend on path consistency, at least over a scale of a round trip time or more.

Looking at the information managed by the forwarding plane, we identify two major open issues:

- What is the level of path property information provided by the network? The optimum situation, consistent with the principle of the reachability plane (that each packet is self-describing), is that each packet is marked with a complete description of forwarding resources available on the path, but this could impose a significant per-packet overhead (in size and forwarding cost). More constrained encodings impose stronger requirements for path consistency and stability for resource control loops to be effective.
- How does the forwarding plane distinguish traffic types? We do, at a minimum, assume that any domain can define certain service classes with different forwarding performance, and that end systems can select between these by embedding information in their packets. However, it is not clear if these service classes can or should be globally defined, or whether they are agreed only at interconnection points – between end systems and networks, and between networks themselves.

2.3. The Transport Services

Transport services are the means by which the packet delivery functions are actually exercised. Re-engineering of the standard transport services is therefore a main method to exploit the functions of the Trilogy architecture. They are implemented in a pure end-to-end fashion and define the communications service offered to applications (email, messaging, file sharing, multimedia, …). Included are functions such as reliability, flow control or message framing, which are totally invisible to the packet delivery service. The identifier spaces involved are also totally separate from the reachability plane's locators; indeed a single node might use several different identifier families, and they may be implicit rather than explicit.

3. Accountability Framework

3.1. From statistical multiplexing to resource accountability

Since the earliest days of telecommunications operators have used statistical multiplexing to maximise the number of customers that can share a backhaul link. This reflects the fact that it is not economically viable to provide each user with an end-to-end dedicated link or circuit at the speed they require. Statistical multiplexing assumes that at any one time only a handful of users will be actively using any given link. This allows telecoms companies to share resources between users deep in the core and thus

take advantage of significant cost savings. The approach has worked well for phone networks. Of course there's a small chance that at a particular moment there isn't enough capacity and so the new phone call is blocked; hence phone operators have to balance the number of customers, their grade of service and the amount of capacity.

Packet switched networks are the logical extension of this drive for efficiency in the network. Statistical multiplexing is done packet-by-packet rather than call-by-call. TCP's job is to perform this multiplexing – the sharing of capacity amongst the users. Historically it has worked well, however it has proved inadequate with the rise of new sorts of application, peer-to-peer being perhaps the most prominent, but also other like voice and business critical VPNs. P2P, for example, undermines the assumption on which statistical multiplexing is built. Nowadays there are some users who download (and upload) content 24 hours a day, 7 days a week. With the growth in P2P TV this can only get worse.

A closer analysis [Briscoe07] reveals that there are actually several interconnected issues:

1. A wider range of users: there are now some users who consume orders of magnitude more bandwidth than others – far more extreme than when everyone just used 'interactive' applications like web browsing. Today less than 1% of customers can generate several 10s% of the traffic.
2. Different types of utility: The TCP algorithm assumes a particular utility function (how much a flow values a byte of information). TCP's utility has the following characteristics:
   - Convexity: twice the bit rate has less than twice the utility
   - Equality: all TCP flows have the same utility function
   - Immediacy: what matters is the bit rate 'now' (within a round trip time); it's irrelevant what the bit rate was or what it will be.

   However, P2P's utility function is very different from TCP's: what matters is how long it takes for the whole film to finish downloading.
3. Non-cooperativeness: Statistical muxing on the Internet also assumes that applications use the TCP algorithm to reduce their rate when there is congestion. But using TCP is voluntary – there are no laws about it! – although luckily (most) application writers build in the use of TCP. However, P2P applications open many TCP flows, typically tens, which of course squeeze the bandwidth available for applications like web browsers that open only one or a few flows. See left hand side of Fig 2.

After the above discussion, the reader may be wondering why they still get a reasonable web browsing experience. Why don't P2P applications get all the bandwidth by opening an ever greater number of flows? Why don't some applications use congestion control that's more aggressive than TCP? (After all, congestion control is a game of chicken; whoever blinks last wins the most bandwidth.) Why don't we have a Tragedy of the Commons?

The reason is that ISPs have introduced DPI. Deep Packet Inspection boxes control the balance of resources between users and between applications. The basic idea is that a device "inspects" each packet and then takes some appropriate action, for example limiting the fraction of bandwidth available for P2P. This leaves more bandwidth available for 'interactive' applications – see centre of Fig 2. ISPs also count how much

traffic each user transfers and cap it according to their fair usage policy. Incidentally, from an operator perspective DPI also allows investment in a capacity upgrade knowing that it won't be "wasted" on P2P.

However, DPI has drawbacks, not least that P2P applications try to disguise themselves as interactive ones, so the DPI has to be cleverer; P2P applications then disguise themselves further (eg using encryption) and then we're in an arms race. So DPI is really a sticking plaster. What is needed is a proper architectural solution for deciding how to allocate resources.

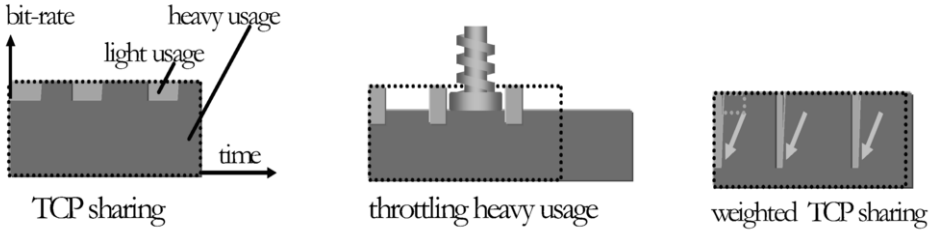

Figure 2: Resource sharing.

Left: TCP sharing (base case). Middle: Limit volume (DPI). Right: Limit congestion volume (this paper)

3.2. Accountability framework for resource allocation

We believe that the architectural answer is "accountability for congestion volume". What does this mean? Congestion is what happens when too much traffic meets too little capacity. It is the adverse impact that your traffic has on others (and vice-versa), ie its externality. Congestion volume is simply the congestion integrated over time and data, so that all the congestion you cause counts (however many flows you create, whatever route, whatever the activity factor etc). For creating an accountability framework [Argyraki07], it is important to know the "rest-of-path congestion" [Laskowski06], as explained below; at present, congestion is only known about on an end-to-end basis (by TCP) or at a single point in the network.

We are developing a framework [Briscoe08] containing the elements required to achieve accountability based on congestion volume:

1. congestion information on each packet. This has already been standardised as ECN and is implemented on all routers.

2. a new mechanism such as re-feedback that gives the ability of any point in the network to calculate the congestion on the rest of the path. 'Rest of path congestion' is the total congestion suffered between a particular point in the network and the destination.

3. a policer at the first ingress, to catch those trying to cause more congestion than they're allowed under their contract. [Jacquet08]

4. a dropper at the last egress, to catch those trying to cheat by under-declaring their rest-of-path congestion

5. border gateways, which count the congestion volume in each direction between two adjacent ISPs; this is a bulk measurement (not per flow). There would be inter-ISP contracts, similar to those today for bandwidth.

6. weighted congestion control. The end host runs an algorithm that reacts to congestion but weighted according to the priority of a particular flow. This achieves true end-to-end QoS.

Some immediate implications of this approach:

1. We envisage that a certain amount of congestion volume would form part of the broadband contract (fair usage policy). The end user will almost certainly not be charged for every single byte of congestion volume that their traffic causes.

2. Software on the user's computer will automatically prioritise their traffic, so that their most important applications tend to use up their congestion volume allowance. If there is no congestion on the path, then you can go as fast as you like, since you don't affect anyone else. If there is congestion, then the user should choose their priorities, since only the user really understands the importance of a particular data flow (DPI makes an educated guess, but it may be wrong). See right hand side of Fig 2.

3. It also enables other types of utility to be taken into account. For instance, a mobile terminal might want to save battery power by transmitting fast in short bursts.

4. It enables everyone to get a better experience, compared to DPI and volume capping. The lower right picture shows that both the interactive applications run faster and the P2P downloads finish earlier.

5. Visibility of rest-of-path congestion enables networks to traffic engineer based on how much congestion there is in other networks. This is analogous to an in-car navigation system learning about hold-ups on the rest of your planned route, and so being able to recommend a diversion.

6. The above points emphasise the wide-ranging impact of the accountability framework; it is a mistake to see it as about "just saving operators money".

4. Resource Pooling and Multi-path Transport

The concept of resource pooling is that separate network resources behave like a single large, pooled resource. Indeed the concept underlies the general principle that resilience should be achieved through redundancy and diversity which led to the use of packet switching. In today's environment, we would like to expand resource pooling across multiple links, because the Internet is much more interconnected than in the past. We are investigating an end-to-end multipath-capable transport protocol as a means to achieve resource pooling [Wischik08]. The concept is simple: enable a single logical connection to use multiple paths simultaneously. The main motivation behind this is to improve the resilience of the network through the use of redundancy.

In the recent discussions, for example [SHIM6], it has been assumed that the sender and/or receiver have multiple addresses associated with different network access technologies or providers. After the initial handshake, the sender and receiver exchange IP addresses and sub-flows can be created using different combinations of source and

destination address. Packets that are lost from one sub-flow may be re-transmitted over any of the other sub-flows.

Reasons for the recent renewed interest [Handley08] may include the fact that recent theoretical work has studied how the congestion response should be managed in such a multi-flow environment [Kelly05], [Massoulie07]. This shows that the congestion response of each sub-flow should be coupled together. One implication is that whilst rate (strictly the TCP window) increase is as normal, decrease is more aggressive. This has the result of allocating rates efficiently between the sub-flows. The correct response also ensures network stability. This is important when a path fails or appears; then data should not be moved suddenly from the failed path (or onto the new path), but instead the window should gradually increase over several round trips, whilst communication can meanwhile continue along existing paths.

The primary benefit of such a scheme is improved resilience in the case of failure. However the benefits are broader because such a mechanism also makes the network better able to handle localized surges in traffic and maximizes the utilization of the network [Handley08]. It automatically achieves a load balanced network. These features arise because multi-path transport effectively pools the network's resources. The upper part of Fig 3 shows how gains are made if resources – bandwidth – can be shared in a single pipe. This is realized today with packet switching. The lower part of Fig 3 shows how this concept is extended to a multi-path environment.

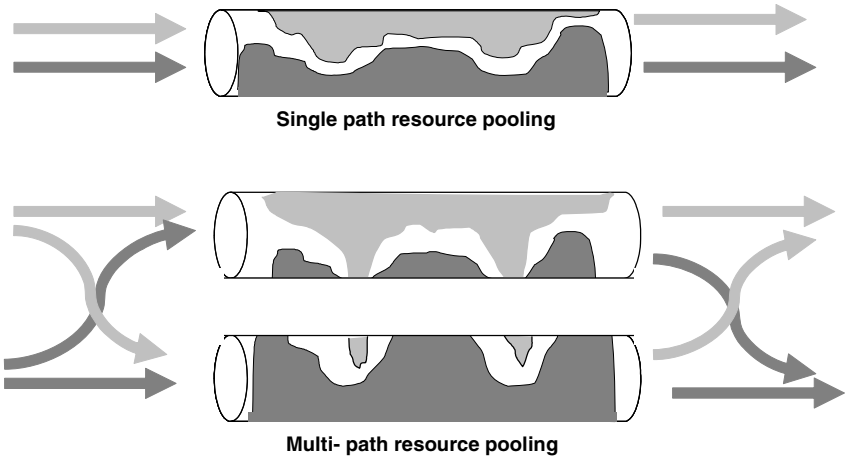

Figure 3: Resource pooling. over single or multiple bit pipes

However, it is important to consider how end-to-end multipath resource pooling will interact (and possibly conflict) with other resource pooling mechanisms that are already in use today. For example, networks may multi-home for resilience and then traffic engineer over these multiple links to utilize the resources properly. The mechanisms used today (essentially BGP routing) unfortunately lead to more scalability and churn problems for the routing system, which decreases the reliability of the whole network, driving more users towards multi-homing; a vicious circle. Peer-to-peer networks also attempt resource pooling in order to maximize their performance –

pooling upstream capacity and also pooling data availability by spreading data over multiple unreliable servers. Content delivery networks also attempt resource pooling, spreading load between multiple servers. Unfortunately the resource pooling of these different entities may be in conflict, for example an ISP lowest cost path choice may well be different from the peer-to-peer user's high bandwidth path choice. Also today we see a much greater range of applications all competing for the same resources – simple data transfers, long-lived bulk flows and voice all co-exist. Users are competing ever more aggressively for a share (fair or otherwise!) of the resources. The costs of this type of conflict can be arbitrarily high [Acemoglu07][Roughgarden02].

Architecturally, it seems that control of multi-path at the transport layer is optimal. When a network operator decides to re-route traffic to achieve better load balancing, the sudden change in traffic patterns could lead to congestion elsewhere in the network (and may in turn lead to traffic engineering downstream attempting to force the traffic back towards the original route). This suggests that control of multipath below the transport layer is too low to ensure network stability and safety. But multi-path management could be offered generically to many applications; having a standardized well understood mechanism may go some way to ensuring that conflicts can be managed. Hence having the functionality in the transport layer seems best.

However, we consider that the network providers need a way to influence the path choice taken – there needs to be some kind of feedback from the network to the users to ensure that the economic stability of the Internet is not lost. One initial idea for how this might be achieved is by the network provider adding ECN congestion marks [Wischik08]; lots of marks will encourage traffic to move away from the non-preferred, expensive link. On the other hand, it seems likely that a multi-path capability that resides with end-users will foster better competition between providers.

There are still outstanding architectural questions:

- This type of resource pooling is most effective for data and bulk transport; it is much less suited to jitter-sensitive applications.
- The end points need to be able to use multiple paths – what is the best mechanism? Provider aggregatable addresses have address management issues which are non-trivial if the access network is multi-homed rather than the end host.
- How many paths do the hosts need to access? The tentative answer is that just a few paths are needed [Mitzenmacher01] – provided they are the right paths [PGB08]! How much benefit could be gained if these were known to be disjoint?

5. Conclusions and next steps

5.1. Conclusions

The original Internet was an academic research network that has proved immensely successful – so successful that managing its rapid growth has been much more important than finishing the research! But increasingly the pressures on the Internet's growth can be attributed to the missing pieces. The missing pieces are well known, and it is hardly surprising that they were not considered critical to the early Internet. The early Internet was not composed of participants with different economic interests; the degree of interconnectivity was low as physical resources were so expensive and scarce that the concepts of having a large number of paths and multiple points of connection were almost unthinkable. The world is a different place today.

We believe that the basic Internet architecture is a very good starting point. Some concepts, like connectionless packet switching, are as good today as ever. Some concepts, resource pooling for instance, need extending and repositioning within the architecture to cover the wider range of resources that are available today. And some concepts, such as resource accountability, need to be added in from scratch. Overall we strive to provide a technical solution that should be able to respond to the changes in society, and that allows different economic models representing different business regimes and indeed different societies to co-exist within the single global network.

We have presented an architecture based on design for tussle. It is radical in that the architecture is in many ways a small, yet defined, step from the current architecture. We have identified the need for separation of reachability (path discovery) and forwarding (path resource management), and also that end hosts and routers should be involved in both processes in a coordinated manner – rather than in the confrontational manner of today. We believe that the addition of accountability should improve user-to-user interactions, network-to-user interactions and network-to-network interactions by enabling all interested parties to actually understand the global network behaviour and understand how their actions influence this behaviour.

5.2. Next steps

As is clear from our open questions, there is still much that needs to be done. We hope that exposing the architecture at this early stage will help foster debate. Meaningful validation of architectural work is always a challenge and we believe that we have reached a point where further progress will now depend on evaluation of concrete case studies or new proposals. We are developing candidate technical proposals in the areas of reachability and resource control, and hope to study their interactions practically. We will evaluate these both for architectural compatibility on the one hand, and simplicity and performance in realistic environments on the other. This combination of engineering evaluation, mathematical analysis, and simulation will be used to refine both the solutions and the architecture within which they fit.

Because the architecture is close to the current system, we hope that migration is plausible. The 'multipath TCP' protocol is very similar to existing transport protocols and so does not require any changes to the network. It depends on end hosts deploying

software pair-wise, and since the end systems directly benefit this migration is very plausible. Migration towards the accountability framework is subject to on-going study.

One of our key motivations is try and incorporate flexibility, so that the architecture can adapt for local business and operational needs. Our aim is to ensure that the future network can be more tolerant of the demands of society - by adding in design for tussle, specifically resource accountability, we hope this architecture will be valid for another 30 years.

References

[Acemoglu07] Acemoglu, R. Johari, and A. Ozdaglar, "Partially optimal routing. IEEE Journal of selected areas in communications", 2007

[Argyraki07] K Argyraki, P Maniatis, O Irzak, S Ashish & S Shenker, "Loss and Delay Accountability for the Internet," In Proc. IEEE ICNP'07 (Oct 2007)

[Briscoe07] B Briscoe, T Moncaster & L Burness, "Problem Statement: We Don't Have To Do Fairness Ourselves" IETF I-D draft-briscoe-tsvwg-relax-fairness-00.txt (work in progress, Nov 2007)

[Briscoe08] Bob Briscoe, A Jacquet, T Moncaster & A Smith, "Re-ECN: Adding Accountability for Causing Congestion to TCP/IP" IETF I-D draft-briscoe-tsvwg-re-ecn-tcp-05.txt (work in progress, Jan 2008)

[Clark88] D. Clark, "The design philosophy of the Darpa Internet protocols," In Proc. ACM SIGCOMM, Vancouver, BC, Canada, Sept. 1988.

[Clark05] D. Clark, J. Wroclawski, K. Sollins, R. Braden, "Tussle in Cyberspace: Defining Tomorrow's Internet", IEEE/ACM Transactions on Networking, 13(3), p. 462-475, June 2005.

[Del3] Trilogy Project Initial overall architecture Report, available from http://trilogy-project.org/publications/deliverables.html , August 2008

[FORCES] http://www.ietf.org/html.charters/forces-charter.html

[Handley08] "Multipath TCP and the resource pooling principle, Mark Handley, Damon Wischik and Marcelo Bagnulo Braun, IETF Dublin TSVAREA presentation, 2008

[Jacquet08] "Policing Freedom to Use the Internet Resource Pool", Jacquet, Briscoe and Moncaster; Re-Arch CoNEXT workshop, Dec 2008

[Kelly05] F. Kelly , T. Voice, "Stability of end-to-end algorithms for joint routing and rate control", ACM SIGCOMM Computer Communication Review, v.35 n.2, April 2005

[Laskowski06] P Laskowski & J Chuang, "Network Monitors and Contracting Systems: Competition and Innovation," In Proc. SIGCOMM'06, ACM CCR 36(4)183--194 (2006)

[Massoulié07] P. B. Key, L. Massoulié, D. F. Towsley, "Path Selection and Multipath Congestion Control", INFOCOM 2007, 143-151

[Mitzenmacher01] M. Mitzenmacher, "The Power of Two Choices in Randomized Load Balancing", IEEE Transactions on Parallel and Distributed Systems, Volume 12 , Issue 10 (October 2001)

[PBG08] P. Brighten Godfrey, "Balls and bins with structure: balanced allocations on hypergraphs", Symposium on Discrete Algorithms, San Francisco, 511-517, 2008

[Roughgarden02] T. Roughgarden and E. Tardos, "How bad is selfish routing?", Journal of the ACM, 2002

[RRG] Routing research group

[SHIM6] http://tools.ietf.org/wg/shim6/

[Thaler00] D. Thaler, C. Hopps, "Multipath Issues in Unicast and Multicast Next-Hop Selection", RFC 2991, November 2000

[Wischik08] "The Resource Pooling Principle", Damon Wischik, Mark Handley and Marcelo Bagnulo Braun. ACM/SIGCOMM CCR, Oct 2008

Acknowledgements

The research results presented herein have received support from Trilogy (http://www.trilogy-project.eu), a research project (ICT-216372) partially funded by the European Community under its Seventh Framework Programme. The views expressed here are those of the author(s) only. The European Commission is not liable for any use that may be made of the information in this document.