

Network Working Group
Internet-Draft
Intended status: Informational
Expires: September 02, 2014

E. Lear, Ed.
Cisco Systems GmbH
March 01, 2014

Internet Technology Adoption and Transition
draft-lear-iab-itat-report-00

Abstract

The Internet Technology Adoption and Transition (ITAT) Workshop took place on December 4th and 5th of 2013. The focus of the workshop was on models to facilitate adoption of Internet protocols, with particular emphasis at the waist of the hourglass. This report summarizes contributions and discussions. As the topics were wide ranging, there is no single set of recommendations for IETF participants to pursue at this time. Instead, in the classic sense of early research, we note areas that deserve further exploration.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on September 02, 2014.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must

include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
1.1.	Organization of This Report	3
2.	Motivations and Review of Existing Work	4
3.	Economics of Protocol Adoption	5
3.1.	When can bundling help adoption of network technologies or services?	5
3.2.	Internet Protocol Adoption: Learning from Bitcoin	5
3.3.	Long term strategy for a successful deployment of DNSSEC - on all levels	6
3.4.	Framework for analyzing feasibility of Internet protocols	6
3.5.	Best Effort Service as a Deployment Success Factor	7
4.	Innovative / Out There Models	7
4.1.	On the Complexity of Designed Systems (and its effect on protocol deployment)	7
4.2.	Managing Diversity to Manage Technological Transition	8
4.3.	On Economic Models of Network Technology Adoption, Design, and Viability	8
5.	Making Standards Better	8
5.1.	Standards: A love/hate relationship with patents	8
5.2.	Bridge Networking Research and Internet Standardization: Case Study on Mobile Traffic Offloading and IPv6 Transition Technologies	9
5.3.	An Internet Architecture for the Challenged	9
6.	Other Challenges and Approaches	9
6.1.	Resilience of the commons: routing security	10
6.2.	Getting to the next version of TLS	10
7.	Outcomes	10
7.1.	Work for the IAB and the IETF	11
7.2.	Potential for the Internet Research Task Force	11
7.3.	Opportunities For others	11
8.	Security Considerations	12
9.	Acknowledgments	12
10.	Attendees	12
11.	References	12
11.1.	Normative References	12
11.2.	Informative References	13
Appendix A.	Changes	14
	Author's Address	14

Lear

Expires September 02, 2014

[Page 2]

1. Introduction

[Ed: this is adapted from our call for papers.]

The Internet is a complex ecosystem that encompasses all aspects of society. At its heart is a protocol stack with an hourglass shape, and IP at its center. Recent research points to possible explanations for the success of such a design and for the significant challenges that arise when trying to evolve or change its middle section, e.g., as partially evident in the difficulties encountered by IPv6. We have a number of other key examples to consider, including the next generation of HTTP and WebRTC. The eventual success of many if not all of these protocols will largely depend on our understanding of not only what features and design principles contribute lasting value, but also on how (initial) deployment strategies can succeed in unlocking that value to foster protocol adoption. The latter is particularly important in that most if not all Internet protocols exhibit significant externalities that create strong barriers to adoption, especially in the presence of a well-established incumbent. Taking into account [RFC 5218](#) on what makes a protocol successful, this workshop sought to explore how the complex interactions of protocols design and deployment affect their success. One of the workshop's goals was, therefore, to encourage discussions to develop an understanding of what makes protocol designs successful not only in meeting initial design goals but more importantly in their ability to evolve as these goals and the available technology change. Another equally important goal was to develop protocol deployment strategies that ensure that new features can rapidly gain enough of a foothold to ultimately realize broad adoption. Such strategies must be informed by both operational considerations and economic factors.

Participants this workshop consisted of operators, researchers from the fields of computer science and economics, as well as engineers. Contributions were wide ranging. As such, this report makes few if any recommendations for the IETF to consider, although there are some.

1.1. Organization of This Report

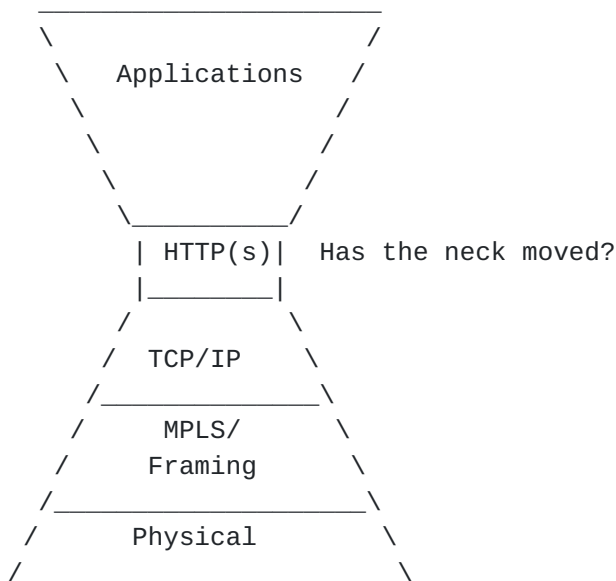
This report records our discussions. At the end of the workshop we reviewed potential follow-up items. These will be highlighted at each point during the report, and a summary is given at the end.

[Section 2](#) discusses the economics of protocol adoption. [Section 3](#) delves into an examination of recent operational challenges and some success stories. [Section 4](#) examines different views of success factors. Finally [section 5](#) summarizes views of the participants, and identifies a few key insights.

2. Motivations and Review of Existing Work

Our workshop began with an introduct that asks the question: is the neck of the Internet hourglass closed for business? We have numerous instances where progress has been slow, the three biggest that come to mind being IPv6 [[RFC2480](#)], SCTP [[RFC4960](#)], and DNSSEC [[RFC4034](#)]. The impact of DNSSEC is of particular interest, because it is relied upon for the delivery of other services, such as DANE [[RFC6698](#)] as well as application discovery mechanisms[refneeded]. Thus slowdown at the neck of the glass can have an impact closer to the lip.

Even when we consider the classic neck of the hourglass to be IP and transport layers, it was suggested that the hourglass might extend as high as the application layer.



This idea was rebutted by the argument that protocols do continue to evolve, that protocols like SMTP and IMAP in the applications space have continued to evolve, as has the transport layer.

We moved on to a review of [[RFC5218](#)] which discusses protocol success factors. This work was presented in an IETF plenary some time ago, and was the basis for this ongoing work. There were two clear outcomes from the discussion. The first was that successive Internet Architecture Boards should review and consider that document in the context of evaluating birds of a feather (BoF) session proposals at the IETF, so that any working group proposal is carefully crafted to address a specific design space and provide positive net value. Another aspect was to continue work on tracking the value specific works in terms of success, wild success, or failure. On that last point, failure remains difficult to judge, particularly at the neck of the hourglass.

[3.](#) Economics of Protocol Adoption

Several papers were submitted that looked at economic aspects of protocol adoption.

[3.1.](#) When can bundling help adoption of network technologies or services?

Economics of bundling is a long studied field, but not as applied to protocols. It is relevant to the IETF and inherent to two key notions: layering and "mandatory to implement". Two current examples include DANE atop DNSSEC and WebRTC atop SCTP. The workshop reviewed a model [[Weber13](#)] that examines the concept that bundling of technologies can lead to several possible outcomes, which includes more or less adoption of both. This will depend on a number of factors, including the costs, benefits, and externalities associated with adopting each. The work we considered involved two independent technologies. One question was what happens where one technology depends on the other. That is directly tied to "mandatory to implement" discussions within the IETF. That is a matter for follow-on work. IETF participants can provide researchers anecdotal experience to help improve models in this area.

[3.2.](#) Internet Protocol Adoption: Learning from Bitcoin

We considered an examination of protocol success factors in the context of Bitcoin[Bohme13]. Here, there were any number of barriers to success, including adverse press, legal uncertainties, glitches and breaches, previous failed attempts, and speculative attacks, amongst others. Bitcoin has thusfar overcome these barriers thanks to several key factors:

- o First, there is a built-in reward system for early adopters. Participants are monetarily rewarded at an exponentially declining rate.

- o There exist exchanges or conversion mechanisms to directly convert bitcoin to other currencies.
- o Finally, there is some store of value in the currency itself.

The first two of these factors may be transferrable to other approaches. That is- one key protocol success factor is direct benefit to the participant. Another key protocol success factor is ability to interface with other systems for mutual benefit.

A key message from this presentation is that if a protocol imposes externalities or costs on other systems, find a means to establish incentives for those other players for implementation. As it happened we had a limited example of how to do this that is directly relevant to the IETF.

3.3. Long term strategy for a successful deployment of DNSSEC - on all levels

We reviewed the approach Sweden's .SE registry has taken to improving deployment of DNSSEC[Lowinder13]. .SE has roughly 1.5 million domains. IIS manages the ccTLD. They made the decision to encourage deployment of DNSSEC within .SE. They began by understanding who their stakeholders were, and examined financial, legal, and technical aspects to deployment. As they began their rollout, they charged extra for DNSSEC. As they put it, this didn't work very well.

They went on to fund development of OpenDNSSEC to remove technical barriers to deployment at end sites, noting that tooling was lacking in this area. Even with this development, more tooling is necessary, as they point out a need for APIs between the signing zone and the registrar.

To further encourage deployment, the government of Sweden provided financial incentives to communities to see that their domains were signed. .SE further provided an incentive to registrars to see that their domains were signed. In summary, .SE examined all the players and provided incentives for each to participate.

The workshop discussed whether or not this model could be applied to other domains. .SE was in a position to effectively subsidize DNS deployment because of their ability to set prices. This may be appropriate for certain other top level domains, but it was pointed out that the margins of other domains do not allow for a cost reduction to be passed on at this point in time.

3.4. Framework for analyzing feasibility of Internet protocols

One of the goals of the workshop was to provide ways to determine when work in the IETF was likely to lead to adoption. We considered an interactive approach that combines value net analysis, deployment environment analysis, and technical architecture analysis that leads to feasibility and solution analysis[Leva13]. This work provided an alternative to [RFC 5218](#) that had many points in common. The case study examined was that of MPTCP. Various deployment challenges were observed. First and foremost, increasing bandwidth within the network seems to decrease attractiveness of MPTCP. Second, the benefit/cost tradeoff by vendors was not considered attractive. Third, not all parties may agree on the benefits.

Solutions analysis suggested five separate approaches to improve deployment, including open source, lobbying of various implementors, proxy deployment, and implementation by parties where they own both ends of a connection.

[3.5.](#) Best Effort Service as a Deployment Success Factor

When given the choice between vanilla and chocolate, why not choose both? We considered an approach that became a recurring theme throughout the workshop, which was to examine when it was necessary to make a choice between technologies, but rather to implement multiple mechanisms to achieve adoption[Welzl13]. The workshop discussed the case of Skype, where it will use the best available transport mechanism to improve communication between clients, rather than to tie fate to any specific transport. The argument goes that such an approach provides a means to introduce new transports such as SCTP. This would be an adaptation of "Happy Eyeballs" [[RFC6555](#)].

[4.](#) Innovative / Out There Models

There were several approaches presented that examined how we look at protocol adoption.

[4.1.](#) On the Complexity of Designed Systems (and its effect on protocol deployment)

We reviewed a comparison between the hourglass model and what systems biologists might call the bowtie model[Meyer13]. The crux of this comparison is that both rely on certain building blocks to accomplish a certain end. In the case of our hourglass model, IP sits notably in the center, whereas in the case of systems biology, as adenosine triphosphate (ATP) is the means by which all organisms convert nutrients to usable energy.

We also examined the notion of "robust yet fragile", which examines the balance between the cost of implementing robust systems versus

their value. That is, highly efficient systems can might prove fragile in the face of failure or hard to evolve.

The key question asked during this presentation was how we could apply what has been learned in systems biology or what do the findings reduce to for engineers? The answer was that more work is needed. The discussion highlighted the complexity of the Internet in terms of predicting network behavior. As such, one promising area to examine may be that of network management.

4.2. Managing Diversity to Manage Technological Transition

We considered the difference between planned versus unplanned technology transitions[Kohn013]. They examined several transitions at the link, IP, and application layers in Japan. One key claim in the study is that there is a phase difference in the diversity trend between each layer. The statistics presented show that indeed HTTP is the predominant substrate for other applications. Another point made was that "natural selection" is a strong means to determine technology.

Along these lines there were two papers submitted that examined the formation and changes to the hourglass in the context of evolutionary economics. Unfortunately the presenter was unable to attend due to illness. The work was discussed at the workshop, and there were different points of views as to the approach.

4.3. On Economic Models of Network Technology Adoption, Design, and Viability

We considered how network protocol capabilities enable certain sorts of services that are beneficial to consumers and service providers. This model looks at smart data pricing (SDP) in which some behavior is desired and rewarded through a pricing model.[Sen13] The example given was use of time-dependent pricing (TDP) and demonstrated how a service provider was able to load shift traffic to off-peak periods. Explicit Congestion Notification (ECN) and RADIUS were used by the project alongside a simple GUI. This sort of work may prove useful to service providers as caching models evolve over time. The question within the room is how will protocol developers consider these sorts of requirements.

5. Making Standards Better

There were several papers that focused on how standards are produced.

5.1. Standards: A love/hate relationship with patents

One of the biggest barriers to deployment is that of the unseen patent by the non-practicing entity (NPE).[\[Lear13\]](#) While this problem is relatively well understood by the industry, the discussion looked at patents as a means to improve interoperability. Those who hold patents have the ability to license them in such a way that a single approach is the result.

[5.2.](#) Bridge Networking Research and Internet Standardization: Case Study on Mobile Traffic Offloading and IPv6 Transition Technologies

Not for the first time there was a presentation and discussion about the gap between the research community and standards organizations. Two cases were examined: mobile offloading and IPv6 transition technologies.[\[Ding13\]](#) In the case of mobile offloading, a mechanism was examined that required understanding of both 3GPP and IETF standards. Resistance in both organizations was encountered. In the 3GPP, the problem was that the organization already had an offloading model in play. In the IETF, the problem was a lack of understanding of the interdisciplinary space. The researchers noted that in the case of the IETF, they may have taken the wrong tact by having jumped into the solution without having fully explained the problem they were trying to solve. In the case of IPv6 transition technologies researchers encountered a crowded field and not much appetite for new transition technologies.

The workshop discussed whether the standards arena is the best venue or measurement of success for researchers. The IRTF is meant to bridge academic research and the IETF. As we will discuss below, several avenues for continued dialog are contemplated.

[5.3.](#) An Internet Architecture for the Challenged

We held a very provocative discussion about whether the existing Internet architecture serves the broadest set of needs. Three specific aspects were examined: geographic, technical, and socio-economic. Researchers presented an alternative hourglass or protocol architecture known as Lowest Common Denominator Networking (LCDNet) that re-examines some of the base assumptions of the existing architecture, including its "always on" nature.[\[Sathiaseelan13\]](#)

The workshop questioned many of the baseline assumptions of the researchers. In part this may have been due to constrained discussion time on the topic, where a fuller explanation was warranted.

[6.](#) Other Challenges and Approaches

We held a number of other discussions about different approaches to technology adoption. We should highlight that a number of papers were transmitted to the workshop on routing security, two of which were not possible to present.

6.1. Resilience of the commons: routing security

We discussed a presentation on the tragedy of the commons in the context of global inter-domain routing[Robachevsky13]. The "Internet Commons" is a collection of networks that we depend on but do not control. The main threat to the commons in the context of BGP is routing pollution, or unwanted or unnecessary routing entries. The Internet Society has been working with service providers to improve resiliency by driving a common understanding of both problem and solution space, and developing a shared view with regard to risk and benefits, with the idea being that there would be those who would engage in reciprocal cooperation with the hopes that others would do similarly in order to break the tragedy.

What was notable in discussion was that there was no magic bullet to addressing the resiliency issue, and that this was a matter of clearly identifying the key players and convincing them that their incentives were aligned. It also involved developing approaches to measure resiliency.

6.2. Getting to the next version of TLS

Originally the workshop had planned to look at the question of whether we could mandate stronger security. This evolved into a discussion about getting to the next version of Transport Layer Security (TLS), and what challenges lie ahead. It was pointed out that there were still many old versions of TLS in existence today, due to many old implementations. In particular, it was pointed out that a substantial amount of traffic is still encrypted using triple DES.

One concern about the next generation is that perfect could become the enemy of good. Another point that was made was that perhaps a testing platform might help interoperability. Finally, there was some discussion about how new versions of TLS get promoted.

7. Outcomes

This wide ranging workshop discussed many aspects that go to the success or failure of the work of the IETF. While there is no single silver bullet that we can point to for making a protocol successful, the workshop did discuss a number of outcomes and potential next steps.

7.1. Work for the IAB and the IETF

The IAB's role in working group formation consists of providing guidance to the IESG on which birds of a feather sessions should be held, review of proposed working group charters, and shepherding some work so that it can reach a suitable stage for standardization. In each of these stages the IAB has an opportunity to apply the lessons of [RFC 5218](#), as well as other work such as the notion of bundling choices, when members give advice.

In addition to working group creation, the IAB has an opportunity to track and present protocol success stories, either through wikis or through discussion at plenary sessions. For instance, there is much interest at the moment of this report in Bitcoin, its success, and what parallels and lessons can be drawn. Specifically it would be useful to track examples of first mover advantages.

Finally, one area that the IETF may wish to consider, relating specifically to DNSSEC, as raised by our speakers was standardization of the provisioning interface of DNSSEC (DS keys) between parent and child zone. Contributions in this area would be welcome.

7.2. Potential for the Internet Research Task Force

There are at least two possible activities that the IRTF might wish to consider. The first would be a research group that considers protocol alternatives and recommendations that might be useful in areas where environments are constrained, due to bandwidth or other resources. Such a group has already been proposed, in fact.

The second possibility is a more general group that focuses on economic considerations relating to Internet protocol design. In particular there were a number of areas that were presented to the working group that deserve further investigation, and could use collaboration between researchers, engineers, and operators. Two examples include work on bundling as well as systems biology.

7.3. Opportunities For others

Incentive models often involve many different players. As we considered work in the workshop, our partners such as ICANN, and the RIRs can continue to play a role in encouraging deployment of protocols through their policies. Their members can also participate in any activity of the IRTF that is related to this work.

Specifically, RIRs have a specific role to play in encouraging security for the routing system, and ICANN has a specific role to play in securing the domain name service.

The suggestion was made that the IETF working groups could leverage graduate students in many Universities around the world in helping review documents (drafts, RFCs etc). This would serve as a source of education in real world processes to students, and would engage the research community in IETF processes more thoroughly, as well as providing a scale-out resource for handling the IETF review workload. Several attendees who have such students were prepared to try this out.

8. Security Considerations

This document does not discuss a protocol. Security for the workshop itself was excellent.

9. Acknowledgments

The IAB would like to thank the program committee, who consisted of Roch Guerin, Constantine Dovrolis, Hannes Tschofenig, Joel Halpern, Eliot Lear, and Richard Clayton.

A special debt of gratitude is owed to our hosts, Ross Anderson and Richard Clayton for arranging an excellent venue for our discussions.

10. Attendees

The following people attended the ITAT workshop:

Aaron Yi Ding, Adrian Farrel, Andrei Robachevzsky, Arjuna Sathiseelan, Bjoern Zeeb, Dave Meyer, Dave Thaler, Dongting Yu, Eliot Lear, Elwyn Davies, Erik Nordmark?, Hannes Tschofenig, Joel Halpern, Jon Crowcroft, Lars Eggert, Martin Stiernerling?, Michael Welzl, Michiel Leenaars, Miyo Kohno, Rainer Bohme, Richard Clayton, Roch Guerin, Ross Anderson, Russ Housley, Sam Smith, Sean Turner, Soumya Sen, Spencer Dawkins, Steven Weber, Tapio Levae, Toby Moncaster, Tony Finch

11. References

11.1. Normative References

[RFC5218] Thaler, D. and B. Aboba, "What Makes For a Successful Protocol?", [RFC 5218](#), July 2008.

11.2. Informative References

- [Bohme13] Bohme, R., "Internet Protocol Adoption: Learning from Bitcoin", December 2013.
- [Ding13] Yi Ding, A., Korhonen, J., Savolainen, T., Kojo, M., Tarkoma, S., and J. Crowcroft, "Bridge Networking Research and Internet Standardization: Case Study on Mobile Traffic Offloading and IPv6 Transition Technologies", December 2013.
- [Kohno13] Kohno, M., Asaba, T., and F. Baker, "Managing Diversity to Manage Technological Transition", December 2013.
- [Lear13] Lear, E. and D. Mohlenhoff, "Standards: a love/hate relationship with patents", December 2013.
- [Leva13] Leva, T. and H. Soumi, "Framework for analyzing feasibility of Internet protocols", December 2013.
- [Lowinder13] Eklund Lowinder, A.M. and P. Wallstrom, "Long term strategy for a successful deployment of DNSSEC - on all levels", December 2013.
- [Meyer13] Meyer, D. M., "On the Complexity of Engineered Systems (and its effect on protocol deployment)", December 2013.
- [RFC2480] Freed, N., "Gateways and MIME Security Multiparts", [RFC 2480](#), January 1999.
- [RFC4034] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Resource Records for the DNS Security Extensions", [RFC 4034](#), March 2005.
- [RFC4960] Stewart, R., "Stream Control Transmission Protocol", [RFC 4960](#), September 2007.
- [RFC6555] Wing, D. and A. Yourtchenko, "Happy Eyeballs: Success with Dual-Stack Hosts", [RFC 6555](#), April 2012.
- [RFC6698] Hoffman, P. and J. Schlyter, "The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security (TLS) Protocol: TLSA", [RFC 6698](#), August 2012.
- [Robachevsky13] Robachevsky, A., "Resilience of the commons: routing security", December 2013.

[Sathiaseelan13]

Sathiaseelan, A., Trossen, D., Komnios, I., Ott, J., and J. Crowcroft, "An Internet Architecture for the Challenged", December 2013.

[Sen13]

Sen, S., "On Economic Models of Network Technology Adoption, Design, and Viability", December 2013.

[Weber13]

Weber, S., Guerin, R., and J. C. Oliveira, "When can bundling help adoption of network technologies or services?", December 2013.

[Welzl13]

Welzl, M., "The "best effort" service as a deployment success factor", December 2013.

[Appendix A.](#) Changes

This section to be removed prior to publication.

o 00 Initial Revision.

Author's Address

Eliot Lear (editor)
Cisco Systems GmbH
Richtistrasse 7
Wallisellen, ZH CH-8304
Switzerland

Phone: +41 44 878 9200

Email: lear@cisco.com

