

Internet Architecture Board (IAB)
Request for Comments: 7305
Category: Informational
ISSN: 2070-1721

E. Lear, Ed.
July 2014

Report from the IAB Workshop
on Internet Technology Adoption and Transition (ITAT)

Abstract

This document provides an overview of a workshop held by the Internet Architecture Board (IAB) on Internet Technology Adoption and Transition (ITAT). The workshop was hosted by the University of Cambridge on December 4th and 5th of 2013 in Cambridge, UK. The goal of the workshop was to facilitate adoption of Internet protocols, through examination of a variety of economic models, with particular emphasis at the waist of the hourglass (e.g., the middle of the protocol stack). 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, the workshop noted areas that deserve further exploration.

Note that this document is a report on the proceedings of the workshop. The views and positions documented in this report are those of the workshop participants and do not necessarily reflect IAB views and positions.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Architecture Board (IAB) and represents information that the IAB has deemed valuable to provide for permanent record. It represents the consensus of the Internet Architecture Board (IAB). Documents approved for publication by the IAB are not a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7305>.

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.

Table of Contents

1. Introduction	4
1.1. Organization of This Report	5
2. Motivations and Review of Existing Work	5
3. Economics of Protocol Adoption	7
3.1. When can bundling help adoption of network technologies or services?	7
3.2. Internet Protocol Adoption: Learning from Bitcoin	7
3.3. Long term strategy for a successful deployment of DNSSEC - on all levels	8
3.4. Framework for analyzing feasibility of Internet protocols	9
3.5. Best Effort Service as a Deployment Success Factor	9
4. Innovative / Out-There Models	10
4.1. On the Complexity of Designed Systems (and its effect on protocol deployment)	10
4.2. Managing Diversity to Manage Technological Transition	10
4.3. On Economic Models of Network Technology Adoption, Design, and Viability	11
5. Making Standards Better	11
5.1. Standards: a love/hate relationship with patents	11
5.2. Bridge Networking Research and Internet Standardization: Case Study on Mobile Traffic Offloading and IPv6 Transition Technologies	11
5.3. An Internet Architecture for the Challenged	12
6. Other Challenges and Approaches	12
6.1. Resilience of the commons: routing security	12
6.2. Getting to the Next Version of TLS	13
7. Outcomes	13
7.1. Work for the IAB and the IETF	13
7.2. Potential for the Internet Research Task Force	14
7.3. Opportunities for Others	14
8. Security Considerations	14
9. Acknowledgments	15
10. Attendees	15
11. Informative References	15

1. Introduction

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. The workshop had a number of other key examples to consider, including the next generation of HTTP and real time web-browser communications (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 how 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 strong externalities that create strong barriers to adoption, especially in the presence of a well-established incumbent. That is, factors beyond the control of the end points (such as middleboxes) can limit deployment, sometimes by design.

The Internet Architecture Board (IAB) holds occasional workshops designed to consider long-term issues and strategies for the Internet, and to suggest future directions for the Internet architecture. This long-term planning function of the IAB is complementary to the ongoing engineering efforts performed by working groups of the Internet Engineering Task Force (IETF), under the leadership of the Internet Engineering Steering Group (IESG) and area directorates.

Taking into account [RFC5218] 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 in this workshop consisted of operators, researchers from the fields of computer science and economics, and engineers. Contributions were wide ranging. As such, this report makes few recommendations for the IETF to consider.

1.1. Organization of This Report


This report records the participants' discussions. At the end, workshop participants 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 reviews the motivations and existing work, and Section 3 discusses the economics of protocol adoption. Section 4 covers innovative models for protocol adoption. Section 5 delves into an examination of recent standards issues and some success stories. Section 6 examines different views of success factors. Finally, Section 7 examines potential next steps.

2. Motivations and Review of Existing Work

Our workshop began with an introduction that asks the question: is the neck of the Internet hourglass closed for business? There are numerous instances where progress has been slow, the three biggest that come to mind being IPv6 [RFC2480], the Stream Control Transmission Protocol (SCTP) [RFC4960], and DNS Security (DNSSEC) [RFC4034]. The impact of DNSSEC is of particular interest, because it is relied upon for the delivery of other services, such as DNS-Based Authentication of Named Entities (DANE) [RFC6698], and it could be used for application discovery services through DNS (specifically where security properties are part of that discovery). Thus, slowdown at the neck of the glass can have an impact closer to the lip.

Even when one considers 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.



HTTP(s) as the new neck?

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.

The workshop moved on to a review of RFC 5218, which discusses protocol success factors. This work was presented in the IETF 70 plenary and was the basis for this ongoing work. There were two clear outcomes from the discussion. The first was that the Internet Architecture Board 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 presented 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 explores how bundling of two technologies may lead to increased or decreased adoption of one or both. This will depend on a number of factors, including costs, benefits, and externalities associated with each technology. (Simply put, an externality is an effect or use decision by one set of parties that has either a positive or negative impact on others who did not have a choice or whose interests were not taken into account.) Bundling of capabilities may provide positive value when individual capabilities on their own do not provide sufficient critical mass to propel further adoption. Specifically, bundling can help when one technology does not provide positive value until critical mass of deployment exists, and where a second technology has low adoption cost and immediate value and hence drives initial adoption until enough of a user base exists to allow critical mass sufficient for the first technology to get positive value. 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

The workshop considered an examination of protocol success factors in the context of Bitcoin [Boehme13]. Here, there were any number of barriers to success, including adverse press, legal uncertainties, glitches and breaches, previous failed attempts, and speculative attacks. Bitcoin has thus far 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, e.g., people find intrinsic value in it.

The first two of these factors may be transferable to other approaches. One key protocol success factor is direct benefit to the participant. Another key protocol success factor is the ability to interface with other systems for mutual benefit. In the context of Bitcoin, there has to be a way to exchange the coins for other currencies. The Internet email system had simpler adaption mechanisms to allow interchange with non-Internet email systems; this facilitated its success. Another more simply stated approach is "IP over everything".

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 happens, there is a limited example that is directly relevant to the IETF.

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

The workshop reviewed the approach Sweden's .SE registry has taken to improving deployment of DNSSEC [Lowinder13]. .SE has roughly 1.5 million domains. IIS (<<https://www.iis.se>>) manages the ccTLD (Country Code Top Level Domain). They made the decision to encourage deployment of DNSSEC within .SE. They began by understanding what the full ecosystem looked like, who their stakeholders were, and the 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. The workshop considered an interactive approach that combines value net analysis, deployment environment analysis, and technical architecture analysis that leads to feasibility and solution analysis [Leval13]. This work provided an alternative to RFC 5218 that had many points in common. The case study examined was that of Multipath TCP (MPTCP). Various deployment challenges were observed. First and foremost, increasing bandwidth within the network seems to decrease the 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 several approaches to improve deployment, including using open-source software, lobbying various implementers, deploying proxies, and completing implementations by parties that 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? The workshop considered an approach that became a recurring theme throughout the workshop -- to not 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 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)

The workshop reviewed a comparison between the hourglass model and what systems biologists might call the bow tie 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, adenosine triphosphate (ATP) is the means by which all organisms convert nutrients to usable energy, and thus resides centrally within the biological system.

The workshop 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 prove fragile in the face of failure or may prove 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

The workshop considered the difference between planned versus unplanned technology transitions [Kohn13]. 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 view as to the approach.

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

The workshop 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 was 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 towards standardization is the result (e.g., they get to decide the venue for their work).

5.2. Bridge Networking Research and Internet Standardization: Case Study on Mobile Traffic Offloading and IPv6 Transition Technologies

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 (Third Generation Partnership Project) 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 tack 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

The workshop engaged in a very provocative discussion about whether the existing Internet architecture serves the broadest set of needs. Three specific aspects were examined: geographic, technical, and socioeconomic. 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

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

6.1. Resilience of the commons: routing security

The workshop 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 by 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 the IETF 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, reviewing 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, at the time of writing, there is much interest 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 are work on bundling and 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 Regional Internet Registries (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 of 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 (Internet-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; it would also provide 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, as well as Bernard Aboba and Dave Thaler. Their earlier work provided a strong basis for this workshop.

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 Robachevsky, Andrew Sullivan, Arjuna Sathiaselan, 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, Miya Kohno, Rainer Boehme, Richard Clayton, Roch Guerin, Ross Anderson, Russ Housley, Sam Smith, Sean Turner, Soumya Sen, Spencer Dawkins, Steven Weber, Tapio Leva, Toby Moncaster, Tony Finch

11. Informative References

- [Boehme13] Boehme, R., "Internet Protocol Adoption: Learning from Bitcoin", December 2013, <http://www.iab.org/wp-content/IAB-uploads/2013/06/itat-2013_submission_17.pdf>.
- [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, <http://www.iab.org/wp-content/IAB-uploads/2013/06/itat-2013_submission_6.pdf>.
- [Kohno13] Kohno, M., Asaba, T., and F. Baker, "Managing Diversity to Manage Technological Transition", December 2013, <http://www.iab.org/wp-content/IAB-uploads/2013/06/itat-2013_submission_7.pdf>.
- [Lear13] Lear, E. and D. Mohlenhoff, "Standards: a love/hate relationship with patents", December 2013, <http://www.iab.org/wp-content/IAB-uploads/2013/06/itat-2013_submission_11.docx>.

- [Leval13] Leva, T. and H. Soumi, "Framework for analyzing feasibility of Internet protocols", December 2013, <http://www.iab.org/wp-content/IAB-uploads/2013/06/itat-2013_submission_4.pdf>.
- [Lowinder13] Eklund Lowinder, A. and P. Wallstrom, "Long term strategy for a successful deployment of DNSSEC - on all levels", December 2013, <http://www.iab.org/wp-content/IAB-uploads/2013/06/itat-2013_submission_5.docx>.
- [Meyer13] Meyer, D., "On the Complexity of Engineered Systems (and its effect on protocol deployment)", December 2013, <http://www.iab.org/wp-content/IAB-uploads/2013/06/itat-2013_submission_9.pdf>.
- [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.
- [RFC5218] Thaler, D. and B. Aboba, "What Makes For a Successful Protocol?", RFC 5218, July 2008.
- [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, <http://www.iab.org/wp-content/IAB-uploads/2013/06/itat-2013_submission_12.pdf>.
- [Sathiaseelan13] Sathiaseelan, A., Trossen, D., Komnios, I., Ott, J., and J. Crowcroft, "An Internet Architecture for the Challenged", December 2013, <http://www.iab.org/wp-content/IAB-uploads/2013/06/itat-2013_submission_3.pdf>.

- [Sen13] Sen, S., "On Economic Models of Network Technology Adoption, Design, and Viability", December 2013, <http://www.iab.org/wp-content/IAB-uploads/2013/06/itat-2013_submission_101.pdf>.
- [Weber13] Weber, S., Guerin, R., and J. Oliveira, "When can bundling help adoption of network technologies or services?", December 2013, <http://www.iab.org/wp-content/IAB-uploads/2013/06/itat-2013_submission_2.pdf>.
- [Welzl13] Welzl, M., "The "best effort" service as a deployment success factor", December 2013, <http://www.iab.org/wp-content/IAB-uploads/2013/06/itat-2013_submission_8.pdf>.

Author's Address

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

Phone: +41 44 878 9200
EMail: lear@cisco.com

