

Network Working Group
Internet-Draft
Intended status: Informational
Expires: May 6, 2021

J. Arkko
Ericsson
T. Hardie
Google
November 2, 2020

**Report from the IAB workshop on Design Expectations vs. Deployment
Reality in Protocol Development
draft-iab-dedr-report-01**

Abstract

The Design Expectations vs. Deployment Reality in Protocol Development Workshop was convened by the Internet Architecture Board (IAB) in June 2019. This report summarizes its significant points of discussion and identifies topics that may warrant further consideration.

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 <https://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 May 6, 2021.

Copyright Notice

Copyright (c) 2020 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 (<https://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	2
2.	Workshop Agenda	4
3.	Position Papers	4
4.	Discussions	6
4.1.	Past experiences	6
4.2.	Principles	7
4.3.	Centralised deployment models	8
4.4.	Security	10
4.5.	Future	10
5.	Conclusions	11
5.1.	Summary of discussions	11
5.2.	Actions	12
5.2.1.	Potential architecture actions and outputs	13
5.2.2.	Potential other actions	13
5.3.	Other publications	13
5.4.	Feedback	13
6.	Informative References	13
Appendix A.	Particant List	17
Appendix B.	Acknowledgements	18
	Authors' Addresses	18

[1.](#) Introduction

The Design Expectations vs. Deployment Reality in Protocol Development Workshop was convened by the Internet Architecture Board (IAB) in June 2019. This report summarizes its significant points of discussion and identifies topics that may warrant further consideration.

The background for the workshop was that during the development and early elaboration phase for a number of protocols, there was a presumption of specific deployment models. Actual deployments have, however, often run contrary to these early expectations when economies of scale, Distributed Denial-of-Service (DDoS) attack resilience, market consolidation, or other factors have come into play. These factors can result in the deployed reality being highly concentrated.

This is a serious issue for the Internet, as concentrated, centralized deployment models present risks to user choice, privacy, and future protocol evolution.

On occasion, the differences to expectations were almost immediate, but they also occur after a significant time has passed from the protocol's initial development.

Examples include:

Email standards, which presumed many providers running in a largely uncoordinated fashion, but which has seen both significant market consolidation and a need for coordination to defend against spam and other attacks. The coordination and centralized defense mechanisms scale better for large entities, which has fueled additional consolidation.

The DNS, which presumed deep hierarchies but has often been deployed in large, flat zones, leading to the nameservers for those zones becoming critical infrastructure. Future developments in DNS may see concentration through the use of globally available common resolver services, which evolve rapidly and can offer better security. Paradoxically, concentration of these queries into few services creates new security and privacy concerns.

The Web, which is built on a fundamentally decentralized design, but which is now often delivered with the aid of Content Delivery Networks. Their services provide scaling, distribution, and Denial of Service prevention in ways that new entrants and smaller systems operators would find difficult to replicate. While truly small services and truly large ones may operate using only their own infrastructure, many others are left with the only practical choice being the use of a globally available commercial service.

Similar developments may happen with future technologies and services. For instance, the growing use of Machine Learning technology presents challenges for distributing effective implementation of a service throughout a pool of many different providers.

In [[RFC5218](#)] the IAB tackled what made for a successful protocol. In [[RFC8170](#)], the IAB described how to handle protocol transitions. This purpose of the workshop was to explore cases where the initial system design assumptions turned out to be wrong, looking for patterns in what caused those assumptions to fail (e.g., concentration due to DDoS resilience) and in how those failures impact the security, privacy, and manageability of the resulting deployments.

While the eventual goals might include proposing common remediations for specific cases of confounded protocol expectations.

The workshop call for papers invited the submission of position papers which would:

- o Describe specific cases where systems assumptions during protocol development were confounded by later deployment conditions.
- o Survey a set of cases to identify common factors in these confounded expectations.
- o Explore remediations which foster user privacy, security and provider diversity in the face of these changes.

A total of 21 position papers were received, listed in [Section 3](#). On site or remote were 30 participants, listed in [Appendix A](#).

2. Workshop Agenda

After opening and discussion of goals for the workshop, the discussion focused on five main topics:

- o Past experiences. What have we learned?
- o Principles. What forces apply to deployment? What principles to take into account in design?
- o Centralised deployment models. The good and the bad of centralisation. Can centralisation be avoided? How?
- o Security. Are we addressing the right threats? What should we prepare ourselves for?
- o Future. What can we do? Should we get better at predicting, or should we do different things?

3. Position Papers

The following position papers were submitted to the workshop (in alphabetical order):

- o Jari Arkko. "Changes in the Internet Threat Model" [[Arkko2019](#)]
- o Vittorio Bertola. "How the Internet Was Won and Where It Got Us" [[Bertola2019](#)]
- o Carsten Bormann. "WiFi authentication: Some deployment observations from eduroam" [[Bormann2019](#)]

- o Stephane Bortzmeyer. "Encouraging better deployments" [[Bortzmeyer2019](#)]
- o Brian Carpenter and Bing Liu. "Limited Domains and Internet Protocols" [[Carpenter2019](#)]
- o Alissa Cooper. "Don't Forget the Access Network" [[Cooper2019](#)]
- o Stephen Farrell. "We're gonna need a bigger threat model" [[Farrell2019](#)]
- o Phillip Hallam-Baker. "The Devil is in the Deployment" [[HallamBaker2019](#)]
- o Ted Hardie. "Instant Messaging and Presence: A Cautionary Tale" [[Hardie2019](#)]
- o Paul Hoffman. "Realities in DNSSEC Deployment" [[Hoffman2019](#)]
- o Christian Huitema. "Concentration is a business model" [[Huitema2019](#)]
- o Geoff Huston. "The Border Gateway Protocol, 25 years on" [[Huston2019](#)]
- o Dirk Kutscher. "Great Expectations: Protocol Design and Socioeconomic Realities" [[Kutscher2019](#)]
- o Julien Maisonnewe. "DNS, side effects and concentration" [[Maisonnewe2019](#)]
- o John Mattsson. "Privacy, Jurisdiction, and the Health of the Internet" [[Mattsson2019](#)]
- o Moritz Muller. "Rolling Forward: An Outlook on Future Root Rollovers" [[Muller2019](#)]
- o Joerg Ott. "Protocol Design Assumptions and PEPs" [[Ott2019](#)]
- o Lucas Pardue. "Some challenges with IP multicast deployment" [[Pardue2019](#)]
- o Jim Reid. "Where/Why has DNS gone wrong?" [[Reid2019](#)]
- o Mohit Sethi and Tuomas Aura. "IoT Security and the role of Manufacturers: A Story of Unrealistic Design Expectations" [[Sethi2019](#)]

- o Andrew Sullivan. "Three kinds of concentration in open protocols" [[Sullivan2019](#)]

These papers are available from the IAB website [[CFP](#)] [[POS](#)].

4. Discussions

4.1. Past experiences

The workshop investigated deployment cases from certificate authorities for web connections (WebPKI) to DNS Security (DNSSEC), from Border Gateway Protocol (BGP) to Network Address Translators (NATs), from Domain Name System (DNS) resolvers to Content Data Networks (CDNs), and from Internet of Things (IoT) systems to instant messaging and social media applications.

In many cases there was either a surprise in how technology was deployed, lack of sufficient adoption, or the business models associated with chosen technologies were not in favor of broader interoperability.

In general, the protocol designers cannot affect market forces but must work within them. But there are often competing technical approaches or features that are tailored for a particular deployment pattern. In some cases it is possible to choose whether to support, for instance, a clear need for an established business, a feature designed to support collaboration among smaller players, or some kind of disruption through a more speculative new feature or technology.

Lessons learned include:

- o Feedback from those who deploy often comes too late.
- o Building blocks get re-purposed in unexpected ways
- o User communities come in too late.
- o The web is getting more centralised, and counteracting this trend is difficult. It is not necessarily clear what technical path leads to distributed markets and de-centralized architectures, for instance.
- o There are also many forces that make it easier to pursue centralised models than other ones. For instance, deployment is often easier in a centralised model. And various business and regulatory processes work best within a small, well-defined set of entities that can interact with each other. This can lead to, for instance, regulators preferring a situation with a small number of

entities that they can talk to, rather than a diverse set of providers.

- o It is important but hard to determine how useful new protocols are.
- o It is difficult for the IETF community to interact with others, e.g., specific business sectors that need new technology (such as aviation or healthcare) or regulators.

4.2. Principles

Several underlying principles can be observed in the example cases that were discussed. Deployment failures tend to be associated with cases where interdependencies make progress difficult and there's no major advantage for early deployment. Despite persistent problems in the currently used technology, it becomes difficult for the ecosystem to switch to better technology. For instance, there are a number of areas where the Internet routing protocol, BGP [[RFC4271](#)], is lacking, but there has been only limited success in deploying significant improvements, for instance in the area of security.

Another principle appears to be first mover advantage. Several equally interesting technologies have fared in very different ways, depending whether there was an earlier system that provided most of the benefits of the new system. Again, despite potential problems in an already deployed technology, it becomes difficult to deploy improvements due to lack of immediate incentives and due to the competing and already deployed alternative that is proceeding forward in the ecosystem. For instance, WebPKI is very widely deployed and used, but DNSSEC ([RFC4033](#)) is not. Is this because the earlier commercial adoption of WebPKI, the more complex interdependencies between systems that wished to deploy DNSSEC, or some other reason?

The definition of success in [[RFC5218](#)] appears to a part of the problem. The only way to control deployments up front is to prevent wild success, but wild successes are actually what we want. And it seems very difficult to predict these successes.

The workshop also discussed the extent to which protocol work even should be controlled by the IETF, or the IESG. It seems unproductive to attempt to constrain deployment models, as one can only offer possibilities but not force anyone to use a particular possibility.

The workshop also discussed different types of deployment patterns on the Internet:

- o Delivering functionality over Internet as a web service. The Internet is an open and standardised system, but the service on top may be closed, essentially running two components of the same service provider's software against each other over the browser and Internet infrastructure. Several large application systems have grown in the Internet in this manner, encompassing large amounts of functionality and a large fraction of Internet users. This makes it easier for web applications to grow by themselves without cross-fertilisation or interoperability.
- o Delivering concentrated network services that offer the standard capabilities of the Internet. Examples in this category include the provisioning of some mail services, DNS resolution, and so on.

The second case is more interesting for an Internet architecture discussion. There can, however, be different underlying situations even in that case. The service may be simply a concentrated way to provide a commodity service. The market should find a natural equilibrium for such situations. This may be fine, particularly, where the service does not provide any new underlying advantage to whoever is providing it (in the form of user data that can be commercialized, for instance, or as training data for an important machine learning service).

Secondly, the service may be an extension beyond standard protocols, leading to some questions about how well standards and user expectations match. But those questions could be addressed by better or newer standards. But the third situation is more troubling: the service are provided in this concentrated manner due to business patterns that make it easier for particular entities to deploy such services.

The group also discussed monocultures, and their negative effect on the Internet and its stability and resistance to various problems and attacks.

Regulation may affect the Internet businesses as well. Regulation can exist in multiple forms, based on economic rationale (e.g., competition law) or other factors. For instance, user privacy is a common regulatory topic.

4.3. Centralised deployment models

Many of the participants have struggled with these trends and their effect on desirable characteristics of Internet systems, such as distributed, end-to-end architecture or privacy. Yet, there are many business and technical drivers causing the Internet architecture to become further and further centralised.

Some observations that were made:

- o When standardising new technology, the parties involved in the effort may think they agree on what the goals are, but often in reality are surprised in the end. For instance, with DNS-over-HTTPS (DoH, [RFC8484](#)) there were very different aspirations, some around improvements in confidentiality of the queries, some around operational and latency improvements to DNS operations, and some about shifting business and deployment models. The full picture was not clear before the work was completed.
- o In DNS, DDoS is practical reality, and only a handful of providers can handle the traffic load in these attacks.

The hopeful side of this issue is that there are some potential answers:

- o DDoS defenses do not have to come through large entities, as layered defenses and federation also helps similarly.
- o Surveillance state data capture can be fought with data object encryption, and not storing all of the data in one place.
- o Web tracking can be combatted by browsers choosing to avoid the techniques sensitive to tracking. Competition in the browser market may help drive some of these changes.
- o Open interfaces help guard against the bundling of services in one large entity; as long as there are open, well-defined interface to specific functions these functions can also be performed by other parties.
- o Commercial surveillance does not seem to be curbed by current means. But there are still possibilities, such as stronger regulation, data minimisation, or browsers acting on behalf of users. There are hopeful signs that at least some browsers are becoming more aggressive in this regard. But more is needed.

One comment made in the workshop that the Internet community needs to curb the architectural trend of centralization. Another comment was that discussing this in the abstract is not as useful as more concrete, practical actions. For instance, one might imagine different DoH deployments with widely different implications for privacy or tolerance of failures. Getting to the specifics of how a particular service can be made better is important.

4.4. Security

This part of the discussed focused on whether in the current state of the Internet we actually need a new threat model.

Many of the communications security concerns have been addressed in the past few years, with increasing encryption. However, issues with trusting endpoints on the other side of the communication have not been addressed, and are becoming more urgent with the advent or centralised service architectures.

Further effort may be needed to minimise centralisation, as having only few places to tap increases the likelihood of surveillance.

There may be a need to update [[RFC3552](#)] and [[RFC7258](#)].

The participants in the workshop agreed that a new threat model is needed, and that non-communications-security issues need to be treated.

Other security discussions were focused on IOT systems, algorithm agility issues, experiences from difficult security upgrades such as the DNSSEC key rollover, and routing security.

The participants cautioned against relying too much on device manufacturers for security, and being clear on security models and assumptions. Security is often poorly understood, and the assumptions about who the system defends against and not are not clear.

4.5. Future

The workshop turned into a discussion of what actions we can take:

- o Documenting our experiences?
- o Providing advice (to IETF or to others)
- o Waiting for the catastrophe that will make people agree to changes? The participants of course did not wish for this.
- o Work at the IETF?
- o Technical solutions/choices?

The best way for the IETF to do things is through standards; convincing people through other requests is difficult. The IETF needs to:

- o Pick pieces that it is responsible for.
- o Be reactive for the rest, be available as an expert in other discussions, provide Internet technology clue where needed, etc.

One key question is what other parties need to be involved in any discussions. Platform developers (mobile platforms, cloud systems, etc.) is one such group. Specific technology or business groups (such as email provider or certificate authority forums) are another.

The workshop also discussed specific technology issues, for instance around IOT systems. One observation in those systems is that there is no single model for applications, they vary. There are a lot of different constraints in different systems and different control points. What is needed perhaps most today is user control and transparency (for instance, via Manufacturer Usage Descriptions (MUDs, [RFC8520](#))). Another issue is management, particularly for devices that could be operational for decades. Given the diversity of IOT systems, it may also make more sense to build support systems for the broader solutions that specific solutions or specific protocols.

There are also many security issues. While some of them are trivial (such as default passwords), one should also look forward and be prepared to have solutions for, say, trust management for long time scales, or be able to provide data minimization to cut down on potential for leakages. And the difficulty of establishing peer-to-peer security strengthens the need for a central point, which may also be harmful from a long-term privacy perspective.

5. Conclusions

5.1. Summary of discussions

The workshop met in sunny Finnish countryside and made the non-surprising observation that technologies sometimes get deployed in surprising ways. But the consequences of deployment choices can have an impact on security, privacy, centralised vs. distributed models, competition, surveillance. As the IETF community cares deeply about these aspects, it is worthwhile to spend time in analysis of these choices.

The prime factor driving deployments is perceived needs; expecting people to recognise obvious virtues and therefore deploy is not likely to work.

And the ecosystem is complex, including for instance many parties: different business roles, users, regulators, and so on, and

perceptions of needs and ability to act depends highly on what party one talks to.

While the workshop discussed actions and advice, there is a critical question of who these are targeted towards. There is need to construct a map of what parties need to perform what actions.

The workshop also made some technical observations. One issue is that the workshop identified a set of hard issues that affect deployment and for which have no good solutions. These issues include, for instance, dealing with distributed denial-of-service attacks or how to handle spam. Similarly, lack of good solutions for micropayments is one factor behind a lot of the Internet economy being based on advertisements.

One recent trend is that technology is moving up the stack, e.g., in the areas of services, transport protocol functionality, security, naming, and so on. This impacts how easy or hard changes are, and who is able to perform them.

It was also noted that interoperability continues to be important, and we need to explore what new interfaces need standardisation -- this will enable different deployment models & competition. Prime factor driving deployments is actual needs; we cannot force anything to others but can provide solutions for those that need them. Needs and actions may fall on different parties.

The workshop also considered the balancing of user non-involvement and transparency and choice, relevant threats such as communicating with malicious endpoints, the role and willingness of browsers in increasing the ability to defending the users' privacy, and concerns around centralised control or data storage points

The workshop also discussed specific issues around routing, denial-of-service attacks, IOT systems, role of device manufacturers, the DNS, and regulatory reactions and their possible consequences.

5.2. Actions

The prime conclusion from the workshop was that the topic is not completed in the workshop. Much more work is needed. The best way for the IETF to make an impact is through standards. The IETF should focus on the parts that it is responsible for, and be available as an expert on other discussions.

5.2.1. Potential architecture actions and outputs

The documents/outputs and actions described in the following were deemed relevant by the participants.

- o Develop and document a modern threat model.
- o Continue discussion of consolidation/centralisation issues.
- o Document architectural principles, e.g., (re-)application of the end-to-end principle.

The first receiver of these thoughts is the IETF and protocol community, but combined with some evangelising & validation elsewhere.

5.2.2. Potential other actions

- o Pursue specific IETF topics, e.g., work on taking into account reputation systems in IETF work, working to ensuring certificate scoping can be appropriately limited, building end-to-end encryption tools for applications, etc.
- o General deployment experiences/advice, and documenting deployment assumptions possibly already in WG charters
- o A report, and a short summary article will be produced from the workshop.

5.3. Other publications

The workshop results have also been reported at [[ISPColumn](#)] by Geoff Huston.

5.4. Feedback

Feedback regarding the workshop is appreciated, and can be sent to the program committee, the IAB, or the architecture-discuss list.

6. Informative References

[Arkko2019]

Arkko, J., "Changes in the Internet Threat Model",
Position paper submitted for the IAB DEDR workshop , June
2019.

[Bertola2019]

Bertola, V., "How the Internet Was Won and Where It Got Us", Position paper submitted for the IAB DEDR workshop , June 2019.

[Bormann2019]

Bormann, C., "WiFi authentication: Some deployment observations from eduroam", Position paper submitted for the IAB DEDR workshop , June 2019.

[Bortzmeyer2019]

Bortzmeyer, S., "Encouraging better deployments", Position paper submitted for the IAB DEDR workshop , June 2019.

[Carpenter2019]

Carpenter, B. and B. Liu, "Limited Domains and Internet Protocols", Position paper submitted for the IAB DEDR workshop , June 2019.

[CFP]

IAB, ., "Design Expectations vs. Deployment Reality in Protocol Development Workshop 2019",
<https://www.iab.org/activities/workshops/dedr-workshop/> , April 2019.

[Cooper2019]

Cooper, A., "Don't Forget the Access Network", Position paper submitted for the IAB DEDR workshop , June 2019.

[Farrell2019]

Farrell, S., "We're gonna need a bigger threat model", Position paper submitted for the IAB DEDR workshop , June 2019.

[HallamBaker2019]

Hallam-Baker, P., "The Devil is in the Deployment", Position paper submitted for the IAB DEDR workshop , June 2019.

[Hardie2019]

Hardie, T., "Instant Messaging and Presence: A Cautionary Tale", Position paper submitted for the IAB DEDR workshop , June 2019.

[Hoffman2019]

Hoffman, P., "Realities in DNSSEC Deployment", Position paper submitted for the IAB DEDR workshop , June 2019.

[Huitema2019]

Huitema, C., "Concentration is a business model", Position paper submitted for the IAB DEDR workshop , June 2019.

[Huston2019]

Huston, G., "The Border Gateway Protocol, 25 years on", Position paper submitted for the IAB DEDR workshop , June 2019.

[ISPColumn]

Huston, G., "Network Protocols and their Use", <https://www.potaroo.net/ispcol/2019-06/dedr.html> , June 2019.

[Kutscher2019]

Kutscher, D., "Great Expectations: Protocol Design and Socioeconomic Realities", Position paper submitted for the IAB DEDR workshop , June 2019.

[Maisonneuve2019]

Maisonneuve, J., "DNS, side effects and concentration", Position paper submitted for the IAB DEDR workshop , June 2019.

[Mattsson2019]

Mattsson, J., "Privacy, Jurisdiction, and the Health of the Internet", Position paper submitted for the IAB DEDR workshop , June 2019.

[Muller2019]

Muller, M., "Rolling Forward: An Outlook on Future Root Rollovers", Position paper submitted for the IAB DEDR workshop , June 2019.

[Ott2019] Ott, J., "Protocol Design Assumptions and PEPs", Position paper submitted for the IAB DEDR workshop , June 2019.

[Pardue2019]

Pardue, L., "Some challenges with IP multicast deployment", Position paper submitted for the IAB DEDR workshop , June 2019.

[POS]

IAB, ., "Position Papers: DEDR Workshop", <https://www.iab.org/activities/workshops/dedr-workshop/position-papers/> , June 2019.

[Reid2019]

Reid, J., "Where/Why has DNS gone wrong?", Position paper submitted for the IAB DEDR workshop , June 2019.

[RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", [BCP 72](#), [RFC 3552](#), DOI 10.17487/RFC3552, July 2003, <<https://www.rfc-editor.org/info/rfc3552>>.

[RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", [RFC 4033](#), DOI 10.17487/RFC4033, March 2005, <<https://www.rfc-editor.org/info/rfc4033>>.

[RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", [RFC 4271](#), DOI 10.17487/RFC4271, January 2006, <<https://www.rfc-editor.org/info/rfc4271>>.

[RFC5218] Thaler, D. and B. Aboba, "What Makes for a Successful Protocol?", [RFC 5218](#), DOI 10.17487/RFC5218, July 2008, <<https://www.rfc-editor.org/info/rfc5218>>.

[RFC7258] Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an Attack", [BCP 188](#), [RFC 7258](#), DOI 10.17487/RFC7258, May 2014, <<https://www.rfc-editor.org/info/rfc7258>>.

[RFC8170] Thaler, D., Ed., "Planning for Protocol Adoption and Subsequent Transitions", [RFC 8170](#), DOI 10.17487/RFC8170, May 2017, <<https://www.rfc-editor.org/info/rfc8170>>.

[RFC8484] Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", [RFC 8484](#), DOI 10.17487/RFC8484, October 2018, <<https://www.rfc-editor.org/info/rfc8484>>.

[RFC8520] Lear, E., Droms, R., and D. Romascanu, "Manufacturer Usage Description Specification", [RFC 8520](#), DOI 10.17487/RFC8520, March 2019, <<https://www.rfc-editor.org/info/rfc8520>>.

[Sethi2019]

Sethi, M. and T. Aura, "IoT Security and the role of Manufacturers: A Story of Unrealistic Design Expectations", Position paper submitted for the IAB DEDR workshop , June 2019.

[Sullivan2019]

Sullivan, A., "Three kinds of concentration in open protocols", Position paper submitted for the IAB DEDR workshop , June 2019.

Appendix A. Particant List

The following is a list of participants on site and over a remote connection:

- o Arkko, Jari
- o Aura, Tuomas
- o Bertola, Vittorio
- o Bormann, Carsten
- o Bortzmeyer, Stephane
- o Cooper, Alissa
- o Farrell, Stephen
- o Flinck, Hannu
- o Gahnberg, Carl
- o Hallam-Baker, Phillip
- o Hardie, Ted
- o Hoffman, Paul
- o Huitema, Christian (remote)
- o Huston, Geoff
- o Komaitis, Konstantinos
- o Kuhlewind, Mirja
- o Kutscher, Dirk
- o Li, Zhenbin
- o Maisonneuve, Julien

- o Mattson, John
- o Muller, Moritz
- o Ott, Joerg
- o Pardue, Lucas
- o Reid, Jim
- o Rieckers, Jan-Frederik
- o Sethi, Mohit
- o Shore, Melinda (remote)
- o Soininen, Jonne
- o Sullivan, Andrew
- o Trammell, Brian

[Appendix B](#). Acknowledgements

The authors would like to thank the workshop participants, the members of the IAB, and the participants in the architecture discussion list for interesting discussions. The notes from Jim Reid were instrumental in writing this report. The workshop organizers would also like to thank Nokia for hosting the workshop in excellent facilities in Kirkkonummi, Finland.

Authors' Addresses

Jari Arkko
Ericsson

Email: jari.arkko@piuha.net

Ted Hardie
Google

Email: ted.ietf@gmail.com

