



HAL
open science

The Internet Engineering Task Force and The Future of Internet

Emmanuel Baccelli, Thomas Heide Clausen, Philippe Jacquet

► **To cite this version:**

Emmanuel Baccelli, Thomas Heide Clausen, Philippe Jacquet. The Internet Engineering Task Force and The Future of Internet. ERCIM News, 2009, pp.20 - 21. hal-00651466

HAL Id: hal-00651466

<https://hal.science/hal-00651466>

Submitted on 14 Dec 2011

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The Internet Engineering Task Force and The Future of Internet

E. Baccelli¹, T. Clausen¹, and P. Jacquet¹

¹INRIA HiPERCOM Project, www.inria.fr

If one wants to identify where ideas and initiatives regarding the Internet are being confronted, the *Internet Engineering Task Force* (IETF [1]) is unavoidable. Created in 1986 by US government agencies (DoD, Department of Energy, NASA, NSF) to supervise the design and deployment of Internet protocols, it was initially open only for US government funded researchers. Early 1987 saw a dozen of industry representatives invited, and in a matter of months, the IETF was opened to anyone interested. In 2008, IETF meetings were attended by roughly 1300 engineers and researchers from all over the world.

The IETF is an R&D forum where network engineers define, describe, review and discuss network protocols – published as Request For Comments (RFC) – which industry then may, or may not, implement and use. IETF meetings are tri-annual, with business in between these meetings being conducted on open mailing-lists.

1 The Organizational Structure of the IETF

Work within the IETF is organized into *working groups* (WG), each of which in charge of a specific problem (e.g. mobile ad hoc routing). Typically, a WG is supervised by two *chairs*.

WGs within the same general field are assembled in a so-called *area* (e.g. the routing area). Each area is supervised by two *area directors* (AD), whose task is to shepherd the creation, activity and eventual demise of WGs in the area. In early 2009, the IETF had 8 areas and 120 working groups.

The assembly of area directors forms the *Internet Engineering Steering Group* (IESG). The IESG, together with the *Internet Architecture Board* (IAB), ensure the overall coherence of the Internet protocols "corpus". IESG and IAB members are periodically replaced, potentially by any other competent IETF

participant.

2 The IETF and Decision Making

The IETF differs from standardization bodies such as IEEE or ETSI in that in the IETF, individuals represent themselves. There is no *de facto* company representation. People from the same company/institution may make conflicting contributions, while people from different companies/institutions may contribute together to a standard without the necessity of a formal agreement. Proposals must be open for other potential contributors without any copyright restrictions. Moreover, the IETF's fundamental motto is: "*We reject kings, presidents and voting. We believe in rough consensus and running code.*"

Rough Consensus Rather than voting (as in the IEEE or ETSI) decisions in the IETF are made based on "rough consensus" in a WG on one hand, as gauged by the WG chairs, and in the IETF as a whole on the other hand, as gauged by the IESG. Well understood by IETF participants, this procedure allows any good idea from any origin to be discussed, bringing contributions from individuals and small institutions on equal footing with those from big companies.

Working Code A proposed protocol can generally not be promoted as potential standard without thorough experimentation (protocols can be experimented via working code without much hardware investment, often none). Furthermore, to avoid artifacts due to internal bugs, several working code bases developed independently following the proposed specification must demonstrate their full compatibility before the standard can be validated.

3 The Pertinence of the IETF

The ability of an R&D forum to meet the positive evolution of a technology depends on how it manages the four following parameters: vision, legacy, luck and necessity.

Vision The IETF clearly has the right focus. While its vision is fuzzy, since initiatives generally come from the bottom, top level directions are very clear. Currently, for instance: mobility, scalability to encompass the Internet of objects, or IPv6. Introduced in the 1990s to address the scarcity of available addresses with IPv4 (four bytes format), IPv6 upgrades IP to a flexible address management scheme over 16 bytes, potentially identifying 10^{38} elements. While

transition from IPv4 to IPv6 is slower than expected due to the generalization of CIDR, NAT, and DHCP, experts predict the allocation of the last IPv4 address to take place in 2010.

Legacy An R&D forum is the meeting place for dreams and possibilities. However, one may present the most brilliant idea in the world in vain, if it is incompatible with existing technology — i.e. "*A good idea is not always a good idea.*" Nevertheless, the IETF is very careful not to bypass any innovative idea, and manages to this end a parallel forum, called the *Internet Research Task Force* (IRTF), where new paradigms (e.g. Delay Tolerant Networking) are trained to fit legacy.

Luck The most important issue in an R&D forum is the ability to manage an unexpected breakthrough. With a culture of ideas beginning at the bottom, even the most crazy idea is welcomed – if it fits legacy and addresses a concrete issue. I.e. "*A good idea can become an extremely good idea.*" A striking example is TCP. In the late 1980's, the challenge was to cope with brutal capacity reduction when data traffic had to cross long-haul networks. Failing to address this issue caused the demise of a concurrent system — ATM. The IETF produced a surprisingly simple, but innovative, solution: with TCP, a source terminal tunes the file transmission pace according to feedback from the destination terminal. Experts consider the strength of TCP (supporting variations of network capacity ranging more than twelve orders of magnitude) as the main reason for the success of the Internet.

Necessity The IETF mandates itself to solve certain problems. For example in the late 1980's, the current routing protocol RIP was detected to fail when a set of routers was brutally removed from the network. This bug, called "*count to infinity*", created a sustained loop which caused an avalanche of disruptions – the Internet was down for two full days. A failure for a system designed with resilience as its core tenant! RIP had to be replaced by a new protocol, specified in emergency: *Open Shortest Path First* (OSPF), widely used nowadays. Less elegant than RIP, OSPF is far more robust, based on an exhaustive mapping of network links, thus letting routers compute new routes and react in real time to disruptive topology changes.

4 The IETF and the Future

The IETF was the birthplace of the Internet of today. Understanding its activities is part of the necessary tasks for individuals and institutions, aiming at anticipating the future of the Internet. As it appears, this necessity is not likely to fade any time soon.

References

- [1] The Internet Engineering Task Force (IETF), www.ietf.org