will help consolidate debate in areas where consensus is beginning to emerge, while fostering and moderating it in areas where disagreement prevails.

Careful consideration of these issues has led us to consider a Wikipedia-like online platform as the first step for this debate. This platform is charged with gathering opinions and consolidating the debate of a growing community of researchers, and will be a repository for relevant work in areas that are critical to the making of the future Internet. The platform will most importantly serve as a tool with which to formulate, debate and finally resolve disagreements. The use of collaborative editing tools like wikis will allow this debate to be recorded and structured. It is important to note that this online platform will differ from a conventional Wikipedia-like encyclopedia in having a clearly defined editorial team, soliciting contributions from identified (registered) users (while being open to be read by anybody), and allowing individual contributors to retain the copyright of their contributions.

At the time of writing, an initial team is in the process of setting up such a platform. It is expected to launch during the spring of 2009, seeded with results from the initial EIFFEL think tank meetings but also incorporating material from experts worldwide. While a particular timeframe will not be enforced, we do expect the debate to pick up quickly, leading to early results.

The positions presented in this paper are the outcome of discussions among EIFFEL think tank members and therefore includes the viewpoints of many individuals involved in these discussions.

Links:
http://www.fp7-eiffel.eu
US initiatives FIND and GENI:
http://www.nets-find.net/
http://www.geni.net

Please contact:
Dirk Trossen
BT Research, UK
E-mail: dirk.trossen@bt.com

The Internet Engineering Task Force and the Future of the Internet

by Emmanuel Baccelli, Thomas H. Clausen and Philippe Jacquet

The Internet Engineering Task Force was the birthplace of today's Internet. Understanding its activities is necessary for individuals and institutions who wish to anticipate the future of the Internet. As things stand, this necessity is not likely to fade any time soon.

For those with a stake in the ideas and initiatives that will drive the Internet in the future, the Internet Engineering Task Force (IETF) 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 to US government-funded researchers. Early 1987 saw a dozen industry representatives invited, and in a matter of months, the IETF was opened to all interested parties. In 2008, IETF meetings were attended by roughly 1300 engineers and researchers from all over the world.

The IETF is an R&D forum in which network engineers de?ne, describe, review and discuss network protocols, which are published as Requests For Comments (RFC). These then may or may not be implemented and used by industry. IETF meetings are triannual, with business in the interim being conducted on open mailing lists.

Organizational Structure of the IETF

Work within the IETF is organized into working groups (WGs), each of which is in charge of a speci?c problem (eg mobile ad hoc routing). Typically, a WG is supervised by two chairs.

WGs within the same general ?eld are assembled in a so-called 'area' (eg 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 eight 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), ensures the overall coherence of the Internet protocols 'corpus'. IESG and IAB members are periodically replaced, potentially by any other competent IETF participant.

The IETF and Decision Making

In contrast to standardization bodies such as IEEE or ETSI, individuals represent themselves to the IETF: 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 for 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 this is gauged by the WG chairs, and in the IETF as a whole it is 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 an equal footing with those from big companies.

Working code: generally a proposed protocol cannot be promoted as a potential standard without thorough experimentation. Experiments on protocols can be performed using working code and minimal hardware investment, often none. Furthermore, to avoid artefacts 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.

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, its 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 1038 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, the most brilliant idea in the world may be presented in vain if it is incompatible with existing technology: “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 (eg 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 ?ts legacy and addresses a concrete issue: “A good idea can become an extremely good idea”. A striking example is TCP. In the late 1980s, 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 ?le transmission pace according to feedback from the destination terminal. Experts consider the strength of TCP (supporting variations of network capacity ranging over more than twelve orders of magnitude) to be the main reason for the success of the Internet.

Necessity: the IETF mandates itself to solve certain problems. For example in the late 1980s, the current routing protocol RIP failed when a set of routers was brutally removed from the network. This bug, called ‘count to infinity’, created a sustained loop that caused an avalanche of disruptions: the Internet was down for two full days. A failure indeed for a system designed with resilience as its core tenet! 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 that allows routers to compute new routes and react in real time to disruptive topology changes.

The IETF and the Future

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

Link:
http://www.ietf.org

Please contact:
Emmanuel Baccelli
INRIA, France
Tel: +33 169334101
E-mail: Emmanuel.Baccelli@inria.fr

Invited Article

Network Virtualization - An Enabler for Overcoming Ossification

by Anja Feldmann, Mario Kind, Olaf Maennel, Gregor Schaffrath and Christoph Werle

While the Internet is currently viewed as widely successful for some of its participants, namely the users and service providers such as Google, it also suffers from ossification in the underlying infrastructure. The ossification has multiple causes, among them the fact that since the Internet works quite well as it is, Internet Service Providers (ISPs) have no incentive to change their ways. Moreover, ISPs suffer from a lack of business perspective due to the predominant charging modi for Internet access: flat rates for users and a combined price model consisting of a base rate and usage-based component for content providers. An additional complication is that traffic grows at a higher rate than that at which the network equipment costs decrease.

As a consequence, there is debate within the ISP community over whether ISPs should become pure bit-pipe providers or should offer value-added services. In addition, some service providers have found that some application support inside the network can help their applications and are considering investing in network infrastructure, eg Google. Indeed, over the last twenty years almost all innovation, eg novel applications, has taken place at the edge of the network, while the core remains almost untouched. However, the time has come to support novel applications with services inside the network (eg via network-based enablers), and to revisit the Internet architecture to add native support for security, mobility and manageability.

For circumventing the difficulty of changing successful networks, the concept of overlays has proven to be very useful. For example, the Internet got started as an overlay on top of the phone network. One of the key insights is that