

The IETF, the IRTF, and the Networking Research Community

Aaron Falk

Chair, Internet Research Task Force^{*}
University of Southern California, Information Sciences Institute
4676 Admiralty Way, Marina Del Rey, California 90292
falk@isi.edu

ABSTRACT

Some in the research community perceive the Internet Engineering Task Force (IETF) as ossified and see the difficulty in standardizing new protocols as an indication that there is no room for research in the Internet or researchers in the IETF. The truth of the matter is a growing number of critical research problems are facing the Internet today, and the IETF needs help from the network research community. Research Groups in the Internet Research Task Force (IRTF) are working on some of these important problems. The networking research community is invited to collaborate on an effort which will have great impact on current and future network functions.

1. THE PROBLEM

Time was when the IETF was largely driven by researchers. In collaboration with operators and vendors, the researchers defined the agenda, proposed solutions, did prototype implementation of new protocols, performed experimental deployments and analyses, revised the specs, and did it again. These days participation in the IETF is dominated by equipment vendors, stack developers, and network operators. Correspondingly, the participation of researchers in the IETF has dropped precipitously. I call it the 'loss of the curmudgeons' :) and I think it is a reason for concern. Researchers, through their participation in working groups, helped the IETF maintain an intellectual consistency – an architecture, if you will – that the more product-driven members of the community often under-value.

The Internet's technology base has grown more complex, making it more fragile, and more demanding. Many demands find their origin in the success of the Internet: regulation, internationalization, privacy, and accessibility over more diverse media. This complexity makes it harder to solve known problems without creating new ones. The need for a consistent architecture is greater than ever, so that development and standardization of protocols (and there are lots of protocols!) will lead to an overall coherent network architecture.

2. SOME INTERNET RESEARCH QUESTIONS

With their diminished participation, the research community is losing awareness of the greatest challenges faced by

IETF community. Ironically, many of these problems are caused by the scale of the Internet and its extension into areas unimagined by its creators.

Here are some examples of research problems faced by today's Internet:

- Should the overloading of identity and location in the IP address be maintained? Many in the research community have already accepted that some sort of an identifier-locator split is necessary. If this is the right path, how and when should it be effected? What is the roadmap? The advent of IPv6, and the need to support multihoming, present an opportunity to develop and deploy some mechanisms. However, this split also brings new design issues. If hosts have two forms of names, rather than simply an IP address, how will the mapping between them be managed – especially if there are multiples of each?
- What notions of end-user identity allow flexible and reliable user and end-system authentication and subsequent authorization? TLS and passwords are in wide use today for authenticating human-user actions, but are not appropriate when one device, like a laptop, wants to authenticate and then authorize another device, like a printer. There exist a variety of protocols and authentication techniques, yet the users and end-systems do not have the credentials to use the more secure mechanisms. How can devices and human users be easily enrolled in the authentication systems? In a home environment, where there is little or no infrastructure, how do devices obtain the credentials that are needed to perform authentication and subsequent authorization checking?
- Given the Internet was originally designed as an open network, how can the network be protected from unwanted traffic, such as denial-of-service or email spam? Much of the IETF's effort in security has been on protecting user data. Some of the proposed mitigations for undesired traffic also have complex trade-offs, some requiring significant deployments and operational changes. An architecture and roadmap need to be developed for the community that balance risk against performance, complexity, and deployment challenges.
- What congestion control architecture for the Internet allows good performance for high-speed flows on fat

^{*}<http://www.irtf.org/chair>

pipes, fair sharing between data and real-time applications, and responsiveness as paths change? There are many proposals for congestion control, each having its own tradeoffs among performance, complexity, deployment challenges, and fairness. These require objective evaluation that we can then use to drive the development of a roadmap.

- How can we evolve the Internet architecture and protocols so that end-users' privacy can be preserved? For example, it is already now relatively easy to correlate the geography of IP address allocations with user identity, e.g., gleaned from web cookies, and to acquire, sell, and exploit information about a user's location. The adult entertainment industry is already doing this. In the future, as networked devices become more ubiquitously connected service providers, application providers, and even eavesdroppers may perform such correlation with much better accuracy than today, potentially leading to complete lack of privacy in our online life. There is a tussle space here, there are legitimate uses for this information along these lines of some granularity - the challenge isn't to strip it out entirely but to provide it in a privacy-preserving fashion.
- How can the network maintain – *and restore* – end-to-end connectivity and freedom to innovate in this time of middleboxes and unwanted traffic? Is it possible to design general mechanisms that allow any host to request communication, using any transport protocol, with any other host, without the network being able to prevent it? What are the semantics of these mechanisms? What are the implications on trust in the network? Can end-users explicitly authorize agents in the network to act on their behalf? Can, for example, end-users acquire increased ability to configure – or at least understand – firewall behavior?

The reality of finding solutions for use on the Internet means that these topics need to factor in economics and incentives, regulatory constraints, and deployment and migration. These topics are often inadequately addressed in network research.

3. THE INTERNET RESEARCH TASK FORCE

The answers to the above questions are being decided today without sufficient input from you, the research community, often in an ad-hoc manner within the IETF. The Internet Architecture Board (IAB) recognizes that a consistent application of design principles will help create long-term solutions necessary to preserve the cohesion of the Internet architecture. The research community is a valuable resource in achieving this goal and the IAB is expected to propose Research Groups within the Internet Research Task Force (IRTF) to address some of the most critical problems [1].

The IRTF is an organization that fosters collaboration on topics related to 'Internet protocols, applications, architecture, and technology.' [2] The work of the IRTF is performed by Research Groups. 'IRTF Research Groups are formed to encourage research in areas of importance to the evolution of the Internet. Clearly, anyone may conduct such

research, whether or not they are members of a Research Group. The expectation is that by sponsoring Research Groups, the IRTF can foster cross-organizational collaboration, help to create "critical mass" in important research areas, and add to the visibility and impact of the work.' [2] Research Groups are well-positioned to enable interaction between researchers, implementers, and operators on hard topics – such as those above – not yet baked enough for standardization. Additionally, they can consider topics much more wide ranging than those appropriate for standardization in the IETF. For example, the Delay Tolerant Networking Research Group considers networking over astronomical distances and in environments where portions of the network are never concurrently connected.

4. YOU ARE INVITED

Longer lasting, more architecturally sound solutions can only be developed with the active input of the research community. However, research results need to be adopted by the Internet to have real impact. For this reason, collaboration between researchers and the engineering community is *vital*.

Your solutions will have impact. Research Group output often feeds into the IETF, where it will evolve the architecture and continue the expansion of the Internet. Getting the answer wrong could mean a technical or regulatory stifling of innovation on the Internet.

On behalf of the IRTF, *I invite you to participate*. Look for new research groups and workshops created in the coming months on topics such as those above. See the IRTF home page (<http://www.irtf.org>) for pointers to the active IRTF research groups, all of which maintain open mailing lists. I welcome proposals for new Research Groups, too. *Please: bring your ideas – and your energy!*

5. ACKNOWLEDGMENTS

Many thanks to Pekka Nikander, Vern Paxson, Lixia Zhang, John Wroclawski for their assistance in preparing this note.

6. REFERENCES

- [1] S. Floyd and V. Paxson. IAB Thoughts on the Role of the Internet Research Task Force (IRTF). Internet-Draft Version 1, IETF, May 2005.
- [2] A. Weinrib and J. Postel. IRTF Research Group Guidelines and Procedures. RFC 2014 (Best Current Practice), Oct. 1996.