

How to Contribute Research Results to Internet Standardization

Marcelo Bagnulo, Philip Eardley, Lars Eggert, Rolf Winter
UC3M, Madrid; BT, Ipswich; Nokia, Helsinki; NEC, Heidelberg

marcelo@it.uc3m.es, philip.eardley@bt.com, lars.eggert@nokia.com, rolf.winter@neclab.eu

This article is an editorial note submitted to CCR. It has NOT been peer reviewed.

The authors take full responsibility for this article's technical content. Comments can be posted through CCR Online.

ABSTRACT

The development of new technology is driven by scientific research. The Internet, with its roots in the ARPANET and NSFNet, is no exception. Many of the fundamental, long-term improvements to the architecture, security, end-to-end protocols and management of the Internet originate in the related academic research communities. Even shorter-term, more commercially driven extensions are oftentimes derived from academic research. When interoperability is required, the IETF standardizes such new technology. Timely and relevant standardization benefits from continuous input and review from the academic research community.

For an individual researcher, it can however be quite puzzling how to begin to most effectively participate in the IETF and - arguably to a much lesser degree - in the IRTF. The interactions in the IETF are much different than those in academic conferences, and effective participation follows different rules. The goal of this document is to highlight such differences and provide a rough guideline that will hopefully enable researchers new to the IETF to become successful contributors more quickly.

Categories and Subject Descriptors

K.1 [Computing Milieux]: THE COMPUTER INDUSTRY—Standards.

General Terms

Standardization.

Keywords

IETF, research.

1. INTRODUCTION

In telecommunications, standards are essential. More often than not, technology interoperability requires an agreement on a single standard for a given problem. However, unlike most research, standards developments are driven by particular real-world problems and require solutions that are not only theoretically correct, but need to be implementable with state of the art technology in a cost-effective manner, and must be incrementally deployable in the actual Internet by the involved stakeholders. In other words, standards should be both theoretically correct and practically applicable.

Unlike in the academic world, the latter is often more important than the former! A practically applicable solution that has some well-defined and acceptable deficiencies trumps a theoretically complete and optimal solution that cannot be deployed. Likewise, a solution to an interesting theoretical problem that does not exist in the deployed Internet at large does not require urgent standardization. Finally, standardization oftentimes focuses on piecemeal improvements to existing technology in order to enhance secondary aspects, which does not excite an academic researcher looking to solve juicy problems.

These differences between academic research and Internet standardization are the main reason why many researchers initially struggle when they begin to participate in the IETF. Symptoms of this struggle occur, for example:

- for ideas that are too far outside the IETF's areas of current work
- for ideas that are too high-level for the IETF to begin protocol-level work on
- proposals that solve problems that are not expected to arise for a very long time
- when giving others a say in how a research idea is being made concrete, or giving over change control entirely
- feeling that the IETF “does not listen” to them or does not have “the right people”
- there seems to be no working group or other venue to bring the work to
- the process is too time consuming
- the researchers do not have the resources to keep the IETF effort active for an extended period of time
- simulation is not a convincing enough argument for the IETF to start working on something
- the research idea is just not implementable in today's Internet

This document attempts to give some basic advice that researchers might want to take into account when deciding to approach the IETF with their ideas, in order to improve their success probability. It is intended to complement the more general advice in [1] about “How to gain prominence and influence in standards organizations”.

The authors have been involved in several research projects, including collaborative ones, which have sought to standardize some of their results at the IETF, and we hope to pass on some advice (sometimes that we have learnt the hard way!). The advice is split into three groups: before you approach the IETF; how to get the IETF to start work on your proposal; and finally how to increase the chances of success once work has begun.

2. IS THE IETF THE RIGHT VENUE?

A researcher should consider whether the IETF is the right venue before bringing a proposal to it. A good idea is to imagine that the IETF has standardized your proposal and it has been deployed, and ask yourself two questions:

1. How would the Internet be better?
2. What Internet nodes would have been upgraded?

It is very important to have a clear explanation about the motivation for your proposal - What would its benefits be? What problem does it solve? Many ideas do not bring a clear benefit to the Internet in the near term (of course they may still be fine pieces of research!). In the past the IETF has often developed protocols that ended up not being used, so it now thinks harder about the benefits before starting new work and makes sure that it solves a current, significant problem rather than one that may theoretically arise in the future. It is best to be specific about what improvement your proposal would make and the use cases in which this would be seen.

It is also important to have a simple description of what additions or changes are needed and to which nodes (be they end hosts, routers, middleboxes etc). Is it substituting for an existing IETF protocol or supplementing one? Again, it is best to be specific - Do both ends need to adopt the new protocol? Can it fall-back or interoperate with the existing IETF protocol? Do the “first movers” (the first nodes that include your protocol) get an improvement, or do the “last movers” gain most? What assumptions do you make about the network or host (perhaps that the host is multi-homed or there are no middleboxes on the path)?

If it is hard to answer these questions, it may indicate that the idea is too high-level or abstract for the IETF. Then it may be better to approach the IRTF (the research arm of the IETF); the IETF needs a specific protocol-level proposal before it can begin work, whilst the IRTF considers work that is not yet mature enough for standardization. Another danger is that the IETF is the wrong standards body, as a different one would need to standardize your proposal.

If your idea involves replacing several IETF protocols and/or upgrading several types of node simultaneously, it is probably best to re-think: the IETF finds it almost impossible to handle radical, “clean slate” proposals that change lots of things at once. Perhaps you can trim off a subset of your idea that's a smaller initial step requiring only an incremental change to an existing protocol, but is still useful?

3. HOW TO GET THE IETF TO START WORK ON YOUR PROPOSAL?

Having decided that the IETF is the right venue, you now need to persuade the IETF to start work on your idea. We discuss three steps that should help - they can be done in parallel - and then briefly how to form a new WG, if that is necessary.

3.1 Identify the right part of the IETF

The IETF is a large organization; therefore you need to communicate with the right part of it. The IETF consists of over 100 working groups (WGs), each responsible for a specific topic. So a good step is to identify whether there is already a WG where your work would fit.

If yes, then join the WG's mailing list and send email and perhaps write an internet draft. A WG's current set of specific items is defined in its “Charter”; be aware that if your proposal falls outside the WG's current charter, then it would have to be extended before formal work could begin. Most WGs think about re-chartering every year or two, although most are OK for some limited discussion on items outside their current charter.

If there isn't a relevant WG, then you should identify the right Area. The WGs are clustered into “Areas” with a common theme such as security, with one or two Area Directors in charge of each Area. You may have to get a new WG created within the most relevant Area; this is a significantly difficult step (see below).

3.2 Build a community

Standards require agreement and approval by a wide range of people. Therefore you need to persuade others of the merits of your idea. In practice you need to go further and persuade others to do work - at a minimum this will be to thoroughly review your proposal and preferably it will be to develop and test it with you. The IETF community needs to see evidence of wider support, interest and commitment - a lack of reaction means work will not go forward (silence is not consent!). At an early stage support could be demonstrated through comments on the mailing list. It is a very good idea to have some internet drafts jointly authored with people from beyond your research team, perhaps an industry player - for example, you could develop a “use cases” document with a “user”, such as an operator.

Working with others has the extra benefit that it will help to clarify your idea and explain better its benefits and how it works. There are many experts at the IETF who can help stress test the idea technically and advise about process and culture. You need to get some of them involved as early as possible.

It may well be worth trying to hold an informal session at an IETF meeting - this can help build a community of interest for your idea; see the advice in [2].

3.3 Outline your protocol

You also need to describe your proposal in a way that others can understand. Your initial document should outline the protocol - it is counter-productive to detail every aspect, unless the protocol is incredibly simple. Firstly, too much detail swamps people with information that they cannot process - most people understand things by learning about them several times at increasing levels of detail. Secondly, providing only an outline makes people feel that they have a chance of making worthwhile suggestions and changes, so they are more likely to actively engage with you. Thirdly, working out details is generally something that a wider group of people is better at than an isolated individual. Fourthly, in order for the IETF to start work, it is more important to convince the IETF that there is a problem that it needs to solve than to convince it about the merits of your solution.

A good idea is to document a “protocol model”, as described in [3]: “a short description of the system in overview form ... to answer three basic questions: 1. What problem is the protocol trying to achieve? 2. What messages are being transmitted and what do they mean? 3. What are the important, but unobvious, features of the protocol?”

3.4 Establish a new WG

You only need to establish a new WG if the idea falls outside the scope of existing WGs. Establishing a new WG nearly always requires a specific session, called a “BoF” (Birds of a Feather), at one of the IETF’s face-to-face meetings. Here the pros and cons of the proposed WG are debated. As part of the preparation for the BoF you need to:

- Build a community (see above)
- Document the benefits – for example, a problem statement and/or use cases
- Document the architecture – for example covering assumptions and requirements on a solution.
- Suggest specific work items for the proposed WG - typically the protocol to be standardized and the supporting informational documents.

Getting approval to hold a BoF and running a successful BoF meeting are both quite difficult. It is highly recommended to work with someone experienced and to read the guidance in [4].

4. HOW TO INCREASE THE CHANCES THAT THE IETF SUCCESSFULLY STANDARDISES YOUR PROPOSAL

Congratulations, you have got the IETF to agree to start working on your proposal. Now it only remains to do the actual work! In this section we give some advice about ways of working that will increase the chances that the standardization runs smoothly.

4.1 Commit enough time, energy and perseverance

Those new to standards bodies may be surprised how long and how much effort it takes to standardize something.

Success at the IETF requires active participation - to convince others your idea is worthwhile, to build momentum, to gain consensus. Although in theory the IETF progresses mainly through mailing lists, in practice face-to-face time is critical, especially for new or substantial work - if possible go to the three IETF meetings a year.

It takes quite a long time for a proposal to turn into an IETF standard - even if the proposal is mature when it is first presented. There are many steps: building a community of interest, convincing the IETF to start work, working through suggestions from technical experts and incorporating their improvements, getting detailed reviews, going through the formal IETF approval process and so on. Even if you can work full time on the proposal, effort is required from other people who can’t. Also, the IETF tends to work in intensive bursts, with activity concentrated in the run-up to and then at the IETF meetings, with lulls of low activity in-between.

The IETF proceeds by “rough consensus” - unlike some other standards bodies, there is no voting and no top-down process from requirements to architecture to protocol. The downside of this is that the IETF is not good at making decisions. Hence you need to persevere and guard against decisions unwinding. On the other hand, if the consensus is to reject your proposal or there is little interest in it, persevering is likely to be a waste of time - probably you should give up or re-start at Section 2.

All this means that it takes a considerable length of time to complete something at the IETF. Two years is probably a minimum. So, although a typical 3 year research project sounds like plenty of time to do standardization, if you haven’t already raised the idea within the first year, you’re probably too late to complete before your project ends. Therefore, since it’s quite likely that the IETF won’t be finished when your project ends, it is particularly important to convince others to help, so that the work is more likely to complete afterwards.

4.2 Be Open and focus out

It is helpful to come to the IETF with an open mind-set.

Co-authorship is good. Some standards bodies value trophy authors, who indicate their support but don’t actually do any work. In the IETF, it is much better if co-authors are actually investing cycles on developing the proposal, whereas simple indications of support can be made on the mailing list or at the meetings.

In particular, if the IETF is going to standardize something, then in effect it takes ownership of it - it is no longer “yours”. Indeed a good milestone of success is when your individual document becomes a WG draft, as then it is owned by the WG. The research mentality is a bit different, as it prizes authorship and confidentiality-until-publication.

It is very important to be open to working with others. Collaborative research projects sometimes find this difficult for two reasons. Firstly, such projects typically have a consortium agreement about confidentiality - it must not prevent you engaging properly day-to-day with people outside the project. Secondly, you may have to spend considerable effort on intra-project coordination - but an individual researcher only has so much energy and enthusiasm for collaborating, so if you spend a lot of time liaising between Work package 2 and Work package 3, then you have little left for working with the IETF.

4.3 Seek resolution not perfection

The research mind-set is often to investigate very thoroughly all possible details about an idea - to seek perfection - sometimes with no particular deadline. The IETF mind-set is to get something done and out there that works, albeit imperfectly; if people find it useful, then there’ll be another iteration to improve it, probably to meet needs that only become apparent on widescale deployment. The philosophy is to find a reasonable solution to the problem that currently exists - time spent over-optimizing may simply mean that the solution has been superseded (perhaps the problem has been solved in some other way, or perhaps the problem was so significant that a different approach had to be found to avoid the problem).

4.4 Implement

The IETF is very impressed by actual implementations: “running code”. It helps smooth the standards process, it helps people believe it really works, and it helps you and others discover any issues.

An implementation that others can download and try is extremely helpful in getting your protocol actually deployed - and presumably that is your real objective, not simply to get an IETF standard! In the longer term, you may need to think how to get it incorporated in the Linux kernel, for instance.

Overall it is very hard to get a protocol in actual widespread use. There are far more IETF protocols on paper than in use.

5. EXAMPLES

In this section, we include some examples where the authors have been deeply involved and have managed (we believe) to bring the research output of a collaborative research project successfully into the IETF.

5.1 Multipath TCP

Multipath TCP enables a regular TCP connection to use multiple paths simultaneously. It extends TCP to allow the use of multiple IP addresses by each endpoint. This work is one output of the Trilogy research project which was brought to the IETF for standardization and it is currently making good progress. We provide a brief overview of the steps taken.

The first stage was doing some early socialization of the main ideas of MPTCP. Presentations were made in several relevant WGs: the Routing Research Group (July 2008) and the Transport Area Open meeting (July 2008 and March 2009). In addition, a mailing list was created, open to anyone who was interested in discussing Multipath TCP related issues in the IETF context, and a public web page was created containing Multipath TCP related material, including papers, Internet Drafts, presentations and code. The feedback received was encouraging enough to continue with the effort of bringing the work to the IETF.

Once we had verified that the proposed ideas had potential traction in the IETF, the next step was to identify the proper venue for the proposed work. There were two choices, namely, to go for a BoF, with a view to a new WG, or to try to add additional work items to an existing WG, in particular TCPM seemed a good candidate. After talking to the Area Directors, it seemed that having a BoF was the right approach, at least for the initial discussion stage. So, a BoF proposal was submitted to the Transport ADs for the IETF 75 meeting held in Stockholm on July 2009. The initial BoF proposal was crafted by Trilogy people, but was sent to the open mailing list for discussion and modification from the rest of the community. The BoF request was approved and the MPTCP BoF was held at the IETF 75 meeting.

The general feedback received during the BoF was that there was enough interest and energy in the community to do this work within the IETF. A first charter draft was posted on the mailing list for comments a couple of months after the BoF. After a month or so of charter discussion on the mailing list, the MPTCP Working Group was created in October 2009. The charter includes deliverables due to March 2011.

The MPTCP working group has, so far, made significant progress and most of the milestones have been delivered on schedule [5].

5.2 Congestion Exposure

Congestion Exposure enables sending end-hosts to inform the network about the congestion encountered by previous packets on the same flow. This allows the network devices to act upon the congestion information and the perceived user behaviour. Like the MPTCP work, it is an output of the Trilogy research project and has been successfully brought to the IETF. We next describe the steps followed to do so.

In this case, early socialization included presentations at the Internet Congestion Control Research Group and the Internet Area meeting at the IETF 75 meeting in July 2009, the creation of an open mailing list to discuss Congestion Exposure related issues in the IETF and posting the related materials such as papers, Internet drafts, and code in a public web page. In addition, an informal, open meeting (sometimes called a Bar-BoF in IETF parlance) was held during the IETF 75 meeting.

After processing the feedback received in the Bar-BoF, a BoF proposal was submitted to the Internet Area ADs for the IETF 76 meeting in November 2009. The BoF was accepted and was held as planned. While the feedback received in the BoF was positive, the IESG was uncertain about chartering a Working Group on this topic. (The IESG is the IETF's management body and consists of all the Area Directors.) In order to address the remaining concerns of the IESG, another BoF was held at the following IETF meeting.

After much debate, the CONEX WG was approved by the IESG but the scope of its charter was limited compared with the original proposal. The CONEX WG [6] held its first meeting at the IETF 78 meeting in July 2010. Its charter contains deliverables up to November 2011.

6. ACKNOWLEDGMENTS

The authors are partly funded by the Trilogy Project [7], a research project supported by the European Commission under its Seventh Framework Programme.

7. REFERENCES

- [1] Eastlake, D. 2005. How to Gain Prominence and Influence in Standards Organizations. RFC4144. <http://tools.ietf.org/html/rfc4144>
- [2] Eggert, L. 2010. Considerations for Having a Successful "Bar BOF" Side Meeting. (Work in progress) <http://tools.ietf.org/html/draft-eggert-successful-bar-bof>
- [3] Rescorla, E. 2005. Writing Protocol Models. RFC4101. <http://tools.ietf.org/html/rfc4101>
- [4] Narten, T. 2009. Considerations for Having a Successful Birds-of-a-Feather (BOF) Session. RFC5434. <http://tools.ietf.org/html/rfc5434>
- [5] Multipath TCP working group. <http://tools.ietf.org/wg/mptcp/>
- [6] Congestion Exposure working group. <http://tools.ietf.org/wg/conex/>
- [7] Trilogy Project. <http://www.trilogy-project.org>