

Into the Future with the Internet Vendor Task Force A Very Curmudgeonly View or Testing Spaghetti — A Wall's Point of View

Randy Bush
Internet Initiative Japan

It is said that many researchers think the Internet Vendor Task Force (IVTF, nee IETF) has become irrelevant. They are either measuring the real internet as a behavioral phenomenon, which is a bit scary if you think about it, or they are wanting to do research 'beyond' the internet. Neither involves the IETF.

This is partly because of the research community's inability to get deployment traction via the IETF path. The '90s poster children, QOS, DiffServ, IntServ, Self-Serv have a long history of attempts at relevance and deployment via the IETF, none successful in the face of a bandwidth glut and lack of end-to-end signaling in their designs. Whether economics will change sufficiently to give them legs is not clear; I would not bet on it.

But do not think that the direct researcher/operator interface is in the best of shape. For example, as it is widely believed that the majority of congestion is on customer access links, why is WRED not enabled on these links? Why have operators not asked the vendors to make it the default for some types of interfaces?

Simple ideas for which there was clear need, ECN, pushback, ... have been poorly deployed at best, and never finished in the IETF community at worst, vendor reluctance being a major block. This inability to get deployment has been exacerbated by the disconnect between the IETF and the operational community - and researchers want their work to be used. And the IETF community has a real problem thinking it's too smart. Despite all the attention to cryptographic algorithm agility, when it came to hash functions, the IETF got it wrong every time [0]. Only IKEv2 was even close to correct.

Given that the IETF has spent the last decade devaluing researchers almost as much as it has operators, is it that there are few if any really serious thinkers left? Or maybe it has become so difficult to have serious new thoughts that the added bureaucratic path through the IETF makes it all just not worthwhile.

But my perspective of the IETF is poisoned by the devaluation of operational realities.

- The IETF has grown so large and so enamored of complexity and featuritis [1] that it is a full-time job to participate. Who can afford to spend full-time on the IETF? Vendors with more features to add to sell more baroque systems. It's not an evil conspiracy, but rather a consequence of being enamored of complexity and the vendors' need to keep selling 'new' 'better' products. This is the path the telcom industry took and we know how much radical technology it develops these decades and how profitable local and long distance minutes are.
- The IETF's vendor/market approach has engendered a 'let the market decide' culture. Instead of hard-thought, rigorous, and simple designs, every possible feature gets added and many competing proposals are approved. This last is like throwing spaghetti at the wall to see what sticks, an amusing tactic to everyone but the wall.

The operators are the wall. And they pay capital cost and operational expense to deploy complex features which vendors market as needs to the users. And then everyone wonders why the margins went down and the prices stayed up.

I think it was Professor Wirth who once said that there are two kinds of standards: unions, where every feature is included so everyone is happy, and intersections, where only the features which everyone agrees are absolutely critical are included. His languages have been elegant examples of the latter, until standards groups got hold of them [2]. We are recreating PL/I and Ada.

- Rigor and excellence have been replaced by a social vision of openness, fairness, and inclusion. Combine this with the above issues, and you get a culture of mediocrity. It becomes more important to make people happy and to be liked than to do rigorous design and development. This forms a cycle, and the downward spiral has been in progress for some years.

One has only to look at the long and painful histories of (not) deploying IPv6, dnssec, various QOS schemes, ... to see the results of the disconnects between research, the

IETF, process and the realities of operational deployment. And no one has bad intent. It's the physics of growth, industry maturity, 42 cases of second system syndrome, and the resulting but anti-functional complexity.

It is easy to lose perspective, become preoccupied with the trees, and think one must be making progress through the forest. When there are as many moving parts as the IETF culture of complexity seems to tolerate, worse encourage, the concept of forest and direction are lost in organizational processes, social welfare, and administrative politics.

A famous Mexican painter, David Siqueros, had a wonderful piece on a church ceiling I visited as a child, The Man who was so Open Minded his Brains Fell Out. As a culture, we are becoming so open minded and

accepting of complexity, both in our product and in our process, that our real engineering vision has fallen out. So, we can look at other matured scientific/engineering disciplines, e.g., aerospace, telco, ... and get a feel for where we are headed, not fun places. Think ITU, ISO, etc. But today's vision of the IETF will take us there, slowly, openly, and with true team spirit.

[I wish to thank a few researchers and operators who read and very seriously improved this broadside. It is probably best if they remain anonymous.]

[0] - <http://www.cs.columbia.edu/~smb/papers/new-hash.pdf>

[1] - <http://rtg.ietf.org/~fenner/ietf/deps/viz/mpls.pdf>

[2] - <http://www.scifac.ru.ac.za/cspt/sc22wg13.htm>