

IETF 111 Plenary Minutes

28 July 2021, Online

[All Slides: IETF 111 Plenary](#)

1. Welcome

Lars Eggert welcomed the community to IETF 111

2. Host Presentation Juniper Networks

The IETF thanked Raj Yavatkar, CTO of Juniper Networks, for hosting IETF 111.

3. Brief Updates

Note: Reports available in the [Datatracker](#)

3.1. IETF Chair Update

Slides: [IETF Chair Report](#)

3.2. IAB Chair Update

Slides: [Internet Architecture Board \(IAB\) Report](#)

3.3. IRTF Chair Update

Slides: [Internet Research Task Force Plenary Report](#)

3.4. RFC Editor Update

Slides: [RFC Editor Update](#)

3.5. NomCom Update

Slides: [NomCom 2021, Update IETF 111](#)

3.6. IETF LLC Board Update

Slides: [LLC Briefing - IETF 111](#)

4. IETF Community Survey

Slides: [IETF Community Survey](#)

5. Internet Engineering Steering Group (IESG) open mic session

Andrew Campling: I notice there have been quite a few issues raised on some of the different mailing lists over the last few weeks on the use of GitHub vs mailing lists and whether the different working groups are getting the balance right. It seems to me from the comments that for whatever reason some of the current working practices are not satisfactory without making judgments as to in what direction. And I wondered if you thought this is something the IESG needs to reevaluate to give more consistent guidance across the working groups.

Lars Eggert: Thanks for the question. This came up at least once in GENDISPATCH and it was my individual draft. We have the GitHub working group that closed which defined some guidelines on how working groups can use GitHub in the IETF. And it sort of leaves us in a position where each individual working group decides for themselves if they are using GitHub and how they are using GitHub. What I can't recall is if we had a discussion on whether a working group that doesn't make an active decision on what goes for them. So GENDISPATCH, for example, in this case as far as I know doesn't have a stated policy for how GitHub is to be used, and so in this particular case I have an individual draft that I use GitHub for, and I posted a link with a pull request that contained a proposed change to the document. The document discussion happened on GitHub and there was some confusion on whether this was wanted or not, and why it wouldn't be an email. I'll point out that for working groups that basically configure GitHub as an organization, it's possible to set up a mailing list such that all the discussion on Github gets reflected there. QUIC did this for the QUIC-ISSUES mailing list. You've seen it mentioned earlier because we excluded it from the stats that Alvaro showed because it is by far the busiest IETF list, and it's sort of automatically generated email mostly, but at least it allows people that don't want to interact with GitHub to at least follow the discussion. But it requires a working group to actively have a GitHub Organization and be configured in a certain way. So it's not something you can do for an individual draft, which is a little bit of a problem. So I'm wondering if other people have an opinion on this one thing that I've seen proposed whether GitHub Working Group (GIT) should be reopened with an extended charter. That's something we could do. But I'm not sure how much consensus there would be for doing something more than we did in the existing GitHub Working Group a while ago because the community is quite split on the question of whether GitHub should be used, and how much.

Branden Moran: I wonder if this would be less of an issue if there were an IETF-hosted issue tracker and Git host available.

Lars Eggert: It has been discussed, and I think, for example, we have a GitLab license for the IETF that we could take into action. One of the arguments that has been brought up is that being on GitHub as the sort of large place where a lot of open source gets developed makes it easy for people who are already on GitHub to participate on IETF documents, and we've totally seen this for the QUIC working group where individuals that hadn't participated in the IETF found the repository somehow and opened issues and made comments and then we never saw them again. So, just sort of drive-by participation is something that GitHub enables, and if we would host our own set up that wouldn't be possible. We can certainly have a discussion about how frequently that happens and how important it is for your organization, but it's been brought up and discussed. I don't think there was consensus on it but GitLab would be something that's available and other Git hosting solutions exist too.

Branden Moran: I see. Thank you.

Geoff Huston: I'm bringing up a question that arose in the SIDROPS working group earlier today (or yesterday, depending on where you live in this world), and it is about what "running code" actually means. The SIDROPS working group is certainly grappling with that, taking the lead from other working groups that appeared to say you really need implementation reports in order to advance a draft. But in some ways that is not quite good enough these days, and what we find, of course, is that the Internet is working on a scale that it has its own stasis and resistance to incorporating what we call experimentation. Now when you say, "this has implementation," do you mean this spec is good enough for independent implementations to interoperate full stop? In other words, the spec is a reasonable guideline for code, but as to whether you should deploy it or not, or it's a good thing eh, someone else's problem, or you could say this specification creates independent implementations that operate at scale and makes the internet better and you should build it. There are a whole bunch of different opinions out there, the DNS certainly appears in those working groups to take the position that Internet-Drafts are purely ephemeral, I'll implement it when it is an RFC and not before. Because you know it is hard to chase butterflies. Whereas others appear to be saying all the rigor, and indeed even some element of acceptance of this, is a good idea and should come at the working group and draft level and what should come out of the IETF are RFCs that bless history. This is a tough question, and I think punting this to individual working groups to actually resolve severally, diversely, and independently is entirely wrong for the IETF. These mixed signals are not only confusing. I tend to suggest they create such a mixed quality of outcome that it actually devalues the RFCs. Lowest common denominator rules tend to apply. What's the IESG think about this, and are they willing to actually reopen what "running code" actually means for the IETF and its working groups? Thank you.

Lars Eggert: Thanks Geoff. Anyone who participated in the working group want to speak to this? I see Alvaro.

Alvaro Retana: This has come up before, and I think that ideally we would have running code for many things before we publish. There are some working groups that Geoff, I know you're familiar with, IDR for example that requires two independent implementations. We've had a ton of discussions with the chairs and others about what does an implementation mean in general. You know we want an actual vendor who provides backbone software for example when we're talking about BGP, right. The last thing we want to do is break BGP. However, and it goes to your point about leaving independent working groups to make the decisions, in my mind at least, I do think that's the right way to do it. The reason is that different working groups have different markets. Different segments, and not all of them are going to go do implementations before things are standardized. Many working groups that work with governments, for example, these manufacturers are going to wait until there is an RFC until the spec has been completely approved before they do implementations. So I think that part of the issue is that one size doesn't fit all now that we have, I don't know, for 120 working groups, it is very very hard to come up with one solution that is going to fit everyone.

Geoff Huston: To Alvaro here, in making the observation that individual working groups have to interpret what they see as running code, in their own at times unique way, I don't think that flexibility to that degree serves the IETF in a general sense. I actually think you need to refine what we mean by running code and work out the difference between a spec that is implementable and a spec that is a good idea for customers to use and work out how running code relates to that because at the moment as I said, I think that degree of variation is actually harmful to the process and the quality of its output and simply saying it's somebody else's problem, working groups should make their own way through this mess, isn't exactly the degree of helpful guidance that normally I would expect from you worthy folk on the IESG. Thank you.

Roman Danyliw: Geoff, I would say you're asking a very important question which is what are the means we can arrive at to have confidence in the specs that we're going to publish. I think certainly code is one of those, but one of the other things I would kind of emphasize with the enthusiasm that you have is that for example, formal verification in the Security Area has been transformative to make sure that the spec or even the code we're producing that's those reference implications has really changed the game for us and largely table stakes for the introduction of new protocols for us and to the question of scale I know prior working groups have also elicited the academic community for things like modsim to make sure that we understand those protocols at scale as well without the deployment so I think there is a number of levers we should be pushing to make sure the quality is high. Thanks.

John Klensin: Geoff, I am very sympathetic to the argument you are raising and I think it is important, but it seems to me that the argument you're making in the distinctions you're asking for is exactly why we started down the path of a three-step standards process about a century ago. And I know that for better or for worse, we've never been able to make that work. I wish we could. Thank you.

Toerless Eckert: So having been on both the development deployment and the standardization side I think I'd like to chime into what Alvaro said with respect that; you can't simply say it is one size fits all, right. So depending on where you are actually working in the standardization, the overhead to just do an experiment in implementing something that ultimately turns out not to be useful has been shown to be very, very expensive in cases of you actually may want to have standardization pre-go the implementation to really work out quirks before your investments. So obviously that's not true everywhere. The more something is easy to implement software, the more the opposite is true, but I certainly don't think that you know one size fits all.

Lars Eggert: Thank you Toerless. Any further questions? No. Okay next up is the IAB Open Mic.

6. Internet Architecture Board (IAB) open mic session

Andrew Campling: I have two questions, if I may. The first is, I think the excellent RFC 8890 Informational document on "The Internet is for End Users" talked about the importance of multistakeholder consultation, and I wondered if there'd been any subsequent thoughts on actually developing a process to do this, because I think there are certainly a number of Working Groups where there's a need for multistakeholder consultation, and I think it potentially becomes quite a barrier to doing it if each develops its own process. So, some guidance on how, I think, would be welcome. And then secondly, to think about my other question, I think [at IETF 110] I raised a question around concerns around centralization, and how some of the proposed standards working through the system seem to be in danger of making that worse, not better, and whether there are any plans to assess the centralization impacts of new standards before they're introduced, alongside security impact. I don't think I've seen progress on that, and would welcome views on what's going to be done to actually take centralization risk seriously. Thank you.

Mirja Kühlewind: Let me start on your first part. I don't think we have a specific process for consulting multistakeholder organizations or parties, but we do discuss a lot, on the one hand, how to better manage our liaisons, how to [do] better outreach, how to better utilize them and how to have a better overview of what's happening at all, and we've also had some discussions about how outreach, and there are multiple dimensions of this. This is like, advertising what the IETF is doing in other organizations, getting feedback, getting people involved, and getting new people involved also. But these are all different aspects with different answers, and we are just starting the discussion, so I don't have an answer to you, but it's still on our radar, but it's also not like [there's] an easy solution on how to improve the situation here. It's more in the assessment phase. I don't know if someone else from the IAB wants to add something?

Tommy Pauly: Can I turn to the second half of the question here? Just about the point about what can we do for concerns about centralization, et cetera, it is definitely something we've discussed as the IAB at retreats et cetera. This is one of the practical items that we've been going through an exercise (that we're still working on), is a list of things, centralization just being one small part of that, that the IAB should be reviewing during BoF chartering, Working Group

chartering, any time we have new work coming in, just as a way to kind of look at the bigger picture of how do these pieces fit together, and are there things that may not come up in other types of review that we think should be on the list that the IAB comes back to. So, that's definitely one of the items that's in there, and at some point we'll formalize this list, and then it'll be public, but we're already thinking about it as we are looking at new work coming into the IETF.

Jari Arkko: Yeah, so Tommy said basically what I wanted to say, but maybe I'll add a few things. One is that many of us actually do care quite a bit in the IETF and the IAB about centralization, consolidation issues, but it's sort of a difficult item to deal with. You have to discuss the specifics of these cases. So, yesterday, for instance, we had a big discussion on impacts of discovery or not in one of the BoFs that we had this week, and you know, whether there's an issue depends on exactly what the use cases are and how we end up using a particular technical design. So, I think it's at the top of our minds, or at least it's one of the things that we care about. But we have to discuss the specific cases in detail and there's no direct golden rule that says do this or don't do that. Thank you.

Wes Hardaker: One thing to note about the Internet being for end users is that one of the things that we are doing right now is actually going to be holding a workshop on measuring network quality for end users, and that is right in line with our goal to figure out how do we get a better sample rate for how things are impacting end users. The call for papers for that is open right now, it's linked off of the IAB website [<https://www.iab.org/activities/workshops/network-quality/>], and I'd strongly encourage people to submit papers for it. The due date is actually August 2nd, which is coming up quickly, but they don't need to be long papers, just what are your thoughts around that, and you can go read the call for papers. I'll drop it in chat, too. Thanks.

Mirja Kühlewind: So as Andrew just mentioned in the chat that he would like to have more guidance, maybe send us a separate email to the IAB and we can figure out what you're actually looking for, and if that is covered somehow by our liaison management, or if you're actually looking for something else that would be helpful. Thanks.

7. IETF Administration LLC open mic session

Dominique Lazanski: Going back to the survey, as I was mentioning, a question on what is being done to potentially address the gender issues and also the issue that was brought up in parallel about feeling welcome or inclusivity. Thanks.

Jay Daley: Shall I go first? At this stage, Dominique, we're just getting the results together. The response obviously needs to have part IESG response and part LLC response. I can't speak for the IESG, Lars will do that. Within the LLC, when the consultation is finished and we've got the final report out there, we'll be having some form of conversation about that. Jason, did you want to add anything? [Jason shakes head no.] It is something that we recognize and do take seriously. It's something we've been talking about for quite some time, which is why the

questions appeared the way they did in the survey, which was to make sure we had the data to validate what for many of us was an instinctive understanding. Taking that forward now I think we have a very clear mandate across the organization, though we have different roles in the organization as to how we do that, but basically a mandate to understand that and do something about it. Lars, did you want to add anything?

Lars Eggert: I can only say that the results came in a few weeks ago and the IESG was in the ramp up to the IETF meeting so we've looked at it and discussed it but we simply haven't had the time to sit down on a chat to dig down and decide what we're going to do with the results. I agree with Jay and the others that the results are eye opening and they certainly motivate an action by the organization or by the community to improve the scores we have in some of those categories. I was quite surprised by the outcomes. Are there any further questions for the LLC Board? It doesn't look like it.

Jason Livingood: One quick comment; there was some discussion in the chat I just wanted to make sure folks took note of. It was about the mechanism by which people can provide feedback to the board on the IASA2 Retrospective. As noted in the blog, we're happy to take it via email, on admin-discuss@ietf.org, on Github, the upcoming second webinar, or office hours in Gather this week. Any of those are fine and if you have some other method or you'd like to have a one on one, we're happy to do that as well.

Lars Eggert: That brings us to the end of the plenary, thank you very much.