

# IETF® Journal



A report from IETF 95, April 2016, Buenos Aires, Argentina. Published by the Internet Society in cooperation with the Internet Engineering Task Force.\*

## INSIDE THIS ISSUE

From the Editor's Desk.....	1
Declaring IPv4 Historic: One Issue, Two Sides.....	1
Message from the IETF Chair.....	2
Words from the IAB Chair.....	3
Meet the IETF Systems.....	8
Turning 30, the IETF Seeks Ideas for Growth.....	8
Using I2NSF for Overlay Network to Avoid Distributed Denial of Service Attacks.....	10
Intelligent Transportation Systems and the IETF.....	11
Things Talking to Other Things about Things.....	15
Internet Regulators, Technologists Seek Ongoing Dialogue.....	16
WG Update: L3SM.....	17
WG Update: TAPS.....	19
TRON Workshop Connects IETF TLS Engineers and Security Researchers.....	20
Hackathon Breaks New Ground in Buenos Aires.....	21
Integrating the Worlds of Technology Design and Policy.....	22
IRTF Update.....	24
ANRP Winners Announced.....	25
IETF Ornithology: Recent Sightings.....	26
IETF 95 At-A-Glance.....	30
Calendar.....	31

## FROM THE EDITOR'S DESK

By *Mat Ford*

**T**HE IETF BEGAN ITS 30TH ANNIVERSARY YEAR WITH A MEMORABLE MEETING IN Buenos Aires. This was the second time the IETF has met in the southern hemisphere. In addition to drawing a large number of local participants, it had a relatively high number of remote attendees.

Our cover article was prompted by a presentation and debate that took place in the sunset4 Working Group on the subject of declaring IPv4 Historic. Read Lee Howard's and Geoff Huston's perspectives and make up your own mind!

In addition, we have several Working Group updates, observations from one of the Internet Society Fellowship to the IETF Programme participants, a view from the pre-IETF Hackathon, and an article about a potential new technology direction for the IETF: Intelligent Transportation Systems. Our regular columns from the IETF, Internet Architecture Board, and Internet Research Task Force chairs, and coverage of hot topics discussed during the plenary meetings wrap up the issue.

Finally, I'm thrilled to announce that our redesigned website has been launched. Check it out at [ietfjournal.org](http://ietfjournal.org). I hope you find the IETF Journal content more accessible and interactive.

We are hugely grateful to all of our contributors. Please send comments and suggestions for contributions to [ietfjournal@isoc.org](mailto:ietfjournal@isoc.org). You can subscribe to hardcopy or email editions at [www.ietfjournal.org/subscribe](http://www.ietfjournal.org/subscribe).

## DECLARING IPV4 HISTORIC: ONE ISSUE, TWO SIDES

By *Lee Howard and Geoff Huston*

**D**URING IETF 95, IN A MEETING OF THE SUNSETTING IPV4 WORKING GROUP, LEE Howard presented on a proposal that recommends that IP version 4, or to be specific, the technical protocol specification documented in RFC 791, be declared Historic.

In the context of the Internet Standards Process, the term Historic has a particular meaning. RFC 2026 defines Historic to mean:

*Continued on page 4*



\*The articles published in the IETF Journal are not intended to reflect the opinions or the position of the IETF or the Internet Society. See <http://www.ietf.org>.

## MESSAGE FROM THE IETF CHAIR

By Jari Arkko

**W**E HAD A GREAT IETF 95 MEETING IN BUENOS AIRES WITH A LOT OF TOPICS AND MANY participants. We had approximately 500 people following the event remotely and more than 50 presentations offered remotely. We even had a steering group member participate remotely. And this is as it should be: while face-to-face meetings are very important for networking, people should also be able to attend over the Internet. After all, we are the *Internet* Engineering Task Force.

Another remarkable aspect of IETF 95 was that it was our first meeting held in South America. We saw slightly more than 1,000 participants on-site, about 140 people from the region. I was very happy to see such strong and active local participation.

The meeting was cohosted by LACNIC and the Internet Society—thank you for stepping up to support this meeting! I was happy to see many local sponsors, too, including IPLAN, CABASE, .AR, and NIC.BR. And my thanks to the other sponsors as well: Neustar, Level 3, Comcast–NBC Universal, Huawei, A10 Networks, and ICANN.



Jari Arkko, IETF Chair

**Another remarkable aspect of IETF 95 was that it was our first meeting held in South America.**

### Technical Topics

Two meetings on the growth of encrypted traffic stood out: (1) LURK, which was on building a distributed system that allows Content Data Networks (CDNs) to employ HTTPS/TLS while not releasing a copy of the private keys to the CDNs; and (2) ACCORD, which was about whether better queuing algorithms or more information about traffic flow priority would be useful for better scheduling of radio resources in mobile networks.

The Internet of Things is another active and interesting area. Low-power wide-area networks were discussed in the LPWAN BoF, and some IoT-related Working Groups, including CORE and ROLL, have completed their initial batches of work and are now looking at new work.

This meeting also saw the first official meeting of the Thing-to-Thing Research Group (T2TRG) at the IETF. This active Research Group is focusing on device-to-device communications and has met twice before the meeting. At the technical plenary, we heard a report from the recent IAB workshop on semantic interoperability problems (page 8).

There was also plenty of work on Internet security. One of the most interesting topics was the work on TLS 1.3, specifically the discussions about its super-efficient 0-roundtrip initialization mode and under what conditions replay attacks can be avoided in that mode.

*Continued on page 6*

The mission of the Internet Engineering Task Force is to make the Internet work better by producing high-quality and relevant technical documents that influence the way people design, use, and manage the Internet. See <http://www.ietf.org>.

### Recent IESG Document and Protocol Actions

A full list of recent IESG Document and Protocol Actions can be found at <https://datatracker.ietf.org/iesg/ann/new/>

## WORDS FROM THE IAB CHAIR

By Andrew Sullivan

I BECAME IAB CHAIR IN DALLAS, TEXAS, IN 2015. IT ISN'T QUITE TRUE THAT AS I BOARDED the plane to Buenos Aires for IETF 95 that I couldn't believe it had been a year; it was more like the year had evaporated while I wasn't noticing.

Naturally, because of the Internet Architecture Board's (IAB's) role in looking after the IETF's relationship with the Internet Assigned Numbers Authority (IANA) and the Internet Corporation for Assigned Names and Numbers (ICANN), a lot of that time went into the IANA transition. But, in a way, the time pressure turned out to be healthy for the IAB. More on this follows.



Andrew Sullivan, IAB Chair

### The IAB Reports

As I noted in the previous edition, we heard positive reactions to the idea that the IAB would make most of its report in email and devote some time to just a few highlights in the meeting. So, we did that again. The report is available at <https://www.iab.org/2016/04/04/report-from-the-iab-before-ietf-95/>. We plan to keep working this way as long as it is useful. Don't forget, you can discuss with the IAB anything you see in that report or in this note, or anything else you want the IAB to attend to. If you want to do it in public, send mail to [architecture-discuss@iab.org](mailto:architecture-discuss@iab.org). If you want to talk to the IAB without causing a public discussion, send mail to [iab@iab.org](mailto:iab@iab.org).

**The IAB annually appoints its chair. I am flattered by my IAB colleagues' trust in me in selecting me for another year. I hope this one doesn't go as fast!**

### The IAB Changes

The first meeting of the calendar year is the time when appointment terms end and new appointments begin. At IETF 95, the IAB had to say goodbye to two departing colleagues: Mary Barnes and Marc Blanchet. At the same meeting, we welcomed Lee Howard and Martin Thomson. It is always difficult to accept that valued colleagues will no longer be available in the same capacity as before. Yet the changes bring fresh perspective and renewal, and that renewal is what ensures the IAB can be of service to the IETF and the Internet. The Internet does not sit still. Neither should we.

The IAB annually appoints its chair. I am flattered by my IAB colleagues' trust in me in selecting me for another year. I hope this one doesn't go as fast!

### Technical Plenary Discussions

After IETF 95, we received some expressions of disappointment that there was no technical topic at the plenary. When we decided that less plenary time is better—and we got a lot of feedback to that effect—we had to acknowledge that every year one plenary needs to include more administrative detail. New IESG, IAOC, and IAB members get

*Continued on page 7*

The Internet Architecture Board is chartered both as a committee of the IETF and as an advisory body of the Internet Society. Its responsibilities include architectural oversight of IETF activities, Internet Standards Process oversight and appeal, and the appointment of the RFC Editor. See <http://www.iab.org>.

*Declaring IPv4 Historic: One Issue, Two Sides, continued from page 1*

A specification that has been superseded by a more recent specification or is for any other reason considered to be obsolete is assigned to the Historic level.

While it looks simple on paper, actually acting on RFC 2026 is anything but. Lines have been drawn and supporting arguments on both sides show significant merit. To shed light on the pros and cons of declaring IPv4 Historic, the *IETF Journal* invited Lee Howard and Geoff Huston to share their thoughts.

Lee Howard

Yes!

The original document defining IPv6 says that “IP version 6 (IPv6) is a new version of the Internet Protocol, designed as a successor to IP version 4 (IPv4)” [RFC 1883]. A “designed successor” may not instantly supersede its predecessor, but that sounds like the intent, and the successful deployment of IPv6 means the time is near.

IPv4 is historic. It has enabled new means of communicating that have changed the world. IPv4 is also historical: it belongs to the past. Like stone knives, it is a technology that enabled other life-improving innovations, but whose time has passed.

A Historic protocol can still be used. Early discussion in the SUNSET4 Working Group shows that it is too soon to deprecate IPv4. “Deprecating” would be saying it should be avoided because it is harmful or not recommended. IPv4 does have inherent limitations that cannot be mitigated: primarily, the length of the address space. IPv6 does not have this limitation.

IPv4 may still be perfectly viable for communication in some circumstances. Other historic protocols are still in use, when administrators understand the risks, usually when both end points and the network between them are under single

administrative control. Network operators are free to continue using IPv4 as long as it suits their needs.

Declaring an Internet Standard to be Historic does have implications. When RFC 791 is moved to Historic status, any Standards Track RFC with a Normative reference to RFC 791 becomes Historic. Over one hundred RFCs cite RFC 791, but not all of them are normative references, and not all of them would reasonably be obsolete. For instance, RFC 7676 defines “IPv6 Support for GRE” and even though it includes RFC 791 as a normative reference, there’s nothing in it that fails if IPv4 is declared Historic.

Some RFCs define IPv4 options, which would seem to make them Historic. Most, such as RFC 1035 “Domain Names - Implementation And Specification” which defines A records and the IN-ADDR.ARPA zone, will be updated by this document, but are not Historic. Other documents with incidental references to RFC 791 should not be affected. Documents requiring updates should be included in [draft-ietf-sunset4-gapanalysis].

#### **Why go to all this trouble?**

It’s not just housekeeping. Although a tidy house is appealing, there’s plenty of clutter in the RFC series. Some clutter doesn’t matter, though—there’s no need to tell people to ignore the disused ashtray under the sofa, they’re ignoring it already, and fewer and fewer people need ashtrays. IPv4, however, is significant, with new transition mechanisms, new optimizations, and new Network Address Translation (NAT) workarounds still being introduced. Developing consensus on that work distracts people, whose time could be better spent developing IPv6 features or optimizing performance or security in IPv6.

It is therefore important to stop working on IPv4. This tool is becoming more fragile (or brittle) over time with patchwork like NAT and its workarounds. An IETF consensus

declaring IPv4 to be Historic will signal to future IETF contributors that we are done with it. The process of considering, discussing, and building consensus on that declaration is how we as a community determine how we want to spend our precious time. We can decide that we no longer want to support those who refuse or have taken too long to upgrade to IPv6.

**Like stone knives,  
[IPv4] is a technology  
that enabled other life-  
improving innovations,  
but whose time has  
passed.**

For those who choose to continue using IPv4, there are some considerations. The IETF may not normally update Historic RFCs. This doesn’t mean that the IETF can never update IPv4, but the bar is set higher, requiring scrutiny from the IESG. Maybe we continue optimizing transition technologies. As described in RFC 6540, IPv6 support is required, and some documents may be confusing as to whether “IP” means IPv4 plus IPv6, IPv6-only, or IPv4-only.

We can’t declare IPv4 Historic tomorrow. “Standards track specifications normally must not depend on other standards track specifications which are at a lower maturity level” [RFC 2026]. Therefore, any RFC depending on IP must have IPv6 at full maturity before declaring IPv4 Historic. Since the IETF IPv6 Maintenance (6MAN) Working Group is in the process of promoting IPv6 to Full Standard, we would have to wait. Being dependent on that work does not mean it’s too soon to discuss it and to work on building consensus.

It is possible that bugs inherent to IPv4 may yet be discovered. This seems unlikely, given the extent of testing and production use it has. RFC 791 has only been directly updated three times:

1. RFC 1349, "Type of Service in the Internet Protocol Suite"
2. RFC 2474, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers"
3. RFC 6864, "Updated Specification of the IPv4 ID Field"

Still, it is conceivable that an inherent flaw will be found, and if IPv4 is Historic, it will be easier to update IPv6 than IPv4. Therefore, for security reasons, the use of IPv6 only is recommended; IPv4 should be used only as needed for backward compatibility.

To quote Leonardo da Vinci, "Art is never finished, only abandoned." It's time to stop working on our previous master work, and show how much farther we have come than our original stone knives. The combined focus of the IETF on IPv6 can give it an even greater resilience, flexibility, and security than we did with IPv4.

**Geoff Huston**

**Not Yet!**

The rationale for the proposed redesignation of IPv4 was that the protocol has been superseded by a more-recent specification, IP version 6. Furthermore, it is thought that this action would add to the impetus to deploy IPv6. The take-up of IPv6 is not overly uniform. While some service providers are enthusiastic proponents of IPv6, many more providers are at best hesitant and at worst ignoring it and hoping that it will go away. As a message to both the laggards and potentially their customer bases, it was argued that the IETF should clearly indicate that it is time to move to IPv6. One way to do that is to declare IPv4 a Historic technical specification.

If that were all there is to it, then perhaps a case can be made that the IPv4 technical specification should be declared Historic. But there are additional aspects to consider. As pointed out by RFC 2026:

Not Recommended: A Technical Specification that is considered to be inappropriate for general use is labeled "Not Recommended". This may be because of its limited functionality, specialized nature, or historic status.

This text seems to imply that a status of Historic also suggests Not Recommended, which may send the wrong signal to the existing user base that relies on IPv4.

The proposed redesignation also would throw the IPv4 specification out of the Internet Standards set.

Specifications that are not on the standards track are labeled with one of three "off-track" maturity levels: Experimental, Informational, or Historic. The documents bearing these labels are not Internet Standards in any sense.

So maybe this is a bigger step than just observing that IPv6 supersedes IPv4. As one commenter in the Working Group session pointed out, declaring IPv4 Historic would likely backfire and serve no better purpose other than exposing the IETF to ridicule. Certainly there is some merit in wondering why a standards body would take a protocol specification

**Declaring IPv4 Historic would likely backfire and serve no better purpose other than exposing the IETF to ridicule.**

used by more than 3 billion people and approximately 10 billion devices every day and declare it to be Historic. In any other context, such adoption figures for a technology would conventionally be called outstandingly successful!

Perhaps we can put this into a broader context by looking at other Historic specifications. Unfortunately, the IETF does not have an obviously consistent story when declaring technical specifications Historic. Some very old and now unused services as described in Request for Comments (RFCs) are not declared to be Historic. For example, we can go a long way back in time to RFC 162, and find a specification that the completely forgotten protocol, Netbugger3, is not Historic. If there are any extant implementations of this curiously-named protocol, I would be keen to learn of them. Similarly, Gopher, a specification that enjoyed a brief moment in the sun in the early 90s before the juggernaut that is today's Web superseded it, is not, according to the IETF, Historic.

So if some pretty obviously defunct protocols are not Historic, what is? A browse of the Historic RFCs reveals a collection of TCP extension specifications, including RFC 1072 and RFC 1106. They were declared Historic by RFC 6247 with the rationale that they "have never seen widespread use". That's not applicable to IPv4 by any stretch of the imagination! Browsing the Historic RFC list at <https://www.rfc-editor.org>, it's evident that there are more than a few RFC documents that never even saw the light of day as current specifications, as they were declared Historic at the outset. Again, quite obviously, this is not applicable to IPv4.

What about other Internet Standards? There is a reasonable case to be made that Internet Standards numbers 23 and 24 (Quote of the Day and Finger, respectively) have long passed out of common use. Historic status seems to be entirely

*Continued on next page*

*Declaring IPv4 Historic: One Issue, Two Sides, continued*

applicable for both of those quite venerable standards, as their previously widespread use has waned to the point of almost complete invisibility.

So we appear to treat Historic status somewhat whimsically. We could be a little more consistent, and in that vein there is a case to be made to push Finger, Quote of the Day, Gopher, Netbugger3, and their like into Historic. The world has moved on and these protocols are stuck in an older world. But not IPv4. And not just because it is used by an unprecedented number of people and devices and we all still rely on it.

It's not just that.

It's because we probably haven't finished with IPv4 yet.

While many folk, including myself, would dearly like to see an all-IPv6 Internet today, I'd like to think I'm pragmatic enough to understand that we're stuck with a dual-stack Internet for many years to come. And that pragmatic observation has its consequences. So far, we have managed to cram some 10 billion unique devices into the Internet. The silicon industry is not going to stop and wait for us to complete this IPv6 transition, and, in the meantime, we can readily imagine a near future that crams every new computer, smartphone, personal pad, television, new car, and a whole heap of other applications, on this dual-stack Internet. We may well need to push the IPv4 Internet to encompass 20 billion or so devices on this strange and protracted dual-stack journey to IPv6.

One immediate change so far is the semantics of IPv4 addresses. Increasingly, IPv4 addresses are ephemeral short-term elements of conversation stream identifiers. They have lost any concept of being a stable endpoint identifier. The more devices we push into the network, the more we change the way IPv4 behaves. As we try and make IPv4 stretch just that little bit farther, we may need to make more subtle

changes, which may or may not impact the current specification for IPv4. We just don't know yet. What we do know is that right now the story is by no means over for IPv4.

As a standards body, it may sound like a good idea for the IETF to send a strong signal to the industry about the need to take this transition seriously by declaring IPv4 to be Historic. But if this is what the IETF does, then the work on IPv4 will probably continue. The risk is that it will continue without the benefit and support of the acknowledged Internet Standards organization, the IETF. And we have some prior experience with what happens then.

The last time the IETF turned its back on a technology specification was the development of the specification of Network Address Translation (NAT). The result was that implementers could not rely on a complete and coherent specification, and they were forced to make it up as they went along. Every NAT product had subtle behavioural differences with every other NAT product. The losers in this scenario were application developers and, ultimately, the users. Applications had to work across NATs and negotiate functionality across a diverse set of undocumented and, at times, inconsistent behaviours. The resulting environment can be brittle and fail in unanticipated ways.

Standards help us understand how to interoperate and how to rely upon predictable ways to interoperate. It takes a lot of the guesswork out of technology specification, and can eliminate a large amount of complexity in implementing technology. I would be horrified if we managed to repeat this mistake at this point in IPv4's life history. Oddly, it is now, while we continue to work through the dual stack phase of the transition to an IPv6 Internet, that the IPv6 part of the Internet needs a consistent and relevant IPv4 specification the most!

IPv4 Historic? No. We haven't finished with it yet! 🎉

*Message from the IETF Chair, continued from page 2*



For a summary of the meeting in video form, visit <https://www.youtube.com/watch?v=pdjunL22WZA>.

### **IETF Hackathon**

This was a wonderful experience, both in terms of what got worked on and the people who participated. There were more than 30 new participants, including more than 10 who were new to the IETF. See Charles Eckel's article on page 21. And thank you to Huawei, our new sponsor for all of this year's IETF Hackathon events.

I very much enjoyed Ole Troan's new project on adding source address routing to Vector Packet Processing (VPP). There was also work on DNS privacy, big data, and many other things.

### **Admin Stuff**

During the meeting, we announced that the IETF is creating an ombudsperson team to help handle any harassment concerns. For more about this team and how to report harassment, see <https://www.ietf.org/ombudsteam>.

Finally, Alia Atlas gave a talk at the plenary on challenges and opportunities associated with the IETF's changing environment. For example, our participation and funding models are changing as more participants attend remotely.

¡IETF 95 se concluyó! ¡Gracias a LACNIC, Internet Society, Buenos Aires, y a todos participantes! 🎉

*Words from the IAB Chair, continued from page 3*

introduced. Once a year we simply must go through a detailed outline of the accounts in public, lest basic transparency be lost. For the same transparency reasons, we cannot cut the open mic. Under these constraints, it was necessary that something be cut from the program, and the technical topic had to be it. But fear not! We expect to continue technical topics at the other plenaries in the year. Look for one in Berlin.

### **Time Demands and Making the IAB Work**

I noted at the beginning that, since I've been chair, the IANA stewardship transition has taken a lot of time. This has been frustrating for me because there are lots of other things that I wanted to do. But it has likewise been inspiring, and has reminded me how effective we are when we divide up the work.

**When the IAB spreads out its work among many different collaborating people, the work happens and we don't have a single choke point.**

Because I've had to devote so much time to the transition, the IAB as a whole has had to work harder to do what otherwise might be done by the chair. The programs, as you have seen from our reports, have become both more effective and more subject to regular review. The good consequences are, I think, seen in the workshops the IAB is holding to address pressing issues and the way that programs are producing topics that inspire IETF work. But to me, the other lesson is just important:

when the IAB spreads out its work among many different collaborating people, the work happens and we don't have a single choke point.

There is still work to do in this direction. The IAB chair automatically inherits certain jobs by virtue of being the chair. Why? Especially in an organization like the IETF, any IAB member is as able to speak as unilaterally for the IAB as the chair is. You may have noticed that we now sometimes send out notices from different members of the IAB saying "for the IAB". We think this is correct: the IAB speaks as one when it does speak, regardless who the mouthpiece is. What matters is that it reflect the IAB's view. We'll probably always need to have a chair to make other organizations think we work like they do. But we don't have to work that way for real.

### **Retreating to Advance**

Every year, the IAB holds a retreat, usually not too long after the new IAB is seated. The goal is to try to ensure that each IAB member has a clear understanding of what others' priorities are for the year, and to ensure that we have a common direction so that we work effectively together. This year, we met in Cambridge, Massachusetts, USA, on 17 and 18 May.

Inevitably, some of this is IAB members talking to each other—the goal, after all, is partly to ensure we're aligned. But the IAB tries to ensure that the discussions in our retreat are also responsive to factors impinging on the Internet. This year, our topics about those external factors included the ongoing influence of the so-called Internet of Things on the architecture of the Internet; this discussion led directly to the IAB's comments to the United States National Telecommunications & Information Administration in response to their request ([https://www.iab.org/documents/correspondence-reports-documents/2016-2/iab-comments-to-ntia-request-for-comments-the-benefits-challenges-and-potential-](https://www.iab.org/documents/correspondence-reports-documents/2016-2/iab-comments-to-ntia-request-for-comments-the-benefits-challenges-and-potential-roles-for-the-government/)

[roles-for-the-government/](https://www.iab.org/documents/correspondence-reports-documents/2016-2/iab-comments-to-ntia-request-for-comments-the-benefits-challenges-and-potential-roles-for-the-government/)). We talked about cross-organization workshops: what has worked, what could use improvement, and what more of this we need to do. By the time you read this, the IAB-cosponsored Internet of Things Software Update workshop will have happened (<https://www.iab.org/activities/workshops/iotsu/>). We also discussed developments in Internet architecture that tend to promote the power or control of the network operator. And we had the good fortune of welcoming Danny Weitzner, Taylor Reynolds, and Dave Clark from the Massachusetts Institute of Technology Internet Policy Research Initiative. They discussed with us the ways that the IAB can and cannot interact effectively with policy makers. Our goal always is to identify those issues that are relevant to the Internet as a whole, and to find the people who are interested in the topic and can help make it better.

### **Let's Advance**

By the time you read this, the IETF will be meeting in Berlin for IETF 96, and the IAB will be pressing ahead on its issues: keeping the different parts of the Internet working as a coherent whole, while remaining faithful to the core design of a network of networks. If you want help understanding how different parts might fit together, or want another point of view on an issue you're trying to sort out, feel free to ask us for it. Send us mail at [iab@iab.org](mailto:iab@iab.org) or discuss your topic on [architecture-discuss@iab.org](mailto:architecture-discuss@iab.org). Or, if you're in Berlin, you can just talk to us. We have red dots, and we're all friendly. Even me. 

**We'll probably always need to have a chair to make other organizations think we work like they do.**

## MEET THE IETF SYSTEMS

By Allison Mankin

**D**URING NEARLY EVERY IETF meeting since 1993, an informal gathering of women participants, the Systems, has taken place. We chose the name *Systems* as an answer to the late Anita Borg's call for women in computer science and engineering to support and celebrate each other.

In 2013 and 2014, gifts by Comcast, EMC, and Verisign Labs established a lunch fund for the gathering. For most of the participants, the Systems gathering is a chance to catch up with friends across all areas of the IETF, to employ an informal mentoring and information gathering forum, and to encourage each other in a largely male-dominated field.



IETF 95 Systems

The Systems IETF list, [systems@ietf.org](mailto:systems@ietf.org), offers this kind of support before and after the face-to-face meetings. The list is for Systems involved in IETF topics—both technical and specific to women. Traffic is typically light with some discussions, but mostly for organizing per meeting gatherings. The list is open to any woman interested in the IETF, whether she participates only by mail or also in person.

If you are interested in learning more about the Systems or contributing to our fund, please contact [systems-admin@ietf.org](mailto:systems-admin@ietf.org).



## TURNING 30, THE IETF SEEKS IDEAS FOR GROWTH

By Carolyn Duffy Marsan

**A**S THE IETF CELEBRATES ITS 30TH BIRTHDAY, THE GROUP'S LEADERSHIP team is looking for ways that the standards body can remain influential and effective for the next 15 years. It invites everyone to participate in this ongoing discussion.

During the plenary session at IETF 95 in Buenos Aires, the Internet Engineering Steering Group (IESG) discussed Internet trends and observations affecting how the IETF operates. Routing Area Director Alia Atlas gave a report from a design team that has developed a draft entitled, IETF Trends and Observations. At issue is how the IETF should continue evolving to meet its goal of making the Internet work better.

"We have changed the world, and we all know that," Atlas said. "The Internet is critical to public and private life. People who are under 30 [years old] can't imagine what the world was like before the Internet... Now the IETF is living in the world that we helped to create, and that creates more opportunities for us."

In particular, the IESG is looking for ways that working groups can do more of their business online using cutting-edge collaboration tools. In addition, the IESG hopes to enhance remote participation, develop local hubs of activity, and reduce its financial dependency on hosting three large meetings per year.



Alia Atlas, panelist and routing area director

**"The Internet is critical to public and private life. People who are under 30 [years old] can't imagine what the world was like before the Internet."**

—Alia Atlas  
Routing Area Director

"We're going to keep changing because we're going to keep taking advantage of the technology and collaboration abilities that we enabled," Atlas said. "We need to continue expanding our community, expanding our social circle. We need to add more people to our meetings... but we also need to keep the operators, developers, and researchers comfortable participating in the low-volume way that they want to."

Atlas pointed out that at its first meeting, the IETF had only 30 people with "the simple idea of rough consensus and running code. The question as we look ahead is: How can the IETF continue to be true to our roots, thrive in this world, and create the future Internet that we need?"

Atlas explained that the Internet Society is reorganizing its support for the IETF and seeking more global sponsors for what it now calls The IETF Endowment. She said the IETF must transition its funding structure away from in-person meetings to a more sustainable structure that will support an increase in remote participation.

She said companies are interested in sponsoring the IETF because it “is a trusted technical authority. People respect the work we do. They know that we understand the technology and that we care that it makes the Internet keep going and get better.”

In addition to increasing remote participation, the IESG hopes to create local hubs with active communities engaged in technical sharing, Hackathons, and social activities.

“We’re going to spread the idea of the IETF and grow the community,” she said. “There are two things that tie us together: one is our love of technology and finding a good practical solution, and the other is finding someone else—one or five other people—to have an awesome technical discussion with... The community is where we have our strength.”

Atlas said the IETF must do a better job of communicating what it is doing and to be more outward-focused, rather than exclusively inward-focused. “New communities may be joining us because they want a technology standardized. We need to be welcoming,” she added.

Atlas concluded by asking IETF participants to read the draft and participate in the mailing list discussion at [ietf@ietf.org](mailto:ietf@ietf.org). “What we really want is your ideas on how the IETF should adapt and improve,” she said. “We’re looking for community

discussion that converges and sets new direction.”

Also at the plenary session, the IAB described its recent Internet of Things Semantic Interoperability Workshop and its Names & Identifiers Program.

IAB member Dave Thaler said the IoT workshop was “extremely productive”, attracting almost 40 attendees to discuss the many different definitions and schemas emerging for various objects in the evolving IoT area.



Dave Thaler, panelist and IAB member



Suzanne Woolf, panelist and IAB member

Suzanne Woolf gave an overview of the IAB Names & Identifiers Program, which has several drafts about the history and semantics of domain names, what an idealized naming system might look like,

and how to look at names and naming in context. Additionally, the IAB held a Birds-of-a-Feather session in Buenos Aires aimed at looking beyond DNS and default context for Internet names.

In other news, Scott Bradner was given a standing ovation after it was announced that he would retire in June. Bradner has been active in the organization since IETF 16 in 1990. He has published 44 Request for Comments (RFC) documents and is the author of the most-cited RFC (2119), which outlines key words for use in RFCs to indicate requirement levels. He was area director in four different areas and is a current Internet Society Board Member and IETF Administrative Oversight Committee member. Bradner was the second person to win the IETF’s highest honor, the Postel Service Award, after Jon Postel himself.

“For me, you have been the person to look up to. It was very easy to work with you. You always have an intelligent answer, and you always go all the way thinking about topics and trying to do the right thing,” said IETF Chair Jari Arkko as he presented Bradner with an award for being the “Mother of Consensus”.

“In general, it’s been positive for me and hopefully for the organization. But the time does come, and it has. Thank you very much,” Bradner said. 🍀



Scott Bradner accepts an award from Jari Arkko for being the “Mother of Consensus”.

# USING I2NSF FOR OVERLAY NETWORK TO AVOID DISTRIBUTED DENIAL OF SERVICE ATTACKS

By Linda Dunbar

IN TODAY'S WORLD, WHERE EVERYTHING IS CONNECTED, PREVENTING unwanted traffic has become a key challenge. More and more networks, including various types of Internet of Things (IoT) networks, information-centric networks (ICN), content delivery networks (CDN), and cloud networks, are in some form of overlay network with their paths (or links) among nodes that are provided by other networks (aka, underlay networks). These paths are considered a single hop by the virtual networks. The approach of overlay networks having their own security solutions cannot prevent various attacks from saturating the access links to the overlay network nodes, which may cause overlay nodes' CPU/links to become too over-utilized to handle their own legitimate traffic.

Very much like traditional networks placing a firewall or an intrusion prevention system (IPS) on the wire to enforce traffic rules, Interface to Network Security Functions (I2NSF) can be used by overlay networks to request that certain flow-based security rules are enforced by underlay networks. By this mechanism, unwanted traffic, including

Very much like traditional networks placing a firewall or an IPS on the wire to enforce traffic rules, I2NSF can be used by overlay networks to request that certain flow-based security rules are enforced by underlay networks.

DDoS attacks, doesn't appear on the physical links and ports to the overlay network nodes, thereby avoiding excessive or problematic overlay node CPU/storage/port utilization.

I2NSF has two types of interfaces: a service layer and a capability layer. The service layer specifies how a client's security policies may be expressed to a security controller. The capability layer specifies how to control and monitor flow-based security functions (NSFs) at a functional implementation level.

The policies over the Service Layer Interface don't care which NSFs are used to enforce the policies. There could be multiple NSFs to enforce one Service Layer policy. The policies over the Capability Layer Interface are to specific NSFs.

The I2NSF Working Group was chartered in October 2015. From IETF 94 to 95, I2NSF WG has adopted four Internet-Drafts: draft-ietf-i2nsf-problem-and-use-cases-00, draft-ietf-i2nsf-framework-00, draft-ietf-i2nsf-gap-analysis-01, and draft-ietf-i2nsf-terminology-00. The Working Group has agreed to use the Event-Condition-Action paradigm for the Flow Based Security Rules polices over both the Service Layer Interface and the Capability Layer Interface.

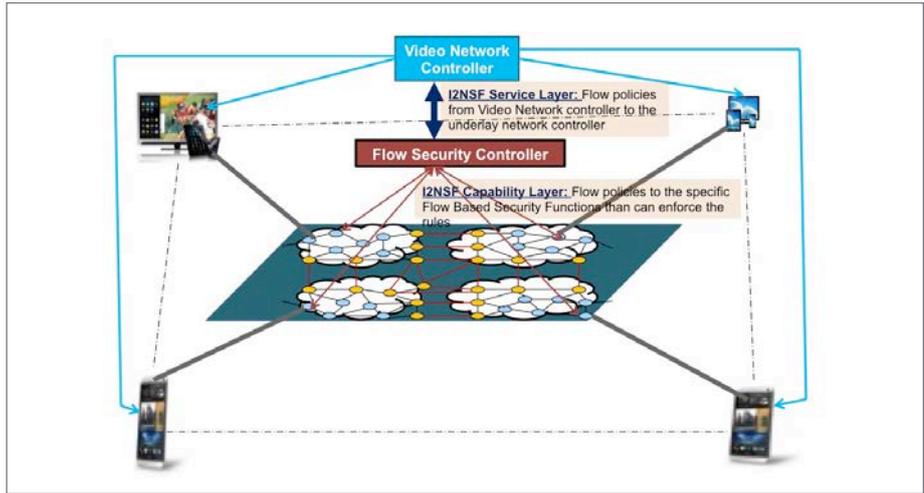


Figure 1. The I2NSF Service and Capability Layers

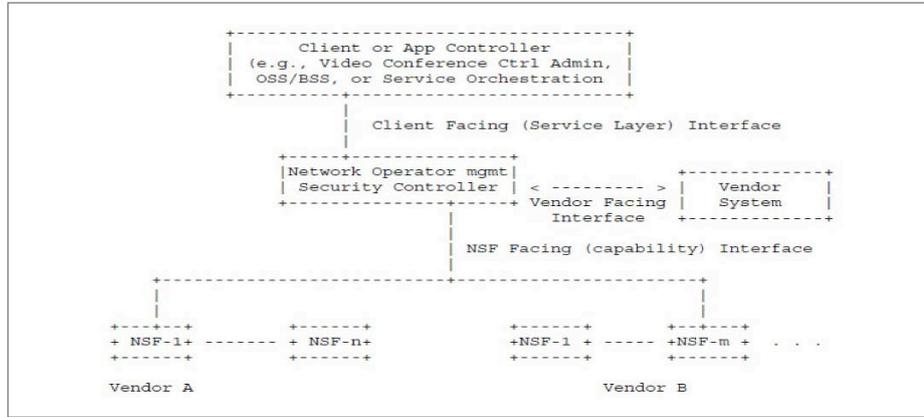


Figure 2. The I2NSF Framework

# INTELLIGENT TRANSPORTATION SYSTEMS AND THE IETF

By *Alexandre Petrescu and Carlos Pignataro*

**I**NTELLIGENT TRANSPORTATION SYSTEMS (ITS) IS A GENERIC NAME FOR USING any of a wide range of information technologies to move people, freight, and other devices across roads, waterways, air, and space. Uses may include Internet access in cars, trains, and planes; multimodal itinerary planning across smart cities; high-speed multioperator road and park tolling; goods-delivery tracking; traffic supervision and management; self-driving cars and car platooning; emergency calls; and highly-improved safety of traffic of automobiles and trucks. To support such a wide range of uses, applications need reliable communication capabilities across complex systems involving a variety of mobile and fixed devices with disparate wireless and wired links. To that end, technical committees at the International Standards Organization (ISO) and the European Telecommunications Standards Institute (ETSI), an ITU collaboration effort, and a connectivity-focused US government programme are but a handful of organizations and activities carrying an ITS acronym in their names (e.g., ISO TC204 ITS, ETSI TC ITS, ITU Collaboration on ITS Communication Standards, USDOT ITS JPO).

Acknowledging that many design requirements of the early-stage Internet were related to safety, reliability, and heterogeneity, and considering the successful Internet deployments of recent decades, it is tempting to contemplate the use of the TCP/IP family of protocols to support applications in ITS use-cases. At IETF meetings, participants are versed in the design and deployment for such requirements. For example, one illustration of reliability is the best-effort nature of IP packet delivery on a path through a maze of routers worldwide, with an available alternative if a known path fails, across heterogeneous networks.

The communication systems currently used in transportation are satisfactory in some use-cases. For example, multimodal itinerary ticketing, 50km/h toll passage, and incumbent car platooning rely extensively on dedicated communication systems; they are all successful in trial phases, even though the penetration of the current<sup>1</sup> IP family of protocols in these use-cases is relatively limited, if not outright absent.

ITS discussions at the IETF offer hope of TCP/IP protocols in vehicular communications in the near future. TCP/IP protocol

stacks are already present in many cars that are connected to the Internet with a cellular modem, typically LTE. In addition, widespread in-car technologies like Mirror-Link, Apple CarPlay, and Android Auto use IP. Current demonstrators of security mechanisms for out-of-car DSRC<sup>2</sup> need an ancillary Internet connection, often LTE with TCP/IP support, to realize the transfer of security material (certificates, revocation lists, and so on).

**One illustration of reliability is the best-effort nature of IP packet delivery on a path through a maze of routers worldwide, with an available alternative if a known path fails, across heterogeneous networks.**

Recent demonstrators featuring vehicle-to-vehicle (V2V) use-cases (e.g., platooning), rely on ITS-G5 link-layers and CAM application layers. Within these demonstrators, the platoon size limits (number of cars in a platoon) have been exhibited due to the radio range of ITS-G5. It is believed that the involvement of IP protocols (with ETSI ITS applications) featuring IP subnet structures between cars may lead to arbitrary-size platoons (like the size of the Internet can be arbitrary), where packets are IP-forwarded rather than broadcasted.

**Recent lessons learned at a European demonstration event illustrated the necessity of sending vehicle position corrections over IPv6.**

Adopting TCP/IP would also help address the fact that the pairing operation of vehicles is independently being developed by organizations that often overlook fundamental interoperability requirements. Furthermore, recent lessons learned at a European demonstration event illustrated the necessity of sending vehicle position corrections over IPv6.

DSRC and ITS-G5 messages, such as BSP or CAM, are broadcast periodically (IP is not used); there is a need for a mechanism to allow the sender to learn whether or not a message was received with a certain degree of reliability. Using the request-response semantics of some IP protocols may help achieve improved

*Continued on next page*

*Intelligent Transportation Systems and the IETF, continued*

reliability when necessary (e.g., ICMP Neighbor Solicitation/Advertisement, or TCP SYN/ACK).

Arguably, where smartphone-to-server TCP/IP is the preferred method of mobile interaction, rarely, if ever, does a deployed multimodal travel planning application run on IPv6. More importantly, too often application-glued-on-link communication protocols (i.e., protocols without networking layers<sup>3</sup>) are involved in communications between automobiles. In this context, further involvement of applications relying on TCP/IP and of IP forwarding mechanisms is expected to result in significant improvements to the security of communications (IPsec), and orders of magnitude more interactions between numerous directly reachable devices in vehicles (IPv6). Applications that are unimaginable today will be possible, when every car can talk to every other car around the world, as computers do via the Internet. And since the value of the network grows with the number of connected parties, it is expected that the Internet's value and reach will increase even more when cars are connected. The potential growth can be further illustrated by vehicles forming an independent network on a road linking smart cities; to some extent the question whether to connect the network of vehicles to the Internet may be turned the other way around.

**The ITS BoF in Buenos Aires**

Participants at the IETF have published Internet-Drafts that are explicitly or implicitly related to ITS use-cases on many occasions in recent years. In April 2016, at the IETF meeting in Buenos Aires, a Birds-of-a-Feather (BoF) meeting, chaired by

**A tutorial on the use of IP in vehicle networks exposed the advantages of a narrow-waist networking layer, including the support of link layers, such as 802.11-OCB with a variety of modulation methods.**

Carlos Pignataro and Russ Housley, was held specifically on the topic of ITS. Problem statements were discussed about the use of IP in vehicle networks: IP for vehicle-to-vehicle and vehicle-to-infrastructure communications, IPv6-over-foo, IP path establishment, and naming. The presented use-cases involve direct communications

between vehicles (V2V), for example, vehicle platooning. Additional use-cases involve communications between a server situated along or near the road (Road-Side Unit) and vehicles passing by. A tutorial on the use of IP in vehicle networks exposed the advantages of a narrow-waist networking layer (compared with network layer absence or with other link-specific or application-specific networking layers), including the support of link layers, such as 802.11-OCB (also known as DSRC or 802.11p), with a variety of modulation methods (e.g., WiFi, LTE, and VLC). Other aspects of using packetised data exchange principles were described as comparing favorably to the use of bouncing-signal principles of communication between vehicles, such as when Light Detection And Ranging (LIDAR) or cameras are used (Figure 1). These two aspects raised a number of comments from the audience; together with previously expressed security and privacy concerns, these comments can be found in the meeting minutes.

The establishment of a Category A liaison between the IETF and the ISO Technical Committee TC204, ITS, was announced during the ITS BoF in Buenos Aires. One liaison statement from ISO/TC204 was announced, with a slide set from the ISO/TC204 liaison officer.

**Charter Work and Interim**

The initial text of a charter<sup>4</sup> for an ITS WG was presented in Buenos Aires. In the initial phases of charter writing, a significant

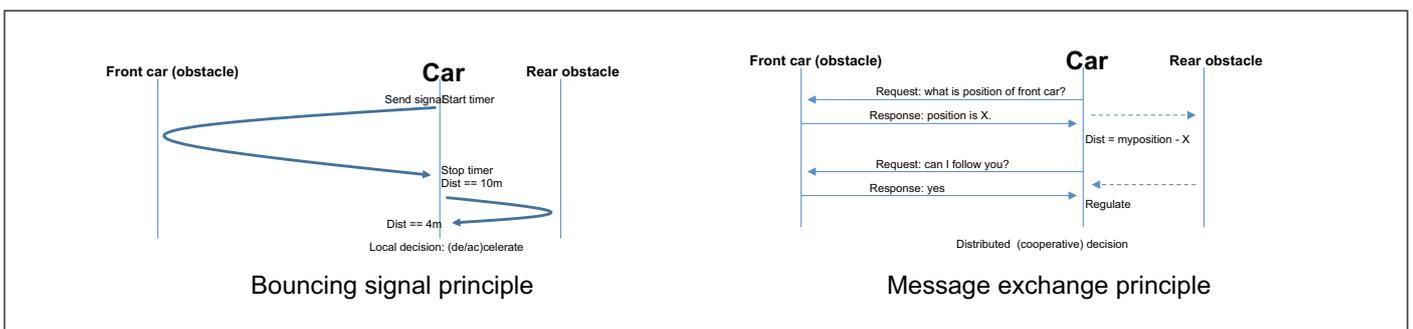


Figure 1. Vehicle-to-vehicle communication principles

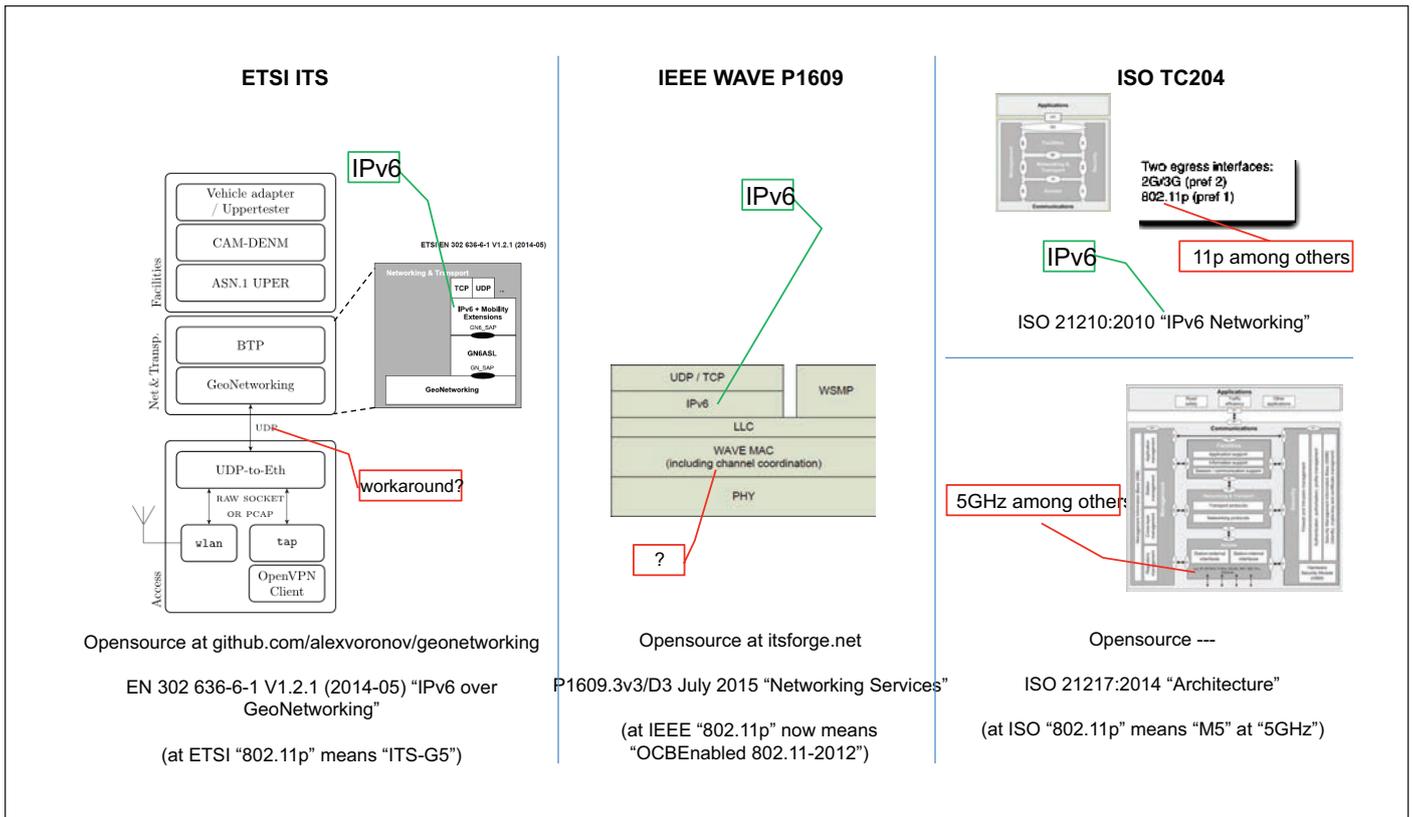


Figure 2. Models of Running IPv6 over 802.11p

number of work domains were suggested: communications between automobiles (V2V, V2I), space, airline, and unmanned aerial vehicle communications, information- and content-centric networking applied in vehicular communications, alternative mobility protocols and locator-identifier split for networked vehicles (AERO, LISP) and more. Facing this potentially enormous scope, extensive discussions led to more than just improving the text, it helped narrow down the number of deliverables: two Informational documents on the context and the problem statement for the use of IP in vehicular communications, and one Standards Track document on “IPv6-over-802.11p”. This charter structure was further finalized during the virtual interim ITS BoF held on 31 May 2016 via audio-conference with remote slide presentation. The details are described in the minutes of the virtual interim meeting.

The item “IPv6 over 802.11p” is regarded as a typical IETF “IPv6 over foo” doc-

ument, based on “IPv6 over Ethernet” RFC 2464. The model of an IPv6-over-802.11p layered stack of protocols can be compared against other models of running IPv6 over 802.11p (DSRC) found at pertinent Standards Development Organizations. Three such models are illustrated in Figure 2.

Prior to the BoF in Buenos Aires, the topic of vehicular networks was presented at the IETF 93 technical plenary in Prague, Czech Republic. Presentations from academic and industry experts in vehicle networks, security, and standardization were discussed. For more on the plenary, see “Vehicular Networks Are Expected to Save Lives But Carry Privacy Risks,” *IETF Journal*, Vol. 11, Issue 2 (<https://www.internetsociety.org/publications/ietf-journal-november-2015/vehicular-networks>).

**A Proposal for an ITS WG**

The goal of the proposed ITS WG is to standardize and/or profile IP protocols for

establishing direct and secure connectivity between moving networks.

**Definitions**

The Working Group defines the terms V2V, V2I, and V2X as follows:

V2V (vehicle-to-vehicle communications). The communications can be direct (without requiring an access point or relay by the road-side), or indirect (relying on one or multiple relays along the road-side).

V2I (vehicle-to-infrastructure communications). Data flows happen between a mobile vehicle and a server in the fixed infrastructure nearby. Sometimes V2I stands for vehicle-to-Internet communications, referring to a server anywhere in the Internet.

V2X (vehicle-to-‘any other’ communications). In some contexts it is a handy term to mean both V2V and V2I at the same time, e.g., “V2X technology enables

*Continued on next page*

Intelligent Transportation Systems and the IETF, continued

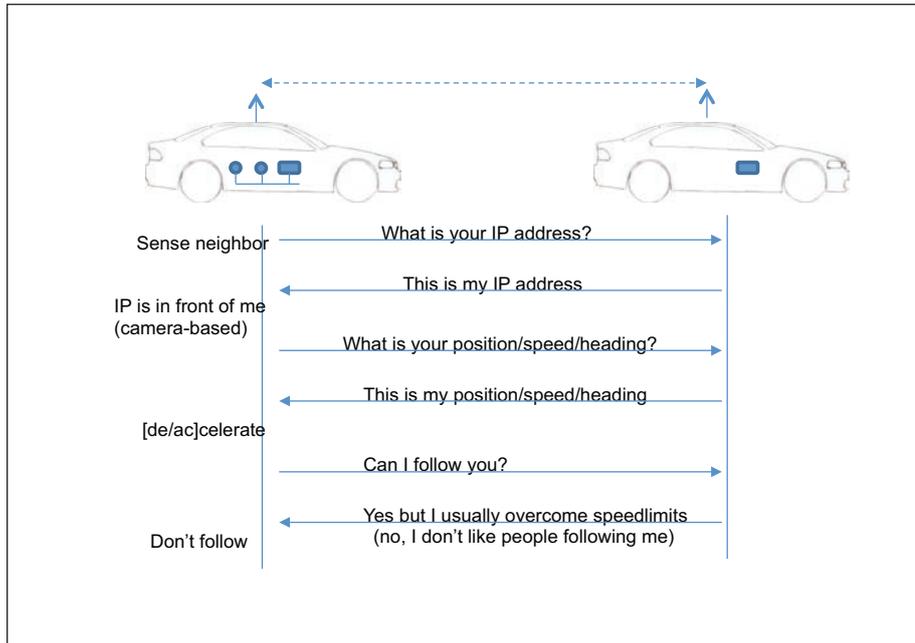


Figure 3. Example Application Using IP Messages for C-ACC

a vehicle to stay connected to both the Internet and other cars". In other contexts it means vehicle-to-'something other than vehicle or infrastructure, most notably a human' communications, e.g., V2P (vehicle-to-pedestrian), V2N (vehicle-to-nomadic pedestrian), and V2D (vehicle-to-device of pedestrian).

**Models and Use-Cases**

Two use-cases were discussed at the BoF: cooperative adaptive cruise-control (C-ACC) and platooning. These communication models are illustrated in Figures 3 and 4.

**Looking Forward**

The charter text is now stable, and work has started on the initial work items. Four Internet-Drafts have been identified as good candidates for the first three work items. More people have joined the email list and some of those have expressed interest in submitting Internet-Drafts to address the goals in the current proposed charter.

If you are interested in the use of IP protocols in vehicular communications, please subscribe to the email list <https://ietf.org/mailman/listinfo/its> and submit an Internet-

Draft targeting one of the three proposed work items: ITS General Problem Area, IPv6 over 802.11p, or Problem Statement. You are also invited to read the existing Internet-Drafts in this group, review them, and make comments. 

**Footnotes**

1. Ahead of general thinking, the IETF considers the current version of the IP family of protocols to be IPv6. Work is being considered to declare IPv4 as "Historic", or "Restricted Standard", and, simultaneously, work is ongoing to promote IPv6 as "Internet Standard". See page 1.
2. The term Dedicated Short-Range Communications (DSRC) is used with multiple meanings. An earlier IETF Journal article defines DSRC as "802.11e for quality of service; 802.11j-2004 for half-clocked operations, which are a more robust form of communication; and 802.11p for operation in the 5.9 GHz band and a new mode called OCB for Outside the Context of a Basic Service Set." Also, DSRC MAC and PHY layers are defined by ASTM E2213 - 03(2010), which may refer to IEEE documents. The DSRC application layer is defined by SAE J2735\_201603. In Europe, DSRC is defined by CEN/TC278. An additional standard used in Europe for 5 gigahertz, in lieu and place of DSRC, is defined by ETSI as "ITS-G5".
3. An early example of an application protocol glued onto the link layer (without a network layer) is the Wireless Application Protocol (WAP). It was used in the initial deployments of interactive applications in the first smartphones, only to be phased out by the arrival of TCP/IP. Today, WAP has largely disappeared, yet similar tendencies persist to develop such protocols within and outside ITS.
4. See the latest proposed charter at <https://tools.ietf.org/wg/its/trac/>.

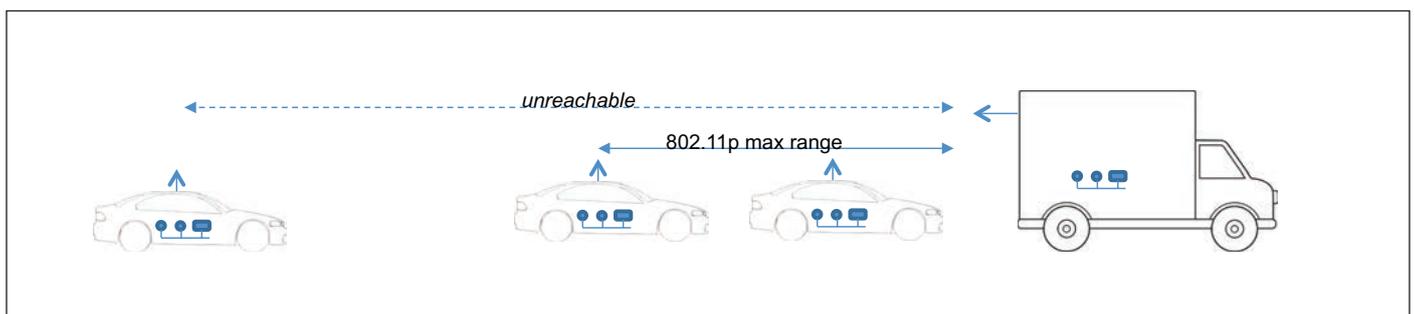


Figure 4. Scalability and Interoperability Issue of Initial Non-IP Platooning Demonstrators



## THINGS TALKING TO OTHER THINGS ABOUT THINGS

By Dave Thaler and Andrew Sullivan

**T**HE IETF IS ABOUT INTEROPERATION. *ROUGH CONSENSUS AND RUNNING code* are all about making diverse things work together as much as possible. One of the places that things—in this case, “Things”—need to line up is in the application layer.

For the Internet of Things (IoT) to become the reality many popular accounts would suggest, various kinds of Things need to be able to talk to one another, and not only at the lowest levels. For example, one promise of the Internet of Things is that the lights and the thermostat and the garage door can all collaborate to make your house more comfortable. And the whole system is likely to be better overall if each part works together, no matter who made each device—just the way the Internet has grown and succeeded.

A key theme of Dave Thaler and Hannes Tschofenig’s talk at the IETF 92 Technical Plenary was the duplication and gratuitous differences arising from many organizations independently defining data models, or schemas, for each type of IoT device. For example, there were already many different definitions of what a light bulb was! As a follow-up to help tackle this problem, the Internet Architecture Board organized the Internet of Things Semantic Interoperability (IoTSI) workshop held 17–18 March 2016.

Facing this issue brought many people together, including, but not only, those who participate in the IETF, World Wide

Web Consortium (W3C), Open Mobile Alliance (OMA), AllSeen Alliance, Open Connectivity Foundation (OCF), National Institute of Standards and Technology (NIST), CableLabs, ZigBee, and European

**One of the very encouraging items from the workshop is that people from many different sectors of the industry all agree that there is a serious problem to be solved.**

Telecommunications Standards Institute (ETSI). We convened at the IoTSI workshop in the Ericsson offices in Santa Clara, California. For two days, we tried to work out ways to improve semantic interoperability.

How can diverse systems interoperate? Are better standards in information models or data models needed? Is a single framework necessary or is some sort of mapping possible? What can you do when frameworks are formally incompatible? What do we do about end-to-end security when intermediate security models are incompatible?

One of the very encouraging items from the workshop is that people from many different sectors of the industry all agree that there is a serious problem to be solved. Some groups had already started developing common solutions for some things, and the level of information sharing across the group was quite remarkable. This is how interoperation works best: not by trying to impose a single model, but by people with different interests all recognizing a common problem.

Of course, recognition is just a first step. Work still needs to be done to move from recognition to results. While a workshop report is in progress, more important are the follow-on activities. We agreed to start with a wiki to provide pointers to schema repositories, with further developments to follow. We in the IETF, in other standards developing organizations, and in industry have an opportunity to make interoperability in the Internet of Things the positive force that earlier Internet innovations were. Interoperation is what we do, so let’s do it again.

For more information, visit <https://www.iab.org/activities/workshops/iotsi/>. 

# INTERNET REGULATORS, TECHNOLOGISTS SEEK ONGOING DIALOGUE

By Carolyn Duffy Marsan

**T**HE INTERNET SOCIETY SHOULD CONTINUE TO FOSTER THE DIALOGUE between Internet policymakers and IETF technologists with the goal of creating a more open and secure Internet. That was the consensus of an Internet Society-sponsored panel discussion held in Buenos Aires during IETF 95 entitled, Public Policy and Internet Technology Development.

Moderated by Olaf Kolkman, the Internet Society's chief Internet technology officer, the panel featured policymakers from Chile, the Dominican Republic, and Fiji, along with long-time IETF members familiar with regulatory issues.

Dilawar Grewal, vice president of the University of the South Pacific in Fiji, said IETF engineers working on the building blocks of the Internet don't always think about how policymakers will interpret the technology that they are developing.

"A direction from the policy perspective to the technology makers and from the technology makers to the policymakers would be something that could actually change our environment quite dramatically in the future," Grewal said. "That's the kind of bridge we ought to build between what the IETF does and what policymakers do. Right now, there is no connection."

Nelson Guillén, a regulator with the Dominican Republic's INDOTEL, said policymakers are interested in the IETF's work "to make a better and open and stable Internet." Of specific interest are the IETF's work on cybersecurity and quality of service, he added. "We don't have certain information to defend the user's rights to receive the quality level that they have contracted with the company to receive, so Internet service measurements and how to measure them and what's the best tool" are of interest, Guillén said.

Raúl Lazcano, head of the regulatory division of Chile's SUBTEL, said regulatory agencies like his don't have technical

professionals on staff who understand the inner workings of IETF standards.

"The IETF rules are deeply technical, and we have a lack of professionals with the technical proficiency to analyze them," Lazcano said. "The question for the community is: How can you help us be nearer to the IETF, to analyze and to participate more in the technology development?"

Tim Polk, assistant director for cybersecurity at the US Office of Science and Technology Policy, said his agency favors multistakeholder organizations like the IETF for standards development.

"We believe that US companies and others will do better if the cybersecurity standards they use are internationally accepted," Polk said. "From a technologist point of view, I see the IETF position is to deliver

the best technical standards to support the broadest range of policy options. Not all policies in all countries are the same."

However, Polk pointed out that the IETF's cybersecurity standards haven't been created in such a way that people are motivated to implement them. "We have the standards out there that we know if everyone implemented them, the Internet would be a better place. The challenge I see in the policy space is how to motivate adoption," he said.

Cisco Fellow Fred Baker said that often the conversation between Internet technologists and regulators breaks down because they each have distinctive expertise. "Often times, when I'm talking with policy regulators, I'm dealing with misconceptions and baggage and trying to help them understand the world in which they are living so they can make better regulatory decisions," he said.

Polk agreed, noting that an important role for the IETF community is to educate policymakers on what is and isn't possible with the Internet. "Often regulators have a completely different idea of what the Internet is and how it works... It isn't that they aren't well intentioned, but they don't understand how it all works," he said. "They have a responsibility to get educated to reflect the



Raúl Lazcano, panelist and head of the regulatory division at SUBTEL, speaks at the Internet Society-sponsored plenary.

**Polk pointed out that the IETF's cybersecurity standards haven't been created in such a way that people are motivated to implement them.**

real state of the Internet, and we have a responsibility to help them understand what can and cannot be done."

As a rule of thumb, Grewal encouraged IETF engineers to design open standards to ensure that the Internet brings economic development opportunities to all, not just to some. "When technical people focus more on interoperability in standards, they open up a whole lot of doors for policymakers and for the users," he said.

Polk added that another benefit of open standards is that they foster innovation. "If standards are done poorly, they can become an inhibiting factor," he added.

Guillen said it's important to foster interaction between Internet technologists and policymakers because they both want a better Internet for all, and IETF meetings are a great place for these conversations.

"It's good that we know the process of [standards] creation that the IETF uses. That might help us make better regulations," he said. "We should also regulate, but trying to keep things as open as possible in order to let people keep developing and designing new opportunities to use the technology."

Lazcano agreed that it is important for the IETF and policymakers to keep working to bridge the gap between technical issues and political rules through continuous dialogue. "If I have the advisory of a good technical community, I can do rules in the correct way. So these kind of meetings are very important for me," he said. 

## WORKING GROUP UPDATE: L3SM Service models in the IETF

By Adrian Farrel and Qin Wu

**T**HE LAYER THREE VPN SERVICE MODEL (L3SM) WORKING GROUP HAS succeeded in its goal of describing an Internet service as a data model by documenting it in the YANG modelling language.

Driven by the huge interest in Software-Defined Networking (SDN), the use of YANG-based data models is now very popular in open-source projects, throughout the IETF, and in many other standards bodies and forums. A recent count showed more than 200 active Internet-Drafts and Request for Comments (RFCs) describing YANG data models. However, the focus of that work has been

"southbound" of the network controller to configure and monitor network devices and protocols.

The purpose of a service model is to formalize the description of a service on the interface between a network operator and that operator's customers. In this context, an operator may use the data model to describe and discuss services that the operator can supply, and the customer may request a service using the data model transmitted either on paper or via an Information Technology (IT) system. Increasingly, with the growth in SDN projects and products, consideration is given to automating service request and delivery via dynamic software systems.

Ultimately, the intent is that a network operator will convert a customer's request into configurational and operational parameters that control network resources to

**A recent count showed more than 200 active Internet-Drafts and Request for Comments (RFCs) describing YANG data models.**

*Continued on next page*

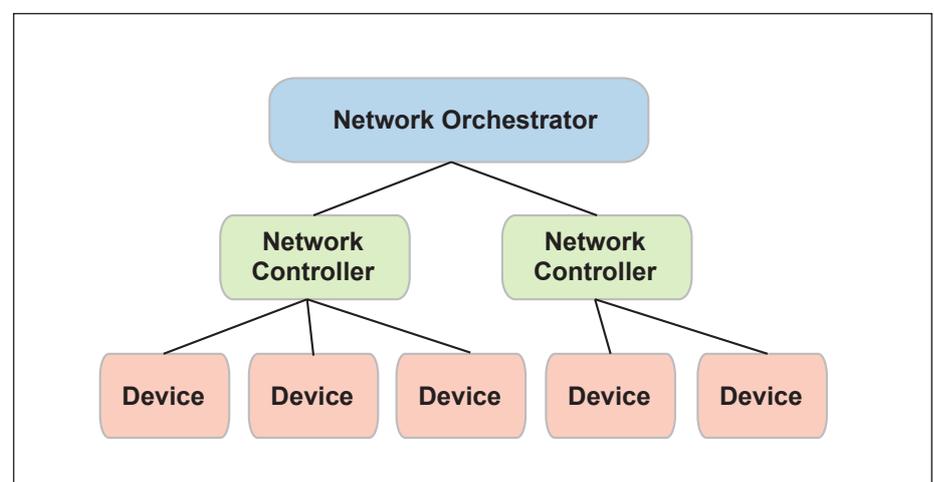


Figure 1. A Simplified SDN Architecture

Working Group Update: L3SM, continued

**Because network operators are the consumers of the data model, it was essential to both involve them and have them control the work.**

deliver the service. In an SDN system, the control and configuration of network resources and protocols are under the control of software systems that determine how best to utilize the network. Figure 1 shows a simplified view of a common representation of the SDN architecture: the network orchestrator plans how the network should be used and communicates with the network

controllers that configure and program the network devices.

The service model applies a level of abstraction so that it contains only the questions operators would ask their customers in order to activate the service (versus including all possible configuration knobs for the devices). That is, because a service request is network agnostic, it must be mapped onto the network orchestrator’s view of the network. This can be achieved by introducing a service orchestrator, as shown in Figure 2. The service orchestrator receives service requests from the customer and maps them to the correct network orchestrator of the operator’s network (or networks) that was chosen to deliver the service.

The L3SM Working Group took a different approach to working from the usual way. First, the work was driven entirely by network operators, not equipment vendors. Participation in the IETF by operators is a precious resource that can focus our work on real problems that need to be solved. Because network operators are the

consumers of the data model, it was essential to both involve them and have them control the work. In addition, it has been a challenge for operators to agree on a common set of parameters to describe the L3VPN service that they each offer in their own unique way. Achieving this agreement was one of the ways the Working Group measured its success.

L3SM also is unusual in that it was created by Area Director Benoit Claise without holding a Birds-of-a-Feather (BoF) meeting. Instead, as soon as the team of operators had written a relatively early, but stable version of the data model and posted it as an Internet-Draft, Benoit chartered the Working Group. This is an example of how the IETF can be relatively quick at handling and progressing new work: the L3SM Working Group expects to last call this data model around IETF 96—14 months after the Working Group was chartered and only 15 months after the first version of the draft was posted.

The L3SM Working Group and draft can be found at <https://datatracker.ietf.org/wg/l3sm>. 

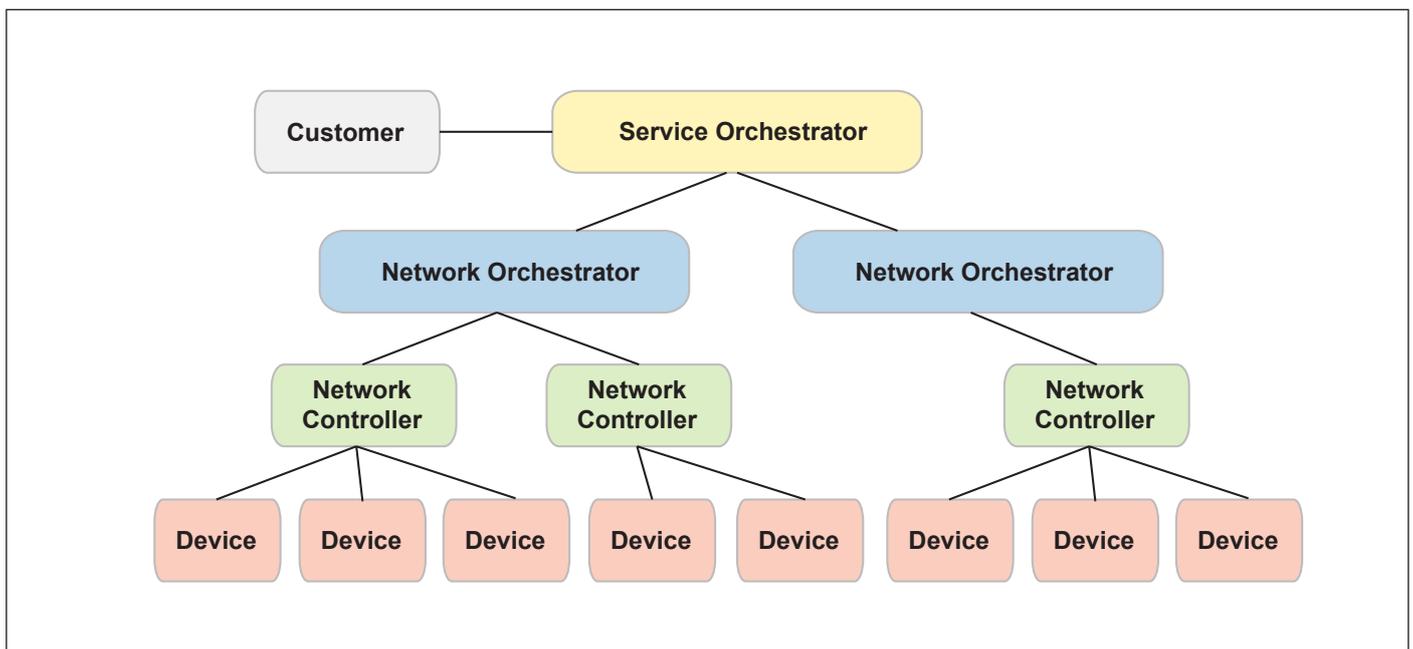


Figure 2. An SDN Architecture Showing the Service Orchestrator

# WORKING GROUP UPDATE: TAPS

## Transport Services can play a vital role in transport protocol evolution

By Zaheduzzaman Sarker and Aaron Falk

THE TRANSPORT SERVICES (TAPS) WORKING GROUP ADDRESSES ONE OF the trickiest challenges that application and service developers face while designing and deploying an application or service: choosing a transport protocol to use. Different applications and services have different requirements. Some, for example, can be tolerant of delay but not of lost data, like file download; others are very sensitive to delay but can accept some data loss, like interactive video. Over time, the IETF has standardized several transport protocols in order to address different application requirements, including TCP for reliability, UDP for unreliable and unordered data, Stream Control Transmission Protocol (SCTP) for multiplexed data streams, and Datagram Congestion Control Protocol (DCCP) for congestion control without reliability. Developers need to understand the capabilities of the various IETF protocols to determine which one is the best fit for their applications and services. If they identify a protocol that is not vanilla TCP or UDP, it may not work end-to-end. Frequently, middleboxes, such as Network Address Translations (NATs), firewalls, and load balancers, block or break unfamiliar protocols. As a result, developers have a strong incentive to choose only between TCP and UDP, and many potential improvements in Internet transport protocols are impeded.

TAPS addresses this situation by defining the relationship between the application and the transport layer in terms of services. In the TAPS model, an application specifies the service features it requires from the transport, and a TAPS mechanism selects the best possible transport protocol to serve the purpose, possibly using probing to verify end-to-end transparency. In this way, the applications can take advantage of modern transport protocols, thereby enabling the network provider and/or stack developer to utilize new protocols and protocol features without breaking the applications.

TAPS does face several challenges. The TAPS Working Group must understand application developers' requirements on transport. The approach used is to create an abstract interface between application and transport, where application preferences and requirements can be expressed. However, there will be situations where an application would like to influence the way transport works in a very protocol- or OS-specific way. One example would be an application that wishes to influence

**Because TAPS creates an abstraction layer between the application and transport, we need to understand the existing interfaces of each transport protocol.**

path selection where there is more than one path available. The transport might not have the knowledge about the cost associated with a certain path that the application may have. Another example is an application capable of handling certain adaptations that might be seen as a job for transport protocol, such as Dynamic Adaptive Streaming over HTTP (DASH) adaptive bitrate for streaming services. These sorts of optimizations are needed,

and in turn, the understanding of both applications and transports are needed in order for the TAPS model to be successful. This combination of expertise can be rather rare in our community.

Because TAPS creates an abstraction layer between the application and transport, we need to understand the existing interfaces of each transport protocol. One of the current issues is the lack of proper interface descriptions in RFCs, as well as the differences between what is implemented in the protocol stack and what is specified in the RFC.

There exists a chicken-and-egg problem with new transport protocols. The lack of transparency inhibits application developers from making use of the new protocols, as they require additional fallback mechanisms for the parts of the Internet where they don't work. However, without the pressure of new protocols being used, there is little incentive for middlebox operators and vendors to fix their behavior and permit protocols other than TCP and UDP. The TAPS Working Group endeavors to break that deadlock by making it easier for application developers to probe and fallback (e.g., by using a TAPS library). This solution would facilitate the use of new protocols where the network permits and would make partial transparency useful.

The IETF recently published the HTTP/2 standard and there are proposals to start working on new transport protocols, such as QUIC and PLUS. While HTTP/1 is still dominant in the Web world and TCP/SCTP is still evolving, the TAPS Working Group can promote that both application and network developers try new transport protocols with new features and help evolve the upcoming transport protocols. It is a unique opportunity for the entire community—including application developers, service providers, protocol stack developers, and network vendors—to collaborate in TAPS to ensure the natural evolution of transport protocol. 

# TRON WORKSHOP CONNECTS IETF TLS ENGINEERS AND SECURITY RESEARCHERS

By Karen O'Donoghue

ON 21 FEBRUARY 2016, THE TLSV1.3 READY OR NOT? (TRON) WORKSHOP WAS held in conjunction with the Network and Distributed System Security Symposium (NDSS 2016) in San Diego, California.

The goal of the workshop was to foster cross-collaboration between the research and standardization communities. The workshop was viewed as an opportunity to get security researchers engaged in the analysis of the Transport Layer Security (TLS) 1.3 specification prior to its publication. The thought was that potential flaws in the specification could be identified and corrected earlier in the process. This would be a big benefit to the Internet in general.

**The newest version of TLS, version 1.3, is currently under development in the IETF. Given the frequency with which flaws are being discovered in security protocols, the earlier we get quality researchers engaged the better.**

TLS is a generic building block that provides confidentiality and integrity in the Internet Protocol suite. It is used to provide end-to-end encryption and authentication for Web, email, and messaging traffic, as



Russ Housley (left) presenting Tibor Jager with an award for the Best Contribution to the IETF.

well as virtually any other conceivable form of Internet communication. The newest version of TLS, version 1.3, is currently under development in the IETF. Given the frequency with which flaws are being discovered in security protocols, the earlier we get quality researchers engaged the better.

The workshop was very successful and included a full day of in-depth presentations and discussions featuring selected published research in this space. In keeping with the overall theme of the workshop, many researchers presented approaches and tools for analysis and verification

of TLS 1.3. The discussion continued with looking at ways to improve the secure implementation of TLS 1.3. Finally, there was a discussion on the topics related to the defense of TLS 1.3 from external factors including the ongoing impact of flaws in Public-Key Cryptography Standards (PKCS) #1 and the issue of metadata leakage and its impact on privacy. The TRON workshop also collected references to related research papers for further analysis.

As a way to further facilitate cross pollination between the two communities, the TRON programme committee presented an award for the Best Contribution to the IETF to Tibor Jager (<http://tiborjager.de/>) for his work on, "On the Security of TLS 1.3 (and QUIC) Against Weaknesses in PKCS #1 v1.5 Encryption". The award was presented to the workshop participant whose work was most likely to have a positive impact on the IETF work in this space. Part of the award includes Tibor attending IETF 96 in Berlin to further the collaboration with IETF security engineers.

For more information, see the workshop programme at <http://www.internetsociety.org/events/ndss-symposium-2016/tls-13-ready-or-not-tron-workshop-programme>.



# HACKATHON BREAKS NEW GROUND IN BUENOS AIRES

*Originally posted by Charles Eckel in the DevNet Open Source Community on 26 April 2016.*

**T**HE IETF 95 HACKATHON IN BUENOS AIRES KICKED OFF WHAT WAS BOTH the first and an extremely rewarding trip by the IETF community to South America. Roughly 100 participants—a record 10% of the total IETF meeting attendees—arrived before the meeting in order to put their talents to use tackling a diverse set of projects aimed at improving the Internet we rely on every day.

The list of projects and teams included many familiar faces, as well as a refreshing set of new participants and challenges. This was the first Hackathon for about one third of the participants, with more than a dozen attending their first IETF meeting. Many first timers were from the host country, including from Buenos Aires and Mendoza; there were two participants from Africa, and others from Europe and the United States. See the story of one IETF and Hackathon first timer at <https://communities.cisco.com/community/developer/opensource/blog/2016/04/25/first-timer-at-ietf-and-ietf-hackathon-shares-his-story>.

Others firsts for this Hackathon include:

- IETF Hackathon materials appeared in the regular meeting proceedings (<https://www.ietf.org/proceedings/95/hackathon.html>).
- An IETF Hackathon github organization was created (<https://github.com/ietf-hackathon>).
- Huawei took over as financial sponsor.

Charles Eckel from Cisco DevNet continues to run the Hackathon in his role as Hackathon chair, and he welcomes Barry Leiba from Huawei as an appreciated and valued Hackathon cochair.

**Participants worked cooperatively and tirelessly, producing fantastic results.**

As at previous Hackathons, participants worked cooperatively and tirelessly, producing fantastic results. Each team summarized their achievements in a brief presentation to the judges and their peers. Top honors and prizes were awarded for especially brilliant accomplishments, including those of the FD.io/VPP team, the TLS 1.3 team (see <https://www.ietf.org/blog/2016/04/ietf-hackathon-getting-tls-1-3-working-in-the-browser/>), and the network-based network analytics team. Some teams demonstrated their work at Bits-N-Bites, including the NETCONF/YANG, I2RS, OpenDaylight teams, DNS/DNSSEC/DANE/DNS-over-(D)TLS teams, and the IBNEMO team. All the presentations and results are available at <https://www.ietf.org/proceedings/95/hackathon.html>.

Mark your calendars now and subscribe to the Hackathon list at <https://www.ietf.org/mailman/listinfo/hackathon> to receive the latest information, including announcements of new projects and the ability to reserve your place in this history-making event.

## HACKATHON HIGHLIGHTS

- 100+ participants, 30+ new participants
- 10+ new IETFers
- New projects: TLS 1.3, VPP, Big Data
- Huawei sponsoring entire year
- Cisco DevNet supporting
- CodeSprint/Hackathon interworking: YANG improved, in draft submit tool
- YANT Pub/Sub client extended – contribute to OpenDaylight
- I2RS findings discussed in Working Group
- Projects on display at Bits-N-Bites
- Hackathon Proceedings in the IETF Datatracker

Following the success in Buenos Aires, the IETF 96 Hackathon in Berlin on 16–17 July should be the biggest Hackathon ever.

For more information, visit <http://www.ietf.org/hackathon/96-hackathon.html>.



Hackathon participant Susan Hares collaborates with others.

# INTEGRATING THE WORLDS OF TECHNOLOGY DESIGN AND POLICY

By Dilawar Grewal, PhD

**M**OST PEOPLE ASSOCIATE THE INTERNET ENGINEERING TASK FORCE (IETF) with all things technical having to do with the Internet. In fact, the IETF's mission statement states, "The mission of the IETF is to make the Internet work better by producing high quality, relevant technical documents that influence the way people design, use, and manage the Internet."<sup>1</sup> It also lists five cardinal principles that it adheres to in pursuit of its mission: open process, technical competence, volunteer core, rough consensus and running code, and protocol ownership. At first glance, the mission statement reads as quite technically oriented. However, the second half of the statement, "...influence the way people design, use, and manage the Internet," alludes to things other than just technical standards. Use and management are often related to things like policies, the human factor, behaviors, rights, acceptability, implementation criteria and protocols, and even politics. It is in this context that this paper explores the IETF experience for participants in the Internet Society Fellowship to the IETF Programme.

The IETF structure and participation protocols are fairly clear to most participants interested in the technical aspects of engineering of and for the Internet. The pathway of in-person or online voluntary participation, discussions, Working Groups, Request for Comments documents (RFC), and so forth, is clear to those interested in discussing, developing, and addressing specific technical issues. There even exists a Working Group on human rights that bridges the gap between the purely technical and human involvement sides of Internet technologies. However, for the most part, IETF meetings lean mostly toward the development of technology and not as much toward the human interaction component. This is in part because the application layer is usually exempt from IETF discussions, and most human interactions occur at that layer. Even though its mission statement refers to people

**For the most part, IETF meetings lean mostly toward the development of technology and not as much toward the human interaction component.**

and use, the integration of policy as an interface between technology and human use is at a nascent stage in IETF activities. It is against this background that the Internet Society invites a variety of users and developers of technology to engage with the IETF community as Policy Fellows.

My first experience as a Policy Fellow was at IETF 94 in Yokohama in November

2015. In parallel to the IETF conference, some end-users and policymakers from around the world were invited to attend a Policy Fellows meeting. The intent of the meeting was to introduce the policymakers and implementers to the structure of IETF, basic Internet technologies, management tools for such technologies, and networks of technologists; and to provide exposure to the mechanisms and issues around development of Internet technologies, standards, and protocols. Even though policy and its role in the development (and sometimes control) of the use of Internet and Internet-related technologies is not directly discussed at IETF conferences, exposing the Policy Fellows to the technologists and the technologies helped bridge the gap from the technology developers side to the policymakers side. In other words, the policy people now were slightly more aware of the world of the technologists and engineers. Meeting activities included presentations by technologists and technology developers, as well as attendance at Working Group meetings. But there were no activities designed to expose the technologists to the world of the policymakers and implementers. Thus, the world of the technologists and the policy people remained separated, despite being connected by a single-lane, one-way bridge. A true development environment



Dilawar Grewal, PhD, also shared his views at the Internet Society briefing panel at IETF 95.

**Meeting activities included presentations by technologists and technology developers, as well as attendance at Working Group meetings. But there were no activities designed to expose the technologists to the world of the policy makers and implementers.**

requires that this bridge be two-way and that there be points of contact between both worlds. The interdependence between technology development and policy development collectively defines the Internet of tomorrow. Generating mechanisms whereby these two worlds become more and more aware of each other is an essential component of true development for the future.

My experience at IETF 95 in Buenos Aires in April 2016 was significantly different in two ways. One, I attended the IETF conference as a technologist in areas of interest to me and not as a policy maker or implementer. Two, I attended some of the sessions organized for the Policy Fellows as an outsider. At IETF 95, the Policy Fellows group mostly consisted of South American regulators and telecommunications' operators. The format of the meetings was similar to IETF 94, but with one major difference: a joint session of technologists and Policy Fellows was organized in which a panel of technologists, policymakers, and regulators discussed the role of and interaction between technology developers and policy developers. The session was filled to capacity and was very well attended on the webcast, as well. It was obvious that the idea of a marriage between future technology protocols and policies was of interest to many.

Awareness of the need for a coalition between core protocols and policies centers around the following three points:

- Acknowledging that even though most policies are executed at the application layer, awareness of policy needs at the protocol design phase is both necessary and useful.
- Efficiency as a design driver for technology is not necessarily the best way to develop technology protocols that require integration with policies for optimal deployment outcomes.
- Envisioning a future where machines talk to each other and impact human lives, sometimes without humans being aware, lays great responsibility on the shoulders of designers to see the good and bad of possibilities. The designs of today are the ones that will either help or hinder much of humanity from progressing and being equal players in the future.

Communications and information may be the two most significant contributors to the development of an individual's opportunities. Control over communications and information is also perhaps the most significant way in which peoples' lives, development, culture, and freedom can be influenced. Technologies and technology protocols together form the basis on which machines store, generate, and communicate information; without proper balance between freedom and control at the base level, it is not possible to enhance either freedom or control at higher levels. Policies are what govern the end-user experience, both as individuals and as

peoples. If the base level of technology design curtails freedoms, there is no policy that can provide it at the end-user level.

Most engineers design from the perspective of achieving efficiencies. Efficiencies mean that certain objectives are met via the shortest route possible. Designing by efficiency produces unexpected consequences as the social context evolves and the technologies of tomorrow become the technologies of today. In order to preserve the human freedoms in the evolving social context, the degrees of freedom in the design phase of technologies need to be carefully increased. This is best accomplished by introducing the slightly inefficient world of policies into the very efficient world of technology design engineers.

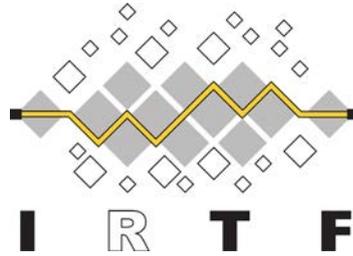
The IETF and Internet Society environments play a vital role in bringing together both technologists and policy people. Exposure to each other's worlds is the linchpin to sustaining awareness of the importance of the marriage between technology protocols and policy. To that end, I envision that IETF and the Internet Society Fellowship to the IETF Programme build on the success of the joint panel discussions of IETF 95 and generate more pathways in future meetings that continue to promote the integration of technology design and policy into the work of its intellectuals, designers, engineers, and implementers. 

**Exposure to each other's worlds is the linchpin to sustaining awareness of the importance of the marriage between technology protocols and policy.**



## IRTF UPDATE

By Lars Eggert



**D**URING IETF 95 IN BUENOS AIRES, Argentina, seven out of the ten chartered Internet Research Task Force (IRTF) research groups (RGs) held meetings:

- Crypto Forum (cfrg)
- Human Rights Protocol Considerations (hrpc)
- Internet Congestion Control (icrg)
- Information-Centric Networking (icnrg)
- Network Function Virtualization (nfvr)
- Software Defined Networking (sdnr)
- Thing-to-Thing (t2trg)

In addition to the meetings of those already chartered research groups, the Proposed Measurement and Analysis for Protocols Research Group (maprg) and the Proposed Network Machine Learning Research Group (nmlrg) met.

At the IRTF Open Meeting at IETF 95, the first two of six winners of the 2016 Applied Networking Research Prizes (ANRP) presented their research. Roya Ensafi presented her examination on how the Chinese “great firewall” discovers hidden circumvention servers, and Zakir Durumeric presented an empirical analysis of email delivery security. The ANRP is awarded for recent results in applied networking research that are relevant for transitioning into shipping Internet products and related standardization efforts. Everyone is encouraged to nominate relevant scientific papers they have recently authored—or read!—for consideration for the award. Please see <https://irtf.org/anrp> for details.

**Join the IRTF discussion list to stay informed about these and other happenings. The website is <https://www.irtf.org/mailman/listinfo/irtf-discuss>.** 

# APPLIED NETWORKING RESEARCH PRIZE WINNERS ANNOUNCED

By Mat Ford

THE APPLIED NETWORKING RESEARCH PRIZE (ANRP) IS AWARDED FOR recent results in applied networking research that are relevant for transitioning into shipping Internet products and related standardization efforts. The ANRP awards presented during IETF 95 went to the following two individuals:



2016 ANRP winner Roya Ensafi

- **Roya Ensafi.** For examining how the Chinese “great firewall” discovers hidden circumvention servers.

See the full paper at <http://conferences.sigcomm.org/imc/2015/papers/p445.pdf>.

- **Zakir Durumeric.** For an empirical analysis of email delivery security.
- See the full paper at <http://conferences.sigcomm.org/imc/2015/papers/p27.pdf>.

Ensafi and Durumeric presented their findings to the Internet Research Task Force open meeting during IETF 95.



2016 ANRP winner Zakir Durumeric

Slides are available at <https://www.ietf.org/proceedings/95/slides/slides-95-irtfopen-0.pdf> and <https://www.ietf.org/proceedings/95/slides/slides-95-irtfopen-1.pdf>.

Thanks to Meetecho, audio/video from the presentations is available at [http://recs.conf.meetecho.com/Playout/watch.jsp?recording=IETF95\\_IRTFOPEN&chapter=chapter\\_1](http://recs.conf.meetecho.com/Playout/watch.jsp?recording=IETF95_IRTFOPEN&chapter=chapter_1) (from 00:05:30).

ANRP winners have been selected for all of the IETF meetings in 2016. The following winners will present at the IETF 96 meeting in Berlin:

- **Samuel Jero**, a postgraduate researcher at Purdue University. Samuel will present a security analysis of the QUIC protocol.
- **Dario Rossi**, a professor at the computer science and networking department of TELECOM ParisTech. Dario will present on characterizing anycast adoption and deployment in the IPv4 Internet. 

**The call for nominations for the 2017 ANRP award cycle is open starting July 2016.** Join the [irtf-announce@irtf.org](mailto:irtf-announce@irtf.org) mailing list for all ANRP related notifications.

# IETF ORNITHOLOGY: RECENT SIGHTINGS

Compiled by Mat Ford

**G**ETTING NEW WORK STARTED IN THE IETF USUALLY REQUIRES A BIRDS-OF-A-FEATHER (BoF) meeting to discuss goals for the work, the suitability of the IETF as a venue for pursuing the work, and the level of interest in and support for the work. In this article, we review the BoFs that took place during IETF 95, including their intentions and outcomes. If you're inspired to arrange a BoF meeting, please read RFC 5434, "Considerations for Having a Successful Birds-of-a-Feather (BoF) Session".

## Babel Routing Protocol (babel)

**Description:** Babel is a loop-avoiding distance vector routing protocol that has good provisions for dynamically computed metrics and remains robust even in the presence of metric oscillations and failure of transitivity. Babel has seen some production deployment, notably in hybrid networks (networks that combine classical, wired segments with mesh segments) and in global overlay networks (networks built with large numbers of tunnels spanning continents). Babel is also considered as part of the IETF Homenet protocol stack. There exist three independent implementations of Babel, all of which are open source.



Long-tailed Meadowlark  
(*Sturnella loyca*)

**Babel has seen some production deployment, notably in hybrid networks (networks that combine classical, wired segments with mesh segments) and in global overlay networks (networks built with large numbers of tunnels spanning continents).**

The core of the Babel protocol is described in detail in RFCs 6126 and 7557, which are both Experimental. While these RFCs have been useful (as indicated by the independent reimplementations of Babel), a number of parties have expressed a desire to have a new specification that clarifies RFC 6126 according to the feedback provided by the independent reimplementations, and to integrate the contents of RFC 7557 without expanding the scope of Babel.

The goal of this BoF was to discuss the value and scope of the work required to create a standards track successor to RFCs 6126 and 7557, including what technical topics need attention as part of advancement. The BoF also discussed the applicability of Babel.

**Proceedings:** <https://www.ietf.org/proceedings/95/minutes/minutes-95-babel>

**Outcome:** This was a well-moderated and productive discussion that focussed on the background to Babel and the work that needs to be done in an IETF working group. There was solid interest in the room from people interested in contributing to the work and reviewing documents. A charter for a WG has been circulated for review.

### Low Power Wide Area Networks (lpwan)

**Description:** Low-Power Wide Area Networks (LPWAN) are long-range low-power lossy networks, many of which operate in license-exempt bands. LPWANs provide low-rate connectivity to vast numbers of battery-powered devices over distances that may span tens of miles. Existing pilot deployments have shown the huge potential and met industrial interest, but the loose coupling with the Internet makes the device management and network operation complex and implementation-specific. As of today, there is little to no use of IETF technologies in LPWANs at large, and there is a need to evaluate their applicability.

**Proceedings:** <https://www.ietf.org/proceedings/95/minutes/minutes-95-lpwan>

**Outcome:** Discussion ranged across a wide variety of L2 technologies that could be classified as LPWANs. The implications for the IP protocol stack of each were also discussed. The group will need to focus more on fewer L2 technologies and engage with external SDOs to make progress.

**As of today, there is little to no use of IETF technologies in LPWANs at large, and there is a need to evaluate their applicability.**

### Alternative Resolution Contexts for Internet Naming (arcng)

**Description:** RFC 819 describes Internet names as a set of directed graphs in an absolute naming context. While that work eventually led to the creation of the Domain Name System, it is important to note that it does not imply that there will be a single resolution system for Internet names. Although the most common Internet names are those which are part of the Domain Name System, that set of names is not the whole.

A number of independent naming and resolution contexts have emerged. In addition to those created for onion routing and multicast DNS, there are large sets associated with the Handle system, Uniform Resource Names (URNs), and Internet scale proprietary names (e.g., Twitter handles). It is apparent that the desire to reuse Internet protocols that default to DNS-based resolution in other resolution contexts has created ambiguities in the resolution context that should be used for individual names. Those ambiguities may result in operational difficulties (queries in the wrong context) and in concerns about limitations implied for DNS-based names.

**Proceedings:** <https://www.ietf.org/proceedings/95/minutes/minutes-95-arcng>

**Outcome:** This meeting was intended to address the questions of whether there are interesting problems here and whether it is possible to provide good guidance on resolving them. There was some support for the idea that a WG-forming BoF could be held during IETF 96 in Berlin. The chairs encouraged attendees to write drafts describing their preferred solutions for this problem. Writing about technology efforts that have these issues now would also be helpful.



*Blue and Gold Macaw  
(Ara ararauna)*

### Limited Use of Remote Keys (lurk)

**Description:** HTTPS in typical use authenticates the server by proving ownership of a private key, which is associated with a public-key certificate. Currently, most trust models assume private keys are associated and owned by the HTTP server and that the server is responsible for both the hosted content and for the network delivery. Although these assumptions were largely true in the past, today the deployment of services on the Internet often relies on multiple distributed instances of the service. Similarly, the delivery of popular content often splits the roles of providing the content and delivering the content. In such architectures, the application, such as a Web browser, expects to authenticate a content provider, but is actually authenticating the node delivering the content. In this case, the confusion mostly results from using a secure transport layer to authenticate application-layer content.

**Proceedings:** <https://www.ietf.org/proceedings/95/minutes/minutes-95-lurk>

**Outcome:** There was a lack of agreement about the problem and how to address it, and a range of potential solutions were identified and discussed. Focussing on the Content Delivery Network use case was identified as a way to make progress, and there was consensus that the IETF was an appropriate place to work on this problem. It is clear that more work to define the problem scope will be necessary before that work can start in earnest. This seems likely to return as a future BoF meeting.

**Currently, most trust models assume private keys are associated and owned by the HTTP server and that the server is responsible for both the hosted content and for the network delivery.**

### Intelligent Transportation Systems (its)

**Description:** The goal of this group is to standardize and/or profile IP protocols for establishing direct and secure connectivity between moving networks (see page 11).

It concentrates on 1-hop moving network to nearby moving network communications. This has immediate applicability in mobility environments such as vehicle-to-vehicle or vehicle-to-infrastructure communications. In some of the moving network applications, the window of opportunity for exchanging data with the immediate infrastructure may be very short. The safety and security requirements are higher in connected mobility environments. The links are very heterogeneous, such as 802.11p/ac/ad OCB, Infrared, VLC, cellular, 802.15.4, and so forth.

The BoF was intended to bring implementers, users, and experts from academia, institutes, IT, the automotive industry, and public authorities together to discuss.

**Proceedings:** <https://www.ietf.org/proceedings/95/minutes/minutes-95-its>

**Outcome:** There was good discussion of the problem space and several comments about the need to address security and privacy considerations. There is a pressing need to engage key SDOs and industry players, in order for the output of any IETF work on this topic to see widespread deployment in vehicular networks. A WG charter will be developed on the mailing list.



*Narrow-billed Woodcreeper*  
(*Lepidocolaptes angustirostris*)

### **Alternatives to Content Classification for Operator Resource Deployment (accord)**

**Description:** Mobile Radio Access Networks (RANs) have historically allowed content-type based classification to associate service descriptions with flows with the goal of efficient use of the often-volatile radio bearer. The increased use of Transport Layer Security (TLS) and other encrypted transports eliminates this metadata from the view of the operator and forces a reexamination of this method. While having endpoints expose metadata to the radio access network (RAN) outside of the encrypted channel would resolve this, it would degrade the confidentiality expected by users and require extensive coordination among application developers, user endpoint manufacturers, and RAN operators. To avoid these disadvantages, the WG will examine both what specific network treatments need to be elicited for the efficient operation of RANs, if any, and what the minimal communication to elicit those treatments would be. This BoF session was part of the follow-on activity stemming from the Internet Architecture Board (IAB) Managing Radio Networks in an Encrypted World (MaRNEW) workshop in 2015 (<https://www.iab.org/activities/workshops/marnew/>).

**Proceedings:** <https://www.ietf.org/proceedings/95/minutes/minutes-95-accord>

**Outcome:** The meeting benefitted from a high-level introduction that helped bring attendees up to a minimal understanding of the relevant parts of mobile network architecture. Discussion revolved around whether progress could be made now by running experiments without any additional protocol machinery (0-bit solution) and whether it was necessary to provide some minimal signalling (1-bit solution). The difficulty of extracting useful data from operators was noted. It was also noted that TCP optimizers seem to have fallen out of favour with many operators as they don't offer any performance improvement. Discussion also included general ways to improve performance in radio networks.

### **IAOC Meeting Venue Selection Criteria & Procedures (mtgvenue)**

**Description:** The IETF has expressed concern regarding the process of selecting a meeting venue. The IETF Administrative Oversight Committee (IAOC) and IAOC Meetings Committee have undertaken to document the process, which has been posted at an IAOC-private site for some time and is being updated, in an Internet-Draft for community discussion. This meeting was to allow community discussion of concerns relating to meeting venue selection and the draft process.

**Proceedings:** <https://www.ietf.org/proceedings/95/minutes/minutes-95-mtgvenue>

**Outcome:** The session provided a lot of background about the way the IAOC Meetings Committee has been operating and the draft set of meeting venue criteria that have been developed. An analysis of the impact of applying the draft meeting criteria to IETF venues for IETF 74 to IETF 100 was also presented. There was some time available for input and discussion from the community, and that continues on the mailing list.

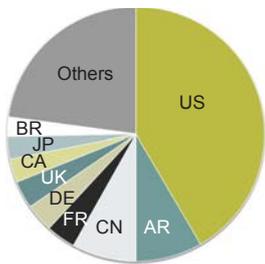


*Toco Toucan  
(Ramphastos toco)*





## IETF 95 AT-A-GLANCE



On-site participants: 1002  
(140 from South America)

Newcomers: 171

Number of countries: 55

### IETF Activity since IETF 94 (01 November–3 April 2016)

New WGs: 5

WGs closed: 7

WG currently chartered: 144

New and revised Internet-Drafts (I-Ds): 1968

RFCs published: 133

- 76 Standards Track, 6 BCP, 7 Experimental, 44 Informational

### IANA Activity since IETF 94 (October 2015–February 2016)

Processed 1581+ IETF-related requests, including:

- Reviewed 113 I-Ds in Last Call and 133 I-Ds in Evaluation
- Reviewed 133 I-Ds prior to becoming RFCs, 68 of the 133 contained actions for IANA

Added 11 new registries since IETF 94 (October 2015–February 2016): mrt-parameters, abfab-parameters, ppsp-tp, emergency-call-additional-data, bgp-ls-parameters, content-security-policy-directives, dncp-registry, ospf-parameters, cdni-parameters, markdown-variants, owamp-parameters

SLA Performance (September 2015–February 2016)

- Processing goal average for IETF-related requests: 97%
- The draft 2016 SLA between ICANN and IAOC for the protocol parameter work is still under review.

IANA and DNSSEC

- As of 28 March 2016, 1093 TLDs have a full chain of trust from the root. [http://stats.research.icann.org/dns/tld\\_report/](http://stats.research.icann.org/dns/tld_report/).

### RFC Editor Activity since IETF 94 (January–March 2016)

Published RFCs: 101

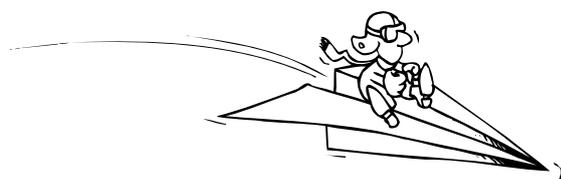
- 63 IETF (10 IETF non-WG), 2 IAB, 1 IRTF, 5 Independent

Improvements to website based on community feedback

- Cluster pages make it easier to check the I-D holding up a cluster.
- Added “Discuss this RFC” (pointing to the WG mailing list), where applicable.

Responded to three legal requests

## Be the First Host on Your LAN to Receive the *IETF Journal*!



Receive the latest edition of the *IETF Journal* as soon as it is available—in hardcopy or via email. Subscribe today at:

[www.ietfjournal.org/subscribe](http://www.ietfjournal.org/subscribe)

Want it faster? Follow the [@ietfjournal](https://twitter.com/ietfjournal) Twitter stream to read the articles as they are published.

# IETF MEETING CALENDAR

For more information about past and upcoming IETF meetings visit [www.ietf.org/](http://www.ietf.org/).

**IETF 97**      **Date** 13–18 November 2016  
**Host** TBD  
**Location** Seoul, South Korea

**IETF 99**      **Date** 16–21 July 2017  
**Host** Comcast–NBCUniversal  
**Location** Prague, Czech Republic

**IETF 98**      **Date** 26–31 March 2017  
**Host** TBD  
**Location** Chicago, IL, USA

**IETF 100**      **Date** 12–17 November 2017  
**Host** Cisco Systems  
**Location** Singapore

Special thanks for hosting IETF 95



The Internet Society Fellowship to the IETF, as part of the Internet Society Next Generation Leaders Programme, is sponsored by



This publication has been made possible through the support of the following Platinum Programme supporters of the Internet Society





# IETF Journal

IETF 95 • Volume 12, Issue 1 • July 2016

Published three times a year  
by the Internet Society.

Galerie Jean-Malbuisson 15  
1204 Geneva, Switzerland

**Editor**

Mat Ford

**Associate Editors**

Megan Kruse • Michelle Speckler

**Contributing Writer**

Carolyn Marsan

**Editorial and Design**

Speckler Creative

**Editorial Board**

Jari Arkko

Mat Ford

Olaf Kolkman

Megan Kruse

Andrew Sullivan

Greg Wood

**Email**

[ietfjournal@isoc.org](mailto:ietfjournal@isoc.org)

**Find us on the Web**

[www.ietfjournal.org](http://www.ietfjournal.org)

**Editor's Note**

*IETF Journal* adheres to the *Oxford English Dictionary*, 2nd Edition.

Unless otherwise noted, photos are

@Internet Society.

## IETF Journal

Internet Society  
Galerie Jean-Malbuisson 15  
1204 Geneva, Switzerland