

A report from IETF 84, July 2012, Vancouver, Canada. Published by the Internet Society in cooperation with the Internet Engineering Task Force*

Inside this Issue

From the Editor's Desk	1
World IPv6 Launch: It's Happening!.....	1
Message from the IETF Chair	2
Words from the IAB Chair	3
Remembering Wendy Rickard	4
IAB Panel Debates Management Benefits, Security Challenges of Software-defined Networking.....	9
IPSO Alliance Successfully Demonstrates Internet of Things Interoperability	11
Software-defined Networking Efforts Debuted at IETF 84	12
Moving Toward a Censorship-free Internet	16
Using JSON in IETF Protocols.....	18
IRTF Update	20
Regulatory Fellows Attend IETF 84 in Vancouver.....	21
Different Experiences, A Common Goal: ISOC Fellows at IETF 84.....	22
IETF Ornithology: Recent Sightings	24
IETF 84 At-A-Glance ...	26
Calendar	27



From the Editor's Desk

By **Mat Ford**

Amongst fireworks and nude bicyclists of coincident festivals, the 84th meeting of the IETF took place in downtown Vancouver, Canada, earlier this year. In this issue of the IETF Journal you'll find a roundup of some of the discussions and people that helped make it such a great meeting.

Our cover article provides detailed insight into the profound impact of World IPv6 Launch on the deployment of IPv6 in networks and services around the world. We also present articles on such diverse topics as, how to build "censorship-proof" networks, using JSON in IETF protocols, smart-object interoperability testing, and software-defined networking.

Also in this issue are our regular columns from the IETF, IAB, and IRTF chairs, coverage of hot topics discussed during plenary meetings, and an introduction to the ISOC Fellows who attended the IETF 84 meeting.

The Internet Area of the IETF is large and diverse—it's welcome news that the chairs of each WG within this area have offered short updates on recent progress. The goal of these updates is not to give a detailed analysis of all the WG activity, but rather to provide highlights of activity that may be of interest to people not actively involved in the WG. The WG summary reports are available on the Int Area wiki: <http://wiki.tools.ietf.org/area/int/trac/wiki/IETF84>.

As always, we are hugely grateful to our contributors. Please send comments and suggestions for contributions to ietfjournal@isoc.org. And remember, you can subscribe in hardcopy or via email at <https://www.internetsociety.org/publications/ietf-journal/ietf-journal-subscription>.

World IPv6 Launch: It's Happening!

By **Erin McGann**

The World IPv6 Launch on 6 June 2012 saw leading web sites, Internet service providers (ISPs), and home-router equipment manufacturers turn on IPv6 by default. At IETF 84 the Internet Society (ISOC) brought together content providers, access providers, and Internet measurement experts to discuss the launch, and to share their findings.

Leslie Daigle, chief Internet technology officer for the Internet Society, kicked off the panel by acknowledging the amazing and collaborative industry event that was World IPv6 Launch.

Continued on page 5

* 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 Russ Housley

IETF 84 was a well-attended, successful meeting. Approximately 1,175 people from 52 countries came to Vancouver and actively engaged in developing the future of the Internet. It was exciting to see so many people collaborating.

Google was the meeting host, the hotel facilities were very comfortable, and the social on Tuesday evening was both fun and educational. Hyatt Regency Vancouver, Telus, and Verisign were meeting sponsors, and they assisted Google in ensuring the event was successful. Thanks to all for your support.

The first-ever Bits-N-Bites reception was held on Thursday evening. The event, modeled after the North American Network Operators' Group's Beer 'n Gear event; it featured exhibit tables for sponsors and free food and drink. Huawei, ICANN, IPSO Alliance, the Internet Society, and Juniper sponsored tables; all sponsors seemed to have visitors throughout the event. Thanks for supporting the first Bits-N-Bites.



Internet Society and IETF cookies at the debut Bits-N-Bites reception.

After each session, the blue sheets are scanned and the scans are included in the proceedings for that session. A notice has been added to the top of the blue sheet to alert people that it will become part of the proceedings. After the blue sheets are scanned, they are discarded.

Bernard Aboba, as the IAB Chair, and Russ Housley, as the IETF Chair, signed an affirmation of the Modern Global Standards Paradigm. This document and more information about the paradigm can be found at <http://www.open-stand.org/>.

IETF 85 will take place in Atlanta, GA, U.S.A., on 4–9 November 2012. The North American Cable Industry will be the meeting host. Scheduling information for the upcoming IETF meetings can be found at <http://www.ietf.org/meeting/>. I look forward to seeing you in Atlanta. 🍪



Russ Housley, IETF Chair

Since IETF 83, six working groups (WGs) have been chartered and four have closed. Between meetings, the Internet Engineering Steering Group (IESG) approved 112 Internet-Drafts for publication as RFCs. The RFC Editor published 140 new RFCs.

Starting with IETF 84, a new blue sheet policy was implemented. The blue sheets no longer ask for email addresses; they now ask for organization affiliation to discern among multiple people with the same name.

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 **Bernard Aboba**

IAB Retreat

The Internet Architecture Board (IAB) held its 2012 retreat 10–11 May in Washington, DC, U.S.A.¹ During the retreat, IAB members reviewed the Programme and Initiative model,² which was developed last year. They concluded that the Programmes proved useful in enabling the IAB to achieve both administrative and technical objectives over a multiyear period, while the benefits of the Initiatives (work bundles of both shorter duration and lower priority) were less clear. As a result, the IP Evolution and Emergency Services Initiatives were promoted to Programmes, a new IAB Processes and Tools Programme was created, and the HTTP/Web Evolution, IPv6 for IAB Business, and DNS Initiatives were closed.



Bernard Aboba, IAB Chair

The IAB now has Programmes in the following areas:

- Administrative: RFC Editor (RSOC), IANA Evolution, ITU-T Coordination, Liaison Oversight, IAB Tools and Processes
- Technical: Privacy, IP Evolution, Internationalization, Emergency Services

IETF 84 Technical Plenary

The IETF 84 Technical Plenary,³ organized by Jon Peterson, focused on Software Defined Networking (SDN). Nick Feamster presented “The Past, Present, and Future of Software-Defined Networking;” David Ward’s presentation, “Programmatic Internet” described some of the issues with the current SDN architecture, and what SDN might mean for the IETF; Ted Hardie presented “Units of Evolution,” a description of how SDN changes the way network elements evolve.

IAB/IRTF Workshop on Congestion Control

The IAB and IRTF held a workshop entitled, “Congestion Control for Interactive Realtime Communication,” on 28 July 2012, in Vancouver, Canada. Papers and materials for the workshop are posted; minutes are under development.⁴

RFC Series

The IAB published “RFC Editor Model (Version 2)” (RFC 6635)⁵ and “Independent Submission Editor Model” (RFC 6548)⁶ as Informational RFCs within the IAB stream. Collectively, these documents obsolete RFC 5620.⁷

IANA Evolution Programme

The IETF and ICANN have executed a supplemental agreement, which describes the service levels expected for 2012.⁸ In May, the IAB submitted an updated ICANN performance evaluation.⁹ On 2 July, the U.S. National Telecommunications and Infrastructure Administration (NTIA) awarded the IANA Functions contract to ICANN.¹⁰

Privacy Programme

The Privacy Programme posted an IPv6 privacy survey.¹¹

Continued on next page

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>.

Words from the IAB Chair, continued

Other Document Actions

The IAB adopted “Technical Considerations for Internet Service Blocking”¹² as an IAB work item. It approved publication of “Design Considerations for Protocol Extensions”¹³ and “IETF and ITU-T Standardization Sector Collaboration Guidelines”¹⁴ as Informational RFCs within the IAB stream; and issued a Call for Comment on “Architectural Considerations on Application Features in the DNS”¹⁵ which was completed on 16 August, and on “Principles for Unicode Code Point Inclusion in Labels in the DNS,”¹⁶ which will be complete on 30 September.¹⁷

Affirmation of the Modern Global Standards Paradigm

At IETF 84, Russ Housley announced that the IETF, IEEE Standards Association, and W3C were working together to develop an “Affirmation of the Modern Paradigm for Standards.” After an initial round of feedback and edits, IETF Last Call was announced

on 10 August and it was complete on 24 August, at which point Russ announced the conclusion of the Last Call.¹⁸ More information can be found at <http://open-stand.org/>, which went live on 29 August.¹⁹ The IAB also agreed to adopt the Affirmation as an IAB document.²⁰

Liaison Relations

The IAB has responded to a liaison from the Open Network Foundation (ONF).²¹

The IESG, IAB, and the IEEE 802 Executive Committee met in Milpitas, California, on 25 July 2012.²² Minutes of the meeting are under development. Next steps arising from the meeting include revising RFC 4441.²³

The IAB approved “IETF and ITU-T Standardization Sector Collaboration Guidelines” as an Informational RFC within the IAB stream, and ITU-T TSAG sent their approval in a liaison message.²⁴ In addition, the IAB and IESG posted a liaison to ITU-T TSAG

on “Observations on Contribution 74 to TSAG Concerning Modification of IETF Protocols Technologies by the ITU-T.”²⁵

Appointments

The IAB has reappointed Eric Burger to the ISOC Board of Trustees,²⁶ and Ole Jacobsen to the ICANN Nomcom.²⁷



References

1. IAB Posts Minutes of 2012 Retreat, <http://www.iab.org/?p=5970>
2. IAB Programs, <http://www.iab.org/activities/programs/>
3. Technical plenary materials, <http://www.ietf.org/proceedings/84/technical-plenary.html>
4. IAB/IRTF Congestion Control Workshop Materials Posted, <http://www.iab.org/?p=6229>
5. RFC Editor Model (Version 2), <http://tools.ietf.org/html/rfc6635>
6. Independent Submission Editor Model, <http://tools.ietf.org/html/rfc6548>
7. RFC Editor Model (Version 1), <http://tools.ietf.org/html/rfc5620>
8. ICANN and IETF Execute Supplemental Agreement, <http://www.iab.org/?p=6045>

Remembering Wendy Rickard

13 August 2012

The Internet Society is deeply saddened at the passing of Wendy Rickard after an extended battle with cancer.

Wendy was a long-time contributor to the Internet Society and the Internet community. As part of her work with the Internet Society, Wendy initiated the *OnTheInternet* magazine (*OTI*) in 1996. *OTI* was published first in print and then online until 2001 and was one of the first periodicals that covered the broad range of technical, policy, and development topics relating to the Internet. She was the initial and long-time editor of the Internet Society Annual Report and contributed to many other Internet Society publications. She also served until very recently as the Associate Editor of the *IETF Journal*.

Wendy worked with Internet Society Members and Chapters across the globe to capture and promote the important Internet development work we were all doing. She was able to capture the difference the Internet made in the lives of individuals around the world.

“Wendy believed in the power of the Internet to transform. She understood its potential and was able to make it real and accessible long before it was so commonplace. She was a remarkable woman and will be missed by all of us at the Internet Society,” said Lynn St. Amour, Internet Society president and chief executive officer.

Beyond her Internet-related accomplishments, Wendy is remembered for her incredible passion and the seemingly endless creative and positive energy she brought to everything she did.

9. IAB submits updated ICANN performance evaluation, <http://www.iab.org/?p=5930>
10. Commerce Department Awards Contract for Management of Key Internet Functions to ICANN, <http://www.ntia.doc.gov/press-release/2012/commerce-department-awards-contract-management-key-internet-functions-icann>
11. IAB posts IPv6 Privacy Survey, <http://www.iab.org/?p=6052>
12. IAB adopts "Technical Considerations for Internet Service Blocking," <http://www.iab.org/?p=6242>
13. IAB approves publication of "Design Considerations for Protocol Extensions," <http://www.iab.org/?p=6041>
14. IAB approves publication of "IETF and ITU-T Standardization Sector Collaboration Guidelines," <http://www.iab.org/?p=6118>
15. Call for Comment on "Architectural Considerations on Application Features in the DNS," <http://www.iab.org/?p=6053>
16. Principles for Unicode Code Point Inclusion in Labels in the DNS, <http://tools.ietf.org/html/draft-iab-dns-zone-codepoint-pples>
17. Call for Comment on "Principles for Unicode Code Point Inclusion in Labels in the DNS," <http://www.iab.org/?p=6306>
18. Affirmation of the Modern Global Standards Paradigm, <http://www.iab.org/?p=6240>
19. Leading Global Standards Organizations Endorse 'OpenStand' Principles that Drive Innovation and Borderless Commerce, <http://www.iab.org/?p=6297>
20. Affirmation of the Modern Paradigm for Standards, <http://tools.ietf.org/html/draft-iab-modern-paradigm>
21. IAB responds to liaison from the Open Networking Foundation (ONF), <http://www.iab.org/?p=6109>
22. IAB and IESG to meet with IEEE 802 Executive Committee, <http://www.iab.org/?p=5793>
23. The IEEE 802/IETF Relationship, <http://tools.ietf.org/html/rfc4441>
24. IAB approves publication of "IETF and ITU-T Standardization Sector Collaboration Guidelines," <http://www.iab.org/?p=6118>
25. IAB and IESG post liaison to ITU-T TSAG, <http://www.iab.org/?p=6002>
26. IAB reappoints Eric Burger to ISOC Board of Trustees, <http://www.iab.org/?p=6018>
27. IAB reappoints Ole Jacobsen as IETF Representative to the 2013 ICANN Nomcom, <http://www.iab.org/?p=6257>

World IPv6 Launch: It's Happening!, continued from page 1



Mat Ford, Internet Society Panel speaker and technology program manager, Internet Society, speaks of the success of the World IPv6 Launch.

"For a long time the criticism for IPv6 deployment has been 'it's only happening in research networks and it's certainly not happening in North America,' and that is simply not true anymore."

—Mat Ford
Technology Program Manager
Internet Society

"For anyone who was doubtful that we could actually do anything collaborative and outside of business interests anymore, this was good proof [that we can]," she said.

Mat Ford, technology program manager at the Internet Society, organizes and leads aspects of ISOC's work in deployment of open standards. Presenting some of the high-level and measurable outcomes from the event, Ford dispelled some of the criticisms around IPv6 deployment: "[The thing] I really want to emphasize from these statistics is that although there are a lot of research networks in this list, they are not at the top. The networks at the top are major broadband ISPs in France, North America, and interestingly, in Romania. For a long time the criticism for IPv6 deployment has been 'it's only happening in research networks and it's certainly not happening in North America,' and that is simply not true anymore."

For web sites, Ford found that the number of web sites serving IPv6 spiked last year on World IPv6 Day in 2011 and then dropped away, but in 2012, there was both a much larger number of web sites participating and a sustained response afterwards. "This was a truly global event, and a major step change in

the levels of IPv6 deployment in access networks and content providers around the world."

George Michaelson, senior research scientist at the Regional Internet Registry serving the Asia-Pacific region, captured data on IPv6 deployment through clever use of online advertising tools. Embedding Flash code in online adverts allowed them to obtain measurements from a large and diverse sample of users that would have been very hard to reach otherwise.

"The distribution of IPv6 to end users is really quite variable by economy," he explained. "It is not as straightforward as saying the G8 are doing it and it's the other economies that aren't. It really depends on the nature of your economic investment, timelines for your mobile plant upgrade, broadband upgrade, what kind of machines [end users] are buying, what their cycles are—it's not an even distribution."

Michaelson's data turned up some surprising results, even within North America. "The numbers in Canada [show] an upward trend, but it's somewhat slow. But if you compare that

Continued on next page

World IPv6 Launch: It's Happening!, continued from page 5

to America, for instance, this is a really compelling story of continuous uptake rate. This is a good story. I think you're used to us standing up and saying 'woe is me, woe is me, [IPv6 isn't happening.] It is happening.'

He cited the estimated IPv6-user population figures, including 3 million in the United States, 2.3 million in China, 2 million in Japan and 2 million in France, as examples.

John Brzozowski, Comcast's distinguished engineer and chief architect for IPv6, noted that to achieve approximately two percent of Comcast's active customers using IPv6, they needed to deploy IPv6 on approximately half of their network. Approximately 70 percent of these customers are using a computer connected directly to a cable modem; 30 percent are using home routers.

"One of the things we feel is important to highlight here is that if we understand what devices are not using v6, we can then use that as an opportunity to have feedback here in this community and to also work more with consumer electronics," explained Brzozowski. "One of the things we feel is pretty significant is that while there is a growing momentum for [IPv6] across consumer electronics in the form of home networking, the other part of consumer electronics—your televisions, your Roku—needs some tender loving care as far as [IPv6] is concerned."

Brzozowski highlighted some of the services that represent the bulk of their IPv6 traffic, namely Netflix, YouTube, and the iTunes App Store. Approximately six percent of Olympics streaming over YouTube to Comcast customers was over IPv6.

Lorenzo Colitti is currently the technical lead for Google's IPv6 efforts, which includes everything from performance metrics to Android



George Michaelson, Internet Society Panel speaker and senior research scientist, Asia Pacific Network Information Centre, cites IPv6-user population figures.

"This is a good story. I think you're used to us standing up and saying 'woe is me, woe is me, [IPv6 isn't happening.] It is happening.'"

—George Michaelson
Senior Research Scientist
Asia Pacific Network Information Centre

development to government outreach. For World IPv6 Launch, Google helped participants prepare, published data, and built a tool that tracked deployments.

"We measure 2.5x growth of IPv6 traffic over the last year, and if anything it's accelerating—this is on top of [IPv4] growth, of course," Colitti explained. "It's basically 2.5x per year that [IPv6] is gaining compared to [IPv4], if you extrapolate that, as the pointy-haired boss-types like to do, you get to 50 percent of traffic in about six years, or 100 percent in seven years."

For Google, a major focus for World IPv6 Launch was also tracking IPv6 breakage, to ensure bad user experience was kept to a minimum. Colitti pointed out some networks are responding to IPv6 by filtering AAAA records at their DNS servers, and that this was not a tactic they recommended. "The

underlying problem doesn't go away, it's just masked. Once you mask the problem it becomes very difficult to measure it, and once you can't measure it, you don't know when to turn the filtering off. It's hard to get out of it."

Google tracks how many search queries there are per second, and on World IPv6 Launch, they witnessed an increase of 75 percent of queries over IPv6.

"One key message from our perspective as a content provider is that we have seen deployments in every part of the world, on every access technology, and these are real deployments," said Colitti. "We've seen a real impact on the whole ecosystem—we've seen web sites, we've seen router vendors, phone vendors, and home router vendors actually enable IPv6, turn it on by default and leave it on. So the next time somebody says to you 'I need the

IETF to standardize this otherwise I can't deploy it', think carefully whether that's true, or whether there are other ways to solve the deployment problem. The thing is, all these networks beg to differ. They say that, actually, IPv6 is deployable right now."

Lee Howard, director of network technology for Time Warner Cable, has primary responsibility for the cable company's IPv6 deployment. He described the relationship between the one percent of their users actively using IPv6 and the scale of their deployment as the product of multiplying fractions. "Half of cable modems have to be enabled to get to one percent of users. Here's why: because those cable modems are evenly distributed across your entire footprint, or close enough to evenly distributed."

He went on to explain the breakdown further: "Only about half of operating systems in residential networks support IPv6, plus or minus five or ten percent.

That half is everything that's not Windows XP... Therefore you have to enable so many devices in order to get



Lee Howard, Internet Society Panel speaker and director of network technology, Time Warner Cable, answers audience questions.

down to that half of devices that are not Windows XP, in order to get something that can actually use the IPv6. So we're still multiplying fractions here—50 percent times 30 percent times 50 percent is down to a fairly small number.

"As John [Brzozowski] pointed out, you have to look at the people who are directly connecting a device into their cable modem, and that's about 15 percent. So that takes us down to a very, very small number," he said.

"One percent doesn't sound like a lot, but since you're multiplying so many fractions, it is actually a huge roll-out. It represents a huge deployment in any network in order to get that [many users] actually using IPv6."

Erik Nygren is chief architect in Akamai's platform infrastructure engineering organization, and leads their IPv6 initiative. "On our network right now, we have IPv6 in more than 53 countries around the world and in 600 of our locations."

Nygren discussed the different factors at play when trying to calculate IPv6 growth, including how much content is available, how many clients have network connectivity over IPv6 and whether there is browser or operating system support.

"When we combined all of these factors together," he explained, "what we saw was very significant growth as part of the World IPv6 Launch. In particular, the number of IPv6 addresses we saw year over year between 2011 and 2012 rose over 400 fold, up to 19 million unique IPv6 addresses. The number of requests we served, even just in that 24-hour period of World IPv6 Launch, was more than three billion requests over IPv6. This combination of more content and more clients does mean more traffic. Although a lot of these [figures] are from taking a bunch of small fractions and multiplying them together, one side effect is that ... taking one of those small fractions and doubling it means you've doubled the overall aggregate number."

Nygren broke down where clients are coming from and how they are



Internet Society Panel on World IPv6 Launch, continued



Internet Society Panel (left to right): Mat Ford, Internet Society; George Michaelson, Asia Pacific Network Information Centre (APNIC); John Brzozowski, Comcast; Lorenzo Colitti, Google; Lee Howard, Time Warner Cable; and Erik Nygren, Akamai.

connecting, noting that there are some areas of the world where there is more 6to4 than native IPv6. “One of the biggest changes from a network perspective over the past year is that some big, top, American ISPs came into the fray and really started making IPv6 available to their end user subscribers. Among the top six U.S. networks, 86 percent of the IPv6 requests we saw came from those. Verizon Wireless alone, with their Android LTE devices, was over a third of the IPv6 traffic during that 24-hour window.”

Addressing the pervasive view that IPv6 was something that only occurred in Asia or in Europe, Nygren pointed out these numbers prove that’s not the case anymore.

“One of the top questions we get from content provider sites is. ‘When I make my site available over IPv6, what percentage of requests will come to my site over IPv6?’” Nygren explained that the answer varied greatly depending on audience. “If you have an audience of global consumer end users, you’re going to have less IPv6 preference, somewhere in the half a percent, to 1.5 percent range. If you’re at the other end, a particular country with lots of IPv6, or you’re a router manufacturer, you may

have an IPv6 preference rate that’s in the two to three percent range, or even higher.”

Nygren highlighted a dramatic shift in the overall IPv6 preference rate in the United States: in the past year, for a sample of consumer web sites [the preference rate] has gone up by a factor of nine. He credited this jump to a number of big U.S. ISPs coming into play. “As this continues to grow in the United



Lorenzo Colitti, technical lead for Google’s IPv6 efforts, speaks during the Internet Society Panel.

States, as other networks keep turning this on, as more content becomes available over the next few years, hopefully we’ll get to the point where IPv6 becomes the dominant Internet protocol.”

Leslie Daigle posed the question to the entire panel of whether someone with IPv6-only connectivity could now say they were on the Internet, or, if not, when could that statement be made. Although several panelists agreed that point had not been reached yet, Leslie went on to explain that if she had asked that question last year, “eyes would have rolled back in their sockets,” and the response would have been different—as the response will be very different next year. “We’re going through the ugly part of transition,” she commented.

Several panelists discussed the possibilities of deploying IPv6 and supporting IPv4 through a legacy, backwards-compatibility approach, described by Colitti as “IPv4 by carrier pigeon.” It was acknowledged that this would be easier for some networks than others to manage.

Russ Housley closed the session by saying, “On Friday morning Vancouver time, the IETF web server was not available because there was a denial-of-service attack against our server. For the first time, all of the traffic taking that server down, was IPv6.” The room erupted with cheering and applause.

Several panelists discussed the possibilities of deploying IPv6 and supporting IPv4 through a legacy, backwards-compatibility approach, described by Colitti as “IPv4 by carrier pigeon.” It was acknowledged that this would be easier for some networks than others to manage.

IAB Panel Debates Management Benefits, Security Challenges of Software-defined Networking

By Carolyn Duffy Marsan

Software-defined Networking (SDN) could radically change how networks are designed and operated, according to panelists at the Internet Architecture Board's Technical Plenary discussion in Vancouver, Canada, which considered whether the IETF should develop protocols to support the emerging network architecture.

SDN separates the control plane from the data plane in network devices such as switches and routers. The leading SDN architecture is OpenFlow, which is a Layer 2 protocol that allows software running on routers to determine the path of packets through the network switches as a way of enabling more sophisticated traffic management.

Dave Ward, a Cisco Fellow, defined SDN as enabling applications not resident in embedded operating systems to extract and program state into networking nodes. He listed the following ways that SDN challenges basic assumptions about network design that have been popular for the last 20 years.

- Using a logically centralized database to keep all the information about the state of network devices, rather than having the devices store that information.
- Adopting new ways for network management systems to interact with network devices, rather than relying on transactions where state is written into configuration files.
- Taking a centralized view of network topology, rather than a distributed view.

The main goal of SDN is to make it easier for network operators to configure networks for specific applications and services, Ward said. The SDN architecture does this by moving the network away from command line interfaces and providing a standard interface that

allows intelligent applications to customize network policies in real time.

Other potential advantages of SDN include faster deployment of computer, storage, and network services, as well as enabling new services. SDN could enable better usage of network capacity, plus faster restoration when outages occur, Ward said.

However, he also pointed out that "it's not all rosy in the arena of SDN."

The current SDN architecture is missing information to understand topology, utilization, delay, loss, or jitter, he said. Other features that are not available include loop detection, the ability to remove duplicates, or fix errors in state. SDN doesn't support horizontal communications between network nodes or controllers to enable collaboration between devices.

Nonetheless, Ward favors SDN. "SDN will augment and add functionality and make the Internet easier to operate, but as it's currently constructed it's missing a bunch of features that are necessary," he says.

Ward says SDN requires standard interfaces to the Internet that could be created by the IETF. "Many of the critical features of networking nodes are not standardized," Ward says. He recommends that new working groups be formed to address the architecture work that will be required by SDN.

Nick Feamster, associate professor at the University of Maryland's computer

science department, said the promise of SDN is the ability to change the network as easily as changing applications. Instead of having closed, proprietary network devices, SDN provides a single software program such as OpenFlow that can control the behavior of entire networks.

Feamster says the benefits of SDN include centralized control of the network, more sophisticated control of network flows, and the ability for operators to more easily evolve network capabilities.



Ted Hardie, panel speaker and network engineer at Google

"What we're trying to do is enable operators to specify much more sophisticated policies and actions than they otherwise would be able to do," Feamster explained.

He described the results of two SDN deployments using so-called Lithium controllers that extend the OpenFlow control model to allow network operators to take actions based on time, history, and user. Possible applications of Lithium event-based controllers include parental controls and usage caps for home networks, as well as access control on campus networks.

In home network deployments, ISPs are placing Lithium controllers outside the home so that residential customers needn't operate their own home networks. "Those people who are unskilled and uninterested in running a

Continued on next page

IAB Panel Debates Management Benefits, Security Challenges of Software-defined Networking, continued

home network—which is most of us—won't have to," Feamster said.

Feamster reported that 225 routers are deployed in home networks as part of an ongoing SDN trial involving several ISPs. The trial will be extended to try denser deployments in apartment buildings and integration with other in-home devices such as phones.

"We're looking at how software-defined networking can simplify network monitoring and management. In our deployments and case studies, we've found that we need new control models to solve some of the problems," Feamster said.

Ted Hardie, a network engineer with Google, is a fan of SDN, which he predicts will change the way networks evolve by allowing devices to change in terms of their function—from being routers to switches to load balancers to firewalls, for example. "SDNs allow a network element to be any of these, or a bit of each, and to change over time," he explained.

Hardie said SDNs would allow networks to be routed differently. So instead of always routing for shortest path with added traffic engineering, networks might be routed to boost utilization, reduce latency, or accomplish some other goal of the network operator. Hardie referenced a data center-to-data center network that Google has built using OpenFlow that improves link utilization and drives down cost.

"What you can do in SDN is create entirely different optimalities," Hardie said. "The network can behave differently for different flows and different optimalities."

SDN is not without its challenges. Audience members identified several of these, including security risks from a centralized network control and state information, and reduced resiliency

as network elements lose the ability to route around outages.

Bob Hinden, a Check Point Fellow, called SDN's security implications "terrifying." He said that the reason network devices are specialized is because they are customized for each function. He warned that a reprogrammable network device that does all functions well is "very optimistic."

But presenters argued that the benefits of SDN outweigh the risks.

"The charter of the IETF is obviously Layer 3 routing protocols, and right now I suggest that some of the work we would want to do is to augment the current routing systems with these programmable interfaces... so we will not add routing loops and so they will work with our routing and signaling protocols," Ward said. "If this work doesn't get taken up because of lack of interest by this community, then it's guaranteed to be done somewhere else. I strongly suggest it be taken up here." 

Open Internet Endowment

Funding for IETF activities primarily comes from three sources: participant registration fees, corporate sponsorships, and an annual contribution from the Internet Society. The IETF leadership believes that the IETF must diversify its revenue sources in order to preserve the important work of the IETF. The Open Internet Endowment (OIE) is being established as one way to provide funds for the IETF as long as the purpose exists. Once the OIE is capitalized, its principal generates a return—the kind of return that can provide a new revenue source for the IETF. The larger the endowment, the larger the support it will provide.

At IETF 84, the Internet Society, as the administrative home of the IETF, and IETF leadership launched the OIE to the IETF community. The endowment's stated purpose: to support open standards initiatives, including the IETF and complimentary activities that promote the health of the Internet.

In addition to support for the IETF, the OIE will also provide funding for initiatives that support an open Internet. At IETF 84, an aspect of the OIE was launched that will directly support the work of the IETF. Thus, the introduction of the OIE was called a *family* launch. The IETF community was invited to be first supporters and give directly before a more public announcement of the OIE takes place, most likely in 2013.

The response was fantastic! During the Vancouver meeting 110 persons gave or pledged support of more than \$28,000. Monies will continue to be raised to capitalize the endowment—look for the OIE booth at IETF 85 in Atlanta. And remember, everyone who gives at IETF 85 will receive donor recognition gifts (including the new OIE t-shirt!) based on their level of their giving.

For more information and to donate, see www.openinternetendowment.org.



IPSO Alliance Successfully Demonstrates Internet of Things Interoperability

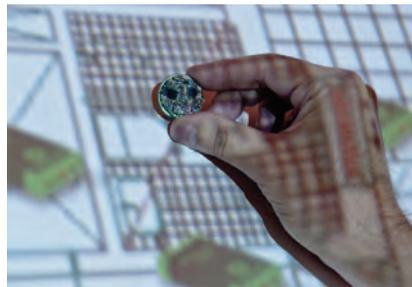
By Geoff Mulligan

During IETF 84 in Vancouver, Canada, the Internet Protocol for Smart Objects (IPSO) Alliance conducted an interoperability event that primarily focussed on the IPSO Application Framework, but also included testing of 6lowpan, Routing Protocol for Lossy Networks (RPL), and Constrained Application Protocol (CoAP)—all IETF developing standards. This was not the first Interop held in conjunction with the IETF; the Alliance organized a successful test of RPL at IETF 77.

Colocating the Interop with the IETF made it possible to provide direct feedback to engineers working on the various Internet Drafts. In addition, we were able to leverage the vast experience, knowledge, and support of the IETF network operations center (NOC) team for network configuration and management.

IPSO has held six interoperability events since its conception in 2008. The purpose of these events is to provide developers an opportunity to come together in a vendor-neutral setting and learn from one another about implementation issues for the various protocols. Afterward the Alliance and the vendors provide feedback to the IETF about inconsistencies in the Internet Drafts or RFCs. A number of issues related to early versions of RPL were communicated back to the Routing over Low Power and Lossy Networks (ROLL) working group in order to improve the developing drafts.

At this latest event, the participants demonstrated Internet of Things interoperability using the current IPv6 and Web standards developed for machine-to-machine (M2M) devices and cloud services that will be used in connected home, building automation, lighting, and smart-energy related applications. The event showcased Web-enabled smart objects from multiple vendors, each exchanging application payloads



IPSO smart-object device

in an interoperable manner based on the IPSO application framework.

The IPSO Application Framework (see www.ipso-alliance.org), defines a representational state transfer (REST)ful design for use in IP smart objects for M2M applications. It identifies a set of REST interfaces that may be used by a smart object to represent its available resources and to interact with other smart objects and backend services. It has been extended to cover a wide range of use cases and to more precisely describe the parameters of the smart objects involved in this Interop during IETF 84. This template enabled participants to interact seamlessly between devices located at the Interop in Canada and remotely in Finland, France, Sweden, and the United States via a cloud-hosted application.

During the three-day event, eight IPSO member companies participated

in the Interop, including ElectroTest Sweden AB, Ericsson, Nivis, Nokia, Proto6, Sensinode, Sensus, and Wateco. Some of the devices tested included electric meters, edge routers, and numerous types of home and building sensors and controllers, and a proximity sensor that monitored access and movement in a participant's hotel room.

There is a lot of preplanning work that goes into bringing together a successful testing event. The Alliance's Interop and Smart Energy committees worked tirelessly to define the goals and subsequent protocol parameters to provide the best opportunity for successful interoperability. In the end, every one of the devices was able to exchange sensor data with the cloud server and some of the devices were able to be controlled from the server. Future Interops are planned that will extend the functionality for direct device-to-device control and look at demonstrating industry-focussed applications for areas such as home automation, appliance control, healthcare, building control, and automotive.

At the end of the member Interop event, the IETF community was invited to see some of the devices and products being developed, including IPv6-enabled light bulbs and an IPv6-connected heart rate monitor, as well as electric meters, routing devices, sensors, and cloud applications.

The IPSO Alliance plans two additional interoperability test events during 2012: the first will be in conjunction with the European Telecommunications Standards Institute (ETSI) event in October and will focus on the CoAP protocol; the second will be in Atlanta at IETF 85, where we will be looking to test IP over 802.15.4g and additional extensions to the IPSO Application Framework. 

Future Interops are planned that will extend the functionality for direct device-to-device control and look at demonstrating industry focussed applications.

Software-defined Networking Efforts Debuted at IETF 84

By Haiyong Xie, Tina Tsou, Diego Lopez, Ron Sidi, Hongtao Yin, and Pedro Andres Aranda

Demonstrations at the Bits-N-Bites session of the IETF 84 meeting in Vancouver, Canada, featured the framework of a software-defined networking (SDN) solution suitable for the telco environment, as well as new proposals such as an SDN northbound interface for the coexistence of SDN with application-layer traffic optimization (ALTO) and the application of ALTO in an east-west interface interconnecting SDN controllers.

SDN is a hot topic these days, as illustrated by Monday's plenary discussion and the incredibly crowded SDNRG BoF at the IETF 84 meeting. SDN allows separation of the control plane and the data plane in network equipment, whereby the control plane is implemented in software in a logically centralized manner, and the data plane is implemented in commodity network equipment. OpenFlow is a leading enabler of SDN architectures.

As a slightly more mature topic, ALTO has an established IETF Working Group of its own. ALTO provides applications with information to, for example, perform better-than-random initial peer selection, and its services may take different approaches, including maximum bandwidth, minimum cross-domain traffic, and lowest cost to users. The ALTO information and services could help network applications to optimize application-layer traffic and to improve the performance of both the network and applications.

The demonstrations were based on two Internet Drafts:

- The SDN+ALTO draft1 (draft-xie-alto-sdn-use-cases) describes architectures for the coexistence of SDN and ALTO.
- The SDNi draft2 (draft-yin-sdn-sdni) proposes SDN interconnectivity and message exchange among multiple SDN controller instances within the scope of a single administrative domain.

The SDN+ALTO draft focuses mainly on interacting with an ALTO server to provide SDN-specific network state information and to provision the network in a more efficient and optimal manner. However, network provisioning through SDN+ALTO only may be difficult due to challenges such as granularity, scalability, and real-time issues. These difficulties could be addressed via the approach proposed in the SDNi draft.

The ALTO information and services could help network applications to optimize application-layer traffic and to improve the performance of both the network and applications.

SDN Domains: Partitioning the Control Plane

Both drafts emphasize the new concept of SDN domains, which was introduced to support the need for partitioning a network control plane among different controllers within an administrative domain. An SDN domain can be defined as the portion of the network (a "subnetwork," although not necessarily in the traditional IP-oriented sense) being managed by a particular SDN controller. Figure 1 shows an example of this concept. Among the reasons for introducing the concept of SDN domains we can emphasize:

1. **Scalability.** Because the number of devices an SDN controller can fea-

sibly manage is likely to be limited, a reasonably large network may need to deploy multiple SDN controllers.

2. **Privacy.** A carrier may choose to implement different privacy policies in different SDN domains because, for instance, an SDN domain may be dedicated to a set of customers who implement their own highly customized privacy policies requiring that some networking information in this domain (e.g., network topology) should not be disclosed to an external entity.
3. **Incremental deployment.** A carrier's network may consist of portions of legacy and nonlegacy infrastructure. Dividing the network into multiple, individually manageable SDN domains allows for flexible incremental deployment and a more-seamless network evolution.

There are many other reasons for partitioning a network using the concept of

SDN domains. For instance, if an SDN controller is used to manage a certain portion of the network, losing connectivity between the controller and the network could result in outages (or reduced service); this mandates locality between the controller and the network under its control in order to minimize the risk of network problems causing more network problems. Additionally, large amounts of mergers among companies introduce the need to interconnect data centers, and in many cases a certain customer will deploy its services on multiple data centers from multiple providers. In these cases, the data centers will be SDN managed in the near future, thereby requiring SDN interconnection.

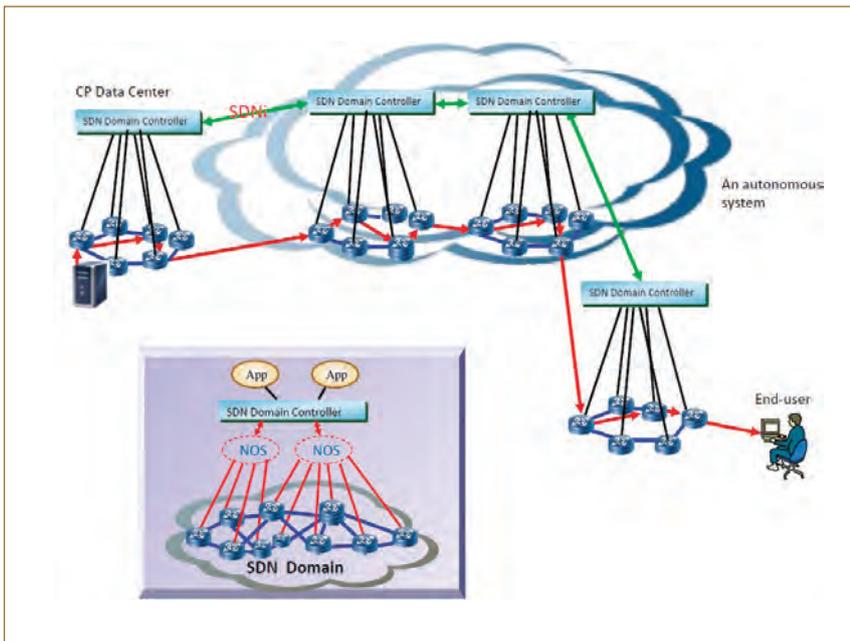


Figure 1. Example of SDN Domains

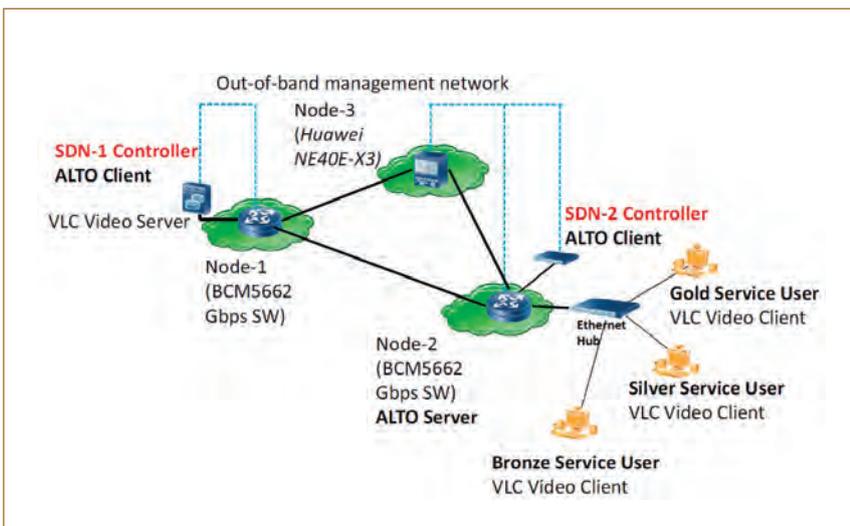


Figure 2. Demonstration Setup

In the demonstrations, a simple network topology was used to illustrate these key ideas. The topology consisted of three infrastructure nodes and two SDN domains, as shown in figure 2. All three nodes supported OpenFlow 1.0 and were SDN-capable. Node-1 and Node-2 were 1 Gbps switches based on the BCM5662 chipset, and Node-3 was a Huawei NE40E-X3 router. The two SDN domains are labeled SDN-1 and SDN-2. Domain SDN-1 consisted of

only Node-1; domain SDN-2 consisted of Node-2 and Node-3. Each of these two SDN domains had its own SDN controller managing the corresponding nodes. An ALTO client was integrated with each of these two SDN controllers, and to simplify the demonstration configurations, an ALTO server was collocated with Node-2, and a VLC video server was collocated with controller SDN-1. The demonstration showed three video streams corresponding to

three network flows, from the video server to three clients connected to Node-2.

SDN+ALTO: A Northbound Interface for the Coexistence of SDN and ALTO

In a network where SDN and ALTO both operate, a northbound interface is needed to support their coexistence. In draft-xie-alto-sdn-use-cases, the authors propose that SDN controllers become a special type of ALTO client in the following ways:

1. SDN controllers feed the ALTO server with SDN-domain-specific and possibly aggregated network information, policed by the specific privacy policies that are enforced by the corresponding SDN domains. The draft does not specify what specific protocol the SDN controllers should use to export information. Some protocol drafts (e.g., BGP-LS3) could be applicable for this purpose, however, the privacy policy, which governs how network information should be aggregated before exporting to the ALTO server, could be difficult to integrate with BGP-LS unless the latter is significantly improved by explicitly allowing the implementation and execution of foreign privacy policies.

2. SDN controllers obtain key SDN-domain-specific information from the ALTO servers in order to manage the whole network. Key information obtained from ALTO must be customized specifically for SDN and SDN domains, since the SDN controllers cannot use the vanilla non-SDN-domain-specific ALTO information directly. For example, an SDN controller needs a network cost map or end-point ranking that is specific to the controller's SDN domain. More

Continued on next page

Software-defined Networking Efforts Debuted at IETF 84, continued

important, the SDN controller may require key information, such as a set of SDN-domain-level paths, in order to correctly and fully utilize the cost map or end-point ranking. In non-SDN environments, ALTO clients likely do not have the ability to choose routes, therefore cost maps need not come with any path information. However, in SDN-enabled networks, SDN controllers are special infrastructural ALTO clients that are able to—and should—choose routes. In this scenario, cost map (or end-point ranking) by itself without any path information is not very useful.

In the demonstration for SDN+ALTO coexistence (see figure 3), three flows—gold, silver, and bronze video streams—were injected into the network one by one. Controller SDN-1 obtained cost maps or ranking information from the ALTO server and decided which path should be set up for each flow. According to the information returned by ALTO, the upper path (Node-1 to Node-3 to Node-2) was set up for the gold flow, while the lower path (Node-1 to Node-2) was chosen for the silver flow and the bronze flow. There are many possible reasons for this flow assignment, including engineering the traffic to minimize the maximum number of link utilizations (putting the silver and bronze flows together minimizes the maximum of aggregated flow rates on the two paths); or the existence of transient flows consuming a significant portion of the capacity on the upper path when assigning the gold flow (necessitating that the silver flow be assigned on the alternative path, however, the transient flow may disappear after the assignment).

In the demonstration, assigning the silver and bronze flows to the lower path caused congestion on Node-2. The congestion may have been a result of

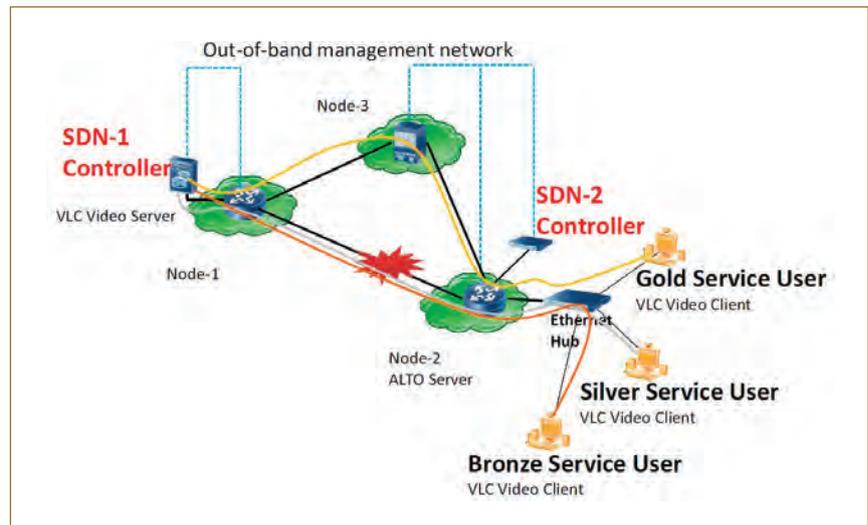


Figure 3. SDN+ALTO Demonstration

background traffic on the link between Node-1 and Node-2. Such background traffic may have existed before or after setting up paths for the two flows. In



Weiqian Dai and Tina Tsou explain the demonstration to Scott Bradner and others.

both cases, controller SDN-1 as a client of ALTO was not able to receive the most up-to-date information from the ALTO server, because ALTO only provides coarse-grain, nonreal-time cost maps or ranking services. As a result, the video quality of both the silver and bronze streams was severely degraded.

This demonstration suggests that with ALTO alone as the source of network information for SDN, network performance may suffer—the coarse granularity of ALTO is not sufficient for SDN to operate on the fine granularity

of network flows. Clearly, we need other mechanisms (e.g., SDNi) to solve such problems.

SDNi: An East-West Control Plane Interface for SDN Controllers

Multiple SDN controllers are likely to be deployed in a reasonably large network administrative domain; interconnecting these controllers in order to share information and coordinate their decisions becomes inevitable—and a key factor for success. In draft-yin-sdn-sdni, the authors propose a protocol framework called SDNi to interconnect SDN controllers. This draft describes the interfaces for exchanging information among multiple SDN domain controllers and the benefits of coordinating network provisioning using SDNi.

The SDNi demonstration shown in figure 4 illustrates this framework. Besides interconnecting multiple controllers, SDNi also allows them to share control-plane network information and coordinate their decision-making processes.

In this case, when congestion happens, Node-2 detects and reports the congestion information to its controller, which again shares this information with controller SDN-1. On receiving

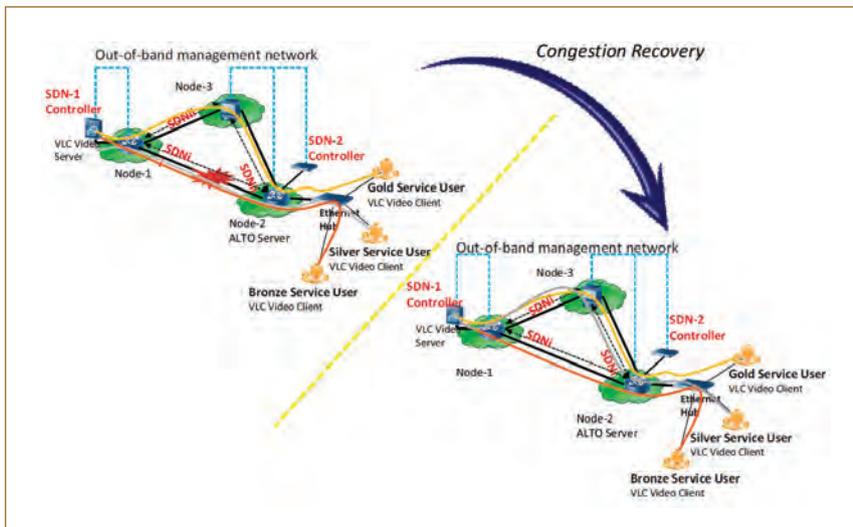


Figure 4. SDNi Demonstration

the congestion report through SDNi, controller SDN-1 can then recover from the congestion by rerouting the flows. More specifically, the silver flow is shifted from the lower path to the upper path. After this shift, the video quality of both silver and bronze flows is significantly improved.

Feedback and Conclusions

The new proposals for SDN evolution described in the two drafts mentioned previously were demonstrated during the Bits-N-Bites session in what we believe is the purest IETF style: running code. Demonstrating via the archetypal video streaming service, the Telco SDN framework, as well as the innovations on

the northbound and east-west interfaces in the SDN architecture, provided a win-win solution for network carriers and Internet content providers. In particular, SDN+ALTO achieved goals such as flexibility, privacy preservation, and independent evolution; SDNi provided an approach to partitioning a network administrative domain and to interconnecting islands of SDN domains to form a coherent network.

The commitment and effort it took to put on the demonstrations were acknowledged and appreciated by many IETF participants. Russ Housley, IETF chair, said to Tina Tsou “I saw people at your table all evening. Thank you very much for your contribution

The new proposals for SDN evolution described in the two drafts mentioned previously were demonstrated during the Bits-N-Bites session in what we believe is the purest IETF style: running code.

to the IETF.” IETF Participants from network carriers such as KDDI commented that it was great to see an initial effort of SDN solution done in IETF while most of the people were still talking about SDN in slides, and that they’d like to see more developments in this area. Said Raquel Morera from Verizon, “we’re interested in SDN and in following developments in the area.”

The contributors to the demonstrated draft and the demonstration itself included (alphabetical within each institution):

- Telefonica: Pedro Andres Aranda, Diego Lopez
- Huawei Technologies: Weiqian Dai, Ritesh Mukherjee, Tina Tsou, Haiyong Xie, Ken Yi, Hongtao Yin
- Contextreme: Ron Sidi
- Alcatel-Lucent: Vijay Gurbani



References

1. H. Xie, T. Tsou, D. Lopez, H. Yin, and V. Gurbani. Use cases for alto with software defined networks. (IETF Internet-Draft), draft-xie-alto-sdn-extension-use-cases-00, 2012.
2. H. Yin, H. Xie, T. Tsou, D. Lopez, P. A. Aranda, and R. Sidi. SDNi: A message exchange protocol for software defined networks (SDNs) across multiple domains. (IRTF Internet-Draft), draft-yin-sdn-sdni-00, 2012.
3. H. Gredler, J. Medved, S. Previdi, and A. Farrel. North-bound distribution of link-state and TE information using bgp. (IETF Internet-Draft), draft-gredler-idr-ls-distribution-02, 2012.



Demonstration personnel (from left to right): Weiqian Dai, Tina Tsou, Ritesh Mukherjee, Ken Yi, Huaru Yang, and Haiyong Xie.

Moving Toward a Censorship-free Internet

By Johan Pouwelse

Internet censorship was the focus of an initiative proposed in Vancouver at an informal bar birds-of-a-feather (BoF) meeting. In this article Johan Pouwelse discusses his motivations for organizing the meeting. An initial draft discussion document is available.¹

Bits moving across the Internet are vulnerable to surveillance and censorship on an unprecedented scale. Today, both Internet providers and governments possess the ability to monitor the moves of their digital citizens from central infrastructure points—an ability that creates significant potential for abuse, and a threat that goes beyond the scope of mere monitoring or filtering.

Lessons Learned from the Arab Spring

Although Internet kill switches do not exist, governments have demonstrated their ability to disable communications networks in times of crisis. During the 2011 Arab Spring, Egyptian authorities demanded that telecommunication companies sever their broadband connections and mobile networks—both local and European operators were forced to comply, and, as a result, digital Egypt vanished. Despite the country's decentralized infrastructure, an Internet blackout was relatively easy to carry out.

The roles—and consequences—of social media (e.g., Facebook and Twitter) during that same period further illustrate the capacity governments have for Internet censorship and the challenges activists face in combating it. The April 6 Youth Movement from Egypt committed digital dissent in full public view. According to *The New York Times*,² the movement “provided a structure for a new generation of Egyptians, who aren't part of the nation's small coterie of activists and opinion makers, to assemble virtually and communicate freely about their grievances.”

But moving protest organizations to social media not accessible to the public-at-large can hold surprising risks. On the ground, the movement's strike organization and protests in Facebook groups, many with thousands of followers, triggered arrest and imprisonment. Protesters in other countries quickly took note of Egypt's lesson and disabled their public Facebook profiles. In response, one government initiated social media searches on incoming, young, plane travelers by forcing them to log in to Facebook upon arrival, thereby revealing online activities and any anti-government sympathies.³

Is there a Role for the IETF?

What happened in Egypt underscores how essential it is that IETF participants fully and comprehensively understand the entire Internet ecosystem when considering the question of a censorship-free Internet. Both government-regulated Internet and public discussions have risks.

Should the IETF engage in this very political aspect of the Internet? Or are there other organizations better equipped to deal with it?

Room consensus quickly moved toward the importance of documenting censorship models and assessing the need for new technology—using the Arab Spring scenario as inspiration. At IETF 85 in Atlanta we hope to present an improved Internet Draft detailing the scenario and how social networks, microblogging, and camera phones proved essential for a new generation of Internet users.

Getting Around Censorship

The onion routing technology in Tor has proved itself over the past 10 years as an effective tool against censorship. The technology's anonymity, unlinkability, and unobservability have made it popular—numerous users are willing to accept its slower browsing experience in exchange for its privacy-enhanced web access.⁴ But onion routing alone cannot overcome the threat of government-imposed Internet shutdown. The challenge is to design a censorship-free Internet sustainable even when an adversary controls the underlying infrastructure.

As early as 2006 it was reported that individuals in wide swathes of the Arab world were using Bluetooth technology to bypass police restrictions. According to news reports,⁵ communication between men and women in this region, in extreme cases referred to as dating, had been made possible by cellphone technology. When Bluetooth-capable phones are in close proximity, they can engage directly in digital and social chatter—no other infrastructure is needed. Moreover, when sharing photos or bandwidth-hungry videos with friends it also pays to be close. Government provided cellphone networks might not be filtering you, but can still be dreadfully slow. It therefore pays to use cell phones' Bluetooth-based, direct file-transfer features—and it comes as no surprise that wireless-transfer apps have seen millions of installs.

A query of Google Trends for the phrase *Bluetooth transfer* reveals a geographical spread of this interesting social phenomenon. Below is a list of countries ranked according to search volume.⁶

1. Philippines
2. India
3. Pakistan
4. Singapore
5. Malaysia
6. United Arab Emirates

7. Hong Kong
8. Indonesia
9. South Africa
10. Iran

It seems millions of mobile phone owners are already employing the social practice of wireless data exchange.

The Musubi smartphone app represents another key, censorship-free, technology advancement. Developed at Stanford University, it offers instant messaging service and media sharing capabilities similar to WhatsApp, Ping, and Blackberry Messenger. What makes it unique is that all data and processing resides on the smartphones, not in the cloud. This decentralization removes the need for central processing and provides significant decoupling from the underlying infrastructure. Exchange of cryptographic keys is integrated in the friending process—Musubi essentially builds a decentralized social graph. But Musubi is also limited—all data transfers go through central servers, as it lacks NAT-traversal capability.

A more general solution is found in delay-tolerant networking (DTN) technology, which uses a simple store-and-forward primitive to communicate over heterogeneous links. Mobile ad hoc networks have been studied within the Internet Research Task Force (IRTF) since 1997 and we hope that much of that knowledge can be reused, despite our scenario differing slightly from DTN (as being investigated by the IRTF [RFC4838]). The DTN focus is on finding routes to an explicitly given destination, usually by maintaining routing tables. As our earlier Bluetooth-based filesharing example showed, dissemination in the Arab Spring scenario is likely to involve an explicit copy between people who trust each other.

Microblogging proved to be a vital tool in the Arab Spring. Several research groups investigated Twitter-like services without the need for central servers. According to one Twitter investor and

former engineer, “Done right, a decentralized one-to-many communications mechanism could boast a resilience and efficiency that the current centralized Twitter does not. Decentralization isn’t just a better architecture, it’s an architecture that resists censorship and the corrupting influences of capital and marketing. At the very least, decentralization would make tweeting as fundamental and irrevocable a part of the Internet as email.”⁷

The Twimight project by ETH-Zurich university shows that decentralized microblogging already exists. Researchers developed an Android application that uses Twitter servers in normal conditions, but switches to a Bluetooth-based disaster mode when Internet connectivity is lost.

Censorship-free technology also have arisen from within the IETF.

- A voice-over-IP protocol using peer-to-peer technology and using a distributed hash table (DHT) for scalability has been standardized. Unfortunately, DHTs are notoriously difficult to secure.
- The peer-to-peer streaming protocol (PPSP) working group in the transport area developed a serverless video streaming protocol, using Bittorrent-like swarming. Pioneer Research UK showed a fully functional set-top box using this new draft protocol with support for both live streaming of BBC feeds and video-on-demand playback at the IETF’s recent meeting in Paris.
- Currently an open source PPSP implementation is available for Android which integrates with Twitter. By tweeting links such as `ppsp://2b2fe5f1462e5b7ac4d7`, it is now possible to augment a Tweet with eyewitness video footage. This architecture has interesting anti-censorship properties, as it is free from DNS, HTTP, or any other server infrastructure. The SHA1

hash part of this URI is used to find peers in a low-latency DHT, which are then used to stream video in peer-to-peer fashion.

A Powerful Adversary

We must assume from the Arab Spring scenario the existence of a powerful adversary. The following threats⁸ have been identified for similar circumstances:

- The adversary can observe, block, delay, replay, and modify traffic on all underlying transport. Thus, the physical layer is insecure.
- The adversary has a limited ability to compromise smartphones or other participating devices. If a device is compromised, the adversary can access any information held in the device’s volatile memory or persistent storage.
- The adversary can choose the data written to the microblogging layer by higher protocol layers.
- The adversary cannot break standard cryptographic primitives, such as block ciphers and message-authentication codes.

The advances listed previously indicate the wealth of experience, related technologies, and available building blocks that an IETF initiative could use to work toward a censorship-free Internet. 

References

1. <http://tools.ietf.org/html/draft-pouwelse-censorfree-scenarios-01>
2. <http://www.nytimes.com/2009/01/25/magazine/25bloggers-t.html>
3. <http://online.wsj.com/article/SB125978649644673331.html>
4. http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5569995
5. <http://www.washingtonpost.com/wp-dyn/content/article/2006/08/05/AR2006080500930.html>
6. <http://www.google.com/trends/?q=bluetooth+transfer>
7. <http://al3x.net/2010/09/15/last-thing-about-twitter.html>
8. <https://fulpool.org/btp.pdf>

Using JSON in IETF Protocols

By Andy Newton

In IETF protocol activities in which a data serialization format is not natural to the problem domain, chances are discussion will be raised about using JavaScript Object Notation (JSON). And chances are that discussion will compare JSON with the Extensible Markup Language (XML). Today, in the IETF, where XML was once the preferred text-based data format, JSON is being used more and more.

This increased use by the IETF mirrors a trend in the broader industry. ProgrammableWeb.com keeps a directory of Internet application programming interfaces (APIs). In 2009 they listed only 191 registered JSON APIs.¹ Today, 3,285 JSON APIs are on their registry (compared to 4,543 registered as XML). That's a pretty good showing, especially considering that XML is more than twice the age of JSON.

The juxtaposition of JSON with regard to XML was probably best put by James Clark, technical lead of the W3C Work Group that created XML 1.0, in a 2010 blog post:²

“I think the Web community has spoken, and it's clear that what it wants is HTML5, JavaScript and JSON. XML isn't going away but I see it being less and less a Web technology; it won't be something that you send over the wire on the public Web, but just one of many technologies that are used on the server to manage and generate what you do send over the wire.”

In other words, JSON will not replace XML, but a lot of the structured data communications in which XML was once the only consideration will eventually migrate to JSON.

While a lot can be said for the heft of XML complexities contributing to this trend, a better explanation of JSON's popularity is that it is simply simpler than XML. After all, the abundance of XML technologies did not inhibit XML adoption before JSON came on scene—by any measure, XML is a success. But JSON is a simpler solution to a subset of

the problems solved by XML—and that subset is a rather large one.

JSON's simplicity comes in two forms: the rules for serializing JSON data, and the interfaces for parsing JSON data, particularly in today's more-popular, dynamic languages.

JSON's simplicity comes in two forms: the rules for serializing JSON data, and the interfaces for parsing JSON data, particularly in today's more-popular dynamic languages.

Unlike XML, JSON serialization and parsing rules are fairly simple. JSON data is composed of objects, which, in turn, have members, and arrays, and the more primitive types, string, boolean, number, and null. They possess neither mixed content nor metadata structures. And they contain no processing instructions or character entity references—no facility at all for commenting.

Thanks to these basic differences, the API in many programming languages can also be quite simple. Because there is no need to account for mixed content, processing instructions, and the like, many dynamic languages easily map JSON data into their native object and array constructs—resulting in seamless access to both JSON-rendered data and the other program's data. For many programmers, outside of a parse command, there appears to be no JSON API.

The Devil Is in the Details

As with any technology, the lure of simplicity and the need for features is a function of the problem being solved. Any specification activity looking at JSON should evaluate the following issues.

JSON Is Not JavaScript. Though widely known, sometimes it bears repeating. JSON is a language-independent data format. That said, understanding JSON's heritage is important for two reasons.

- The constructs of JSON objects and arrays follow JavaScript semantics. Members of JSON objects are considered unordered; no member

should be repeated within an object (this may not produce an error, but the behavior is dependent on the parser's implementation). Arrays, however, are considered ordered collections of values.

- Although JSON is not JavaScript, it is valid JavaScript syntax. Therefore, for security reasons, JSON parsers should take precautions against JavaScript embedded in JSON data (e.g. JavaScript engines should not simply parse JSON using the eval function), and serializers should prevent the injection of JavaScript into JSON data.

Character Set Support. JSON is strictly Unicode; protocols relying on non-Unicode encodings will be incompatible. This is not generally a problem, but does contrast with XML, in which US-ASCII, ISO-8859-1, Shift_JIS, Big-5, GB, and other character sets can be explicitly specified.

Mixed Content and Markup. Another contrast with XML is JSON's lack of mixed-content support, whereby data-format syntax can be intermingled with message text. This is what makes XML ideally suited for markup and annotating text with metadata. For those purposes, JSON is a poor choice.

Schemas. There is no one standardized schema language for JSON, although several are presently in the works (including one by this author). The need for a JSON schema language is controversial—JSON is regarded by most as simple enough on its own. Indeed, there is no shortage of JSON-based interchange specification making due without schema formalism.

The inherent complexity of a generalized schema language may also be unnecessary. The Application-Layer Traffic Optimization (ALTO) working group has created a simple schema language for JSON based upon the C programming language's struct syntax—an approach that sufficiently covers all ALTO needs.

But the significance of JSON Schema should not be understated. While JSON Schema is not yet an IETF proposed standard, it has been used in the specification of many private, JSON-based protocols.

Namespaces. JSON does not (yet) have a standardized, namespace specification. A few have been proposed, but none have gained significant traction.

Many JSON specifications have no need for a namespacing system—nor for the complexities that come with a federated namespacing system. One example is the Web Extensible Internet Registry Data Service working group, which has developed a simple prefix rule for JSON names to avoid collisions between names in standardized specifications and custom extensions.

Performance. The needs of a protocol are often dictated by performance, of which there are multiple aspects. Common

wisdom holds that JSON is less verbose than XML, resulting in smaller payload sizes and needed memory.

But JSON is also considered to be faster than XML. A study conducted at Montana State University³ that modeled the transfer of data over a network found that in some scenarios JSON could be 41 times faster than XML—and consume noticeably less CPU time and memory.

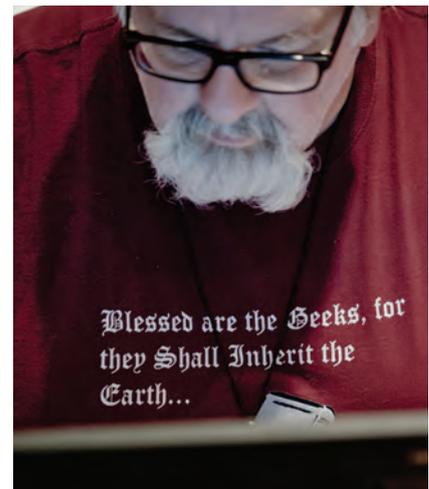
The types of APIs available can also impact performance. Particularly for applications with large datasets in which an in-memory, tree-node API would consume too many resources. For years, XML benefited from the quasi-standard SAX specification, an event-based API for consuming XML. JSON parsers with similar concepts are now available: Jackson, LitJSON, YAJS (its various language bindings), and GSON, to name a few.

Normative Reference. No discussion about using JSON in IETF protocols is complete without noting that the JSON-defining document, RFC 4627, is published as an Informational RFC. Because JSON plays such a pivotal role in IETF standards, there are discussions of moving JSON into the standards track. Until then, RFC 4627 is listed in the DOWNGREF registry⁴ and can be used as a normative reference.

References to ECMA's ECMAScript Language Specification also are unnecessary as a normative reference, as the ECMAScript 5.1 Edition⁵ normatively refers to RFC 4627 with slight differences called out.

An Eye to the Future

With the aforementioned considerations in hand, one must recognise those IETF activities in the area of JSON infrastructure: the Applications Area working group (APPSAWG) has discussions from time to time regarding JSON Schema (and other JSON schema languages) and is currently working



An IAB/IRTF workshop participant exhibits the combination of irreverence and humility we've come to expect at IETF meetings.

on JSON Patch and JSON Pointer, specifications for referencing values in a JSON document and for applying changes to a JSON document with HTTP PATCH. In addition, the JavaScript Object Signing and Encryption (JOSE) working group is tasked with developing cryptographic integrity and encryption specifications for JSON.

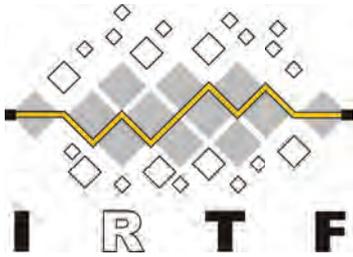
As it is the broader trend in the computing industry, one should expect to see more JSON RFCs from the IETF—both as client standards and as building blocks for JSON-based stacks. 

References

1. "The Stealthy Ascendancy of JSON," Lori MacVittie, F5 DevCentral, <https://devcentral.f5.com/weblogs/macvittie/archive/2011/04/27/the-stealthy-ascendancy-of-json.aspx>
2. "XML vs The Web," James Clark, James Clark's Random Thoughts, http://blog.jclark.com/2010/11/xml-vs-web_24.html
3. "Comparison of JSON and XML Data Interchange Formats: A Case Study," Nurzhan Nurseitov, Michael Paulson, Randall Reynolds and Clemente Izurieta, <http://www.cs.montana.edu/izurieta/pubs/caine2009.pdf>
4. DOWNGREF Registry, <http://trac.tools.ietf.org/group/iesg/trac/wiki/DownrefRegistry>
5. "ECMAScript Language Specification, 5.1 Edition," ECMA-262, June 2011, <http://www.ecma-international.org/publications/files/ECMA-ST/ECMA-262.pdf>

IRTF Update

By Lars Eggert



During IETF 84 in Vancouver, Canada, two chartered Internet Research Task Force (IRTF) research groups (RGs) held meetings:

- ICCRG—Internet Congestion Control research group
- ICNRG—Information-centric Networking research group (new!)

The ICNRG was chartered following IETF 83, and members held their first official meeting in Vancouver. Information-centric networking is an approach to evolve the Internet infrastructure by introducing uniquely named data as a core Internet principle. Data becomes independent from location, application, storage, and means of transportation, enabling in-network caching and replication.

Expected benefits include improved efficiency, better scalability with respect to information and bandwidth demand, and increased robustness in challenging communication scenarios. The ICNRG charter is available at <http://irtf.org/icnrg>.

In addition to the meetings of chartered research groups, a new proposed research group on software-defined networking held an initial, very well attended meeting. Despite a quite broad problem space that will require further refining, the group seems to be on a good trajectory to be chartered. A second, side meeting discussed proposing a research group on energy efficiency.

Since IETF 83, two research groups have closed due to lack of energy and participation: IP Mobility Optimizations (MOPOPTS) and Host Identity Protocol (HIPRG).

Since IETF 83, no new RFCs were published on the IRTF RFC Stream.

The IRTF Open Meeting at IETF 84 was the venue for the first of three Applied Networking Research Prize (ANRP) winners of 2012 to present their research. Alberto Dainotti presented his research into Internet communication disruptions due to filtering. Two additional ANRPs were awarded in 2012, and winners will present their work at IETF 85 in Atlanta, Georgia, U.S.A.

The 2013 nomination and selection cycle of the ANRP has already begun. The ANRP is awarded for recent results in applied networking research that are relevant for transitioning into shipping Internet products and related standardization efforts. It is supported by the Internet Society in coordination with the IRTF. See <http://irtf.org/anrp> for details on the award and instructions for nominating researchers for the prize. The deadline for nominations for the 2013 cycle is 30 November 2012.

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



IETF 84 participants

Regulatory Fellows Attend IETF 84 in Vancouver

By Karen Mulberry

After the success of ISOC's Pilot Fellowship Programme for Regulators (IETF 83 Paris, France), the Internet Society formally launched the programme in Vancouver at IETF 84 by inviting another five public policy guests to the meeting. In addition to attending IETF Working Group sessions, the Fellows spoke with experts within the IETF community on topics such as IP routing, security, and management of key Internet resources. These discussions enabled the Fellows to gather both expert viewpoints and a better understanding of the kind of solutions that have been developed through the IETF standards process.

Fred Clark, one of the five Regulatory Fellows selected to attend IETF 84, is manager of telephony at the Superintendence of Telecommunications in Guatemala. His technical background is systems engineering, and his main areas of work are telephony, Internet, and universal service. Clark brought more than 20 years of experience in the implementation of software and technology projects in urban and rural environments to his experience as a Fellow. Following, he shares his expectations and the rewards of attending his first IETF meeting.

My goal for IETF 84 was to learn all I could about the transition from IPv4 to IPv6, and how Internet standards are made. During the six days of the conference, I learned all that time and my

knowledge allowed me to about IPv6—in that regard, my expectations were met 100 percent. In addition, I went there with a point of view that got wider and wider as the days went by and as I was exposed to different meetings and conversations with very fascinating people. I can say now that I came in with a view from a spyglass, and came out with a very rich panorama of the Internet. My time at IETF 84 was an invaluable learning experience that at the end was larger than my original expectations.

The opportunity to [attend the meeting in person] was invaluable. It's very important to learn and experience from the inside how the Internet works, what organizations make it happen, and what is the role of the Internet Society in keeping it the way we know it and use it—in complete freedom, open to everyone, and without any discrimination or prejudice.

Without a doubt, those are the characteristics I would like to see and keep for the Internet in years to come. All I would like to add is a wider penetration of the Internet throughout *all* of our countries, so that we may all enjoy the possibility of learning and growing in an open-Internet world. 🌐



Internet Society Public Policy Fellows and Internet Society staff, IETF 84, (left to right) Sally Shipman Wentworth (Internet Society), Fred Clark, Jesada Sivaraks, Kila Gulo-Vui, Nigel Cassimire, Karen Mulberry (Internet Society), and Markus Kummer (Internet Society).

Different Experiences, a Common Goal: ISOC Fellows at IETF 84

By Erin McGann

On 30 July 2012, 11 network engineers and entrepreneurs from nine emerging nations met at the Hyatt Regency in Vancouver, Canada, as part of the ISOC Fellowship to the IETF Programme. While most were only familiar with the IETF online, participation as a Fellow to the IETF was their opportunity to meet working group chairs and other IETF members face-to-face and to share their experiences in person.

Offered to technology professionals, advanced IT students, and other qualified individuals from emerging economies, the IETF Fellowship programme increases the diversity of inputs to, and global awareness of, the IETF's work. Every year since its inception in 2006 the programme has connected growing numbers of new Fellows with mentors and other members with similar interests within the IETF community. Although Fellows represent a wide range of backgrounds and experiences, they share a commitment to the betterment of their region and a passion for the collaborative process that defines the IETF.

Nomsa Muswai, a network engineer at Zimbabwe Online, learned about the IETF Fellowship programme last year and was determined to participate. "I met Steve Conte at a United States Telecommunications Training Institute workshop, and he told me about it," she says. "I told my company I was going to come and although they didn't send me here, they supported me. I believe that unless you take giant steps by yourself, you don't get chances."

Muswai views her position as a network engineer as a way to give back. "It gives me a platform from which to impact society," she said.

Paul Muchene, builds websites and mobile web programs in Kenya. "In Kenya, [the web] is a very emerging industry—more and more people are beginning to see the benefits of going online," he says. "We are where America was, perhaps, during the dot-com bubble. There's a bubble of people wanting web sites and domain names—so that's what I'm doing."

In addition to running his company, Muchene is also the network lead at iHub_, an open space in Nairobi for innovators, developers, venture capitalists, entrepreneurs, and graduate students to meet and work. Muchene keeps the network up and running, including recently deploying IPv6.

Shabbir Ahmed is an associate professor in computing science at the University of Dhaka. "We have a huge information gap in Bangladesh," he says. "We usually teach according to what textbooks are available, so in most cases we're not

conversant with cutting-edge technology. Attending a meeting like IETF 84 definitely enriches my knowledge—which I can then pass on to my students."

"It was a surprise when I learned that the man who was sitting beside me was a very well-known researcher in my field. I'd read lots of his papers—to suddenly find this man sitting beside me was very exciting!"

—Shabbir Ahmed
IETF 84 ISOC Fellow

Dorcas Muthoni, founder of Nairobi-based open source consulting firm OpenWorld, marks her third visit as an IETF Fellow this year. "OpenWorld builds web and mobile applications targeted towards enterprise and government. I work from Nairobi, in the East African region. That's what brings food to the table," she says. "But I also do other things. In 2004, I started AfChix, a mentorship programme for tech women, because there are so few of us in the region. My goal is to make sure that more young women can get training, feel technically up-to-date, and ultimately present themselves for leadership roles."

ISOC Fellows to the IETF are paired with mentors who share their interests. Mentors help Fellows navigate the meeting process, introduce them to other members, and answer questions.

Muswai immediately connected with her mentor. "Fred Baker is awesome," she says. "I'm fortunate to have him as my mentor. Through him I'm getting to know a lot of people whose names I knew previously only through their work."

She admits that the intensity of the meetings is definitely something to contend with. "Here [in person], I get so much information, sometimes I need to go back into a quiet space to go over it again."

"To benefit from the meetings you must be able to follow them," says Muthoni. "I find it best not to attend *every* session on *every* topic that interests me. I pick one, master the art of how the meeting runs, and then read the archives after that."

Many of the Fellows agree that meeting other IETF members in person makes following the mailing lists easier once they return home.



IETF 84 Fellows (left to right): Dorcas Muthoni, Vinayak Hegde, Baasan Sivaraks, Fernando Gont, Nigel Cassimire, Shabbir Ahmed, Xiaohong Deng, Saurabh Srivastava, Albert Kamaga, Nomsa Muswai, Mamodu Faty, Paul Muchene, Jesada Sivaraks, and Fred Clarke.

“Initially the mailing lists can be overwhelming,” says Muchene. “I get quite a bit of mail and the discussions and drafts can be a lot. But once I attended a meeting physically, it ceased to be overwhelming—I became more excited and started appreciating the work that was done in the mailing list. I’m much more focussed now—I know exactly which drafts to read and which ones to comment on.”

As a professor of computing science, Ahmed will benefit from the contacts he made at his first IETF meeting. “My institution doesn’t have access to the hardware needed to run simulations,” he says. Now he’s connected with IETF members who can help guide his research despite his institution’s limited equipment.

“My goal is to make sure that more young women can get training, feel technically up-to-date, and ultimately present themselves for leadership roles.”

—Dorcas Muthoni
IETF 84 ISOC Fellow

Muswai found the depth of specialized technical knowledge at the IETF to be different from what she was used to in Zimbabwe. “In my country, as in most of the southern part of Africa, there is a lot of development but everything is still a bit vague,” she says. “There’s a lot of generalization. [In other parts of the world] people concentrate on the finer details of a network. It’s quite frustrating because sometimes we have to outsource the implementation, although we do most of the maintenance. We don’t get much chance at home to learn deeper things. [At the IETF], they really know what they’re talking about. It’s not like learning stuff from school and then

trying to implement it. When I hear IETF members talk—it’s clear that they’re actually involved in the making of things.”

Fellows’ interests span from DNS to applications to MPLS, but all share a conviction to ensure open standards.

“If it weren’t for [the IETF], which pushed for open standards, it would have been impossible for me to start my business,” says Muthoni. “It helps us access things that otherwise would have been [out of our reach]. It also promotes more cooperation. If you look at this forum you’ll find people from different companies, from different backgrounds, all thinking: how can we get things working in a standardized way, so people can publish whatever they want to publish? That’s a really a big thing—working together. Otherwise we create only small ecosystems, like boundaries of countries that everyone needs visas to cross.”

Muswai sees how open access directly affects people in Zimbabwe. “Because it’s for everyone, the Internet makes it possible for a lot of things to occur—economically, politically, and socially. When people who have had no access to the Internet start using it, you see a lot of progression in those things that they couldn’t do before—even things as simple as communication.”

In emerging economies, Internet and IT industries can bring much-needed jobs. “More reliance needs to be made [on industries] besides tea, coffee, and tourism—these are very shaky,” says Muchene. “It’s very important to have an open Internet, a place for people to express ideas, to innovate on top of the IP network, and to perform transactions.”

“Participating in the IETF’s work to make the Internet work better through an open process feels great,” says Muswai. “It’s how I can make a difference. It’s going to change a few things in my life and I’m totally happy about it.”

For more information on the IETF Fellowship, please visit <http://www.internetsociety.org/fellows-ietf>.

If you know a great candidate for the IETF Fellowship programme, let us know! Visit <https://www.isoc.org/leaders> or email leaders@internetsociety.org for more information and an application.

To learn more about the IETF and its work on Internet standards, visit the IETF website at <http://www.ietf.org>.

IETF Ornithology: Recent Sightings

Compiled by Mat Ford

Getting new work started in the IETF usually requires a birds-of-a-feather (BoF) meeting to discuss goals for the work, to assess the suitability of the IETF as a venue for pursuing the work, and to identify the level of interest in and support for the work. In this article, we'll review the BoFs that took place during IETF 84, including their intentions and outcomes. If you're inspired to arrange a BoF meeting, please be sure to read RFC 5434: Considerations for Having a Successful Birds-of-a-Feather (BoF) Session.

Data Set Identifier Interoperability (DSII)

Description: This BoF was not intended to form a working group at this session. The purpose was to discuss how to achieve interoperability among persistent identifiers for data sets made available on the Internet. The initial use case of interest is scientific data sets produced by different research teams; other use cases might include media developed by different sources and combined into a common collection. Access policies based on identifiers, discovery, association of meta-data, and data integrity are expected to be later topics, but these will likely be covered in follow-on mailing list discussion. The meeting reviewed existing methods such as DOI, URN, PURL, and then discussed core requirements.

This BoF was not intended to form a working group at this session. The purpose was to discuss how to achieve interoperability among persistent identifiers for data sets made available on the Internet.

Proceedings: <http://www.ietf.org/proceedings/84/minutes/minutes-84-dsii>

Outcome: The meeting reviewed a conceptual framework for work in this area, core issues for data sets and their identifiers, and several current data set identifier systems. The group briefly discussed interoperability mechanisms. The chairs concluded by thanking the group for the discussion and asked folks to continue the discussion on the mailing list.



Red-tailed Hawk
(*Buteo jamaicensis*)



Canada Goose
(*Branta canadensis*)

RFC Format (rfcform)

Description: The BoF reviewed the current requirements for RFC formatting, and then discussed the RFC format proposals that have been published as Internet-Drafts. For a summary of the more contentious issues relating to RFC format, see <http://www.rfc-editor.org/rse/wiki/doku.php?id=formatsummary>.

Proceedings: <http://www.ietf.org/proceedings/84/minutes/minutes-84-rfcform>

Outcome: The meeting recapped the progress of the discussion to date and reviewed the current proposals.

RTP Media Congestion Avoidance Techniques (RMCAT)

Description: Delivery of interactive, real-time media over the Internet is often in the form of sets of media flows using RTP over UDP. There is no generally accepted congestion-control mechanism for this kind of data flow. With the deployment of applications using the RTCWEB protocol suite, the number of such flows is likely to increase, especially nonfixed-rate flows, such as video and adaptive audio. There is therefore some urgency in specifying one or more congestion-control mechanisms that will meet general acceptance. This was a WG-forming BoF meeting to discuss chartering a WG to initially, and amongst other things, find or develop candidate congestion control algorithms, verify that these can be tested on the Internet without significant risk, and publish one or more as Experimental RFCs.

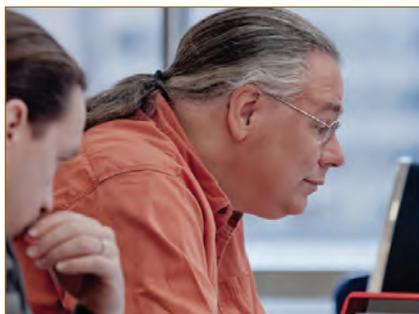
With the deployment of applications using the RTCWEB protocol suite, the number of such flows is likely to increase, especially nonfixed-rate flows, such as video and adaptive audio. There is therefore some urgency in specifying one or more congestion-control mechanisms that will meet general acceptance.

Proceedings: <http://www.ietf.org/proceedings/84/minutes/minutes-84-rmcat>

Outcome: The chairs wrapped-up the discussion by asking three questions:

- Do you think that the problem is clear, well scoped, solvable, and worth solving?
There was a hum in favour.
- Do you support forming a WG with the charter outlined?
There was a strong hum in favour.
- Would you be willing to work on one or more of the drafts outlined?
There was a significant constituency of people willing to work on the drafts.

The proposed charter will be updated based on the discussion and circulated on the mailing list for review. The chairs and the area directors will then move to form a working group. 🍷



IETF 84 At-A-Glance

Registered attendees: 1174

Newcomers: 195

Number of countries: 52

IETF Activity since IETF 83 (March–July 2012)

New WGs: 6

WGs closed: 4

WG currently chartered: 117

New or revised Internet-Drafts: 1607

IETF Last Calls: 102

Internet-Drafts approved for publication: 106

RFCs published: 140

- 81 Standards Track and 4 BCP
- 41 Informational and 11 Experimental

IANA Activity since IETF 82 (March–July 2012)

Processed 1366 IETF-related requests, including:

Reviewed 115 I-Ds in Last Call and reviewed 123 I-Ds in IETF Evaluation

Reviewed 105 I-Ds prior to becoming RFCs, 62 of the 105 contained actions for IANA

Processing goal average for IETF-related requests: 87%

Projects and Deliverables

- Phase 1 of integration of tools testing complete, Phase 2 in planning
- XMLization of registries 86% complete
- Added designated-expert names to individual registries
- Implementing improvements to multicast address and media types request forms

RFC Editor Activity since IETF 83 (March–July 2012)

Published RFCs: 150

Internet-Drafts submitted for publication: 142

- 96 IETF WGs

Blue Sheets

New policy implemented

Bits-N-Bites Debut

Modelled after attendance-optional NANOG Beer 'n Gear

Further experiments at IETF 85 planned

Modern Global Standards Paradigm

Affirmed by Bernard Aboba and Russ Housley

Draft available at <http://www.ietf.org/proceedings/84/slides/slides-84-iesg-opsplenary-4.pdf>

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:

<http://www.internetsociety.org/ietfjournal>

IETF Meeting Calendar

IETF 85

4–9 November 2012
 Host: North American Cable Industry
 Location: Atlanta, GA, USA

IETF 87

28 July–2 August 2013
 Host: TBD
 Location: Berlin, Germany

IETF 86

10–15 March 2013
 Hosts: NBC Universal, Comcast
 Location: Orlando, FL, USA

IETF 88

3–8 November 2013
 Host: TBD
 Location: Vancouver, BC, CA

For more information about past and upcoming

IETF Meetings

<http://www.ietf.org/meeting/>

Special thanks to



for hosting IETF 84

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



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

Bernard Aboba

Leslie Daigle

Mat Ford

Russ Housley

Lucy Lynch

Greg Wood

Editor's Note

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

Unless otherwise noted, photos are
the property of the Internet Society.

Find us on the Web

www.internet-society.org/ietfjournal

ietfjournal@isoc.org

Email

IETF Journal
IETF 84 • Volume 8, Issue 2 • October 2012



IETF Journal

Internet Society
Galerie Jean-Malbuisson 15
1204 Geneva, Switzerland