Volume 10, Issue 2 • November 2014

# the IETF® Journal

Inside this Issue



*A report from IETF 90, July 2014, Toronto, Canada. Published by the Internet Society in cooperation with the Internet Engineering Task Force**



## From the Editor's Desk

***By Mat Ford***

The beautiful Canadian city of Toronto was the venue for the 90th meeting of the Internet Engineering Task Force hosted by Ericsson. As always, *The IETF Journal* presents an interesting sample of the events, discussions, and people that contributed to another great IETF meeting.

Our cover article discusses the emerging field of autonomic networking—a potentially important step forward for plug-and-play networking. We also have an article about the Internet-of-Things Plugfest that took place during IETF 90, information about the Stack Evolution in a Middlebox Internet (SEMI) workshop coming in January 2015, and an article concerning routing on content locators that describes the technology demonstrated during the IETF 90 Bits-n-Bites event.

We also celebrate the most recent winner of the Applied Networking Research Prize, and document an Internet Society panel event on the prospects for Internet security and privacy.

Our regular columns from the IETF, IAB, and IRTF chairs, and coverage of hot topics discussed during the plenary meetings wrap up this issue. For more details of the Internet Area of the IETF in particular, a Working Group summary report is available at http://wiki.tools.ietf.org/area/int/trac/wiki/IETF90.

We are hugely grateful to all of our contributors. Please send comments and suggestions for contributions to ietfjournal@isoc.org. You can subscribe to hardcopy or email editions at www.internetsociety.org/publications/ietf-journal/ietf-journal-subscription.

## Autonomic Networking

***By Brian Carpenter***

IETF 90 included a well-attended Birds of a Feather (BoF) meeting entitled, "Use Cases for Autonomic Networking (UCAN)." The associated mailing list is called ANIMA, which stands for Autonomic Networking Integrated Model and Approach. So what exactly is autonomic networking?

The dictionary recursively defines *autonomic* as "relating to, affecting, or controlled by the autonomic nervous system," which doesn't help much. The autonomic nervous system is an important aspect of an animal's body—it takes care of vital functions, such as breathing and swallowing,

*Continued on page 4*

# Message from the IETF Chair


Jari Arkko, IETF Chair

***By Jari Arkko***

IETF 90 was a great meeting with a good turnout—1,231 people on site from 54 countries. My personal highlights for the week include discussions around the Internet of Things, Internet security and privacy, and the transition of Internet Assigned Numbers Authority (IANA) oversight.

I would like to thank all of our participants, on site and remote, including that participant who sent a robot to attend on his behalf! I would also like to thank our sponsors, specifically Ericsson, our host.

## The Internet of Things

This is a big topic for the IETF—we seem to add more work at every meeting! Three new items arose this time.

- The low-power and lossy networking Plugfest, where participants tested their implementations against each other. These kinds of tests are a big part of the IETF mode of operation. While formally outside the meeting, implementers often gather at the IETF meeting to run such tests.
- The ACE working group (WG), which is focusing on how to bootstrap security and authorisation in a network of smart objects.
- The Bits-n-Bites event, which debuted new format and a focus topic. This time, ten organisations demonstrated Internet of Things solutions to a large audience of interested participants. We will continue the Bits-and-Bites event series at future IETF meetings—please propose focus topics that you would like to see.

One of the difficult tradeoffs discussed this meeting was how increased use of encryption affects caching and other network functions.

## Security and Privacy

Earlier this year we concluded that the IETF needs to do its part to ensure that Internet technology provides better tools against mass surveillance activities. Improving the security of the Internet is no easy task, but we are working hard on several fronts, including updating the Transport Layer Security (TLS) and Hypertext Transfer Protocol (HTTP) protocols (see TLS and HTTPBIS working group efforts).

One of the difficult tradeoffs discussed this meeting was how increased use of encryption affects caching and other network functions. Although this continues to be a challenge, it is clear that HTTPS remains as an end-to-end security solution. Various caching and secure tunneling solutions may arise for other traffic, however.

*Continued on page 5*



**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 Russ Housley***


Russ Housley, IAB Chair

## Highlights from the IAB Retreat

For three days in May 2014, in Cancun, Mexico, the IAB met in conjunction with the Latin America and Caribbean Network Information Centre (LACNIC). One day of the three was spent meeting jointly with the Internet Engineering Steering Group (IESG); we also reached out to the LACNIC community—an effort that will hopefully bring more engineers from Latin America and the Caribbean into the IETF community.

The IAB restructured its programmes, thereby enabling the IAB to focus on those topics most relevant to today's Internet.

The IAB restructured its programs, thereby enabling the IAB to focus on those topics most relevant to today's Internet. Specifically, as a reaction to revelations about pervasive monitoring, the Privacy and Security Programme represents a significant effort for the IAB in 2014.

Following are the current IAB programmes:

- Emergency Services
- IANA Strategy
- Internationalization
- IP Stack Evolution
- ITU–T Coordination
- Liaison Oversight
- Name Resolution
- Privacy and Security Programme
- RFC Editor (includes RSOC)

More information on each of these programmes is available at www.iab.org/activities/programs/.

## IAB Encourages NIST to Use Open and Transparent Processes

The IAB sent comments to the US National Institute for Standards and Technology (NIST) encouraging transparency and openness in their standards development processes. These comments are particularly relevant as it seems that at least one of the NIST security standards provides unexpected access. The entire comment can be seen at www.iab.org/wp-content/IAB-uploads/2014/04/IAB-NIST7977-20140407.pdf.

*Continued on page 6*

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.

*Autonomic Networking, continued*

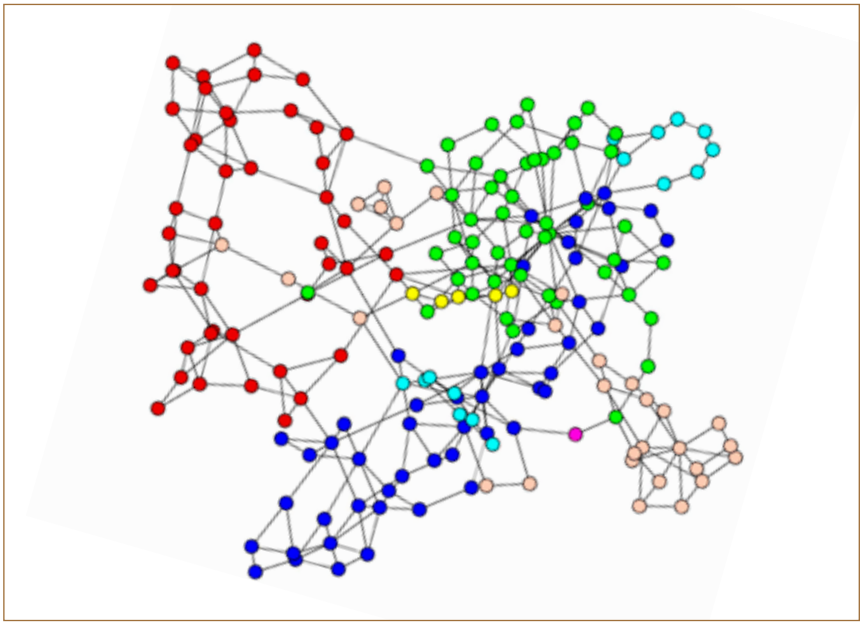

One way to look at autonomic networking is "plug and play for the ISP" or "plug and play for the enterprise network."

without conscious control. This is what inspired the concept of autonomic computing, which was introduced by IBM in 2001 with the intention of making computing systems as self-managing as possible. Wikipedia explains it in more detail: "Autonomic computing refers to the self-managing characteristics of distributed computing resources, adapting to unpredictable changes while hiding intrinsic complexity from operators and users." Autonomic networking, which has been an active area of research in recent years, refers to the application of such ideas to networks. One forum for these discussions has been the Network Management Research Group of the Internet Research Task Force (IRTF).

## Autonomic Networking (AN)

One way to look at autonomic networking is "plug and play for the ISP"or "plug and play for the enterprise network." This is a step forward from the original concept of plug and play for home networks, which has long been recognised as a vital requirement (see, for example, the work of the IETF HOMENET working group).

The goal of self-management includes self-configuration, self-optimization, self-healing, and self-protection. AN puts operational intelligence into algorithms at the node level in order to minimize dependency on human administrators and central management. Nodes that participate in AN discover information about the surrounding network and negotiate parameter settings with their neighbours and other nodes. Ideally, autonomic nodes use stable, closed-loop control methods to achieve self-management, instead of using more traditional top-down network configuration and monitoring tools to set and verify their parameters. Nodes may also have learning and cognitive capability, including the ability to self-adapt decision-making processes based on information and knowledge sensed from their environment. In the most sophisticated cases, advanced data analytics may be part of the input to the autonomic mechanisms.

[Autonomic networking] puts operational intelligence into algorithms at the node level in order to minimize dependency on human administrators and central management.

## More than Science Fiction

Many aspects of small networks have been self-configuring for years, including unmanaged home and small office networks. And numerous existing protocols have a flavour of autonomic properties (e.g., the spanning-tree algorithm needs no manual configuration in order to operate, and some routing protocols require very little configuration). Recently, prototypes and initial products of explicitly autonomic protocols have emerged from some of the major networking equipment vendors. However, it is clearly necessary to have some basic standards in place if AN is to become relevant to large multi-vendor networks.

## Why Now

The main motivation is not new: large network operators, both Internet service providers (ISPs) and enterprises, have been increasingly suffering from the problems and difficulties caused by the central configuration of hundreds or thousands of network elements. Now, after years of research and discussion, ideas about how to achieve autonomic networking are becoming concrete. Fortunately, it is now also economic to provide enough computing power and memory in network elements to support AN. The time is therefore ripe for a standardisation effort.

## Network Parameters

A number of use cases for large networks were proposed at the UCAN

BoF: network address and prefix management, optimisation of mobile backhaul links, risk-aware routing, and detection of service-level agreement (SLA) violations. Other examples are starting to emerge as well, including monitoring and reporting, and others will certainly follow. Two very fundamental aspects of AN can be viewed as use cases in themselves: securely bootstrapping new devices, and creating a secure autonomic control plane for use by specific AN applications.

### Keeping Control

While it is obviously desirable to reduce the need for tedious human interventions, it is essential that network managers can ensure that the network does what is needed and remains fully secure—even if many nodes are configuring and managing themselves. For this reason, the model for AN must include a mechanism for communicating the *intent* of human managers to all self-managing nodes for matters such as resource control, service requirements, and security policy. At the same time, in real networks, AN mechanisms will need to coexist with traditional top-down management and monitoring tools for many years, so it must be possible to introduce AN technology in small steps.

[I]t is essential that network managers can ensure that the network does what is needed and remains fully secure—even if many nodes are configuring and managing themselves.

### Next Steps

At the time of this writing, an ANIMA working group is under discussion in the IETF. A complete solution for autonomic networking would be a very ambitious goal. The scope of the proposed effort is much more modest: define a minimum set of specific reusable infrastructure components to support autonomic interactions between devices, and specify the application of these components to one or two elementary use cases of general value. The main goal is therefore to develop common infrastructure components for distributed functions. The infrastructure should be capable of providing the following services to those distributed functions:

- A common way to identify nodes
- A common security model
- A discovery mechanism
- A negotiation mechanism to enable closed-loop interactions
- A secure and logically separated communications channel
- A consistent autonomic management model

Some important topics are intentionally not included in these initial goals as they are considered separate matters that should be considered later:

- Mechanism for distributing policy intent to autonomic nodes
- Use of data analytics by autonomic nodes
- Other external information sources
- Systemwide integration of autonomics

### Further Reading

draft-irtf-nmrg-autonomic-network-definitions and draft-irtf-nmrg-an-gap-analysis

Mailing list: anima@ietf.org

### Acknowledgements

Several IRTF and IETF drafts were raided for text and ideas, and useful comments on this article were made by Leo Dorrendorf, Sheng Jiang, and Alexandre Petrescu.

*Message from the IETF Chair, continued*

The new TCPINC working group had it first meeting on developing a new layer of opportunistic security. This is mainly for applications, such as the Web, that don't use current transport layer security.

### IANA

The IETF has been actively discussing the transition since the announcement from the US government in March. I am happy about it, but we at the IETF also see it as a part of longer-term evolution already in place with regards to how we deal with the oversight of IANA. In the last 15 years, we have developed the contracts, oversight mechanisms, and processes on which our part of IANA runs.

Our meeting confirmed that the IETF community believes these mechanisms are sufficient going forward. In the coming weeks and months, we will document how these mechanisms address the oversight requirements. I feel optimistic about the process. A few weeks after the meeting, we created the IANAPLAN WG as a forum to discuss the topic. (See "IANA Transition Update," on page 22 for more on this topic.)

### Next Up

IETF 91 is slated for 9–14 November in Honolulu, Hawaii. I would like to welcome everyone to the meeting!

Between meetings, the work of the IETF runs on mailing lists. What can we expect in the coming months? The major projects, such as WebRTC, HTTP 2.0, and so on will continue. Key milestones ahead include the publication of the final HTTP 2.0 RFC (this year), as well as the conclusion of our part in the IANA transition work (planned for completion within 2014). Please visit our newcomers page if you would like to join us in this important work.

*Words from the IAB Chair, continued*

### IANA Stewardship Transition

The IAB sent comments to the Internet Corporation for Assigned Names and Numbers (ICANN) on the principles, mechanisms, and process to develop a transition plan for the stewardship of Internet Assigned Numbers Authority (IANA) from the US National Telecommunication and Information Administration (NTIA) to the multistakeholder community. NTIA asked ICANN to facilitate a process to create a proposal, and the IAB provided these comments on the draft plan. The entire comment can be seen at www.iab.org/wp-content/IAB-uploads/2014/04/iab-response-to-20140408-20140428a.pdf.

During IETF 90, participants of the IANAPLAN Birds of a Feather (BoF) set the direction for the creation of an IETF community plan for the IANA protocol parameters. The community decided that the IAB IANA Evolution Programme would take the pen, and that a new working group in the General Area would provide review and comment. Once the working group reaches consensus, there will be an IETF-wide last call. Then, the resulting plan will be sent to the IANA Stewardship Transition Coordination Group (ICG) to be combined with the plans from the names and numbers communities. Finally, the ICG will deliver the combined plan to NTIA after the Internet community has reviewed it.

The community decided that the IAB IANA Evolution Programme would take the pen, and that a new WG in the General Area would provide review and comment.

The ICG comprises 30 members and two liaisons from a vast number of Internet-related organizations.

- The IAB appointed Russ Housley and Lynn St.Amour to the ICG.
- The IESG appointed Alissa Cooper and Jari Arkko to the ICG.
- The ICG selected Alissa Cooper to be its chair.

The ICG membership is listed at www.icann.org/resources/pages/coordination-group-2014-06-17-en.

### Highlights since IETF 89

- The IAB appointed Sarah Banks and Robert Sparks to the RFC Series Oversight Committee (RSOC).
- The IAB appointed Sean Turner to the Internet Society Board of Trustees.
- The IAB appointed Matt Miller as liaison manager to ECMA TC39.
- The IAB published RFC 7288 on "Reflections on Host Firewalls."
- The IAB published RFC 7295 on "Report from the IAB/IRTF Workshop on Congestion Control."
- The IAB published RFC 7305 on "Report from the IAB Workshop on Internet Technology Adoption and Transition (ITAT)."
- The IAB published RFC 7241 on "The IEEE 802/IETF Relationship."

## OVERCOMING THE OBSTACLES TO INTERNET EVOLUTION: ANNOUNCING THE SEMI WORKSHOP

One of the challenges faced by technologists interested in developing new and better ways of transferring data over the Internet is the growing number of devices in the network that cause connections to fail when they see something new or unexpected. These so-called middleboxes also make it hard to add simple refinements to existing ways of delivering Internet application data. It's as if the information superhighway now has automated toll gates at all the major intersections, and if they don't find your registration plate where they expect to find it, you won't be allowed through.

In their revealing study of the problem (http://conferences.sigcomm.org/imc/2011/docs/p181.pdf), Michio Honda and his coauthors found that, "at least 25% of paths interfered with [data] in some way beyond basic firewalling."

In an effort to promote discussion of ways to detect and characterise the behaviour of these middleboxes, and to identify ways to work around the obstacles they present, the Internet Architecture Board (IAB) is organising a workshop on Stack Evolution in a Middlebox Internet (SEMI) in January 2015, in Zürich, Switzerland.

Attendance at the workshop is by invitation only. Prospective participants were invited to submit short position papers outlining their views on one or more topics related to the scope of the workshop. To learn more about the purpose and scope of the workshop you can download the call for papers at www.iab.org/wp-content/IAB-uploads/2014/09/stackevo-cfp.pdf. Position papers will be published on the IAB website (www.iab.org/activities/workshops/semi/) in due course.

The SEMI workshop is sponsored by the Internet Architecture Board, the Internet Society, and ETH Zürich. Mirja Kühlewind and Brian Trammell of ETH Zürich are cochairs.

# Experts Say Economics and Politics Hamper Efficient Routing of Internet Data

***By Carolyn Duffy Marsan***

The Internet engineering community faces many economic and political barriers to creating a more efficient routing system, according to an IETF 90 technical plenary session about how Internet topology interacts with geography.

"There are, of course, people and governments who are keen to link intranetwork traffic to geography, to geopolitical boundaries," said panel moderator Andrew Sullivan, an IAB member and director of architecture at Dyn. "But there are also real issues of geography that affect the way the network operates."

The first presenter was Antonio Gamba-Bari, a PhD student at the University of Toronto's Faculty of Information and a member of its IXmaps project. IXmaps is an Internet mapping tool that enables end users to see how their personal data travels across the Internet and identify privacy risks. Under development since 2009, the IXmaps tool has gained prominence by offering transparency into Internet traffic surveillance by the US National Security Agency (NSA).

"We encourage people from disparate geographic locations and ISPs to install and run our traceroute program to feed our database with traceroutes through hostname parsing, latency comparison and topological analysis," Gamba-Bari said. "We geolocate the intermediate routers for mapping the routes their packets take. We highlight the exchange points where these routes pass through suspected sites of NSA interception."

IXmaps has gathered more than 30,000 traceroutes from more than 250 contributors and tracking data from more than 2,500 URLs.

IXmaps wants to dispel the notion that the Internet is a "cloud," and instead demonstrate that it consists of a few Internet exchange points that route a massive amount of traffic. In the United States, for example, almost all Internet traffic passes through switching centers in 18 cities. Further, traffic that begins and ends in Canada often travels through the United States, a phenomenon IXmaps calls boomerang routing.

> "There are people and governments who are keen to link intranetwork traffic to geography, to geopolitical boundaries. But there are also real issues of geography that affect the way the network operates."
>
> —Andrew Sullivan, Moderator

What worries IXmaps researchers is the fact that the NSA has a wide-ranging surveillance system that intercepts, copies, analyzes and stores all Internet traffic on US networks.

"Our work has shown that NSA interception in just 18 US cities can capture nearly 100 percent of US domestic traffic," Gamba-Bari said. "Foreign traffic that transits the United States is also very likely to be intercepted. From our data, we estimate that 25 percent of domestic Canadian traffic is routed via the United States and, hence, subject to NSA surveillance."

University of Toronto researchers are transforming IXmaps from a prototype into a more widely usable Internet mapping and policy analysis tool, thanks to a grant from the Canadian Internet Registration Authority (CIRA). The goal is for IXmaps to become more reliable and flexible as well as to improve the accuracy of the geolocation component. IXmaps also hopes to expand beyond North America.

"We welcome offers of help internationalizing IXmaps and making it more sustainable," Gamba-Bari said. "We will put it under a free Open Source software license to make it easier for others to take it in their own directions."

Andrew Sullivan, panel moderator, IAB member, and director of architecture at Dyn

Next, the Internet Society's director of development strategy, Jane Coffin, spoke about the group's effort to build local infrastructure, which includes Internet exchange points (IXPs) as well as the human, technical and governance infrastructure around them.

An IXP "is a physical location where different IP networks meet to exchange traffic and to keep local traffic local," Coffin said, adding that it is more than boxes and wires. "Ninety-five percent of this is the human engineering—how we bring the different character sets

*Continued on next page*

*Experts Say Economics and Politics Hamper Efficient Routing of Internet Data, continued*

together of ISPs, network operators, and research and education networks."

By building IXPs in far-flung locations, ISOC is creating local communities of interest. "When you improve the quality of Internet services," Coffin said, "you drive up demand. Latency comes down, and quality of service usually goes up."

In addition, IXPs attract content development. "Content is generated by businesses that have confidence in those infrastructures," she said. "We know this is a catalyst for overall Internet development from our experience and what we've seen."

Jane Coffin, panelist and Internet Society director of development strategy

"When you improve the quality of Internet services, you drive up demand. Latency comes down, and quality of service usually goes up."

—Jane Coffin, Panelist

For example, a new IXP in Kenya resulted in latency reductions from 200 to 600 milliseconds down to a range of 2 to 10 milliseconds. Not only did end users in Kenya see better Internet performance from the new exchange point, but there were cost savings of $1.5 million per year on international transit for local mobile operators. Further, the new IXP facilitated e-government services, with the Kenya tax authority peering there.

Coffin said similar improvements in Internet performance and economics are occurring in Argentina and Brazil. "In Ecuador, before the IXP went in, international transit was $100 per megabit per second," Coffin said. "It is now local traffic that is $1 per megabit per second."

Coffin said developing countries can use these new IXPs to deploy emerging technology, such as public key encryption, IPv6 and top-level domains. "After a content delivery network [CDN] cache was installed in Quito in 2009, traffic went up by 700 percent. This is local traffic," she added.

Coffin said there are more than 350 IXPs around the world. ISOC is not only building new IXPs around the globe, but helping grow IXPs that already exist. ISOC provides equipment, technical assistance, and economic guidance, and cooperates with local government.

She outlined two ongoing projects in the Africa Union, called Axis I and II. So far, they've held 30 best practices workshops and launched four IXPs this year. In addition, they've sponsored five regional meetings across Africa to bring regulators, policymakers, and network operators together to discuss the importance of growing the regional Internet infrastructure.

"There are so many landlocked countries in Africa that it's important for some of those government entities to work together," Coffin said. "There was one instance in Zimbabwe where it took almost two months to string some fiber about a hundred meters, due to the fact that it was over a historic bridge."

Coffin emphasized that IXPs are designed to keep Internet traffic local for better performance and lower costs, not to be centralized locations for government surveillance.

Coffin said ISOC is working with the Regional Internet Registries such as LACNIC in Latin America, as well as individual country network information centers like NIC.br in Brazil. Large corporations are supporting the work by providing grants and equipment.

Coffin emphasized that IXPs are designed to keep Internet traffic local for better performance and lower costs, not to be centralized locations for government surveillance. "It is not set up to be a monitoring facility for deep packet inspection. Or at least that's our philosophy," she said.

Finally, Amogh Dhamdhere described the data regarding network topology and geography available from the Center for Applied Internet Data Analysis (CAIDA) at the University of California, San Diego. CAIDA operates a network measurement infrastructure called Archipelago that consists of 102 monitors, which collect data about IPv4 and IPv6 traffic in 39 countries. Archipelago collects traceroutes from the entire routed IPv4 and IPv6 space, as well as alias resolution measurements for router-level topologies and measurements of interdomain congestion.

Dhamdhere, a researcher at CAIDA, said the Archipelago infrastructure has collected 6 terabytes of compressed data since 2007, all of which is available to network research and operators. CAIDA provides the raw traceroutes in their original form, curated topology data sets and asynchronous number topologies for IPv4 and IPv6.

The IETF 90 technical plenary concluded with a lively Q&A session for audience members.

"One of the goals of a currently funded project that we're working on is to make it easier for researchers and people interested in this kind of analysis to actually access this data and do interesting things with it," Dhamdhere said. "We're building support for rich queries on this traceroute data, and the idea is to put them together with other kinds of data such as geolocation, annotated AS-level topologies, and router-level topologies."

Eventually, CAIDA wants to provide data that can be used for regional analysis such as measuring how many routes for Canada-to-Canada communications exited Canada and traverse through U.S. networking hubs.

"Suppose we predicted that a certain region was going to be affected in the sense of a natural disaster like a hurricane or a storm coming up or political instability. We'd like to know all the paths from our current monitors that traverse that region," Dhamdhere said. "These paths might be rerouted or might even go down when something actually happens."

He said CAIDA is looking for volunteers to host additional Archipelago monitors—Raspberry Pi computers that cost only $35 each. People with Archipelago monitors can take advantage of an interactive topology-on-demand service called Vela that visualizes traceroutes on a map. Another CAIDA service is a DNS-based geolocation service, which gives hints about the geographic location of a domain. Finally, CAIDA offers a repository of tools and data regarding autonomous systems (AS), including a map that uses geolocation data to infer where an AS has deployed infrastructure.

CAIDA is looking for volunteers to host additional Archipelago monitors—Raspberry Pi computers that cost only $35 each. People with Archipelago monitors can take advantage of an interactive topology-on-demand service called Vela that visualizes traceroutes on a map.

"We have an interface where operators can go in and enter corrections to the inferences we've made," Dhamdhere added.

One work in progress at CAIDA is data about which networks peer at IXPs.

"We're trying to expand the set of Internet exchange points from which we can actually infer reliably the set of connected networks," Dhamdhere said. "We've recently done some work on mining historical peering data… to figure out colocation by different networks at IXPs, what kind of peering policies they advertise, how all of this evolves over time, and we can actually find interesting things like geographical expansion of networks just by looking at historical peering data."

CAIDA has used its data to analyze country-level Internet blackouts and outages such as those that happened in the Arab Spring as well as the impact of natural disasters such as earthquakes and hurricanes. CAIDA is "trying to develop metrics and tools to automatically detect outages of this type," Dhamdhere said.

He said most of CAIDA's data, research, and tools are available online to the IETF community. "If you'd like to collaborate on anything or just get access to the data, then we'd love to hear from you," he added.

Sullivan asked the panelists why network operators don't spend money to build IXPs given that it is more efficient to keep traffic local.

Coffin said that in countries such as Chad, the network operator can't afford to build new infrastructure. In other countries, such as Cote d'Ivoire, the incumbent doesn't have an economic incentive to build a community of interest around an IXP. She said ISOC has to explain to governments the importance of IXPs and that the benefits are not always obvious. Being able to donate equipment and staff to install it and train others to run it helps get IXPs built, she added.

With regard to boomerang routing, Gamba-Bari said it may occur because network infrastructure doesn't exist to support a more direct route. For example, traffic from Halifax to Vancouver must travel through the United States. In other cases, inefficient data paths result from networks deciding to peer with some networks and not others. That's why traffic from one Toronto

*Continued on next page*

*Experts Say Economics and Politics Hamper Efficient Routing of Internet Data, continued*

Amogh Dhamdhere, panelist and researcher Center for Applied Internet Data Analysis at the University of California, San Diego

building to another might end up going through the United States.

A commenter from the audience, Jacques Latour of CIRA, pointed out that Canada's network incumbent doesn't want new IXPs built or to peer with local ISPs because it cuts into their revenue stream. He said all of the new IXPs in Canada are bringing in tier one carriers from other countries that are competing with the Canadian incumbent provider and driving down prices, which is good for consumers. He said CIRA is helping set up new IXPs in every province.

"The core of the Internet is the IXP," Latour said. "That's where you generate bandwidth. That's where content providers go. This is where people can get high volume of data for low cost."

Coffin added that it takes years to build a community of interest around new IXPs in developing countries, to explain the economic advantages, and address worries about surveillance being installed in these locations. In Trinidad, it took seven years; Bolivia took three.

"It is very difficult to grow communities of interest," Coffin said. "They are not mushrooms—you don't just sprinkle some water and they come up. It takes a lot of time and energy."

# Internet Society Panel Explores Internet Security, Privacy Outlook

## New challenges, emerging technologies will influence this ongoing tussle over the next decade

***By Carolyn Duffy Marsan***

No technological magic bullet is on the horizon to solve the Internet's security and privacy challenges during the next 10 years, according to a panel discussion sponsored by the Internet Society (ISOC) that was held in conjunction with IETF 90 in Toronto.

Moderator Andrei Robachevsky, technology programme manager at ISOC, noted that the Internet engineering community lacks a good understanding of the security and privacy qualities of the Internet as a whole.

"Some of the fundamental elements have known vulnerabilities. Take for instance, BGP [Border Gateway Protocol] and TLS [Transport Layer Security]. While fixes are underway, they are far from being widely deployed," Robachevsky said. "At the same time, if you look at the Internet, so far it has proven to be very resilient. What holds the Internet together? Is it technology? Is it people? Is it money?"

Robachevsky asked four experts to identify key issues that will shape Internet security and privacy during the next decade.

Lucy Lynch, director of Trust and Identity Initiatives at ISOC, said the main challenge for improving the Internet's security and privacy is scale.

"I think we have in the security domain and in the privacy domain some of the tools we need. Sometimes they work well together, and sometimes they conflict," Lynch said. "What we don't have is a systems view of how you compose those elements at scale… Getting a systems point of view with our current elements that allows us to operate at scale is the end goal 10 years from now."

Moderator Andrei Robachevsky noted that the Internet engineering community lacks a good understanding of the security and privacy qualities of the Internet as a whole.


Andrei Robachevsky, moderator and technology programme manager at the Internet Society (right)

Wendy Seltzer, panelist and policy counsel and Technology & Society Domain lead at the World Wide Web Consortium (center)

> "As we design technology and as we build on it, we need to think of the interfaces for social controls and legal and regulatory controls to make sure the systems we are building have the properties of protecting users."
>
> —Wendy Seltzer, Panelist

Cisco Fellow Dave Oran said the Internet increasingly reflects all the problems of the physical world, including conflicts, politics, money, and criminality.

"Our challenge looking out 10 years is can we from a technology, policy, and overall citizens-of-the-world perspective use the Internet to actually improve the world as a whole," Oran said. "That's a very difficult job, but our leverage is higher than it's ever been. That's one reason I think looking forward to what security technology and what the security environment could be will be critically important."

Wendy Seltzer, policy counsel and Technology & Society Domain lead at the World Wide Web Consortium, said technological solutions alone can't fix the Internet's privacy and security problems.

"As we design technology and as we build on it, we need to think of the interfaces for social controls and legal and regulatory controls to make sure the systems we are building have the properties of protecting users," Seltzer said. "Some of that will be designing usability into the system so that end users—whether technical or less technical—can understand the choices we are asking them to make and can do appropriate risk analysis."

Danny McPherson, senior vice president and chief security officer for Verisign, said additional security systems, such as badge readers and travel itineraries, will help protect the network infrastructure; but these systems also create more data, which could be abused from a privacy perspective. He pointed out that once an IP address or domain name has been accused of being a security threat, there is no process for rehabilitating it.

"As there are more indicators of compromise and more intrusion sets and other things that people use to protect systems, one concern is the scorched earth notion," McPherson explained. "Most of what is shared in security is a number space or name space and maybe some behavioral aspect of the host that appears to be malicious activity. It's interesting that for a namespace, it's hard to get that reputation back. We take a scorched earth approach. If I pick up a domain name or an IP address, how usable are those and how much residue is left from previous activities?"

McPherson expanded on that notion by explaining his concerns about information sharing among security-related groups. Today, there is no process for removing misinformation from the data these groups share, he said.

"Closed security groups that are trying to protect against some aspect of attack or a botnet or deal with cybercrime… are very effective," he said. "But there is not a lot of provenance of where did this data come from, how we get the information back, how we get our reputation back, or what happens if misinformation is injected into this information. In 10 years, we will have a lot of scorched earth in numbers and name spaces. It's going to be hard to find situations where a number doesn't break in some [intrusion detection system] or some sensor or some [intrusion prevention system] or is blacklisted somewhere."

Robachevsky asked the panelists to identify emerging privacy or security approaches that may prove to be revolutionary.

Danny Mcpherson, panelist and senior vice president and chief security officer at Verisign

Oran outlined the pros and cons of an alternative architecture known as information-centric networking (ICN) that he has been working on for three years. While the Internet focuses on securing transport channels through protocols, such as TLS and IPsec, ICN doesn't worry about channels and instead secures content with built-in encryption.

"There are a couple of benefits to this approach. It's simpler to understand the chain of custody of the content, and it allows you to protect the content at rest in the same way as it was protected

*Continued on next page*

*Internet Society Panel Explores Internet Security, Privacy Outlook, continued*

Dave Oran, panelist and Cisco Fellow (center)

while in the communication system," Oran explained. "These systems were designed with integrity and provenance built in."

Oran pointed out that while the source of data in an ICN architecture is anonymous, the name of the content is public. "You're trading off consumer anonymity for content anonymity," he said. "It's not clear if that's the right tradeoff."

While the ICN approach eliminates many types of attacks, it still leaves the Internet open to Distributed Denial of Service attacks. It's unclear how Internet business models would evolve in an ICN architecture, Oran added.

> "There is no magic bullet. The hard problems are still hard. Trust management is an unsolved problem in the ICN world, just like in the IP world."
>
> —Dave Oran, Panelist

"There is no magic bullet," Oran said. "The hard problems are still hard. Trust management is an unsolved problem in the ICN world, just like in the IP world."

Oran mentioned two other promising technologies: functional encryption and privacy-preserving query systems. Functional encryption allows a user to perform computations on encrypted data, which is useful for middleboxes that perform operations on data. But, today's computing technology needs to be orders of magnitude faster in order to make functional encryption practical. Meanwhile, privacy-preserving query systems attempt to improve the data confidentiality of large database systems by conducting limited queries on encrypted data.

"These are just some technologies that may be important in our Internet lives some number of years out," Oran concluded.

Seltzer added that privacy-preserving query systems depend on cooperative protocols where data collectors limit the data that is being shared and participants limit their disclosures.

"Along with these mathematical tools, we will also need the social organizing functions of distributed systems and technologies that build from control by the end user and that enable us as users to exercise some collective action to demand better security and privacy from the systems that we use," she explained.

Lynch said the future will require Internet engineers to keep balancing between security and privacy and that this tussle will not end in the foreseeable future. She said the relationship between data collectors and users is asymmetric today, with data collectors having an enormous amount of power while individual data subjects have little power.

"People don't pay to be part of some of these systems. So they pay with data instead of paying with cash. An economic model that shifts all of that might shift the concern for protecting the individual data subject in a way that's privacy preserving, but it would require the data subject to be willing to put up something, either accurate and anonymized data or cash or something else," Lynch said. "It's more about finding balance points between security, privacy, secrecy, and the public good."

In response to a question from the audience, Oran pointed out that the Internet engineering community is better at security and privacy than it was 10 years ago, and that more improvement can come in the next decade, too.

> [T]he future will require Internet engineers to keep balancing between security and privacy... this tussle will not end in the foreseeable future.

"It's no longer acceptable to design, let alone deploy, a technology without understanding the security properties and consider security as part of the design," Oran said. "I'm somewhat optimistic that we've gone through a phase change. It's unlikely that somebody gets very far with a design before somebody else asks: How secure is it? What is the threat model? What are the vulnerabilities? And how does it change the attack surface?"

**Editor's note:** This article takes a deep dive into some of the new technology developments that were demonstrated during the IETF 90 Bits-n-Bites event in Toronto. It also offers insight into one aspect of the exciting and ground-breaking work encompassed in the Internet Research Task Force's Information Centric Networking Research Group. *See https://irtf.org/icnrg for more on this topic.*

# Information Routing on Content Locators

***By Dimitri Papadimitriou, Sahel Sahhaf, and Wouter Tavernier***

The increase in Internet traffic volume for applications, such as mobile video and cloud computing, has led to various technologies enabling content distribution that rely on caching and replication. As they are proprietary and deployed in silos, these content-distribution technologies do not allow unique and securely identified named information independent of the distribution channel. Moreover, these technologies are usually implemented as an overlay, which leads to needless inefficiency. Two seed models—the overlay model and the name-based routing model—can be identified as the root of most currently proposed approaches to information-centric networking (ICN). The latter aims at enabling data to become independent from their network/physical location, application, and storage support, as well as their transport/distribution channel in support of both user and content mobility.

## The Limits of Name-based Routing

The name of a resource indicates what we seek, an address (locator) indicates where it is. Following this distinction, the two major alternatives currently under investigation consider either routing on content names (with IP addresses keeping the network locator semantic) or introducing another level of indirection (overlay) by extending the IP address semantic to include a location name from where the information can be retrieved.

The first alternative suffers from limited scaling (in terms of memory space) as approaches such as Content-Centric Networking (CCN)[1] are confronted with name spaces that have not been designed to sustain forwarding performance and scaling. Indeed, the routing information corresponding to each content object has to be maintained in the routing table, the number of content objects is very large (between $10^{15}$ and $10^{22}$)[2], and the size of the routing tables becomes a key concern. Assume for instance that routing tables would include one entry per top-level domain, a name-based routing table would have to hold 2 x $10^8$ routes (compared to the 5 x $10^5$ active BGP routes as of mid-2014). Moreover, the resulting size increase of the routing tables and associated processing would worsen over time as the number of domains increases by 10–15 percent per year (as indicated in the Verisign report of April 2013). Finally, the adoption of a new content name space is challenging, but the maintenance of a hierarchical tree-based structure (summarization) is even more difficult. The second alternative raises the same kinds of issues as any network overlay following the aphorism of D. Wheeler that describes the problem of too many levels of indirection: "All problems in computer science can be solved by adding another level of indirection... but that usually will create another problem."

If names become the fundamental structure of the information exchange process, name-based routing becomes a scaling and performance bottleneck. So, unless names get overloaded with locator semantics, there is no means by which the resolution of a name into an address can be avoided.

If names become the fundamental structure of the information exchange process, name-based routing becomes a scaling and performance bottleneck. So, unless names get overloaded with locator semantics, there is no means by which the resolution of a name into an address can be avoided. The implication being that conventional approaches to ICN like CCN or its multiple variants (see "Networking named content"[3] and "Named Data Networking"[4]) shift this operation to a network-layer functionality. On the other hand, this concept exacerbates the memory-scaling limits of stretch-1 shortest path routing already observed in many experimental studies and theoretic investigations.

## Content Locators

With both alternatives, the fundamental problem is as follows: localization and routing refer to distinct functions associated to distinct objects (address vs. route), which cannot be derived from each other using local knowledge. Moreover, the higher the level of information on which the routing decision is performed, the higher the memory

*Continued on next page*

Ideally, the locator space should be as independent as possible from topology changes, while also providing sufficient information to compute distances where it is processed, provided that different receivers compute distances following the same distance function.

cost. On the other hand, lowering the level by providing additional resolution processes increases the communication cost and latency. This observation leads to the consideration that (1) content-based forwarding requires one to keep locators (instead of names) as the information unit for routing/forwarding decisions and, (2) the locator space, say $X$, should also elevate the memory-scaling problem induced by stretch-1 shortest path routing. To this end, the *localization function* performing at $x \in X$ shall, after resolving the content name into the corresponding locator, say $y \in X$, compute the distance $d(x,y)$ and the *routing function* (distributed by nature) shall determine locally and independently for any destination the adjacent node along a loop-free path, such that incoming messages directed to destination y can reach it. Hence, we seek a locator space $X$ that can be processed at end-points by the localization function and at intermediate nodes by the routing function.

It remains to be seen which locator value space would best fit this combined role. An obvious choice would be a so-called topology-dependent label space. However, such a value space is prone to renumbering, even in the case of non-local topological change; hence, it is unsuitable. Another possible choice would be the reuse of the IP address space as existing Internet routing protocols operate using such a locator space. However, IP locators have no associated distance metric; meaning, no distance computation and no selective localization of content are possible when the same content object is available at multiple locations. Ideally, the locator space should be as independent as possible from topology changes, while also providing sufficient information to compute distances where it is processed, provided that different receivers compute distances following the same distance function (we will see later the implications of the model selection when detailing coordinate computation). Instead, by associating data objects with content locators from a (geo)metric space $(X,d)$ where, to each element of the space $X$ corresponds a globally unique geometric coordinate, and $d$ is a distance function (or metric) that verifies "$x,y \in X$, $d(x,y) = d(y,x) > 0$ when $x \neq y$ otherwise $d(x,y) = d(y,x) = 0$ together with the triangular inequality, we satisfy these requirements and prevent renumbering in the case of topology change.

## Geometric Routing on Content Locators

In addition, the selection of this coordinate space must be performed hand-in-hand with the design of a routing scheme capable of addressing the memory space complexity ($O(n \cdot (log(n))$) characterizing stretch-1 routing algorithms such as Border Gateway Protocol (BGP). For this purpose, we introduce in "Geometric information routing"[5] a variant of geometric routing in which distances between nodes are computed from the length of the corresponding geodesic segments drawn out of the hyperbolic plane $H^2$. From a routing-information-distribution perspective, this routing scheme operates following a modified distance-vector algorithm enabling the construction of geodesic segments. These segments are combined by means of procedures similar to pathlet/segment routing. To reduce the number of routing entries, coordinate sets are represented by geometric areas. As geometric coordinates are associated with the locations of content objects, the underlying model is referred to as *geometric information routing*. Coordinates can be computed offline by different procedures: by means of distance preserving graph embedding or a variant of the Vivaldi algorithm in the hyperbolic space represented by the hyperboloid model (also referred to as the Loid model). The distortion introduced by the hyperbolic model is determined by two parameters: the distance function and the curvature, which

[T]he knowledge of the geometric properties of the Internet topology yields a coordinate-computation algorithm of similar computational complexity to the canonical Vivaldi algorithm with Euclidean distance.

represents the amount by which an object in the space deviates from being flat (i.e., Euclidean). Instead of selecting the hyperbolic space curvature that produces the lowest error (thus avoiding the trial-and-error of embedding functions), we make use of the fundamental formula which relates the curvature $\kappa$ of the hyperbolic space to the value $\delta$ of the topology graph: $\delta = \frac{\ln(1+\sqrt{2})}{\sqrt{-\kappa}}$. In other words, the knowledge of the geometric properties of the Internet topology yields a coordinate-computation algorithm of similar computational complexity to the canonical Vivaldi algorithm with

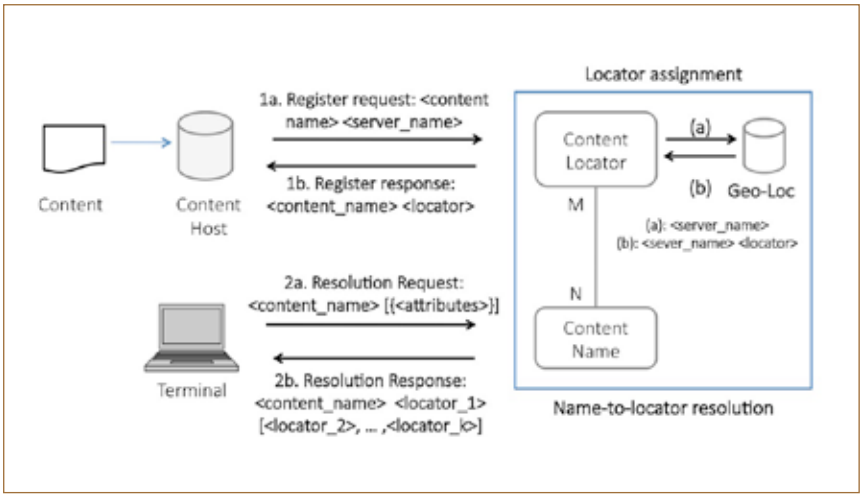


Figure 1. Content Name Registration and Resolution

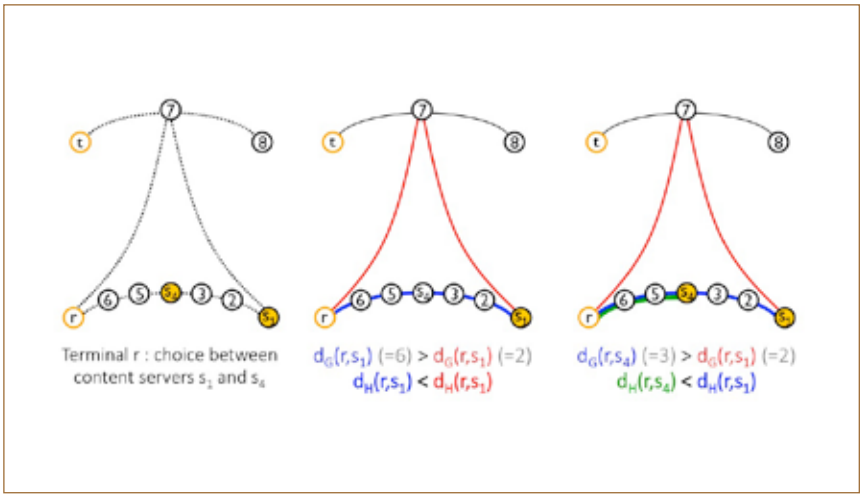


Figure 2. Example for Geometric Information Routing Operation

Euclidean distance. If the coordinate computation procedure could not be unified with the routing-information primitive exchange, a simple extension to the Inverse ARP protocol would enable coordinate allocation.

Compared to the map-based model proposed in "Greedy forwarding in dynamic scale-free networks embedded in hyperbolic metric spaces"[6], the proposed approach does not rely on the construction of a virtual map obtained by applying empirical rules derived from the hidden hyperbolic space underlying the observable Internet router/AS topology, instead it uses the observable topological properties. Compared to Border Gateway Protocol (BGP) with IP geolocation it relies on the exchange of content locators taken out of the IP address space (also known as a push model). However, the IP locator space has no associated distance metric preventing selective localization when the same content object is available at multiple locations.

To illustrate the operation of this information routing model, assume that the requester $r$ associated with coordinate $x$ queries for a given content object and issues a request following the process illustrated in Figure 1. Referring to the middle of Figure 2, the requester receives locator (coordinate) $y_1$ which is associated to the server $s_1$. Since the path following the minimum hop-count distance $d_G(r,s_1)$ (in red) does not correspond to the path with the shortest hyperbolic distance $d_H(r,s_1)$ (in blue), the path-selection process depends on the distance metric. Moreover, as a copy of a given object can be available at multiple servers locations $s_1,s_4$, which is often the case in distributed information systems, the requester may receive a response including multiple locators, i.e., the coordinates set $\{y_1,y_4\}$ (right-hand-side of Figure 2). From this set, it can determine the hyperbolic distance $d_H$. Hence, it can also select the 'nearest' server where a given content object is accessible by computing the minimum distance $d_H$ which is actually different from one obtained if the hop count distance $d_G$ would be used to perform that selection. Indeed, referring to the right-hand-side of Figure 2, one can see that $\min\{d_G(r,s_1),d_G(r,s_4)\} = d_G(r,s_1)$ (in blue) whereas $\min\{d_H(r,s_1),d_H(r,s_4)\} = d_H(r,s_4)$ (in green); thus, the selection of the closest (or nearest) server differs between the metric graph $X_G$ and the topology graph $V(G)$. The reverse operation also applies: by receiving incoming packets that include in their header the coordinate associated to the source

[T]he proposed approach does not rely on the construction of a virtual map obtained by applying empirical rules derived from the hidden hyperbolic space underlying the observable Internet router/AS topology, instead it uses the observable topological properties.

Continued on next page

$x$ (i.e., the locator associated to the requesting terminal $r$), the destination can determine the distance $d(s_1,r)$ without involving additional resolution or translation. This technique avoids requiring discovery of a reverse forwarding path from $y_1$ back to the requester $x$, as is the case when the content object name is used to forward the request towards the content host.

## Conclusion and Next Steps

We propose an alternative routing scheme for ICN in which content names are assigned locators, and geometric routing is performed on those locators. This model—in which routing decisions are performed on content locators—avoids name-to-locator resolution by intermediate nodes, and instead considers that name resolution by end hosts provides them with the information they need to select the destination. The salient feature of this addressing and routing model comes from (1) the property of coordinate-based content locators: these coordinates can be used by the distributed routing function to perform geometric routing decisions, and (2) the fact that as both localization and routing functions are performed on content locators, the result is localization of cached content. Indeed, as there is no distinction between a server and a cache locator—if an intermediate node keeps a local copy of a content object, it can register it as any other object stored on a content server. Subsequently, a requester could receive from the name resolver a (set of) content locator(s) associated to a (set of) intermediate node(s) instead of a server.

To demonstrate the successful operation of the proposed information routing scheme—based on functionality and potential performance gains in memory, bandwidth, and sender-to-receiver delay—we developed new routing and forwarding elements to enable packet processing according to geometric routing operations. We experimentally validated this information-routing model and scheme on the iLab.t virtual wall testbed and compared it with information-centric networking based on BGP with IP geolocation.[7] Against the latter, we gain a factor $n$ (number of nodes) in memory space required to store routing information, and a factor $\sqrt{(n)}$ in the memory space required to store routing tables while limiting the routing path stretch increase to a small additive constant (characterizing the geometric property of the topology). The notion of distance in geometric routing and caching in intermediate nodes affects the capacity utilization and results in more-balanced traffic on the links and a significantly decreased end-to-end delay between terminals and servers. Future work will investigate the possible unification of coordinate computation procedures with routing information exchange primitives. Novel caching strategies exploiting topology properties (such as node degree and centrality) remain for further exploration. More important, investigation of information routing on content locators can also seed new research topics combining addressing and routing algorithms.

Future work will investigate the possible unification of coordinate computation procedures with routing information exchange primitives. Novel caching strategies exploiting topology properties (such as node degree and centrality) remain for further exploration.

## References


1. T. Koponen, M. Chawla, B. G. Chun, A. Ermolinskiy, K. H. Kim, S. Shenker, and I. Stoica, "A data-oriented (and beyond) network architecture," *Proc. of 2007 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications,* ACM SIGCOMM '07, pp.181–192, New York, NY, USA.
2. D. Kutscher (Ed.), "ICN research challenges," Work in progress, February 2014.
3. V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, and R. L. Braynard, "Networking named content," *Proc. of CoNEXT 2009,* pp.1–12, Rome, Italy, December 2009.
4. L. Zhang, A. Afanasyev, J. Burke, V. Jacobson, K. Claffy, P. Crowley, C. Papadopoulos, L. Wang, and B. Zhang, "Named Data Networking," *ACM SIGCOMM Computer Communication Review (CCR),* 44(3):66–73, July 2014.
5. D. Papadimitriou, D. Colle, P. Audenaert, and P. Demeester. "Geometric information routing," *Proc. of IEEE International Conference on Advanced Networks and Telecommuncations Systems (ANTS),* pp.1–8, Dec.2013.
6. F. Papadopoulos, D. Krioukov, M. Bogua, and A. Vahdat, "Greedy forwarding in dynamic scale-free networks embedded in hyperbolic metric spaces," *Proc. of IEEE INFOCOM 2010,* pp.1-9, 2010.
7. S. Sahhaf, D. Papadimitriou, W. Tavernier, D. Colle, and M. Pickavet, "Experimentation of geometric information routing on content locators," To appear in *Proceedings of CNERT 2014,* colocated with International Conference on Network Protocols (ICNP) 2014, October 2014.


# Routing Security on the Internet: Is It Really Worth the Effort?

***By Mat Ford***

During the Internet Research Task Force open meeting in Toronto, the third Applied Networking Research Prize for 2014 was presented to Robert Lychev and his coauthors for studying the security benefits provided by partially-deployed S*BGP.

Many widely used communication protocols on the Internet were not originally designed with security in mind—they were intended for parties that trust each other. As the Internet has evolved, new protocols intended to address specific security vulnerabilities have been developed. Deployment of these protocols can take a long time, therefore questions about the interactions of new secure protocol solutions with legacy insecure protocols are important.

For routing of Internet traffic, Border Gateway Protocol (BGP) is a key technology and much work has been done to address the real security vulnerabilities of BGP via developments like the Resource Public Key Infrastructure (RPKI) and BGP Security Extensions (BGPSEC). Lychev and his collaborators were interested in understanding the security properties of BGPSEC in partial deployment. In particular, what does partially deployed BGPSEC offer over RPKI or, "Is the juice (additional security benefits) worth the squeeze (extra efforts of deployment)?"

In their paper, "BGP Security in Partial Deployment" (*Proc. ACM SIGCOMM*, Hong Kong, China, August 2013), Lychev and his coauthors, Sharon Goldberg and Michael Schapira, report that (1) partially deployed security measures sometimes introduce new vulnerabilities, and (2) partial deployment provides only meagre benefits over RPKI if operators do not prioritise security over all other considerations in their routing policies.

> "I learned quite a bit from this meeting. I met a lot of people, and I hope to start new collaborations with some of them in the near future."
>
> —Robert Lychev, 2014 ANRP Winner

Speaking about the award and his trip to the IETF meeting in Toronto, Lychev said, "I learned quite a bit from this meeting. I met a lot of people, and I hope to start new collaborations with some of them in the near future."

Robert's slides are available at www.ietf.org/proceedings/90/slides/slides-90-irtfopen-1.pdf. Audio is available at http://recordings.conf.meetecho.com/Playout/watch.jsp?recording=IETF90_IRTFOPEN&chapter=chapter_0 (starting at 00:09:00).

Robert Lychev (right) receives his award from Lars Eggert, Internet Research Task Force chair.

**2015 ANRP Nominations**

The nomination period for prizes to be awarded in 2015 closed on 31 October 2014.

# Low-power, Lossy Network Plugfest Demonstrates Running Internet of Things Code

***By Thomas Watteyne, Ines Robles, and Xavier Vilajosana***

During IETF 90, the 6TiSCH, 6lo, and ROLL working groups (WGs) hosted the Low-power Lossy Networks (LLN) Plugfest—an event designed to bring together IETF participants interested in gaining hands-on experience with Internet of Things (IoT) technology. Eight teams presented implementations of Request for Comments (RFCs), Internet Drafts (I-Ds), tools, and hardware related to technology standardized by the 6TiSCH, 6lo and ROLL WGs. The focus of the implementations was on IEEE802.15.4e Timeslotted Channel Hopping (TSCH), the 6TiSCH architecture, IPv6 over Low-power Wireless Personal Area Networks (6LoWPAN)[1], and Routing Protocol for Low-Power and Lossy Networks (RPL)[2]. This article explains the technical challenges and outcomes of the event, the importance of running code, and the federating role of such events across IoT-related WGs.

## Introduction

Several WGs design solutions for wireless networks of constrained devices, also known as low-power and lossy networks, the cornerstone of the Internet of Things.

- The 6lo (WPAN) WGs has standardized an adaptation layer to efficiently carry (long) IPv6 packets over (short) IEEE802.15.4 frames, and other link-layer technologies.
- The ROLL WG has defined the RPL routing protocol, which enables multihop topologies in those constrained networks.
- The 6TiSCH WG is chartered to enable an IPv6-based upper-stack to run on top of the IEEE802.15.4 TSCH link layer. TSCH technology, widely deployed in low-power wireless industrial monitoring solutions, enables ultra-low power and high reliability, and introduces determinism to LLNs.
- Other groups in the LLN space include the CoRE, ACE, DICE, and LWIG WGs.

Each of the aforementioned WGs focuses on a subset of the IoT. For this reason, it is important that we also keep the big picture in mind, that we continuously integrate RFCs and I-Ds from these WGs into one working network to enable the flagging of potential conflicts or missing work. To that end, the 6TiSCH, 6lo, and ROLL WGs cohosted Plugfest[3] at IETF 90, cochaired by Xavier Vilajosana from the Universitat Oberta de Catalunya and Ines Robles from LMF Ericsson.

## Technical Outcomes

Pascal Thubert from Cisco Systems[4] and Thomas Watteyne from Linear Technology[5] presented a joint implementation of the multi-LLN architecture defined by the 6TiSCH WG. They interconnected two SmartMesh IP wireless networksthrough two Cisco i3000 industrial switches. Linear Technology's

[I]t is important that we keep the big picture in mind, that we continuously integrate RFCs and I-Ds from these WGs into one working network to enable the flagging of potential conflicts or missing work.

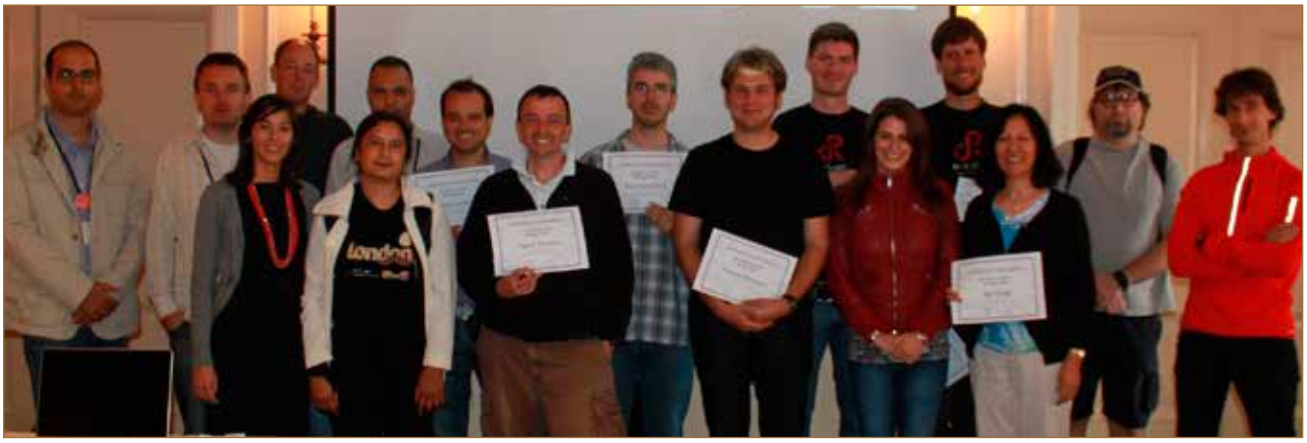

From left to right: Pouria Zand, Alexander Roscoe, Victoria Pimentel Guerra, Jürgen Schönwälder, Samita Chakrabarti, Cedric Adjih, Nicola Accettura, Pascal Thubert, Pere Tuset, Thomas Watteyne, Thomas Eichinger, Ines Robles, Oliver Hahm, Qin Wang, Michael Richardson, Xavier Vilajosana

SmartMesh IP is a commercial product which implements 6LoWPAN and IEEE802.15.4e TSCH, and uses techniques similar to the ones being standardized by 6TiSCH. The Cisco switches play the role of 6LoWPAN Backbone Router[6], federating the two independent SmartMesh IP wireless networks under a single IPv6 prefix.

The University of California, Berkeley's OpenWSN[7] project is an open-source implementation of the protocol stack standardized by the 6TiSCH WG, ported on a variety of hardware and software platforms. Nicola Accettura (University of California, Berkeley), Pere Tuset and Xavier Vilajosana (Universitat Oberta de Catalunya), Qin Wang and Tengfei Chang (University of Science and Technology Beijing), Marcelo Barros and Vitor Garbellini (Universidade Federal de Uberlândia), and Thomas Watteyne (OpenWSN project coordinator) showed a 10-mote network of OpenMote devices implementing the latest I-Ds developed by the 6TiSCH WG[8,9,10,11]. The demonstration consisted of a Constrained Application Protocol (CoAP) client triggering the reservation of link-layer cells along a multihop route.

OpenMote[12] is a startup company that provides an ecosystem of hardware for the Internet of Things. At its core is the OpenMote, a fully programmable and easy-to-use IEEE802.15.4-embedded communication platform.

The University of California, Berkeley's OpenWSN project is an open-source implementation of the protocol stack standardized by the 6TiSCH WG, ported on a variety of hardware and software platforms.

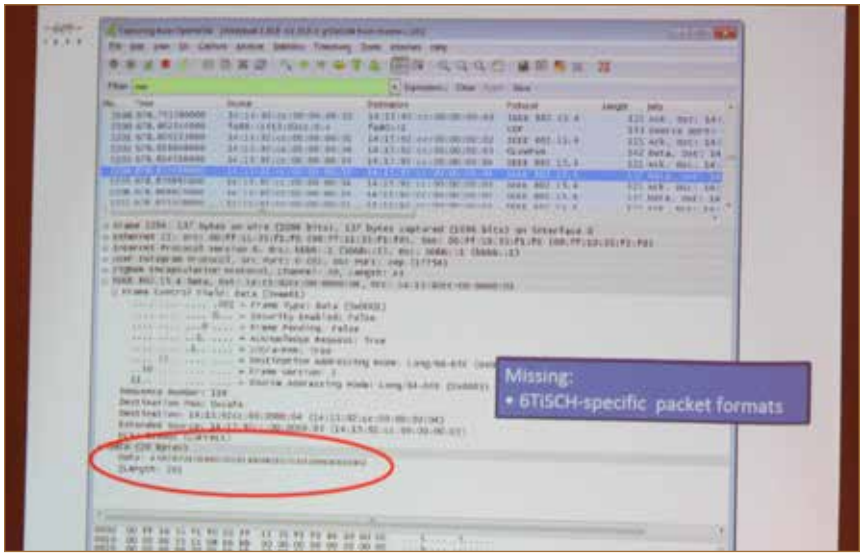


The plugfest included work to extend Wireshark to accurately dissect IEEE802.15.4e TSCH packets.

Multiple open-source implementations, including OpenWSN, fully support the OpenMote. Cofounders Pere Tuset and Xavier Vilajosana demonstrated how an OpenMote can be turned into a wireless packet-capture device for Wireshark, and discussed ongoing work supporting the FreeRTOS operating system in tickless mode.

Cedric Adjih from Inria gave a live demonstration of FIT-IoT's IoT-LAB[13] testbed. This open testbed comprises 2,728 wireless devices deployed on six sites across France. A RESTful web interface enables a user to remotely reserve a number of devices, reprogram them with custom firmware, and monitor their activity. Several open-source implementations, including OpenWSN and RIOT[14], can be used on the IoT-LAB platforms.

Oliver Hahm (Inria) and Thomas Eichinger (FU Berlin) presented RIOT, the friendly operating system for the IoT. They gave a live demonstration of the RIOT operating system running the protocol stack from the OpenWSN project on the IoT-LAB_M3 board.

Jürgen Schönwälder (Jacobs University Bremen) demonstrated an implementation of the 6loWPAN-MIB[15] in Contiki, running on the AVR Raven.

OpenMote is a startup company that provides an ecosystem of hardware for the Internet of Things. At its core is the OpenMote, a fully programmable and easy-to-use IEEE802.15.4-embedded communication platform.

This I-D defines a set of counters for monitoring the behavior of a 6LoWPAN stack implementation for packets, errors, compression, fragmentation parameters, etc. The demonstration involved retrieving those counters both through Simple Network Management Protocol (SNMP) and CoAP on 6lo[16] devices.

Vincent Ladeveze[17] is developing Wireshark dissectors for IEEE802.15.4e TSCH and other I-Ds from the 6TiSCH WG. He is building a 16-channel IEEE802.15.4 sniffer by connecting 16 devices running custom OpenWSN firmware to a computer on which there is software that aggregates these streams of captured packets and forwards them to Wireshark.

*Continued on next page*

The ability to simultaneously sniff 16 IEEE802.15.4 frequencies is necessary for debugging channel hopping solutions, such as IEEE802.15.4e TSCH.

Nestor Tiglao (University of the Philippines) presented Sewio's open sniffer solution, in which multiple wireless devices can be scattered around an area and report, through the Ethernet subnet, the different wireless packets they have captured to a single Wireshark instance. This enables the debugging of geographically spread wireless networks.

### Nontechnical Outcomes

The success of the IETF 90 LLN Plugfest underscores the importance of running code early in the standardization process. In the case of ongoing standardization work in the 6TiSCH WG, having implementations running when the documents are still at the I-D stage enables both verification of what is being proposed and iterative improvement and reconsideration of decisions taken to improve the quality of the documents being produced. Because the ability to formally acknowledge implementations increases the quality of the produced I-Ds or RFCs, we support efforts that do so, such as the CodeMatch outreach program.

> The Plugfest reinforced the importance of continuously integrating the I-Ds and RFCs produced by different WGs, and verified that this kind of integration is a complete and conflict-free solution.

Xavier Vilajosana, cofounder of OpenMote, demonstrates how OpenMote can be turned into a wireless packet-capture device for Wireshark.

Moreover, in the LLN standardization space, multiple WGs are focused on a subset of the standardization space. The Plugfest reinforced the importance of continuously integrating the I-Ds and RFCs produced by different WGs, and verified that this kind of integration is a complete and conflict-free solution. It also suggests that the IETF considers creating a body that oversees the IoT-related WGs to flag potential conflicts early in the standardization process, probably well before implementers could.

> [The Plugfest] also suggests that the IETF considers creating a body that oversees the IoT-related WGs to flag potential conflicts early in the standardization process, probably well before implementers could.

### Acknowledgements

We thank Working Group chairs Michael Richardson, Pascal Thubert, Samita Chakrabarti, and Ulrich Herberg for hosting the event; IETF Chair Jari Arkko and Area Directors Adrian Farrel, Alia Atlas, and Ted Lemon for making the event possible; Stephanie McCammon for helping to organize the event; and particularly acknowledge the teams who participated in the Plugfest, whose sleepless nights made the event a great success.

Thomas Watteyne of Linear Technology presented a joint implementation of the multi-LLN architecture defined by the 6TiSCH WG.

### Footnotes

1. "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks" [RFC 6282]
2. "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks" [RFC 6550]
3. Guidelines and additional details are available at https://bitbucket.org/6tisch/meetings/wiki/140720a_ietf90_toronto_plugfest
4. http://www.cisco.com/
5. http:// www.linear.com/dust
6. 6LoWPAN Backbone Router [draft-thubert-6lowpan-backbone-router-03, work-in-progress]
7. http://openwsn.berkeley.edu/
8. Minimal 6TiSCH Configuration [draft-ietf-6tisch-minimal-02, work-in-progress]
9. 6TiSCH Operation Sublayer (6top) [draft-wang-6tisch-6top-sublayer-01, work-in-progress]
10. 6TiSCH On-the-Fly Scheduling [draft-dujovne-6tisch-on-the-fly-03, work-in-progress]
11. The IPv6 Flow Label within a RPL domain [draft-thubert-6man-flow-label-for-rpl-03, work-in-progress]
12. http://www.openmote.com/
13. https://www.iot-lab.info/
14. http://www.riot-os.org/
15. Definition of Managed Objects for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs) [draft-ietf-6lo-6lowpan-mib-01, work-in-progress]
16. A 6lo stack is based on 6LoWPAN (RFC 4944, RFC 6282, RFC 6775) supporting different Link-layers
17. Presented by Thomas Watteyne, on behalf of Vincent Ladeveze.

# IRTF Update

***By Lars Eggert***

During IETF 90 in Toronto, Canada, four out of the nine chartered Internet Research Task Force (IRTF) research groups (RGs) held meetings:

- Information-Centric Networking (ICNRG)
- Crypto Forum (CFRG)
- Software-Defined Networking (SDNRG)
- Network Management (NMRG)

In addition to the meetings of those already chartered research groups, a proposed research group on Datacenter Latency Control (DCLCRG) held its first public meeting. In addition, a meeting was held to discuss a new research group on Network Function Virtualization (NFVRG). Both proposed research groups are planning meetings during IETF 91 in Honolulu, Hawaii; NFVRG also planned an interim meeting in the San Francisco Bay area on 4 September 2014.

A third proposed research group on Global Access to the Internet for All (GAIA) did not meet at IETF 90, and instead will meet 20–21 October 2014 in Cambridge, UK, and again plans to colocate with the Association for Computing Machinery (ACM) Symposium on Computing for Development in December 2014.

Since IETF 89, three new RFCs were published on the IRTF RFC Stream:

- RFC 7122 on Datagram Convergence Layers for the Delay- and Disruption-Tolerant Networking (DTN) Bundle Protocol and Licklider Transmission Protocol (LTP)
- RFC 7242 on Delay-Tolerant Networking TCP Convergence-Layer Protocol
- RFC 7253 on The OCB Authenticated-Encryption Algorithm

The IRTF Open Meeting at IETF 90 was also where the winner of the third Applied Networking Research Prize (ANRP) of 2014 presented his research. Robert Lychev presented his study on the security benefits provided by partially deployed S*BGP (See "Routing Secruity on the Internet: Is It Really Worth the Effort?" on page 17). A second ANRP presentation was postponed until IETF 91 due to visa difficulties, so now the IRTF Open Meeting in Honolulu will feature the final three award presentations for 2014. In lieu of the second ANRP presentation, local Toronto professor Michele Mosca talked to participants about quantum-safe cryptography.

Stay informed about these and other happenings by joining the IRTF discussion list at www.irtf.org/mailman/listinfo/irtf-discuss.

Robert Lychev, the third Applied Network Research Prize winner of 2014, presents his research.

### 2015 ANRP Nominations

The ANRP is awarded for recent results in applied networking research that are relevant for transitioning into shipping Internet products and related standardization efforts. The nominations period for the 2015 ANRP awards closed on 31 October 2014.

Please see https://irtf.org/anrp for more details.

# IANA Transition Update

***Adapted from a 10 September 2014 IETF blog post by Jari Arkko.***
*See the original post at www.ietf.org/blog/2014/09/iana-transition-update/*

The transition of the National Telecommunications and Information Administration's (NTIA's) stewardship of the Internet Assigned Names Authority (IANA) has been extensively discussed, leading to the creation of both the IANA Stewardship Transition Coordination Group (ICG) (www.icann.org/stewardship/coordination-group) and an IETF Working Group (WG) called IANAPLAN (https://datatracker.ietf.org/doc/charter-ietf-ianaplan/).

With so many interested people, so many opinions, and a role for the NTIA that can be difficult to understand, it can all be very confusing. What is the real story, and what is happening? Where are we in this process? What are the challenges?

The first thing to understand is that the role of IANA is in bookkeeping, not setting policy. They perform a very important role for the Internet, but the actual policy decisions—such as the assignment of new top-level domains or protocol numbers—are done elsewhere. For instance, these decisions take place via the multistakeholder gTLD process at Internet Corporation for Assigned Names and Numbers (ICANN), or via community consensus at the IETF.

Second, the transition is focused on the role of the NTIA, to move oversight to the Internet community. It is not about who performs the IANA function. That is not changing. The oversight role brings a responsibility to track performance and make necessary corrective actions. In the past 15 years, the community has already taken on much of the oversight role, so, in my opinion, the transition is not necessarily the abrupt change some think it is.

The primary responsibility for planning the transition lies with the operational communities that interact directly with IANA, specifically the IETF, the regional Internet registries (RIRs), and ICANN's gTLD and ccTLD communities. These communities are also expected to engage with the broader Internet community, to ensure their plans work well for business, civil society, government, and others.

At the IETF, this work is happening in the IANAPLAN WG chaired by Leslie Daigle and Marc Blanchet. Similarly, the RIRs are already setting up their organisations to develop a plan, and various communities at ICANN are working through a Cross-Community Working Group to develop their plan. Links to all these community efforts are on the ICG's web page at www.icann.org/en/stewardship/community.

Although a year sounds like a long time to develop and agree on a plan, developing a plan requires many steps and a lot of community discussion across the world.

The ICG, comprising 32 individuals from the Internet community, coordinates among these efforts. Work is underway, and with very few exceptions, everyone believes the transition is a good thing for the Internet. However, there are some challenges.

The first challenge is the timeline. Although a year sounds like a long time to develop and agree on a plan, developing a plan requires many steps and a lot of community discussion across the world. Time needs to be reserved (a) for the NTIA to determine if they find it acceptable, (b) to test drive any new mechanisms, (c) to ensure all plans work well together, and (d) for sufficient community comment periods.

Another challenge is accountability. Some viewed the NTIA as a backstop in case something bad happened. While not everybody shares that view, everyone does want to ensure that after the transition there are sufficient safeguards in place. ICANN has started an accountability improvement project to ensure this, but finding a solution that satisfies every concern will not be easy.

The necessary solutions depend on which IANA registries are discussed. For instance, I believe the IETF has good accountability mechanisms for protocol parameters. A contract between the IETF and ICANN governs the duties and roles of both parties. Any difficulties are tracked daily and any serious problem could be raised in the organisations, up to invoking the contract's termination clause. Similarly, if IETF policy decisions are mismanaged, there are last call, appeal, nominations committee, and recall mechanisms that enable the community to correct failures or replace leadership. Accountability improvements at ICANN would be useful, but not absolutely required for the IETF part of the transition. This is probably not true for the names part of the IANA registries, however, and significant work is needed there.

A third challenge is reaching everyone who needs to be involved, and finding an agreement. We all have to discuss with a broader community than usual—a broadly supported transition plan needs buy-in from different corners of the Internet community, from engineers to domain name experts to businesses and governments. I am committed to ensuring that the IETF communicates our plan broadly, and draws in interested participants.

To contribute to the IANAPLAN WG, join the mailing list at www.ietf.org/mailman/listinfo/ianaplan.

# IETF Ornithology: Recent Sightings

***Compiled by Mat Ford***

*Ruby-throated hummingbird (Archilochus colubris)*

Getting new work started in the IETF usually requires a birds-of-a-feather (BoF) meeting to discuss goals for the work, the suitability of the IETF as a venue for pursuing the work, and the level of interest in and support for the work. In this article, we'll review the BoFs that took place during IETF 89, 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.

### Abstraction and Control of Transport Networks (actn)

**Description:** Network operators build and operate multidomain networks. These domains may be collections of links and nodes, each of a different technology, administrative zone, or vendor-specific island. Establishment of end-to-end connections spanning multiple domains is a perpetual problem for operators, both because of operational concerns and because of interoperability issues. Due to these issues, the introduction of new services, often requiring connections that traverse multiple domains, needs both significant planning and several manual operations to interface different vendor equipment and technology.

Establishment of end-to-end connections spanning multiple domains is a perpetual problem for operators, both because of operational concerns and because of interoperability issues. Due to these issues, the introduction of new services, often requiring connections that traverse multiple domains, needs both significant planning and several manual operations to interface different vendor equipment and technology.

The aim of ACTN is to facilitate virtual network operation: the creation of a virtualized environment enabling operators to view and control multiple multi-subnet, multitechnology networks as a single virtualized network. Network abstraction of transport networks is also necessary for operators who consolidate their network services into multitenant virtual transport networks. This will accelerate rapid service deployment of new services, including more dynamic and elastic services, and will improve overall network operations and scaling of existing services. Discussion with operators has highlighted a need for virtual network operation based on the abstraction of underlying technology and vendor domains.

This BoF was not intended to form a working group. It was intended to give operators an opportunity to express their current operational practices, highlighting operational pain points, network virtualization requirements and objectives, short-term goals, and longer-term aspirations.

*Continued on next page*

*IETF Ornithology: Recent Sightings, continued*

**Proceedings:** www.ietf.org/proceedings/90/minutes/minutes-90-actn

**Outcome:** Discussion enabled operators to express their needs and issues, and some common threads were visible, e.g., end-to-end services over multidomain or multilayer networks. It was less clear what that would mean in terms of protocol work. Discussion of use cases, potential solutions, scoping the problem etc. will continue on the mailing list.

### Virtualized Network Function Pool (vnfpool)

**Description:** This is the second BoF meeting on this topic, the previous session having taken place during IETF 89. The main goal of this meeting was to come to consensus on a charter for a vnfpool working group. Since the previous meeting, the proposed charter was updated to focus on pooling within an individual VNF, not reliability for the whole service graph; to clarify the relation of vnfpool to the service function chaining working group, and; to leave service state synchronization out of scope in the initial phase of work.

Since the previous meeting, the proposed charter was updated to focus on pooling within an individual VNF, not reliability for the whole service graph; to clarify the relation of vnfpool to the service function chaining working group, and; to leave service state synchronization out of scope in the initial phase of work.

**Proceedings:** www.ietf.org/proceedings/90/minutes/minutes-90-vnfpool

**Outcome:** It remains unclear whether or not standardization of new technology is needed here, and further discussion with a broader community is required to answer some of the outstanding questions. More work on the mailing list is expected.

### Delay Tolerant Networking Working Group (dtnwg)

**Description:** The meeting investigated interest in transitioning technologies developed in the IRTF DTN research group into standards-track activities through the formation of a new IETF working group. The meeting was presented with a draft working group charter, including work items based on the DTN Bundle Protocol. The goal of the meeting was to discuss the draft charter and present the candidate work items, as well as to determine the level of support for conducting the work in the IETF. The desired end state was the formation of a new working group soon after IETF 90.

*Northern cardinal (Cardinalis cardinalis)*

**Proceedings:** www.ietf.org/proceedings/90/minutes/minutes-90-dtnwg

**Outcome:** Although further work is required to refine the charter, this was a productive and positive session in which several people expressed willingness to work on documents in a DTN working group.

### Use Cases for Autonomic Networking (ucan)

**Description:** See "Autonomic Networking," on page 1 of this issue of *The IETF Journal*.

**Proceedings:** www.ietf.org/proceedings/90/minutes/minutes-90-ucan

**Outcome:** This meeting was not intended to form a working group. Many different use cases were presented and the chairs and other interested parties will continue the work of developing a more focused working group charter on the mailing list.

[The meeting] began by recapping the outcome from the IETF 89 igovupdate session, and then discussing both the NTIA transition in the larger Internet community and the Coordination Group.

### IANAPLAN (ianaplan)

**Description:** This was a working-group forming BoF meeting. It began by recapping the outcome from the IETF 89 igovupdate session, and then discussing both the NTIA transition in the larger Internet community and the Coordination Group. Finally, the impact for the IETF and proposed plans for action were discussed, including a draft working group charter. For more discussion of the IANA transition, see "IANA Transition Update," from IETF Chair Jari Arkko on page 22 of this issue of *The IETF Journal*.

**Proceedings:** www.ietf.org/proceedings/90/minutes/minutes-90-ianaplan

**Outcome:** This discussion went well. There was agreement in the room to support formation of a working group with a slightly tighter charter than was presented during the BoF.

*Common pheasant (Phasianus colchicus)*



# IETF 90 At–A–Glance

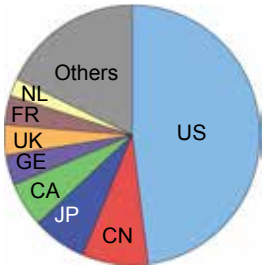


Paid attendance: 1183
(Above forecast by 113 people)

Newcomers: 153

Number of countries: 53

**IETF Activity since IETF 89 (March–July 2014)**

New WGs: 3

WGs closed: 0

WG currently chartered: 123

New and revised Internet-Drafts (I-Ds): 1848

RFCs published: 173

- 105 Standards Track, 7 BCP, 8 Experimental, 50 Informational

Live and Social Media

- YouTube: Tech Plenary streamed 47 live, 301 views
- Twitter: 1186 tweets on #IETF90, exposure via #IETF89=19000, 578 new followers
- Overall, 999.5K impressions on Facebook and Twitter

**IANA Activity since IETF 89 (February–June 2014)**

Processed 1705+ IETF-related requests, including:

- Reviewed 117 I-Ds in Last Call and reviewed 130 I-Ds in Evaluation
- Reviewed 151 I-Ds prior to becoming RFCs, 95 of the 151 contained actions for IANA

Document collaboration with the IETF

- RFC 5226bis is still being reviewed and revised with community members. IETF Last Call expected soon. See https://datatracker.ietf.org/doc/draft-leiba-cotton-iana-5226bis/

SLA Performance (January–June 2014)

- Processing goal average for IETF-related requests: 99%

IANA and DNSSEC

- As of 16 July 2014, 448 TLDs have a full chain of trust from the root. See http://stats.research.icann.org/dns/tld_report/

**RFC Editor Activity since IETF 89 (January–July 2014)**

Published RFCs: 217

- 164 IETF (31 IETF non-WG), 5 IAB, 4 IRTF, 13 Independent

# IETF Meeting Calendar

**IETF 91**
9–14 November 2014
Host: Cisco
Location: Honolulu, HI, USA

**IETF 92**
22–27 March 2015
Host: Google
Location: Dallas, TX, USA

**IETF 93**
19–24 July 2015
Host: TBD
Location: Prague, Czech Republic

**IETF 94**
1–6 November 2015
Host: WIDE
Location: Yokohama, Japan

For more information about past and upcoming

**IETF Meetings**

www.ietf.org/meeting/

**Special thanks to**

ERICSSON

**for hosting IETF 90**

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

Afilias Google Microsoft NBCUniversal

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

Afilias ARIN cable CISCO

comcast ERICSSON HUAWEI

ICANN JUNIPER NETWORKS NBCUniversal RIPE NCC

pub-IETFJv10.2-20141024-en

Published three times a year
by the Internet Society.
Galerie Jean-Malbuisson 15
1204 Geneva, Switzerland

**Editor**
Mat Ford

**Associate Editors**
Megan Kruse • Michelle Speckler

**Contributing Writer**
Carolyn Marsan

**Editorial and Design**
Speckler Creative

**Editorial Board**
Jari Arkko
Mat Ford
Russ Housley
Olaf Kolkman
Megan Kruse
Lucy Lynch
Greg Wood

**Email**
ietfjournal@isoc.org

**Find us on the Web**
www.internetsociety.org/ietfjournal

**Editor's Note**
*The IETF Journal* adheres to the *Oxford English Dictionary*, 2nd Edition.






# IETF Journal

Internet Society
Galerie Jean-Malbuisson 15
1204 Geneva, Switzerland