IETF Journal

Volume 5, Issue 1
June 2009

*A report from IETF 74, March 2009, San Francisco, California. Published by the Internet Society in cooperation with the Internet Engineering Task Force**

Inside this issue

IPv6, Trust and Identity, Key Themes at IETF 74

From the Editor's Desk, by Mirjam Kühne

IPv6 dominated the discussion during several working group and side meetings at IETF 74, culminating in a panel of industry experts and other thought leaders who were brought together to explore the obstacles facing widespread adoption and deployment of IPv6. The discussion is summarized in an article called "The Seven Stages of IPv6 Adoption" (see page 14).

Photo by Peter Löthberg

Cable car in San Francisco not far from the IETF meeting

On a related topic, a BoF (birds-of-a-feather) session looked at a possible solution to the problem of how public IPv4 addresses can be shared among different networks in the event that IPv4 addresses are no longer available to be assigned. A description of that proposal can be found below.

A similarly hot topic these days is trust and identity. In this issue, the *IETF Journal* talks to the cochairs of the OAuth BoF as well as the author of the OAuth specification. The OAuth specification recently was brought into the IETF.

Also in this issue is a summary of the plenary session, including a review of the panel discussion on multiprotocol label switching and an update on what the IETF can learn from the development and deployment of that protocol.

IETF 74 hosted a number of the Internet Society fellows to the IETF and former fellows, who travelled from developing countries for the opportunity to enhance their knowledge and technical skills through involvement with the IETF and to contribute to the work of the IETF (see page 12).

Finally, in an effort to gain a better understanding of our readers, we are embarking on our first *IETF Journal* reader survey. We encourage you to take a minute and fill it out. The survey can be found at http://ietfjournal.isoc.org.

We extend our thanks to those who contributed to this issue. We wish everyone enjoyable reading, and as always, we welcome both comments and contributions for future issues.

Address Sharing— Coming to a Network near You

By Mat Ford, Alain Durand, Phil Roberts, Pierre Levis

Hopefully it is not news to you that allocations of IPv4 addresses from the Internet Assigned Numbers Authority (IANA) are currently forecast to be complete during the first half of 2011 [http://tools.ietf.org/html/draft-ford-shared-addressing-issues-00#ref-IPv4_Report]. Allocations from the Regional Internet Registries are anticipated to be complete a year later, although the exact date will vary from registry to registry. This looming address crunch is causing Internet service providers (ISPs) around

Continued on page 8

Message from the IETF Chair

By Russ Housley


Russ Housley, IETF Chair

In spite of the worldwide economic downturn, IETF 74 was a highly successful meeting. The work of the IETF remains relevant and enthusiastic. Held in San Francisco in March 2009, the meeting drew 1,157 people from 49 different countries. The total number of attendees from California was 273, compared with 143, which is the average number who attended each of the previous three meetings.

Juniper Networks, which hosted the meeting, did a great job, and everyone felt welcome (the T-shirt design drew a number of compliments). The social event was well attended, providing an enjoyable and science-filled evening. The site network was subcontracted to VeriLAN Networks, whose staff, working with dedicated volunteers, made sure the network ran smoothly.

The week was filled with the usual mixture of working group (WG) meetings, birds-of-a-feather (BoF) sessions, research group (RG) meetings, and, as always, many side meetings.

Since IETF 73, one new WG was chartered and six WGs were closed. Currently there are approximately 110 chartered WGs. Between IETF 73 and IETF 74, the WGs and their individual contributors produced 424 new Internet-Drafts and updated 1,013 Internet-Drafts, some of them more than once. The Internet Engineering Steering Group (IESG) approved 106 Internet-Drafts for publication as RFCs. The RFC Editor published 86 new RFCs.

During IETF 74, the IESG passed four seats to new members, and the IETF Administrative Oversight Committee (IAOC) passed two seats to new members. Our thanks to Ed Juskevicius (IAOC), Chris Newman (Applications Area Director [AD]), Jon Peterson (Transport AD and then RAI AD), Jonne Soininen (IAOC), Mark Townsley (Internet AD), and Dave Ward (Routing AD) for their many years of service to the community. We welcome Ralph Droms (Internet AD), Marshall Eubanks (IAOC), Adrian Farrel (Routing AD), Alexey Melnikov (Applications AD), Robert Sparks (RAI AD), and Henk Uijterwaal (IAOC). Thank you for your willingness to serve.

Throughout the week, an IPv6-only network was available for attendees to experience the Internet without IPv4. The discussion of requirements for NAT-PT (network address translation–protocol translation) continues.

In the few weeks prior to IETF 74, there were several intellectual property right (IPR) surprises. To avoid future surprises, a portion of the Wednesday plenary was devoted to a reminder of IETF policies in this area. If you were not there, please review the slides from that session. They are available at http://www.ietf.org/proceedings/09mar/slides/plenaryw-1/plenaryw-1.htm.

The Wednesday plenary also included a discussion of potential changes to the Nominating Committee (NomCom) process. The discussion has continued on the ietf-nomcom@ietf.org mail list. Please join the discussion about the process used for picking IETF leaders.

I look forward to IETF 75 in Stockholm, Sweden, on 26–31 July 2009. The meeting will be hosted by **.se**. I also look forward to seeing you at IETF 76 in Hiroshima, Japan, on 8–13 November 2009. That meeting will be hosted by WIDE. Scheduling information for the next IETF meetings may always be found at http://www.ietf.org/meetings/meetings.html. I look forward to seeing you soon.

New BoF Meetings

Descriptions and agendas for all BoF meetings can be found at http://www.ietf.org/meetings/past.meetings.html.

Applications Area
oauth: Open Web Authentication
mmox: Massively Multi-Player Games and Applications
yam: Yet Another Mail

General Area
pre8prob: Pre-5378 Problem

Internet Area
6ai: IPv6 Address Independence
lisp: Locator/ID Separation Protocol
mif: Multiple Interfaces
netext: Network-Based Mobility Extensions

RAI Area
xmpp2: Extensible Messaging and Presence Protocol 2
atoca: Authority-to-Citizen Alert

Transport Area
shara: Sharing of an IPv4 Address
storm: Storage Maintenance


Olaf Kolkman, IAB Chair

IETF 74 Facts and Figures

Registered attendees 1157
Countries 49
New WG 1
Closed WGs 6
WGs Chartered 110
New Internet-Drafts 424
Updated Internet-Drafts 1013
IETF Last Calls 86
Internet-Drafts approved for publication 106
RFC Editor Actions (November 2008 - February 2009)
86 RFC published of which
- 51 Standards Track
- 3 BCP
- 30 Informational
- 2 Experimental

100 Internet-Drafts submitted for publication
- 79 were submitted by the IETF

IANA Actions (November 2008 - February 2009)
1455 IETF-related requests processed
- 742 Private Enterprise Numbers
- 88 Port Numbers
- 61 TRIP ITAD Numbers
- 30 media type requests

Words from the IAB Chair

By Olaf Kolkman

Spring is in the air . . .

During spring's IETF meeting, the red dots that appear on nametags, which indicate Internet Architecture Board (IAB) membership, are passed from outgoing to incoming members. I had to say good-bye to a number of folk I had not only enjoyed but also had the honour of working with: Loa Andersson, Barry Leiba, Kurtis Lindquist, and Lixia Zhang. Fortunately, people I'm looking forward working with—Marcelo Bagnulo, Vijay Gill, John Klensin, and Jon Peterson—are replacing them.

Springtime is also the time for the IAB to hold its retreat. The main goal of the retreat is for people to get to know each other and to set direction for the IAB's work over the coming year.

This year we met in the Verizon offices in Ashburn, Virginia. Even during spring, the location did not allow for frivolous distractions. It is in the middle of fields that surround Dulles airport and far away from the vices that can be found in Reston, Herndon, and Leesburg. Without those distractions—and others, such as sunlight—we could concentrate on what we'd come to accomplish. The first day, we focused on administrative duties. We discussed some perennial matters but also had a detailed discussion about the various liaison relationships overseen by the IAB. During those discussions we appointed Patrik Fältström to be the liaison to the ITU-T [International Telecommunication Union–Telecommunication Standardization Sector]. Patrik takes over from Scott Bradner, who has been serving the IETF in this role for many years. We also took a significant step forward in approval of the RFC Editor model.

Approval of the RFC Editor model is a milestone in a process that began at last year's retreat. The next major milestone occurs in January 2010, when the RFC Editor functions that are now being executed by the Information Sciences Institute (ISI) of the University of Southern California will be turned over to other people and organizations. The turning over of the RFC Editor function is momentous, particularly when one stops to consider that the RFC series has been edited by ISI for more than 40 years.

The second day we talked about Internet architecture and tried to pick a number of topics to work on in the coming year. The process we used for defining our agenda was to have each IAB member formulate an architectural topic that the member would be willing to lead together with a clear set of milestones and deliverables. We can't be certain that all of the plans we discussed will come to full fruition, so I will not go into detail about the individual efforts discussed. However, most of the topics can be put into three major buckets:

- IPv4 and IPv6 coexistence and how to work toward the best results in IPv6 transition
- Security of the routing data and the routing control plane
- Internationalization issues within the DNS and the application layer, as well as between the DNS and applications

Continued on next page 5

Plenary Report

By Mirjam Kühne

Note: This is not a complete report of the plenary sessions; rather, it is a summary of the highlights of the discussions. All IETF 74 presentations can be found at http:// www.ietf.org/meetings/past.meetings.html.

Mirjam Kühne

Administrative Updates

The IETF Secretariat has been working on a redesign of the IETF Web site, which turned out to be a much bigger job than anticipated. IETF chair Russ Housley thanked the secretariat staff for their hard work.

He also reported on the Code Sprint from early in the week of the IETF meeting, which was a big success. Five releases were developed, including a new version of the datatracker, which improved the authorisation system and which replaced a number of hard-to-maintain legacy scripts.

Finally, Russ thanked all of the sponsors and contributors who had made the meeting possible.

IAOC chair Jonne Soininen and IETF administrative director Ray Pelletier updated participants on issues related to the financial and organisational status of the IETF. In their report, they stated that nearly 50 percent of the IETF's revenues are derived from meeting registration fees, while 14 percent is met by contributions from meeting sponsors and Network Operations Centre (NOC) sponsors, which are secured by ISOC. ISOC provides nearly one-third of the annual funding necessary for IETF operations.

Less than half of the IETF's expenses are allocated for meetings. Day-to-day secretariat expenses—including IT support, RFC Services, and Tools and Administrative costs—total more than USD 2.5 million annually, not including volunteer time and with no direct source of funding.

In 2008, expenses for most activities were under budget. However, while the meeting revenue for IETF 71 and IETF 72 were on target, meeting attendance for IETF 73 was relatively low. At USD 616,000, both IETF host and NOC sponsorships reached a record level.

In light of the current global economic uncertainties, the IAOC has developed financial contingency plans based on 15 percent and 25 percent attendance attritions, which will be evaluated throughout the course of the year. The Internet Society committed to provide a safety net for 2009 should there be an attendee shortfall. This has made it possible for the IAOC to keep the registration fee for IETF 74 at the current rate. Registration fees for future meetings will be reviewed during the year. The IAOC is looking into other opportunities to increase participation at future IETF meetings.

The difficulties associated with visa requirements for entering the United States, especially for people from China, are still concerns. The IAOC requested assistance from the U.S. State Department, without much success. One solution might be to reduce the number of meetings held in the United States, perhaps moving some to Canada instead. The topic was further discussed during the open-mic portion of the plenary, in which a number of people suggested further improvements to remote-participation facilities. Thomas Narten suggested setting up a design team, which would codify rules and behaviour expectations so as to make remote participation easier and more appealing. A mailing list intended to further work on this topic was created following the meeting and can be found at https://www.ietf.org/mailman/listinfo/vmeet.

Photo by Peter Löthberg

Kireeti Kompella (right) of Juniper Networks, host of IETF 74

Ed Juscevicius gave his final report as chair of the IETF Trust prior to turning over the chair position to Marshall Eubanks. The Trust has done work on the legal provisions related to IETF documents. There is also a new frequently-asked-questions document on copyright issues. It is available at http://trustee.ietf.org/docs/IETF-Copyright-FAQ.pdf.

Patents at the IETF

IETF attorney Jorge Contreras and Scott Bradner gave a presentation describing a number of issues related to patents and disclosure obligations. While the rules can be found in RFC 3979, Scott and Jorge highlighted some of the relevant points with regard to the IETF's patent policies. For example, in terms of disclosure obligations, they explained that an IETF participant *must* disclose any known patent that the participant (or the participant's sponsor) controls and that may cover any IETF contribution. However, an IETF participant or anyone else *may* disclose third-party patents the person believes may cover IETF contributions.

Scott and Jorge also said that because IETF participation is by individuals,

disclosure is primarily an individual obligation. Individual engineers must ensure that their employer companies make the required disclosures. If an engineer cannot ensure a disclosure, that engineer should not participate. Similarly, companies that own the patents may be deemed to control the actions of their participating employees.

Disclosure is required "as soon as reasonably possible" after a contribution is published as an Internet-Draft. If this is your own contribution, disclosure should follow relatively quickly, such as within a few days. But Jorge and Scott suggest that you not make the contribution until you're ready to file the disclosure. In addition, if an IETF participant first learns of a patent after publication of the affected Internet-Draft, a disclosure must be made as soon as reasonably possible after the discovery.

Disclosures may be updated voluntarily at any time, but they must be updated when an IETF document changes such that its coverage by a disclosed patent or patent application changes or if the claims of a patent or patent application are amended so that their coverage of IETF documents changes. Failure to comply with patent disclosure requirements is a violation of IETF policy, and the potential legal consequences to companies are considerable.

Both Jorge and Scott strongly recommended that IETF participants refer to RFC 3979 and consult their legal counsels.

IAB Update

Internet Architecture Board chair Olaf Kolkman updated participants on recent IAB activities, including the IAB having finalized two documents since the last meeting: Principles of Internet Host Configuration and Design Choices When Expanding DNS.

Other documents that have been worked on but are still under review include RFC Streams Headers and Boilerplates, the RFC Editor Model, IAB Thoughts on IPv6 Network Address Translation, P2P [Peer-to-Peer] Architectures, Defining the Role and Function of IETF Protocol Parameter Registry Operators, and Evolution of the IP Model.

The IAB continues to follow developments with respect to multiprotocol label switching–transport profile (MPLS-TP). At IETF 74, concerns were expressed about confusion in the marketplace regarding the utility and standardization status of transport multiprotocol label switching (T-MPLS). There have been claims that MPLS-TP is a forward-compatible update from T-MPLS. "This is clearly not the case," said Olaf. The IAB is in the process of reviewing some of the liaison relations between the IETF and other organizations.

At the end of his IAB update Olaf announced the new IAB members and thanked the outgoing IAB members, each of whom received a plaque as an expression of the IAB's gratitude for their service.

MPLS Turns 12

The IETF 74 technical plenary focused on MPLS. The session was intended to analyse MPLS, offering a case study involving the creation and operation of a successful protocol as well as the lessons learned. Four speakers enumerated MPLS's benefits and its impact on the overall Internet architecture, including its effects on higher-

Continued on next page

Comic BoF

Words from the IAB Chair, continued from page 3

In addition, the IAB has expressed a desire to work on certain topics therefrom, and by studying them, we can learn more about the major architectural questions that may face us in the near future. Two areas of interest were (1) technologies for IP in aviation and (2) an Internet populated by large numbers of low-power, autonomous devices.

All of the various projects bring with them different deliverables, some as vague as identifying the actual questions to begin with. This column is not the place to commit to specifics; however, it is clear there are some important issues that need attention.

During spring the seeds have been planted. May the flowers bloom.

Plenary Report, continued from page 5

IETF participants during the plenary session

IETF participants meet in hotel lobby

layer protocol/application operation and delivery. The speakers represented operator, industry forum, and vendor perspectives: George Swallow of Cisco Systems, Tom Bechly of Verizon Business/MCI, and Kireeti Kompella of Juniper Networks. The discussion was introduced and moderated by Loa Andersson and Andrew Malis, who asked the panellists what they would have liked to change in the development of MPLS. The speakers replied that overall, it is a good protocol, even though a few things could be improved. According to Tom, the diagnostic tools "trail the developments." He said that having them sooner would have helped. Kireeti pointed out that a lot of machinery was put in the Label Distribution Protocol that is not used today and that complicates the implementation of the protocol. "There could also be more interoperability," Kireeti said. "Interoperability tests take place in private networks, but there is not a lot of interservice provider interoperability."

When asked about the fact that there are two protocols, Kireeti responded that having two protocols "is painful all around." This is true, he said, for the specifiers, the implementers, and the people who do interoperability tests. It was not a choice we made lightly. Learning from this experience and applying it to the deployment of IPv6, for instance, Kireeti commented that "Seeing so many variants of trying to get IPv4 and IPv6 to interoperate is amazing. I believe we need one solution and should deploy that."

George added that the success of MPLS is to acknowledge that this is an IP world.

"Now we need to recognize that there needs to be a transition to IPv6," he said. "In MPLS we are lucky though: we don't have to deal with hosts."

IAB Open Mic

The open-mic portion of the technical plenary included a discussion about the architecture of the Internet and how it has changed over time. Keith Moore expressed concern that there is no longer a set of shared assumptions. He said he's hearing a lot of proposals that would violate the original set of principles of the Internet, and he's wondering whether there's a process that would bring us back to that state.

To illustrate that the Internet and the underlying assumptions have indeed changed, one participant pointed out

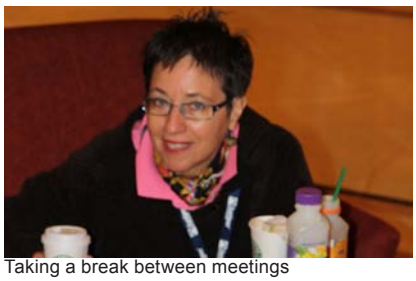

Taking a break between meetings

Interdomain Routing working group presentation

NomCom Changes

Nominations Committee (NomCom) chair Joel Halpern updated IETF 74 participants on recent NomCom activities, after which he welcomed new IESG and IAB members and described problems and questions the NomCom had faced in the most recent round of selections.

Following Joel's presentation, a discussion ensued focusing mainly on the revision of RFC 3777 (IAB and IESG Selection, Confirmation, and Recall Process: Operation of the Nominating and Recall Committees). It was suggested that the document be updated to clarify the role of the liaisons. Joel stated that people who serve on the NomCom should not also serve on any bodies the NomCom serves. "Liaisons should be there to help the NomCom," he said. "They have no vote."

At IETF 74, Spencer Dawkins presented a number of suggested changes to RFC 3777. The most significant change would be one in which the nominees list would be openly published, which could help remove certain inconsistencies, even though doing so could create the potential for lobbying for certain candidates. Though most attendees agreed that the change would be reasonable, Ted Hardie said he felt the problem is symptomatic of a larger problem—that is, the old boys network problem. He cautioned that the way the pool of volunteers is handled right now "eliminates a certain group of people." He suggested changing the NomCom process so that a candidate's absence from an IETF meeting would not disqualify the candidate from serving on the NomCom. The discussion will be continued on the IETF mailing list.

IETF Journal Glossary Now Online
For a complete glossary of terms and acronyms frequently used in the *IETF Journal*, see http://ietfjournal.isoc.org.

that a number of young gamers, who were in San Francisco during the week of the IETF for a gaming convention, were surprised to hear that network address translation (NAT) was not an original assumption of the Net from the start.

Most of those who participated in the open-mic session as well as IAB members agreed with Keith, though most said they felt differently about the seriousness of the problem. Kurtis Lindqvist said he believes that the reason for the changes evolved because the Internet has been so successful. And that, he said, is because the Internet allows people to develop new applications and to make money. "It is an interesting observation, but I can't decide if this is really a problem," he said.

Scott Brim spoke to the fluidity of the Internet's architecture, saying that it is even more fluid now. "There is a lot happening," he said. "These are interesting times." He agreed with Tony Hain, who earlier cautioned participants to "be careful that we can take out the stuff we put into the network once we solved the basic principles."

IETF participants during a working group session

IAB members attend BoF session at IETF 74

Others felt more strongly that a shared set of assumptions is important and that one needs to make sure they stay consistent. "All of us struggle over how many of the original assumptions we can recover," said Dave Oran.

The last part of the open-mic session addressed the work of the IAB and how it can be made more comprehensible and transparent. Most participants agreed that because the IESG is involved in operational issues, its work seems much clearer than that of the IAB. The role of the IAB is more difficult to grasp. Olaf encouraged people to speak up or to send suggestions to the IAB mailing list at iab@iab.org.

It is usually helpful when IAB members attend birds-of-a-feather or working group meetings. IAB statements or documents, too, are seen as useful contributions. "An IAB statement can have a lot of weight in the outside world," said Alain Durand, even though, more often than not, the IAB is simply distilling ideas shared by the community, which is an important part of the work of the IAB.

Photos by Peter Löthberg

Remembering Jim Bound and Steve Coya

Sadly, the Internet Society and the IETF lost two good friends over the past few months.

Photo by Peter Löthberg

On 2 March, **Jim Bound** passed away at the age of 58. In his role as chief technology officer of the IPv6 Forum, Jim was a passionate advocate for the adoption and deployment of IPv6. He was a member of the IETF's Internet Protocol Next Generation Directorate, which, in 1994, selected IPv6 from among several proposals to become the basis of the IETF's work on the next-generation Internet protocol. He was awarded the IPv6 Forum Internet Pioneer Award as IPv6 Lead Plumber. IETF chair Russ Housley expressed sadness at the news, calling attention to Jim's strong support for both the IETF and IPv6.

In announcing his passing to the IETF community, his friend and colleague Yanick Pouffary described him as a man of integrity who made "a profound impact on our industry and everyone who worked with him."

Photo by Ole Jacobsen

Steve Coya, who worked for many years at the Corporation for Research and Network Initiatives, passed away on 3 June. Having spent part of his early career at MCI, Steve is best known for his work at CNRI, where he served as executive director of the IETF, overseeing the IETF Secretariat and organizing IETF meetings during most of the 1990s and early 2000s. Steve was appreciated for his disarming sense of humor and his desire to solve whatever problem was at hand. "It was always nice to have one unflappable optimist in the midst of IETF havoc," commented Paul Mockapetris. "He will be missed."

The IPv6 Forum has announced that those who wish to make donations in Jim's memory send them to the Children's Cancer Research Fund, Jim Bound, c/o Stephen Ellis, PO Box 570, Hollis, NH 03049, U.S.A.

Contributions in memory of Steve Coya can be made to the Alzheimer's National Association, Capital Area Chapter, 11240 Waples Mill Rd., Suite 402, Fairfax, VA 22030.

Address Sharing, continued from page 1

the world to start to question how they will continue providing IPv4 service for IPv4-speaking customers when there are no longer sufficient IPv4 addresses to allocate. Universal IPv6 deployment was originally thought to be the solution to ensure continued global addressability of an ever-expanding network. However, it appears likely that there will be a gap between the demise of the IPv4 free pool of addresses and the arrival of IPv6.

Several possible solutions aimed at bridging that gap are now emerging. In this article we discuss some of the criteria that will help the community evaluate the merits of those choices and we cover the common and potentially serious issues to which address sharing across multiple subscribers may inevitably give rise. In addition, while network operators are busy devising solutions to the addressing problem that is looming on the horizon, content providers are encouraged to consider the impact of shared addressing on their business and operational practices.

Guiding Principles

> *The end-to-end principle is the core architectural guideline of the Internet.* Section 2 of RFC 3724

RFC 3724—The Rise of the Middle and the Future of End-to-End: Reflections on the Evolution of the Internet Architecture—provides a concise history of the end-to-end principle. While the original articulation was concerned with where best to place functionality in a communication system, the growth and development of the Internet have resulted in an expansion of the scope of the end-to-end principle. The principle now encompasses the question of where best to locate the state associated with Internet applications. This expanded principle is well articulated in RFC 1958: Architectural Principles of the Internet:

> *An end-to-end protocol design should not rely on the maintenance of state (i.e., information about the state of the end-to-end communication) inside the network. Such state should be maintained only in the endpoints, in such a way that the state can only be destroyed when the endpoint itself breaks (known as fate-sharing).*

The end-to-end principle is, arguably, the fundamental principle of the Internet architecture. In a sense, the Internet is the embodiment of the principle. By allowing either tacit or explicit erosion of the principle as we apply our understanding to the construction and operation of the global network, we allow the dismantling of the utility itself. Unfortunately, address sharing threatens just such erosion.

Shared-addressing solutions are being proposed as pragmatic responses to the very real problems faced by operators who need to be able to continue providing service for customers who do not have IPv6-capable equipment or who want to access services that are available only via IPv4. However, while we advocate solutions that allow continued operation of the IPv4 Internet and the continued provision of services for IPv4-speaking customers, we do not in any way advocate prolonging the life of IPv4 or of any solution that delays the widespread adoption of IPv6.

Universal IPv6 deployment was originally thought to be the solution to ensure continued global addressability of an ever-expanding network. However, it appears likely that there will be a gap between the demise of the IPv4 free pool of addresses and the arrival of IPv6.

Based on the importance of the end-to-end principle and the ultimate goal of global addressability through widespread IPv6 deployment as discussed earlier, solutions to the problem of how to continue providing IPv4 service post-IPv4-address completion should be judged on two primary criteria:

1. The ability of the end user to readily control the parameters of the solution to minimize the impact of the solution on the end-to-end communication and
2. The extent to which the solution affords a natural progression to widespread IPv6 deployment.

Adherence to these criteria will minimize the impact of address sharing on end-to-end communications, and it will keep the network on track toward the universal deployment of IPv6.

Potential Responses to IPv4 Address Shortages

Assuming ISPs wish to continue growing their businesses in a post-IPv4 world, there are a number of possible avenues they can take:

- Obtain previously allocated IPv4 addresses through some unspecified means;
- Deploy large-scale address translation and allocate customers with private addresses; or
- Deploy a port-shared addressing solution whereby customers would get public addresses with fewer available ports.

Let us deal with each of these in turn.

Obtain Previously Allocated Addresses

Acquisition of previously allocated IPv4 addresses by whatever means is a strategy with currently unknown (but definitely limited) viability. It is also impossible to estimate in advance the cost of such an approach, so it does nothing to minimize business risk. Acquiring previously allocated addresses may provide a short-

term tactical solution whereby a relatively small number of addresses are required urgently to address a specific need. It is not a solution that has the potential for long-term network business growth. It is likely that previously allocated blocks acquired by whatever means will be small and that obtaining lots of contiguous small blocks may be impossible. This would inevitably lead to operational complexity and associated cost for the network operator taking this approach. In other words, except as a short-term solution, it is operationally unsustainable.

Deploy Large-Scale Address Sharing and Allocate Private Addresses

In light of the two criteria for judging solutions to the IPv4 address shortage that we have identified, so-called carrier- grade network address translation (NAT) proposals—otherwise known as CGN (I-D.nishitani-cgn)—raise several issues. Centralization of NAT functionality in the network core may reduce the abilities of end users to deploy applications as they wish without support from the network operator. This means that unadorned CGN solutions may struggle to meet the first criterion. Providing mechanisms for end users to control their treatment by the CGN may go some way toward mitigating that concern; however, those mechanisms would need to be very carefully engineered to avoid raising additional scalability and resilience concerns of their own. CGNs may create a single point of failure for all of their clients, and they may decrease the resilience of the network from an end user's perspective. CGN implementations may also struggle when considering the second criterion, as there is no requirement to make use of IPv6 technology as part of the solution. For these reasons there is a real risk that CGNs will do nothing to advance the state of IPv6 deployment; in fact, they will serve only to degrade the utility of the current network. However, for ISPs that don't have any control over their customer provided equipment (CPE), CGN is an obvious and flexible solution for continuing to provide IPv4 service post-runout.

While the subject of CGN deployment has arisen recently in the context of IPv4 address depletion, some operators, particularly mobile network operators, have long histories of allocating private addresses to their subscribers. Recent discussions have indicated that the increasing sophistication of both mobile handsets and the applications that run on them is driving operators of mobile networks toward public addressing solutions, including IPv6 deployment, to improve scalability and to minimize operating expenses. This suggests that those operators with real-world experience of CGN technology are already choosing to migrate away from it as a solution to their addressing needs.

Improving on CGN

How could we do better? There are proposals currently in the IETF that attempt to address one or both of the criteria identified earlier. These alternative proposals use IPv6 as a transport substrate for the legacy traffic [I-D.durand-softwire-dual-stack-lite]—thereby motivating IPv6 deployment—and may also ensure that control over the fate of end-user applications be kept as close to the end user as possible by distributing the NAT functionality toward the CPE [I-D.ymbk-aplusp]. However, some reduction of utility for IPv4-speaking Internet users is unavoidable in the future. It is inevitable that a reduced number of ports will be available for individual end-user applications. Operation of servers on well-known ports will most probably be an activity that is restricted to users willing to pay premiums for higher tiers of service contract. These may turn out to be good incentives for end users to migrate to IPv6.

Issues with Shared-Address Solutions

A number of proposals that came up for discussion as part of the Sharing of an IPv4 Address (shara) BoF meeting at IETF 74 in San Francisco rely on the concept of address sharing across multiple subscribers in order to achieve their goals. These proposals include carrier-grade NAT [I-D.nishitani- cgn], Dual-Stack-Lite [I-D.durand-softwire-dual-stack-lite], NAT64 [I-D.bagnulo-behave-nat64], IVI [I-D.baker-behave-ivi], Address+Port proposals [I-D.ymbk-aplusp] [I-D.boucadair-port- range], and SAM [I-D.despres-sam]. In many operator networks today, a subscriber receives a single public IPv4 address at the subscriber's home or small business. Within that home or small business there is a NAT function that translates private addresses (RFC 1918 addresses) issued from devices within the home. All of those devices share the single public IPv4 address, and all are associated with a single small set of users and a single operator subscriber account. With the new proposals, a single public IPv4 address would be shared by a number of homes or small businesses, such as multiple subscribers, so the operational paradigm described earlier would no longer apply. All of the previously described proposals share a number of technical or operational issues, and these are addressed in the subsections that follow.

> The end-to-end principle is, arguably, the fundamental principle of the Internet architecture. In a sense, the Internet is the embodiment of the principle.

Continued on next page

Address Sharing, continued from page 9

Fragmentation and Broken Applications

Address sharing has the potential to break a wide range of applications, such as applications that establish inbound communications, carry port information in the payload, carry address information in the payload, use well-known ports, do not use ports (Internet Control Message Protocol, or ICMP), assume uniqueness of IP addresses, or explicitly prohibit multiple simultaneous connections from the same IP address.

In addition, IP fragmentation will require special handling.

Port Distribution, Port Reservation, Port Negotiation

When we talk about port numbers, we need to make a distinction between outgoing connections and incoming connections. For outgoing connections, the actual source port number used is usually irrelevant. But for incoming connections, the specific port numbers allocated to customers matter because they are part of external referrals; in other words, third parties use them to contact services run by the customers. It is desirable to make sure those incoming ports remain stable over time, which is challenging because the network does not know anything in particular about the applications that it is supporting. The network has no real notion of how long an application or service session will remain ongoing and, therefore, how long it will require port stability. According to actual measurements, the average number of outgoing ports per customer is much, much smaller than the maximum number of ports a customer can use at any given time. However, the distribution is heavy tailed, so there are typically a small number of subscribers who use a very high number of ports [CGN_Viability]. This means that an algorithm that dynamically allocates outgoing port numbers from a central pool is much more efficient than are algorithms that statically divide the resource by pre-allocating a fixed number of ports to each subscriber. Similarly, such an algorithm should be better able to accommodate users wishing to use a relatively high number of ports. Early measurements also seem to indicate that, on average, customers use very few ports for incoming connections. However, a majority of subscribers accept at least one inbound connection. That means it is not necessary to pre-allocate a large number of ports to each subscriber. It is possible, however, to either pre-allocate a small number of ports for incoming connections or do port allocation on demand when the application wishing to receive a connection is initiated. The bulk of ports can be reserved as a centralized resource shared by all subscribers using a given public IPv4 address.

> Recent discussions have indicated that the increasing sophistication of both mobile handsets and the applications that run on them is driving operators of mobile networks toward public addressing solutions, including IPv6 deployment, to improve scalability and to minimize operating expenses.

A potential problem with this approach occurs when one of the subscriber devices behind such a port-shared IPv4 address becomes infected with a worm, which then quickly sets about opening many outbound connections in order to propagate itself. Such an infection could rapidly exhaust the shared resource of the single IPv4 address for all of the connected subscribers. The poor network hygiene of one subscriber now threatens the connectivity for all of the immediate network neighbours.

Connection to a Well-Known Port Number

Once a port-address-mapping scheme is in place, connections to well-known port numbers will not work in the general case. Given sufficient incentives, a workaround, such as redirects to a port-specific URL, could be deployed. Some of the existing proposals for application-service-location protocols would provide a means for addressing this problem, but historically, those proposals have not gained much deployment traction.

Universal Plug and Play

Using the Universal Plug and Play (UPnP) semantic, a client asks, "I want to use port number X. Is that OK?" The answer is yes or no. If the answer is no, the client will typically try the next port until either it finds one that works or, after a limited number of attempts, it gives up. To date, UPnP has no way to redirect the client to use another port number, although UPnP IGD 2.0 will most likely fix this for new or upgraded devices. Network addressing translation–port mapping protocol (NAT-PMP) has a better semantic, thereby enabling the NAT to redirect the client to an available port number.

Security and Subscriber Identification with IPv4

Nowadays, a report of abuse is usually in the following form: "IPv4 address x has done something bad at time t0." This is not enough information to uniquely identify the subscriber responsible for the abuse when IPv4 address x is shared by more than one subscriber. This particular issue can be fixed by logging port numbers, but the operations support system will still require updates to deal with service activation, subscriber profile management, and lawful interception.

A number of application servers on the network today log IPv4 addresses in connection attempts to protect themselves from certain attacks. For example, if a server sees too many log-in attempts from the same IPv4 address, it may

decide to put that address in a penalty box for a certain time. If an IPv4 address is shared by multiple subscribers, this would have unintended consequences in a couple of ways. First, it may become the natural behaviour to see many log-in attempts from the same address because it is now shared across a potentially large number of users. Second, and more likely, one user who fails a number of log-in attempts may block out other users who have not made any previous attempts but who will now fail on their first attempt. Moreover, the assumption that a single IPv4 address maps to a single user may be used for other purposes, such as geolocation or counting the number of individual users of a service. All of those things may become more complicated when several subscribers share an IPv4 address at the same time.

Port randomization, which is used for mitigation of blind attacks against established transport connections, will have reduced effectiveness as port entropy gets reduced. In addition, good randomization functions on the operating system may be defeated by nonimplementation on address-sharing CPE.

To some extent, the problems of shared addressing are already with us due to the prevalence of dynamically assigned addresses to domestic broadband subscribers and the use of CPE NAT. However, the point here is that the widespread adoption of port-shared addresses by service providers will make those complications considerably more widespread and more severe.

Concluding Remarks

As we approach the completion of IPv4 address allocations from the IANA, there are various options available to service providers. Of those options, some of the shared-address solutions seem to offer an approach consistent with the long-term goal of IPv6 deployment and to provide maximal preservation of the end-to-end principle. Nevertheless, it must be emphasized that all shared-addressing solutions have a number of common and potentially serious issues. Address sharing among multiple subscribers inevitably will result in a degraded experience of the network for many users, as well as increased operating costs for ISPs. Content providers are encouraged to consider carefully the potential impact of shared addressing on their business and operational practices.

> Operation of servers on well-known ports will most probably be an activity that is restricted to users willing to pay premiums for higher tiers of service contract. These may turn out to be good incentives for end users to migrate to IPv6.

Acknowledgements

This article was largely inspired by conversations that took place as part of an Internet Society–hosted roundtable event for operators deploying IPv6. The participants in that discussion were John Brzozowski, Leslie Daigle, Wes George, Christian Jacquenet, Tom Klieber, Yiu Lee, and Kurtis Lindqvist.

References

[CGN_Viability] Alcock, S., "Research into the Viability of Service-Provider NAT," 2008.

[I-D.bagnulo-behave-nat64] Bagnulo, M., Matthews, P., and Beijnum, I., "NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers," draft-bagnulo-behave-nat64-02 (work in progress), November 2008.

[I-D.baker-behave-ivi] Li, X., Bao, C., Baker, F., and Yin, K., "IVI Update to SIIT and NAT-PT," draft-baker-behave-ivi-01 (work in progress), September 2008.

[I-D.boucadair-port-range] Boucadair, M., ed., "IPv4 Connectivity Access in the Context of IPv4 Address Exhaustion," http://tools.ietf.org/html/draft-boucadair-port-range-01, January 2009.

[I-D.despres-sam] Despres, R., "Stateless Address Mappings (SAMs) IPv6 & Extended IPv4 via Local Routing Domains—Possibly Multihomed," draft-despres-sam-01 (work in progress), November 2008.

[I-D.durand-softwire-dual-stack-lite] Durand, A., Droms, R., Haberman, B., and Woodyatt, J., "Dual-Stack Lite Broadband Deployments Post IPv4 Exhaustion," draft-durand-softwire-dual-stack-lite-01 (work in progress), November 2008.

[I-D.nishitani-cgn] Nishitani, T., Miyakawa, S., Nakagawa, A., and Ashida, H., "Common Functions of Large Scale NAT (LSN)," draft-nishitani-cgn-01 (work in progress), November 2008.

[I-D.ymbk-aplusp] Maennel, O., Bush, R., Cittadini, L., and Bellovin, S., "The A+P Approach to the IPv4 Address Shortage," draft-ymbk-aplusp-02 (work in progress), January 2009.

[IPv4_Report] Huston, G., "IPv4 Address Report," 2009, http://www.potaroo.net/tools/ipv4/index.html.

[RFC1958] Carpenter, B., "Architectural Principles of the Internet," RFC 1958, June 1996.

[RFC3724] Kempf, J., Austein, R., and IAB, "The Rise of the Middle and the Future of End-to-End: Reflections on the Evolution of the Internet Architecture," RFC 3724, March 2004.

Note: For another perspective on address sharing, see the article entitled NAT++: Address Sharing, by Geoff Huston, which appeared in the April 2009 edition of the ISP Column. http://ispcolumn.isoc.org/2009-04/sharing.html.

ISOC Fellows at IETF 74

By Wendy Rickard

Four Internet and network technologists from Africa and South America travelled to San Francisco to attend their first IETF meeting, a trip made possible by the Internet Society as part of their Fellowship to the IETF Programme. Now in its fourth year, the programme helps stimulate Internet growth in developing nations by immersing technologists from those regions in the work being done by the IETF. Fellows are given the opportunity to hone their technical skills and to become more engaged in the standards-development process.

The fellows attending IETF 74 were João Marcelo Ceron of Brazil, Coko Tracy Mirindi Musaza of the Democratic Republic of Congo, Blessings Msowoya of Malawi, and Noah Sematimba of Uganda. Each fellow is paired with a mentor, typically a seasoned IETF participant who can help the fellow navigate the meeting and who can answer questions (see box, this page).

João Marcelo Ceron

As a network manager at the Federal University of Rio Grande do Sul in Brazil, João Marcelo oversees device and protocol configuration as well as management of the Internet exchange point. He also conducts research in the areas of network management and network security, focusing on empirical experiences and problems he encounters in his professional activities. The results of his research are disseminated in presentations and papers, which, he says, helps other network operators who are working with similar issues. In 2008, he presented a paper at the LACNIC (Latin American and Caribbean Network Information Center) conference exploring the limitations of and the potential solutions for BGP 4 management.

João said the RFCs that the IETF generated are important resources for network administrators like himself, mainly because they assist with problems that come up on a daily basis. They also help him understand the characteristics of various protocols.

João is interested primarily in network management issues, such as the work being done within the Inter-Domain Routing working group. He plans to disseminate the knowledge he gained at IETF 74 through presentations to other network managers at his university and to use what he learned in his postgraduate study programme. He also plans to write an article about the experience for publication in his university's journal.

Coko Tracy Mirindi Musaza

In his position as IT consultant at the African Network Operator Group (AFNOG), IETF fellow Coko Tracy specializes in networking and scalable Internet services under FreeBDS, GNU/Linux, and Debian. He also serves as an instructor for AFNOG, teaching the fundamentals of scalable network infrastructures to students and professionals.

His interest in wireless protocol standardization led Coko Tracy in February 2009 to the 2nd Awareness Workshop on Relevance of Low Cost Wireless ICT Solutions. Since then, he has been working on a project to build a wireless communities association in the Democratic Republic of Congo. In addition, he is working with MHDeafCafNet, a wireless network in central Africa, to help connect associations of deaf people and people living with disabilities in the region (see http://www.mhdeafcaf.org). As part of those efforts, Coko Tracy has learned to build antennas that will be used in his projects. "Often, spending money for purchasing antennas makes no sense when I can build them myself and teach others to do so," he said.

Through the IETF, Coko Tracy wants to deepen his understanding of the standardization of wireless protocols. He had subscribed to the Control And Provisioning of Wireless Access Points (CAPWAP) mailing list, but without assistance or a mentor, he found it difficult to understand how the working group was getting things done. Having attended an IETF meeting helped quite a bit. Eventually, he would like to contribute ideas to the working group, as well as to participate in writing a draft for standardization of wireless protocols.

In the meantime, he plans to use what he learned at IETF 74 in his work with the wireless communities association in the Democratic Republic of Congo as well as in his other projects.

Blessings Msowoya

The Malawi Sustainable Development Network is a United Nations Development Programme–supported government Internet Service Provider (ISP) that is executed by the National Research Council of Malawi. The ISP assists in the development of

IETF 74 Fellows and Mentors

Coko Tracy Mirindi Musaza (Democratic Republic of Congo)
Mentor: Margaret Wasserman

Blessings Msowoya (Malawi)
Mentor: Mat Ford

Noah Sematimba (Uganda)
Mentor: John Schnizlein

João Marcelo Ceron (Brazil)
Mentor: Hugo Koji Kobayashi

Returning Fellows

Burmaa Baasansuren (Mongolia)
Mohibul Hasib Mahmud (Bangladesh)
Asim Zaheer (Pakistan)


João Marcelo Ceron


Coko Tracy Mirindi Musaza


Blessings Msowoya


Noah Sematimba

Internet and information services, with an emphasis on sustainable development. As a network engineer for the Malawi Sustainable Development Network Programme (SDNP), Blessings is responsible for supervising network design staff as well as for network implementation and maintenance. Previously, he led a technical team that worked on the Malawi Internet Exchange, and he's on another team that manages the .mw top-level domain.

According to Blessings, it is imperative that the technical section of the Malawi SDNP "be in sync with current best practices," which includes any RFCs that impact them directly. "We provide long-distance wireless connections in remote areas; hence, following best practices and standards is very important for things like deciding which frequencies and bands to use and whether to use public or private address space," he said.

At IETF 74, Blessings was impressed with the discussions that focused on IPv4 and IPv6 issues. The Malawi SDNP has an experimental IPv6 network in which it tests basic Internet services, such as the Domain Name System, firewalls, and mail services. Blessings's attendance at the meeting, he wrote, "will help us increase the services that are running on the IPv6 network." In addition, he will be able to offer guidance to colleagues on current best practices and standards that contribute to the growth of Internet services in Malawi.

Noah Sematimba

IETF fellow Noah Sematimba has been working in the networking and IT fields for the past eight years, including positions with Africa Online, MTN Uganda (the largest telecom in Uganda), and, currently, with Warid Telecom Uganda, where he serves as assistant manager for IT systems. Noah played a key role in setting up the Internet exchange point in Uganda, and he still serves as technical chair of the Uganda Internet exchange point.

Noah says his work with the .ug country code top-level domain drives his interest in the DNSSEC as well as in IP-related developments. He is taking part in his organization's deployment of IPv6, and a lot of the work he is doing is deeply influenced by work being done by the IETF.

Since 2002, when he first began subscribing to the mailing lists, Noah has followed with great interest the activity in the namedroppers and dnsops working groups. Attending an IETF meeting offered him the opportunity to meet many of the people involved in creating standards and to get involved in the process in a way he could not otherwise. He said that building relationships with key people would facilitate his ability to work with those people on future projects.

Since attending IETF 74, Noah says he plans to take a much more active role in the IETF. He also plans to promote the work of the IETF to his peers in Uganda as well as to encourage more active participation in the IETF.

The Internet Society extends its deepest gratitude to its ISOC Fellowship to the IETF Programme sponsors: Afilias, Google, Intel, Microsoft, and Nominet Trust.

Sponsorship to assist future fellowship programme fellows is strongly encouraged. In addition to demonstrating an organization's commitment to technical capacity building and leadership development in less-developed regions, sponsorship affords an organization a range of sponsorship benefits. For information on how to become a sponsor and to learn how sponsorship can benefit your organization, visit http://www.isoc.org/educpillar/fellowship/sponsorship.shtml or e-mail fellow-sponsor@isoc.org.

"We provide long-distance wireless connections in remote areas; hence, following best practices and standards is very important for things like deciding which frequencies and bands to use and whether to use public or private address space."

— Blessings Msowoya

The Seven Stages of IPv6 Adoption

More than 10 years have passed since RFC 1883—the document that outlined the IPv6 specification—was finalized. Yet even with the depletion of IPv4 addresses looming large on the horizon, IPv6 adoption remains surprisingly, if not persistently, low. With the need for adoption and deployment growing more urgent, both the Internet Society and the Internet Engineering Task Force have been working on ways to raise awareness of the importance of IPv6 for the continued growth and functionality of the Internet.


Photo/Internet Society
Attendee asks a question to the IPv6 panelists

In an effort to bridge the engineering and the rest of the IP-address-dependent world, the Internet Society hosted a panel discussion in March 2009 in conjunction with IETF 74 for the purpose of presenting a wide range of perspectives on IPv6 adoption. The panel, entitled The Seven Stages of IPv6, outlined the opportunities made available by IPv6 from the perspective of network citizens who are in the seven stages of dealing with the enormity of change.

Discussion was wide-ranging, but certain key messages emerged:

- IPv6 is ready for deployment, and deployment is as straightforward as any network technology rollout.
- Even as the general uptake is fairly slow, there are important pockets of IPv6 deployment, demonstrating movement.
- The alternative to deploying IPv6 is not "leaving the network as it is"; the nature of the IPv4 network is changing in response to the lack of available addresses.

Where We Are

For more than a decade, the Internet development community has been aware that in the long run, IPv4 will not be capable of providing enough addresses to allow each machine on the network to have its own address. By 1995, work on IPv6 was completed. Today the real task is to facilitate the spread of IPv6 uniformly across the global Internet.


Photo/Internet Society
ISOC staff preparing for the IPv6 session

IPv6, according to moderator Leslie Daigle, isn't the question; it is the answer. "The question is, Do we want to continue to have an Internet that continues to be expanded by innovations? If that's the case, we need to deploy IPv6," she said.

Back when IPv6 was being finalized, it was thought that the transition strategy would be dual stack: a network supporting both IPv4 and IPv6 simultaneously. According to panellist Russ Housley, that meant engineers would start incorporating the new technology onto the old technology, and "once everybody was able to communicate over IPv6, we could start disabling IPv4, and everything would transition smoothly." The strategy didn't work as planned, even though those who implemented IPv6 felt that it worked fine.

Today the Regional Internet Registries (RIRs) are issuing 12 IPv4 /8s per year, and the distribution rate is not slowing down. In fact, according to panellist Richard Jimmerson, it's picking up, particularly in such regions as Asia Pacific, where there are significant numbers of underserved areas and great demand for IPv4 address space. At the end of 2008, there were 34 /8s remaining at the Internet Assigned Numbers Authority (IANA) to be allocated to the RIRs. "At the current rate, it is expected that the remaining /8s will last approximately two years," Richard said.

Even as the time to address pool exhaustion approaches, the IETF continues to develop new tools to bridge IPv4 and IPv6. "When the IETF puts out a specification, we don't just forget about it," panellist Jari Arkko said. "We actually care a lot about continued accuracy and maintenance of the specifications." However, that's only a small part of the overall effort. The real challenges are deployment and working on new features. For the most part, said Jari, these are things like IP diagnostics in both IPv4 and IPv6.

Is IPv6 the Question—or the Answer?

The panellists agreed that the ways people use the Internet today are much different from before—and much different than anyone doing development work then might have anticipated. Today, as Russ pointed out, people expect to be able to carry in their pocket a device that is always on and always connected. "That means we need an address space that allows every device to be always on and always connected," he said. "Within the IETF, there is a working group devoted to low-power, battery-operated devices that are in your home, on your desk, or even on your thermostat." If every house in the world were equipped with such equipment, the number of addresses needed would exceed the space that was ever available in IPv4.

Moreover, there is no way of knowing what new applications are on the horizon. "We didn't know in the 1980s that all these things were going to come along," said Lorenzo Colitti. "So now we have two choices: either we can stay with the

original architecture of the Internet—where you can deploy applications by simply deploying a machine here and a machine there and having them talk to each other, which allows the Internet to continue to operate as a communications medium—or we can choose to deploy NAT [network address translation], which would fundamentally change the architecture of the Internet. We don't know what the future applications will be. The sky's the limit."

IPv6 Resistance

Much of the discussion focused on why motivation to adopt and deploy IPv6 is so low and why denial over the need to do so is so high. Many of the panellists pointed to a lack of economic incentive, which may be true in some respects, but that doesn't mean there aren't benefits. As panellist Alain Durand pointed out, an economic incentive means, "If I deploy something to work with this technology, I will benefit." Every technology that has been successfully deployed in the past 15 years has been incremental, he said.

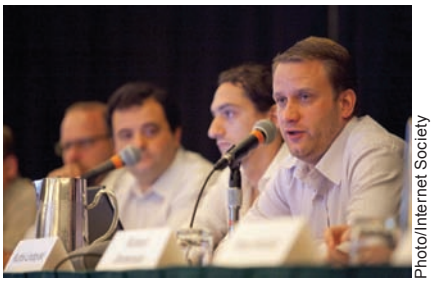

Photo/Internet Society

Kurtis Lindqvist speaking during the IPv6 panel session

nor additional revenue streams for businesses and organizations.

Perception could also be part of the problem. "The Internet of today is much different than it was when IPv6 was first developed," said Kurtis, "and it is being deployed today in much different ways than was originally anticipated." That means the migration from IPv4 is also going to be different from what we might have expected. The Internet, he said, is not homogenous; it is different depending on what part of the world you're in. Similarly, different organisations are in different stages of IPv4 depletion and IPv6 adoption.

> "The question is, Do we want to continue to have an Internet that continues to be expanded by innovations? If that is the case, we need to deploy IPv6."
>
> — Leslie Daigle

Panellist Kurtis Lindqvist looked to other factors that are delaying adoption and deployment of IPv6, such as the time it takes to deploy, the fact that it is not backward compatible, and the reality that it offers neither new features nor additional revenue streams for businesses and organizations.

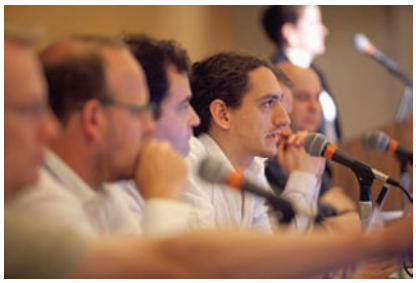

Photo/Internet Society

Lorenzo Colitti speaking at the IPv6 Panel session

Regardless of the ancillary explanations, both the panellists and the audience members kept returning to the idea that the primary obstacle to the adoption of IPv6 is a lack of economic incentive. That assumption is partially supported by the results of a survey recently published by the Internet Society. (See http://www.isoc.org/pubs/2009-IPv6-OrgMember-Report.pdf.) As part of the results, a number of ISOC's Organization Members say they do not see a specific business case for IPv6, although they recognize that their customers are demanding it. "We're hearing a lot of people talking about managing large-scale NAT devices," said Lorenzo. "NATs are expensive and difficult to maintain." In fact, he claims that most of those who have deployed IPv6 would agree that doing so is much simpler than deploying layers of NAT. "It is refreshingly simple to look at a network with only global addresses and have it work the way it should," he said.

The good news is that the depletion of IPv4 address space is not like running out of oil. "It's not as if all of a sudden

Continued on next page

IETF 74 IPv6 PANEL

Jari Arkko
Researcher
Ericsson Research

Sebastián Bellagamba
Manager, Regional Bureau for Latin America and the Caribbean
The Internet Society

Lorenzo Colitti
Network Specialist
Google

Leslie Daigle (moderator)
Chief Internet Technology Officer
The Internet Society

Alain Durand
Director of IPv6 Architecture and Internet Governance
Office of the Chief Technology Officer
Comcast

Russ Housley
Chair
Internet Engineering Task Force

Richard Jimmerson
Chief Information Officer
American Registry for Internet Numbers

Kurt Erik "Kurtis" Lindqvist
Chief Executive Officer
Netnod Internet Exchange in Stockholm

The Seven Stages of IPv6 Adoption, continued from page 15

no cars will be able to run and you can't drive to work the next morning," Alain assured the audience. "Everything that has been deployed still works." However, knowing that your computer will still work when IPv4 addresses are depleted does not mean that adoption of IPv6 isn't a necessary change.

In the developed world, the obstacles to change are directly tied to economics or are consequences of avoidance. In the developing world, the obstacles have more to do with whether governments support an information economy and how much they're willing to invest to make it happen. In those regions, the governments are only just beginning to become aware of the IPv4-to-IPv6 issues. So, when it comes to the seven stages, panellist Sebastián Bellagamba says, in Latin America and the Caribbean, "we have not even begun to enter them."

Throughout the developing world, governments are in nascent stages of understanding the new roles they play with regard to technology development and the future of their countries' economies and societies. As Sebastián said, not only do governments today act as regulators; they also are heavy users and often even service providers. For most of those governments, the Internet and related technologies represent a way to attract development and economic growth to their countries. In Argentina, for example, the Internet has become so important that 75 percent of the internal revenue passes through it. "So in that case, if something happens to the Internet," Sebastián said, "the income of the government is affected."

Photo/Internet Society

IPv6 panellists and audience members mingle

Overcoming Denial

According to Lorenzo, at Google, it began when a couple of people started to deploy IPv6 as a small project and then a pilot network was built. Once the network was up, they saw how the applications followed. "We did it in stages," he said. "The principle that guided us, which I strongly believe is good for deployment, was that it doesn't have to be as capable as your IPv4 stack on Day One. The traffic levels are not comparable. However, it does have to be done properly, and it has to be production ready and supported. It has to be designed according to the same quality standards that you would meet for any other kind of technology infrastructure. Otherwise, it is of no use to anyone."

Lorenzo encouraged those who have a production-ready IPv6 network to talk to Google because the company can provide all Google content and services over IPv6. "That means you not only get to use your IPv6 network; you also get to find out what the problems are, if there are any problems," he said. "Also, you get to find out if other people are implementing it, and you get to be able to say, 'Yes, we do support IPv6 on our network.'"

However, Lorenzo also cautioned the audience to be aware that traffic will appear overnight. "When you do large deployments," he said, "it will just appear out of nowhere. There is no organic growth. We turned IPv6 on for Google maps, and we saw a threefold increase overnight."

For those who are hesitating, the unmistakable message was that, with a few exceptions, IPv6 is fine and ready to be deployed. The challenges, however, are real. From the perspective of the service provider, they are what Alain refers to as "the two long-tail problems of IPv4." The first long tail is what is happening in the home. "It's not only about whether

Photo/Internet Society

Russ Housley presenting during the IPv6 panel

or not Windows supports IPv6," he said. "It's also about the latest gadgets. Today there are cameras with WiFi interfaces that can upload pictures to the Web." Those may be nice services, but all of them are implemented with IPv4, and unfortunately, those devices do not upgrade to IPv6. The same thing is true of the 60-inch television with the cable modem integrated and the software that allows a user to browse the Internet. "That's all IPv4, not IPv6," he said, "and that's a problem."

The second long tail is what is going to happen with content. "My thanks to Lorenzo for getting Google on IPv6," said Alain, "but what about the second tier of Web services, such as news agencies? What about the third tier of Web services, such as small shops?" Eventually, all of those will migrate to IPv6, but it will take some time. Turning on the IPv6-only service is not going to serve the needs of customers who have devices that work only with IPv4, nor will it serve the needs of customers who want to access content that is available only with IPv4. "When dealing with those realities," Alain suggests, "perhaps what we need is a two-pronged approach." The first part involves embracing IPv6 and getting as many endpoint devices on it and as much traffic as possible to it. The second part has to do with realizing that the IPv4 world cannot be abandoned. "It's not as if we move to a new world and the old world becomes lost," he said. "No. We need a bridge between the IPv4 and the IPv6 worlds."

The Road to IPv6 Adoption

In addition to efforts by the IETF and the Internet Society to bring the issue to the fore in both technology and policy venues, the American Registry for Internet Numbers (ARIN) and the other Regional Internet Registries (RIRs) have engaged in awareness campaigns throughout the world. According to Richard, ARIN began going to trade shows in 2006, exhibiting, giving presentations, and talking to people about IPv4 depletion and IPv6 adoption. What the RIRs encountered, according to Richard, really were, as Leslie described, disbelief and denial. "The people we talked with did not believe we would run out of IPv4 addresses, and they were not interested in IPv6," he said.

Recently, there has been a noticeable shift, and Richard says he believes the audience is becoming much more receptive to the IPv6 message. That could be due to the Internet community and the technical community's coming together to work toward widespread IPv6 adoption, which would be one way of assuaging the fears and overcoming denial. "They seem to be working towards acceptance," Richard said. Timing also plays a role. "No one really does anything before they have to," said Jari.

With regard to creating incentives, one journalist asked whether avoidance of future costs would qualify as an economic incentive. "I think people do things if there's a reason to do it,"

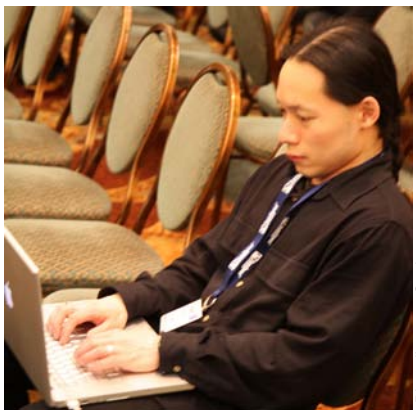

Photo/Internet Society

IETF participant waiting for a session to start

Richard responded. It's important to remember that there is no master plan for deploying IPv6. As Leslie said, for most of the technologists who work in an IETF-like, multistakeholder environment, that's a feature, because it means different pools can develop at their own rate. As the survey made clear, there is no direct or concise business incentive for moving to IPv6, but customers are asking for it. "That *but* part of the comment is important for understanding the entire context," said Leslie. "It goes back to the issue that what customers actually want are applications that work. That means they want continued global addressing in the network, which means, at this point, IPv6."

To understand that role of the marketplace, it might also be important to distinguish between business incentives and business drivers. "I think that, in some ways, if there were better business drivers, we would pay for some implementation of IPv6," said Sebastián, who also said the economic incentives are there, particularly among countries that depend on the Internet for economic growth. In many places in the world, he added, "they need more addresses in order to grow. The only addresses they can get are IPv6 addresses. Therefore, there is some clear economic incentive in that area."

What can governments do to promote the transition to IPv6? According to Sebastián, the first thing they need to do is to address it. Sebastián reminded the audience that IPv6 is not a purely regulatory issue. "It is an issue that has to be addressed by governments together with the private sector," he said.

Moving toward Acceptance

One of the benefits of the way the migration from IPv4 to IPv6 has evolved is how the migration has helped the technical community understand a lot more about how the Internet works, even about how IPv4 networks work. "As we deploy IPv6," Kurtis said, "we need to go back and change some of our original ideas about how to do this. In the meantime, we are acquiring valuable operational experience." Similarly, the technical community needs to continue working on a transition strategy that will be seamless and invisible to the end user. "It shouldn't matter to you whether your Web site or e-mail is sent over IPv4 or IPv6," he said.

Coming to acceptance where IPv6 is concerned may be a struggle, but if the Internet is to continue to evolve, it is a necessary struggle. The shepherding of organizations, service providers, and governments toward acceptance has become a priority to organizations such as the IETF and the Internet Society. As Leslie explained at the beginning of the discussion, the Internet Society is the organizational home of the IETF, but it has a broader mission to promote the evolution and use of the Internet for the benefit of all people throughout the world. "From that perspective, we believe that the IETF has an important role to play with respect to IPv6, and not only in terms of developing the specification," she said.

Beyond the efforts to raise awareness, many of the panellists are simply true believers, and they hope the message will inspire action. "If we want the Internet to be around in three to four years in its current state, then we want to use IPv6," Lorenzo said. "It will allow the Internet to continue to function as we know it. And it will keep the Internet open."

More articles on this topic can be found at http://www.isoc.org/pubpolpillar/docs/ipv6-way-forward.pdf and http://www.isoc.org/pubpolpillar/docs/ipv6-government-role.pdf

Bringing OAuth to the IETF

The IETF Journal meets with OAuth experts Hannes Tschofenig, Blaine Cook, and Eran Hammer-Lahav to discuss the decision to bring OAuth to the IETF and the future of the Internet's latest security specification.

At IETF 74 in San Francisco, *IETF Journal* editor Mirjam Kühne and Trent Adams (Outreach Specialist, Identity Community, at the Internet Society) sat down with OAuth BoF cochairs Hannes Tschofenig and Blaine Cook as well as Eran Hammer-Lahav, who authored the OAuth specification document, to find out more about the decision to bring OAuth into the Internet Engineering Task Force, about how the specification compares with similar resources, and about next steps in its development and application.

IETF Journal: There are a number of resources available today that address the growing need to manage and protect a user's identity on the Internet while making it possible for users to share information from one site to another. How is OAuth related to other work in the identity space?

Eran: There are a few different mechanisms. One is SAML [security assertion markup language], and the other two are OpenID and OAuth. There are a few others in the mix, but those are the main ones in the identity space.

IETF Journal: And what was each of them designed to do?

Eran: SAML was designed mainly for use within business enterprises, and it is perhaps the most complicated of the three because of its use of XML structures. It is also very robust and very powerful. OpenID is fundamentally about single sign-on, and it depends on Web redirections. OpenID is meant exclusively for Web usage and is designed for interactions between human beings. Because of its architecture, OpenID can communicate only what will fit in a URI. It is not capable of handling anything more sophisticated.

"OAuth is a way to delegate both access and permission. It is very simple, and in many ways, it borrows from the culture of OpenID in terms of equal access, while at the same time learning from its mistakes."

— Eran Hammer-Lahav

OAuth is a way to delegate both access and permission. It is very simple, and in many ways, it borrows from the culture of OpenID in terms of equal access, while at the same time learning from its mistakes. It is designed to be a generic access mechanism. So, if you have other authentication mechanisms on the Web, this is just one more option in that stack. However, it does provide more options: it can be used from server to server as well as from client to user.

Right now, OAuth is not really a standard; it is primarily a guide to best current practices. The next phase of its development will need to focus on interoperability, an area in which it is not strong at the moment.

IETF Journal: Is that primarily the reason for bringing this into the IETF?

Eran: There were a number of different reasons for doing so. I believe it started when Mark Nottingham, who is the chair of the httpbis WG [Hypertext Transfer Protocol Bis Working Group] and a world-renowned expert on caching and proxy, gave us feedback on the OAuth specification. He said, "You know, this is a really good start, but it doesn't play well with the actual HTTP stack." We put it through a security review, but we did not have the full skill set of an IETF security review. Instead, we had two or three people from two different companies doing a security review. But two people doing a security review is not quite the same thing as a full-scale IETF security area review.

After that, we decided to publish it as an Internet-Draft, which was a way to get feedback from the IETF. Then we thought, maybe we should go to an IETF meeting and do a Bar-BoF. Then we thought, why not just do a BoF? After that, things happened very fast. When we published the Internet-Draft, we asked what track it should be put on. We were told, "Oh, just put it on the

Trent Adams

Blaine Cook


Eran Hammer-Lahav

Standards track. You can always change it later." So, the original motivation to bring it to the IETF was to get some feedback from experts. That was really what all this was about.

As we started talking about it, we were hearing that the IETF has been trying to develop something like this for a long time—without much success, because it is very hard to develop a brand-new security protocol from scratch. In some ways, it is very difficult to get new things off the ground from within the IETF. A lot of work is developed outside and then brought into the IETF.

Blaine: For me there was another aspect: On the consumer side, the adoption was pretty good, but not so much with enterprise adoption. Businesses want to use it, but they are concerned because OAuth is not yet standardized. We would hear, "You guys are just a bunch of Web guys, so how can I trust you?" And they want to do their own security release, because they don't trust themselves [laughs]. I think with these types of protocols, there is a point at which engaging in a proper process with a standards body behind it is really important, especially when it involves things like being able to delegate access from a specific application to my bank, so that I can gain access to your checking account in order to deposit money.

Hannes: Many of the more conservative groups, such as government agencies, often require a standards-track RFC to exist before agreeing to adopt a particular protocol. And it often creates problems for them to come up with the mechanisms to achieve that.

IETF Journal: It sounds like there were two aspects: The first is the evaluation of the current specification and figuring out if there are any issues that might come up during the review, and then, on the backside, the second one is the issue of legitimacy. Therefore, when enterprises ask who has had a look at it, you can raise their comfort level by pointing to the work that has been done on the protocol within the IETF.

Blaine: Yes.

IETF Journal: It also sounds like what you want is not just to release version 1.0 but also to have something that is supported by a long-term process. Is that also a consideration?

Eran: Yes. On the other hand, we also understand that once something gets fed into a Standards track as an RFC, the community cannot expect a new version in six months.

IETF Journal: But what the community can expect is that people will have their eyes on the ball, and that when review is required—whatever the time frame—it will happen.

Eran: At the moment this is a community-driven specification. There are a small number of people who agree that OAuth is stable for now and that they are not going to do any more work on it. Apart from that, there is nothing to prevent anyone from changing it. That is why bigger companies are skeptical. It will take them six to eight months to incorporate the protocol into their development cycle. If the specification changes after six months, it will have been a waste of their effort and money.

"This idea of perfect security is imaginary. We are trying to enable people to take control over what they are doing and make sure that when they take an action there is not some other action happening that they don't know about."

— Blaine Cook

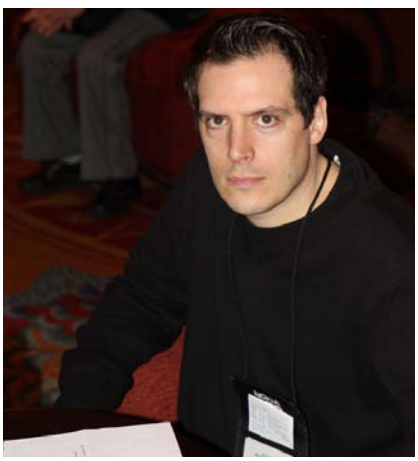
Hannes Tschofenig

I would not say we are bringing OAuth to the IETF; we are attempting to standardize one document as a trial for the community. We are taking this one very specific document to the IETF to have it standardized. We want a more complete security review; we want interoperability; and we want to improve the quality of the document by having more people look at it.

I envision that a lot of the work on OAuth will continue to happen outside the IETF, that fundamentally, it will remain a community-driven process. If at any time the community produces something that is stable and secure enough to be standardized, then we can come back and do the next step. If at that point the WG has completed its work,

Continued on next page

Bringing OAuth to the IETF, continued from page 19

we can see if we want to recharter the group. Or we could start a new WG.

I am really careful of communicating that we are not *moving* OAuth into the IETF. There is the OAuth community and then there is the IETF. But I expect the same people who have worked on it so far will also be the people who will make the RFC happen. It really is not an us-versus-them situation. The OAuth WG will consist of whoever shows up, and I expect that 95 percent of the people who show up will be the same people who have already worked on it.

IETF Journal: With that in mind, are you concerned about the potential for bifurcation of the standard? For example, let's say there is a group of people within the IETF who want to take the standard in one direction, and another group of people in the community who have a use case that they are trying to tackle, which might take them in another direction. How do you decide what direction to go?

Hannes: The decision will be made by rough consensus of the participants of the BoF or the WG.

Blaine: For me it really is about adoption. There are plenty of security standards that haven't achieved adoption because they are designed from within security ivory towers. We are constantly being surveilled, surveilling each other, and surveilling ourselves. This idea of perfect security is imaginary. We are trying to enable people to take control over what they are doing and make sure that when they take an action there is not some other action happening that they don't know about. And designing for that is really hard. So, we already have bifurcation of OAuth, because it is based on a number of different specifications. In the way it is written now, we hope it is pretty neutral and something that everyone can agree on.

IETF Journal: So, one could say that the IETF standardizes OAuth and then people can implement against that?

Eran: The biggest danger is standards shopping. I have no intention of doing standards shopping. I do not intend to standardize OAuth elsewhere if the IETF is taking a direction I don't like. That is also reflected in the way we licensed the work originally and the way we brought it to the IETF: we basically took a snapshot of the document and pretty much said, "Now the IETF can do whatever it wants with it."

Today we already have an OAuth Core 1.0 specification. It may not be a standard, but it is there and companies are implementing it and using it. If the IETF produces a successful standard, it will still be called OAuth. After that, the market will decide if it likes it. It is our hope that the market will like it and that it will move toward adoption. And that the people who have already implemented OAuth Core 1.0 will say they will also want to support the new one. But let's allow the market to decide.

It is also possible that the IETF will move OAuth into a more extreme direction, requiring that all kinds of things be parts of the standard. If that's the case, what will probably happen is that a lot of the original OAuth authors will leave. The only thing we will probably take with us is the name, not the specification or anything else. We will just say, "Do whatever you want, but don't call this OAuth anymore." Why confuse people? You will end up creating something that you think is better, but it will just be different and then we can let the market decide which one it wants.

IETF Journal: Have you given any thought to the intellectual property rights (IPR) issues that could arise from moving OAuth from its original development paradigm to the IETF?

Eran: Yes. I spent six months negotiating a set of very liberal IPR terms with all of the original contributors. So, if you take OAuth and you change it dramatically so that it does not any longer operate within the boundaries of the community-derived OAuth specification, then you're on your own! If that's the case, you need to go and get the IPR requirements you need. But we have actually found the IETF applications area to be very open-source friendly.

IETF Journal: Does that mean that basically there would be no issue if you choose to take the specification down another path at a later time?

Eran: If the IETF is interested in going in a different direction, and if Blaine and I and a few others are the minority voices in the room, and if we cannot raise our voices anymore, then we very likely will just choose to not participate anymore. But nobody is going to take it and bring it to another standards body, such as OASIS. There is no interest in that. If you look at the original goals we stated earlier, creating a standard is not our main objective; it's more like a bonus.

IETF Journal: Do you agree with what Blaine said earlier about how, to some degree, having a standard opens up markets that you otherwise wouldn't be able to reach?

Eran: Yes, sure, but only as long as that doesn't mean selling out on what we were trying to do originally.

> "We are taking this one very specific document to the IETF to have it standardized. We want a more complete security review; we want interoperability; and we want to improve the quality of the document by having more people look at it."
>
> — Eran Hammer-Lahav

Blaine: I don't have any financial interest in OAuth, and as far as I know, nobody in the community is interested in serving as OAuth advisors. No one has any intellectual property claim that would be of any value. I don't need OAuth to go into the IETF. Twitter has OAuth, and Flickr has it, as do others. So, the specific properties I would like OAuth to have are already happening. As far as not having to give out my password on the Web in those applications that I wouldn't really trust with that kind of thing, well, that is done. So, it is really about recognizing that it would be really cool if more enterprises started thinking about these kinds of lightweight but strong security mechanisms.

In terms of fragmentations of the specification and what might happen next, it's important to remember that OAuth could not have happened 10 years ago. The culture on the Net was different. Even two years ago, when we started to write this up, we used MD5, not SHA-1. There were no libraries at the time. Things have changed significantly in this space. Things evolve, and as OAuth gets adopted, we can do more work on it. I fully expect that something better than OAuth will come along in 5 to 10 years, and then we will start to use that. The technology will be better and our comprehension of the problems as the wider community—not just the security community—will be better.

IETF Journal: Was there a moment when you had to decide which community could give you the best security review and the highest level of legitimacy? Were there other players on the table? Or did you just stumble onto the IETF?

Eran: I have been approached by a colleague who is active in the IETF, and then by Lisa Dusseault, who is one of the applications area directors to officially bring the question to the IETF to see if the IETF might be interested. We had a BoF in Minneapolis at IETF 73, where we asked ourselves if this is an area that may interest the IETF. After that, the question was, Is OAuth a good and reasonable solution to use as a starting point? Those were the two main questions we started off with, and then we started working on the charter. People have also approached me informally from both the W3C and OASIS.

"I fully expect that something better than OAuth will come along in 5 to 10 years, and then we will start to use that. The technology will be better and our comprehension of the problems as the wider community—not just the security community—will be better."

— Blaine Cook

IETF Journal: So, it sounds like you did not step back and say, "Here are the three things I need, let's look at which standards organization would be the most appropriate, is that right?

Eran: That's right. However, after I started talking to Lisa, I did look around a bit because this is something you want to do only once and you don't want to discredit yourself. You don't want to go to the IETF and then after the first BoF think, well, I don't like the way this is going in the IETF. I'll go to the W3C and see if they want it.

It was an easy decision based on one criterion—that is, that every single person who has been involved in OAuth until now could continue to work on it in the IETF. Participation at the IETF is open, and there is no membership fee. It is much more difficult to participate in other organizations, such as OASIS and the W3C. It is expensive to become a member.

Morally, I felt that was the right way to approach it. I felt that people might not agree on the need to standardize it or on the value of it or even that the IETF is the best place to do it. But at least they're welcome to join. And all they have to do is join the mailing list. That's the only barrier.

Blaine: If you compare the barriers, for instance, of the W3C or OASIS, it is much more difficult to participate.

Eran: That ruled them out immediately. The W3C membership model is really difficult. And joining OASIS is expensive.

Blaine: To some extent, that's also a concern with the IETF. It may have no membership fees, but the costs of attending the meetings are pretty high. Many people do this in addition to their regular work. Eran is one of the few people who has an employer's approval to be working on it.

IETF Journal: On the other hand, wouldn't it be good to expand the weight of the WG and to broaden the participation? Right now it is pretty U.S.-centric.

Blaine: I agree. I think it's important to meet in Stockholm at IETF 75, which will get more Europeans involved. It would also be a good time to push forward the work on its adoption. At the moment, not many people in Europe know about OAuth.

Hannes: All decisions have to be confirmed on the mailing list anyway. So, the meetings are more like social events.

IETF Journal: Yes, but meeting in person is a good way to overcome potential trust issues with other people.

Hannes: That's true.

Photos by Peter Löthberg


Aaron Falk, IRTF Chair

IRTF Report

By Aaron Falk

What follows are summaries of several updates on the Internet Research Groups (RGs), some of which were reported during the Technical Plenary at IETF 74.

There are three bits of status regarding the Internet Research Task Force (IRTF): First, since IETF 73, the IRTF has not published any new RFCs because of document dependencies that are holding up final establishment of the IRTF RFC publication stream. However, three research group (RG) documents are in the RFC Editor's queue. Additionally, we are finalising the IRTF streams document rights. Our intent is to maximize commonality with the IETF process so as to ease documents' ability to move between the IRTF and the IETF.

Second, there has been activity in the form of the creation of two new research groups: a group of folks, organized by Martin Stiemerling, has been holding BarBoFs to discuss an RG on network virtualization, and Paul Hoffman is developing a draft charter for an RG to discuss alternate public key formats, certificates, and services called PKNG.

And third, most of the IRTF RGs are now fairly active. During IETF 74, four RGs met: DTNRG, RRG, P2PRG, and HIPRG. Most research groups meet at least once a year at an IETF meeting, and several meet more frequently by holding additional meetings at such venues as academic conferences to attract greater research participation.

Recently, I've been giving a very short overview of a couple active research groups during the IETF technical plenary. This is to introduce folks in the IETF to work going on in the IRTF and to encourage more participation. The following two sections are introductions to the Crypto Forum Research Group and the Routing Research Group, as presented during the IETF 74 technical plenary.

Crypto Forum Research Group (cfrg)

The CFRG is a forum for discussing and analysing general cryptographic aspects of security protocols. One of the main goals is to offer guidance on the use of emerging mechanisms and new uses of existing mechanisms in the tradition of RFC 1321 (MD5) and RFC 2104 (HMAC). Another important goal of the work is to create a bridge between theory and practice.

IETF working groups that are developing protocols that include cryptographic elements are welcome to bring questions concerning the protocols to the CFRG.

The CFRG is currently working on a number of important topics. One involves hash functions, wherein the goal is to transition away from MD5 (and SHA-1). In that context, the CFRG is identifying IETF's uses and security goals and is discussing reviving and extending RFC 4270 (Attacks on Cryptographic Hashes in Internet Protocols).

Other topics the RG is working on are Password-Authenticated Key Exchange (currently reviewing draft-sheffer-emu-eap-eke-00, "The EAP-EKE Method"), threshold cryptography (see draft-mcgrew-tss-02, "Threshold Secret Sharing"), and threshold signatures. The last might be a topic of possible future work and is relevant to Domain Name System Security Extensions (DNSSEC) and Public-Key Infrastructure (PKIX).

Routing Research Group (rrg)

The RRG is trying to solve the problem of uncontrolled growth of the routing table. One of the major causes of routing-table growth is multihoming. Multihomed sites inject one prefix or multiple prefixes into the routing system. Routing costs increase with the number of multihomed sites.

The primary goal of the RG is to develop a routing architecture that can provide effective control on routing overhead and that is independent from the number of multihomed sites. Another goal is to avoid the need to renumber when changing service providers. A new routing protocol should also be incrementally deployable and possess equal or better security.

The RRG has continued its efforts to sort out and reassess existing proposals. There are currently nine proposals listed on the RRG wiki. One proposal recently posted argues for an evolution path that will lead to a scalable routing architecture. That same proposal was presented during the RRG meeting at IETF 74. The group is also working to clarify the terminology used in routing scalability discussions. RRG originally planned to offer a recommendation by March 2009, but investigation efforts led to new understandings of the problem and solution space, and now that date has been pushed out by one year, to 2010.

Several members of the RRG participated in a seminar called Naming and Addressing for the Future Internet, which was held in Dagstuhl, Germany, in March 2009. (See http://www.dagstuhl.de/de/programm/kalender/semhp/?semnr=09102.)

For more information about the Internet Research Task Force, visit http://www.irtf.org/.

Recent IESG Document and Protocol Actions

A full list of recent IESG Document and Protocol Actions can be found at http://www.isoc.org/ietfjournal/DocProtoActions0501.shtml.

IETF Meeting Calendar

IETF 75
26–31 July 2009
Host: .SE
Location: Stockholm, Sweden

IETF 76
8–13 November 2009
Host: WIDE
Location: Hiroshima, Japan

IETF 77
21–26 March 2010
Host: TBD
Location: Anaheim, CA, USA

IETF 78
25–30 2010
Host: SIDN
Location: Maastricht, the Netherlands

Register now for

IETF 75

26-31 July 2009

Stockholm, Sweden
http://ietf.org/meetings/75/

Early bird registration: USD 635 (through Friday, 17 July 2009)
Regular registration: USD 785
Full-time students: USD 150 with on-site proof of ID

IETF 75 is being hosted by .se

Special thanks to Juniper Networks for hosting IETF 74

Special thanks to .se The Internet Infrastructure Foundation for hosting IETF 75

The ISOC Fellowship to the IETF is sponsored by

Afilias Global Registry Services, Google, intel, Microsoft, .se The Internet Infrastructure Foundation

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

Afilias Global Registry Services, Alcatel-Lucent, ARIN, Cisco, Comcast, .se The Internet Infrastructure Foundation, Ripe NCC

IETF Journal
IETF 74
Volume 5, Issue 1
June 2009

Published three times a year by the Internet Society

Galerie Jean-Malbuisson 15
1204 Geneva
Switzerland

Managing Editor
Mirjam Kühne

Associate Editor
Wendy Rickard

Editorial and Design
The Rickard Group, Inc.

Editorial Board
Leslie Daigle
Peter Godwin
Russ Housley
Olaf Kolkman
Lucy Lynch

E-mail
ietfjournal@isoc.org
Find us on the Web at
http://ietfjournal.isoc.org

Editor's Note:
The *IETF Journal* adheres to the *Oxford English Dictionary* Second Edition

Internet Society