



A report from IETF 76, November 2009, Hiroshima, Japan. Published by the Internet Society in cooperation with the Internet Engineering Task Force\*

### Inside this issue

- Bandwidth, Transition, Top IETF 76 Agenda .. 1
- The Bandwidth Bandwagon ..... 1
- Message from the IETF Chair ..... 2
- New BoF Meetings .... 2
- Words from the IAB Chair ..... 3
- IETF 76 Facts and Figures ..... 3
- IETF 76 Plenary Report ..... 4
- Google: IPv6 is Easy, Inexpensive..... 12
- ISOC Fellows Enjoy In-Person IETF Experience ..... 13
- RFC Editor in Transition: Past, Present and Future..... 14
- Congestion Exposure: We're All in This Together ..... 18
- The Kantara Initiative for Online Identity: A One-Year Progress Report ..... 20
- IRTF Update ..... 22
- Calendar ..... 24

## Bandwidth, Transition, Top IETF 76 Agenda

**From the Editor's Desk, by Mat Ford**

As regular readers of the *IETF Journal* are aware, Mirjam Kühne moved on during the summer to a new role and new challenges. As managing editor of the *IETF Journal* she will be missed, although her new role will keep her in touch with all things IETF, I have no doubt. Temporarily taking the reins for this edition, I would like to begin by soliciting your input on the future direction of the *IETF Journal*. What works? What doesn't work? What would be of most interest to you and your colleagues? All comments, contributions, suggestions, and feedback can be sent to [ietfjournal@isoc.org](mailto:ietfjournal@isoc.org).

In this issue, the subject of bandwidth on the Internet takes centre stage, with articles reporting on the Internet Society panel event called The Bandwidth Bandwagon (this page) and on the motivations for a very interesting BoF meeting on congestion exposure (see page 18).

Also in this issue is a summary of the plenary sessions, including a review of the informative and entertaining presentation called Internationalization in Names and Other Identifiers given by John Klensin, Stuart Cheshire, and Dave Thaler (see page 4). The plenary also witnessed the presentation of the inaugural Itojun Service Award, and we have a special article commemorating that event (see page 12).

As all contributors to the IETF should be aware, the RFC Editor is in transition as the Information Sciences Institute of the University of Southern California relinquishes a role it has had for 40 years. We could not let this milestone pass unremarked, and Leslie Daigle has provided us with an article based on interviews with some of the key individuals involved in the move (see page 14).



Hiroshima, Japan, site of IETF 76

Photo by Peter Lötberg

*Continued on page 4*

## The Bandwidth Bandwagon

**By Mat Ford**

*A panel discussion at IETF 76 helps shed light on the realities of bandwidth growth, operator responses to a changing landscape, and new, relevant IETF work.*

While the Internet did experience episodes of “congestion collapse” more than 20 years ago, the mechanisms implemented at that time to address the problem have largely stood the test of time. Despite this, rumours of imminent network meltdown are never very far away. In November 2009, the Internet Society organized a panel discussion in Hiroshima, Japan, adjacent to the IETF 76 meeting, for the purpose of making the issues of growing Internet bandwidth accessible to a wider audience—in essence, “pulling the message out of engineering and talking to the real world,” as

*Continued on page 8*



\* The articles published in the IETF Journal are not intended to reflect the opinions or the position of the IETF or the Internet Society.

## Message from the IETF Chair

By Russ Housley

IETF 76, held in Hiroshima, Japan, was a very successful meeting. The work of the IETF remains relevant and energetic, as was clearly demonstrated by the attendance of 1,106 people from 44 different countries. It was a pleasure to see so many enthusiastic people collaborating. Significant progress was made by many working groups.

The WIDE Project was a fantastic host. Led by Jun Murai, the WIDE Project coordinated with 19 other sponsors as well as the city of Hiroshima to provide a warm welcome and an effective meeting venue. Banners were displayed by the city to welcome IETFers. The enormous Christmas light display on the main street had been set up two weeks early so that IETF 76 participants could enjoy it each evening.

The mayor of Hiroshima joined us at the social on Tuesday evening, helping open a large cask of sake for all to enjoy. Wearing an IETF 76 T-shirt, the mayor wished us a successful meeting, which certainly came to pass.

The meeting site network was installed and operated by volunteers from the WIDE Project with assistance from the usual set of dedicated volunteers from the IETF community. The network was robust and reliable. In addition, the WIDE Project installed radio-frequency-identification readers at every microphone in the meeting rooms, which read an RFID tag in each speaker's IETF badge. If the speaker had chosen to participate in the experiment, then the speaker's name and picture were displayed on a screen at the front of the room. Many participants—note takers in particular—found the display helpful.

Since IETF 75, five new working groups (WGs) have been chartered, and two WGs were closed. There are a total of about 115 chartered WGs. Since the last IETF, the WGs and their individual contributors produced 412 new Internet-Drafts and updated 857 Internet-Drafts, some more than once. The Internet Engineering Steering Group approved 87 Internet-Drafts for publication as RFCs. The RFC Editor published 92 new RFCs.

Like at the previous meeting, one of the hot topics during IETF 76 was the development of tools to facilitate the transition from IPv4 to IPv6. The technical plenary, which featured a discussion titled Internationalization in Names and Other Identifiers, generated a lot of lively discussion.

I look forward to IETF 77 in Anaheim, California, on 21–26 March 2010 as well as IETF 78, which is scheduled for 25–30 July 2010 in Maastricht, Netherlands, and will be hosted by SIDN. Looking ahead, IETF 79 is set for 7–12 November 2010 in Beijing, China. The host is Tsinghua University. Scheduling information for the next IETF meetings may always be found at <http://www.ietf.org/meeting/upcoming.html>. I look forward to seeing you at these meetings. 



Russ Housley, IETF Chair

### New BoF Meetings

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

#### Applications Area

hybi	Bidirectional or server-initiated HTTP
IRIs	Internationalized Resources Identifiers
grobj	Generic Referral Object
decade	DECoupled Application Data Enroute
6lowapp	Application protocols for low-power v6 networks

#### Internet Area

aplus	Address plus port
-------	-------------------

#### RAI Area

codec	Internet wideband audio codec
-------	-------------------------------

#### Routing Area

karp	Keying and authentication for routing protocols
------	---

#### Transport Area

homegate	Broadband Home Gateway
conex	Congestion exposure
ppsp	Peer to Peer Streaming Protocol



Olaf Kolkman, IAB Chair

## Words from the IAB Chair

By *Olaf Kolkman*

My first computer was a Sinclair ZX81, a tiny little black box that, as Wikipedia explains, did not use ASCII but had its own character set. Character code 0 was space; codes 1–10 were used for block graphics; and codes 11–63 corresponded to punctuation, numbers, and uppercase characters. Character codes 128–191 were reverse video versions of the first 64 characters. Other codes represented BASIC keywords and control codes such as NEWLINE. There were no lower-case characters.<sup>1</sup> At the time (I was 15 or 16 years old), I did not give that limited character set a lot of thought. And I never considered how its character codes might need to be encoded and mapped if the machine were ever to communicate with the outside world.

Today, billions of users and even more machines are able to communicate over the Internet, and limited character sets like ASCII are perfectly fine protocol elements for intermachine communications not designed for human consumption. When it comes to application content intended for human consumption, however, the IETF, through RFC 2277, required protocols to support UTF-8 encoding of the Unicode character set. The applications and operating systems are expected to translate UTF-8 to and from the user's interfaces.

It is at the intersection where pure machine-to-machine and pure user-to-user communications meet that a lot of the challenges remain. When users need mnemonics to pass to their computers and they want those mnemonics to be written in their own languages and scripts, we will need to answer one particular question: How do we represent what users consider a valid and useful representation of an identifier in an unambiguous manner such that comparison, delivery, and matching can be done in a unambiguous and standardized manner? That underlying question impacts a broad range of work being done within the IETF, and it has been the motivating force behind specific work being done in the two Internationalizing Domain Names in Applications (IDNA) and IDNA-bis working groups, the Email Address Internationalization working group, and the recent Internationalized Resource Identifier BoF. These groups focus on identifiers that are in wide use and based on legacy representation.

At IETF 76 in Hiroshima, Japan, the Internet Architecture Board (IAB) provided an update for the community with regard to internationalization issues, including work being done by the IAB on a document that describes the various encoding schemes that exist for domain names, as well as the problems that occur when assumptions are made about the context in which those names are used.<sup>2</sup> For example, resolver libraries will need to know whether the resolution happens via the DNS, where the encoding of the names is in Punycode, or via alternative encoding schemes, such as UTF-8 DNS name encoding, which is in use in some enterprises. Moreover, within applications, the use of internationalized identifiers can give rise to confusion, as is demonstrated when e-mail addresses (in a user's native script) are included in the header and body of e-mail messages. It is clear that solutions to these problems are not trivial, which is why internationalization of identifiers and names in Internet protocols has been an ongoing area of interest for the IAB and, thus, why the IETF 76 plenary served as a continuation of the technical plenary presented at IETF 66.

That the problem of encoding—and in particular, the mapping of one encoding into another—can bite you at unexpected moments was demonstrated

### IETF 76 Facts and Figures

Registered attendees from 43 countries .....	1133
New WG .....	5
WGs closed .....	2
WG currently chartered .....	115
New Internet-Drafts .....	412
Updated Internet-Drafts .....	857
IETF Last Calls .....	96
Internet-Drafts approved for publication .....	87

#### RFC Editor Actions (July–October 2009)

92 RFC published of which

- 56 Standards Track
- 7 BCP
- 24 Informational
- 5 Experimental

108 Internet-Drafts submitted for publication

- 80 submitted by the IETF WGs
- 19 submitted by IETF individuals
- 9 submitted by IRTF, IAB, and independent submissions combined

#### IANA Actions (July–October 2009)

1600+ IETF-related requests processed

- 784 Private Enterprise Numbers
- 83 Port Numbers
- 131 TRIP ITAD Numbers
- 22 media type requests
- 17 language subtag-related requests

*Continued page 7*

## IETF 76 Plenary Report

By Carolyn Marsan

The IETF 76 Administrative Plenary kicked off with a brief address by IETF chair Russ Housley, who introduced the meeting's host, Jun Murai, of the Widely Integrated Distributed Environment (WIDE) project. Jun thanked attendees for traveling to Hiroshima, a city that he admitted was far away for many participants. Jun pointed out that Hiroshima was an appropriate location for an IETF meeting because it was one of the first cities in Japan to engage in computer networking research. He thanked the city of Hiroshima for giving an Olympic-class welcome to the IETF community.

Jun also thanked meeting attendees for participating in ongoing WIDE-funded research regarding RFID deployment and standards. IETF 76 attendees and speakers were given badges with embedded RFID cards, which were used throughout the meeting for the social check-in, blue sheets sign-up, and open-mic sessions when activated. Jun thanked many of the attendees for their "patience, cooperation, and bravery" in moving past their privacy concerns and participating in WIDE's RFID testing.

### IETF Chair Report

Russ said IETF 76 attracted 1,106 attendees from 44 countries, compared with 937 attendees from 52 countries at the November 2008 meeting in Minneapolis. The largest number of attendees at IETF 76 came from Japan and the United States.

Russ chided attendees for being "procrastinators," pointing out that 70 percent of the new and updated Internet-Drafts published since the group's summer meeting had been posted in

the four weeks prior to the Hiroshima meeting.

In terms of upcoming meetings, Russ pointed out that the March 2010 meeting in Anaheim, California, is without a host. But he said the summer meeting, to be held in Maastricht, Netherlands, will be hosted by SIDN, while the fall 2010 meeting will be held in Beijing, with Tsinghua University serving as host. He reminded attendees that all of the 2010 meetings would feature Friday afternoon sessions.

### Itojun Service Award Presented

Next, Jun presented Google engineers Lorenzo Colitti and Erik Kline as the first annual recipients of the Itojun Service Award, which recognizes extraordinary dedication toward the development and deployment of IPv6.

The Itojun Service Award honours the memory of Dr. Jun-Ichiro "itojun" Hagino, who passed away in 2007, at age 37. The award, established by friends of itojun and administered by the Internet Society (ISOC), recognizes and commemorates the extraordinary dedication

exercised by itojun over the course of IPv6 development. The award includes a presentation crystal, a USD 3,000 honorarium and a travel grant.

Itojun's mother gave a moving tribute, thanking the members of the IETF community who helped establish an award in her son's memory. She said she hoped her son's "passion for IPv6 and his friendships continue forever."

Erik and Lorenzo were honoured for evangelizing IPv6 across Google, for



Jun Murai of WIDE, host of IETF 76

Photo/Peter L  thberg

helping develop IPv6-enabled Google services, and for coordinating Google's IPv6 conferences.

"I went to my first IETF meeting with the express purpose of meeting itojun. That was in December 2007, and I was only able to go to the memorial there," Erik said as he accepted the award. "I didn't meet itojun, but I met many others who helped us do our work and change our culture internally."

### IAOC Report

Bob Hinden, chair of the IETF Admin-

*Bandwidth, Transition, continued from page 1*

Trent Adams and Eve Maler have taken time out from their busy schedules to give us an update on the progress of the Kantara Initiative, which was founded earlier this year (see page 20).

IETF 76 was made richer by the presence of the Internet Society fellows, who travelled from far and wide

for the opportunity to enhance their knowledge and technical skills through involvement with the IETF and to contribute their perspectives during the meeting (see page 13). As Subramanian Moonesamy put it so well, "The ISOC Fellowship Programme provides people from developing countries with the

means to contribute to the IETF and make their voices heard."

My sincere thanks to everyone who contributed to this issue. I hope you find it interesting and, as I mentioned, please send your feedback to [ietfjournal@isoc.org](mailto:ietfjournal@isoc.org). This is your chance to shape the future of the *IETF Journal*. 

istrative Oversight Committee (IAOC), began his report by stating that all of the outstanding issues with the host and the hotel for IETF 79 in Beijing have been resolved. “We’re very confident we can have a normal IETF meeting just like we do everywhere else,” Bob said.

He added that the IETF is close to meeting its 2009 budget despite the economic downturn. The IETF is forecasting revenues of USD 3.2 million, expenses of USD 4.7 million, and a total ISOC contribution of USD 1.5 million.

The IETF took advantage of ISOC stimulus funding so that it could lower registration fees for its meetings and maintain attendance levels. However, the IETF has not tapped into contingency funds that ISOC made available earlier in 2009.

The IETF experimented with one-day passes at the Hiroshima meeting, but it has not decided whether it will continue with that effort at future meetings.

For 2010, IAOC plans to keep registration fees the same, at USD 635 per meeting. The group is projecting USD 3.1 million in revenues, USD 5.2 million in expenses, and an ISOC contribution of USD 2.1 million. The 2010 budget includes investments in new automated tools, including Data Tracker extensions, Secretariat tools, RFC services, and programme management capabilities.

“A lot of our important tools have been developed by volunteers. This has been wonderful, but we’re starting an effort to make these tools more supportable over the long term,” Bob said.

Also in 2010, the IETF plans to begin migrating its RFC Production Centre and RFC Publisher from the University of California’s Information Sciences Institute (ISI) to a new contractor, AMS. The IAOC is still involved in awarding contracts for the RFC Series Editor and Independent Submissions Editor (see the article on page 14 in this issue for more detail and discussion of these developments).

In other IAOC news, Hinden said that the committee was making progress at publishing the minutes from its meetings. All 2008 minutes along with the minutes of 13 out of 19 meetings held in 2009 were published online.

#### Trust Chair Report

IETF Trust chair Marshall Eubanks explained why modifications to the Trust Legal Provisions (TLP) are still in flux.

In October 2009, the trust chair proposed TLP 4.0, which would cover production of IETF, Internet Architecture Board (IAB), Internet Research Task Force (IRTF), and Independent Stream documents. However, the RFC Editor Board has some disagreements with this document. So Marshall said it will be revised again and sent back to the IETF

community for comment in January.

Marshall said the IETF community levied a fair amount of criticism on the trust for the trust’s lack of transparency in its process of modifying the TLP. In response, Marshall said the trust has “tried very hard to

#### NomCom

As of IETF 76, the NomCom was still working on IAB, IAOC, and Internet Engineering Steering Group (IESG) positions opening in March 2010.

Chair: Mary Barnes

#### Voting Members

- Scott Brim
- David Crocker
- Roque Gagliano
- Randall Gellens
- Dorothy Gellert
- Wassim Haddad
- Stephen Kent
- Dimitri Papadimitriou
- Simo Veikkolainen
- Lucy Yong

#### Nonvoting members

- Joel Halpern (Past-year Chair)
- Henrik Levkowitz (Tools Advisor)
- Jon Peterson (IAB Liaison)
- Tim Polk (IESG Liaison)
- Henk Uijterwaal (IAOC Liaison)
- Bert Wijnen (ISOC Liaison)

improve our communications with the community,” and he pointed out that the group’s meeting minutes are now available online.

#### Recognition

Russ and IAB chair Olaf Kolkman recognized and thanked ISI for being the RFC Editor for nearly 30 years. Plaques were presented to ISI employees Bob Braden, Sandy Ginoza, and Alice Hagens. See the article on page 14 in this issue for more detail and discussion of these developments.

#### Open Mic Comments

During the open-mic session, several questions were raised about whether IETF participants would have open, unfettered Internet access while meeting in Beijing in 2010. The IETF leadership said it has been assured by the local host and hotels that there will be no content filtering or blocking of the group’s Internet access.



Photo/Karim Hussein

A performance at the IETF 76 social event in Hiroshima

*Continued on next page*

*IETF 76 Plenary Report, continued*

Additionally, several participants said they would like to see the IETF expand its use of Web conferencing to support remote participation in its meetings, while others complained that Web conferencing is a distraction and reduces the meeting attendees' productivity.

A discussion took place about how best to attract attendees to Friday sessions, with proposals to move the technical plenary or the social to Friday evening.

Finally, IETF participants and leaders debated the purpose of the group's three-step standards process and whether it makes sense for documents to become full standards or to remain as draft standards or proposed standards.

### IETF 76 Technical Plenary

Olaf opened the technical plenary with a welcome to attendees, particularly those following the event remotely through Web conferencing or Jabber.



IETF 76 participants mingle at reception

### IRTF Report

Next on the agenda was IRTF chair Aaron Falk, who said four research groups met in Hiroshima: Host Identity Protocol (hip), Scalable Adaptive Multicast (sam), Delay Tolerant Networking (dtn), and Routing Research Group (rrg). Aaron said the IRTF has six RFCs waiting to be published that

have been held up because of trust issues discussed at the administrative plenary. Aaron highlighted the ongoing work of two research groups—Anti-Spam and Scalable Adaptive Multicast—in the

---

**“A lot of our important tools have been developed by volunteers. This has been wonderful, but we’re starting an effort to make these tools more supportable over the long term.” — Bob Hinden**

---

hopes of generating more interest and activity from IETF participants. For more information, see the IRTF Update on page 22.

### IAB Report

Olaf outlined several IAB documents in his talk. One that is about to be published is the Peer-to-Peer Architecture; another that is under development covers IPv6 Network Address Translators (NATs). The IAB is also working on documents that describe the Internet Assigned Numbers Authority (IANA) function and an update on the IP model.

One piece of news that Olaf shared with attendees: IANA will be signing the .arpa zone using the same mechanism for Domain Name System (DNS) Security Extensions that will be used to sign the DNS root zone.

Olaf said the IAB has created an advisory group that will select candidates to serve as the RFC Series Editor (RSE) and Independent Submissions Editor (ISE) under the IETF's new RFC Editor model. After some delay, the IAB is interviewing candidates for the ISE job and expects to be interviewing candidates for the transitional RSE position.

### Internationalization in Names and Other Identifiers

The IAB is working on a document about internationalization—draft-iab-idn-encoding—that inspired the dis-

cussion led by John Klensin, Stuart Cheshire, and Dave Thaler.

Internationalization “is understood moderately well by a fairly small number of experts, most of whom end up realizing how little we actually understand,” John said at the beginning of the talk. “But it affects a large number of protocols, a large number of people, and it should affect virtually everything we’re doing in the IETF.”

John said internationalization is increasingly common in domain names, e-mail addresses, and URLs because users want to use the Internet in their own languages. Among those pushing for internationalization are users from Arabic-speaking nations, China, South Korea, and Taiwan.

Stuart explained the rationale behind the IETF policy regarding internationalization. The IETF requires that all protocols created after January 1998 should be able to use UTF-8, which is an ASCII-compatible method of encoding Unicode characters as integers. This policy is explained in RFC 2277: IETF Policy on Character Sets and Languages.

Stuart also explained Punycode, which is a method of encoding Unicode characters for use in internationalized domain names. He demonstrated some of the shortfalls of Punycode, particularly that it creates a sequence of bytes that applications can interpret in more than one way.

Photo/Internet Society

Because Internet protocols do not handle international text natively, the number of ways of encoding international text into printable ASCII characters has proliferated, Stuart said, which is creating chaos.

“This can get really crazy,” he said. “Suppose you have a domain name that is part of an e-mail address, which you put in a mailto: URL, which is then appearing on a Web page in HTML text. Is the domain name supposed to be actual rich text as seen to the user? Or is it supposed to be Punycode?”

Next, Dave discussed the difficulties involved with doing matching for internationalized domain names, pointing out how easy it is for users to be confused by what they read on the screen.

“If you have a sufficiently creative use of fonts and you have style sheets from a strange environment, almost anything can look like almost anything else,” John said, pointing out that people tend to see what they expect to see.

He added that the confusion over internationalized domain names could lead to security risks, such as phishing attacks.

“If I can make a string in one script—or partially in one script—look like a string in some other script, I suddenly have an opportunity, especially if I’m what the security people call a bad guy,”

John added, pointing out that such attacks can be deliberate or accidental.

John then talked about some of the difficulties of mapping, including that mapping leads to lost information and differs by language. That’s why the IAB doesn’t recommend developers make up their own mapping systems.

Dave said the situation today is that we have multiple encodings of the same Unicode characters and different encodings, even within DNS. “Because you have all the differences across the protocols, networks, and so on, you can imagine the confusion that results,” he said.

In conclusion, Stuart said the IETF may need to go beyond its current recommendation of supporting UTF-8 as one encoding mechanism to requiring it as the only encoding mechanism for the text that end users see.

Questions from the audience focused on the security problems related to internationalization, such as spoofing one character with another and changing fonts.



John said the Internet engineering community has known for nearly 20 years about the confusion that would be caused in moving from the limited ASCII character set to an environment with tens of thousands of characters.

“I don’t think there are any easy answers,” he said. “But the alternative to this situation is that nobody gets to use their own script, and that answer is completely unacceptable.”

Other questions from the audience revolved around the IAB’s document regarding NATs for IPv6. 

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

in the e-mail that was sent to announce the IETF 76 plenary.<sup>3</sup> The e-mail was composed in an editor and then copied and pasted into a Web form that sends announcements to the IETF list. If you look at the announcement, you can see a few occurrences of `&#8232`, which is XHTML character escape for the Unicode LINE SEPARATOR. Somewhere in the process of cutting, pasting, and CGI script handling, the original encoding present in the editor got translated into the XHTML escaped charac-

ter, which is what ended up in people’s mailboxes.

While most of the casual readers will ignore occurrences of `&#8232` in an e-mail text, those occurrences will become relevant in a comparison of strings (identifiers and names in particular); in other words, the question becomes, Is a line of text that is typed in reverse video on the Sinclair ZX81 actually the same as the one typed using the first 64 characters?

If you want to know more, please read the slide set,<sup>4</sup> the transcript,<sup>5</sup> and the draft,<sup>2</sup> on which the IAB welcomes your feedback. 

## References

1. <http://en.wikipedia.org/wiki/ZX81>
2. <http://tools.ietf.org/html/draft-iab-idn-encoding>
3. <http://www.ietf.org/mail-archive/web/ietf-announce/current/msg06737.html>
4. <http://www.ietf.org/proceedings/09nov/slides/plenaryt-1.pdf>
5. <http://www.ietf.org/proceedings/09nov/minutes/plenaryt.txt>

*The Bandwidth Bandwagon, continued*

moderator Leslie Daigle, the Internet Society's chief Internet technology officer, put it in her introduction.

Leslie said decision makers are increasingly trying to understand the parameters of Internet bandwidth growth and management because these have implications for both network-neutrality debates and business decisions based on predictions of growth and usage. The panel was intended to bring new clarity to the answers to such questions as, What are the bottlenecks? What causes congestion? Is congestion bad? What is the impact? and What is being done about it? The panel was composed of individuals with real data, real network issues to resolve, and real technologies to make Internet bandwidth use more effective and efficient for all.

**Broadband Landscape in Japan**

First up was Kenjiro Cho, a senior researcher at Internet Initiative Japan. Kenjiro presented his research results, which had been based on data collected from six ISPs in Japan starting in 2004 and covering 42 percent of Japanese Internet traffic. As of June 2009, there were 30.9 million broadband subscribers in Japan, and the market is relatively mature, increasing by only 3 percent of households in 2008 to thereby comprise 63 percent of Japanese homes. While growth of cable deployments remains steady, the great majority of households enjoy fibre-to-the-home (FTTH) con-

nections, and existing DSL customers are shifting to FTTH in large numbers. In the Japanese market, 100-Mbps, bi-directional connectivity via FTTH costs USD 40 a month. The relatively high access bandwidth in the Japanese market leads to higher skew in the distribution of per-user bandwidth consumption statistics: there is more variability in bandwidth consumption profiles per user.

ISPs are starting to see the value in sharing traffic growth data as a way to help others better understand their concerns. Of course, ISPs make internal measurements, but measurement methodologies and policies will typically differ from one ISP to the next. By aggregating standardized and anonymized measurements, ISPs can help third parties come to understand the pressures, concerns, and motivations that are shaping their perspective.

Understanding traffic growth on the Internet is critically important, as it is one of the key factors driving investment decisions in new technologies and infrastructure. The balance between supply and demand is crucial. Kenjiro has observed modest growth of about 40 percent per annum since 2005 based on traffic peaks at major Japanese Internet exchanges. For residential traffic, growth rates are similar—around 30 percent per annum. As network capacity is observed to grow at approximately 50 percent per annum, according to various sources, there does not appear to be a problem in catering to Internet traffic growth, at least at the macro scale.

Kenjiro discussed some of the observed shifts in residential user behaviour in the period 2005–09. In 2005, the ratio of inbound to outbound traffic was almost 1:1, suggesting that file sharing was a

**Congestion Collapse**

In the past, when more packets were sent than could be handled by intermediate routers, the intermediate routers discarded many packets, expecting the end points of the network to retransmit the information. However, early TCP implementations had very bad retransmission behaviour. When this packet loss occurred, the end points sent extra packets that repeated the information lost, thereby doubling the data rate sent—exactly the opposite of what should be done during congestion. This pushed the entire network into a congestion collapse, wherein most packets were lost and the resultant throughput was negligible.

Source: [http://en.wikipedia.org/wiki/Congestive\\_collapse](http://en.wikipedia.org/wiki/Congestive_collapse)

very widespread use of the network at that time. In 2009, the outbound traffic (download from a user perspective) was noticeably greater, suggesting a shift from peer-to-peer file sharing to streamed content services. Increases in the mode of download volumes over the period are greater (nearly 4 times: from 32 MB to 114 MB per day) than increases in upload volumes (less than 2 times: from 3.5 MB to 6 MB per day), while average download volumes are now 1 GB per day per user.

In analyzing a scatterplot of in/out volumes per user in 2009, Kenjiro observed that while there are two clusters (client-type users and peer-type, heavy hitters), there is no clear boundary between the two groups. This is an important point to bear in mind when considering the effectiveness of coarse-grained bandwidth management techniques deployed by some ISPs today. Most users make some use of both client-server-style and peer-to-peer-style applications.

Kenjiro concluded with the observation that while the data is interesting, it is nevertheless difficult to predict the future of Internet bandwidth given the



IETF 76 participants listen as panellists discuss bandwidth issues

Photo/Internet Society



Bandwidth Bandwagon panellists (from left) Kenjiro Cho, Danny McPherson, Richard Woundy, and Lars Eggert



Panel moderator Leslie Daigle of the Internet Society



Panellist Kenjiro Cho, of the Internet Initiative Japan, gives a panel presentation

Photos/Internet Society

variety of technical, economic, and political factors at play.

### ISPs Working with IETF

The next presenter was Richard Woundy, senior vice president at Comcast, a large, U.S.-based cable ISP. Richard began by explaining some of the ISP's motivations for congestion management: the need to be responsive to very dissimilar customer application demands; to balance the competing concerns of the Internet community, regulators, invest-

what our customers think of our service," said Richard. "For an ISP it's a balancing act."

For best-effort traffic over the cable network, the Comcast congestion management plan utilizes two different Quality of Service (QoS) levels: Priority Best Effort (PBE), which is the default QoS level, and Best Effort (BE). When levels of traffic on a particular port exceed a set threshold, that port enters a near-congestion state. Customers determined to be contributing disproportion-

congestion management, such as the conex BoF and alto and ledbat working groups (WGs), described in more detail later. Richard said, "It's about making sure—while we're executing a reasonable upgrade schedule—that when flash crowds happen or some new streaming application appears that chews up bandwidth, we can handle all those services gracefully."

### Hypergiants, Port 80, and a Flatter Internet

Danny McPherson, chief security officer at Arbor Networks, then presented the recent results of the ATLAS Internet Observatory, which is collaborative research between Arbor Networks, the University of Michigan, and Merit Network. The ATLAS Internet Observatory utilizes a commercial probe infrastructure deployed at more than 110 participating ISPs and content providers to monitor traffic flow data across hundreds of routers. That commercial probe infrastructure is believed to be the largest Internet monitoring infrastructure in the world, and the observatory's results represent the first global traffic engineering study of Internet evolution.

Major findings from the ATLAS project are, first, the consolidation of content around so-called hypergiants—the 30 companies that now account for 30 percent of all Internet traffic. Content is migrating from the enterprise or network edge to large-content aggregators. Consolidation of large Internet

---

**Understanding traffic growth on the Internet is critically important, as it is one of the key factors driving investment decisions in new technologies and infrastructure.**

---

tors and so on; and the fact that network capacity increases are not instantaneous. Richard noted the daily challenge of having to tune Comcast's network to ensure, for example, that VoIP service providers aren't disadvantaged, followed by the need to then check that, again, for example, another third-party video services provider hasn't gotten unintentionally disadvantaged in the process. For Comcast, the goal of congestion management practice is to ensure consistent performance of Internet applications even in the presence of heavy background traffic, such as from peer-to-peer file sharing. Comcast aims to be both protocol and application agnostic and compatible with current Internet standards. "We're always worried about

ately to the total traffic volume of a port in the near-congestion state will have their traffic marked as BE for a short duration. That marking impacts the traffic of users marked BE only when congestion is actually present; otherwise, PBE and BE traffic are treated identically. In the presence of congestion, traffic marked BE will experience additional latency (on the order of a few microseconds) as it gets queued, while PBE traffic takes priority. Less than 1 percent of Comcast's customer base is impacted by this congestion management plan, said Richard.

Richard also highlighted the work Comcast is doing to collaborate with the IETF on new protocols that could form part of future solutions for end-to-end

*Continued on next page*

*The Bandwidth Bandwagon, continued*

properties has progressed to the point where now only 150 Autonomous Systems contribute 50 percent of all observed traffic.

Second, applications are consolidating around TCP port 80, because the Web browser is increasingly the application front end for diverse content types, such as e-mail and video. For application developers, TCP port 80 works more deterministically due to the presence of middleboxes in the network that filter or otherwise interfere with traffic using different transports and alternative ports.

Third, evolution of the Internet core and economic innovation mean that the majority of traffic is now peered directly between consumers and content. Declining transit prices have not prevented this disintermediation from taking place on a large scale. High-value-content owners are starting to experiment with a paid-peering model and dispensing with transit altogether, meaning that if your ISP doesn't pay to play, then you won't be able to view that content at all—although this phenomenon is difficult to quantify due to the inevitable commercial secrecy surrounding such deals. Disintermediation of the historical tier-1 networks means a flatter Internet with much higher interconnection density.

ATLAS also observed the trend away from peer-to-peer and toward streaming video distribution mentioned by Kenjiro

earlier. Observations of the Internet's size (9 exabytes per month) and growth rate (44.5 percent compound annual growth) also agree with others' analyses. While those numbers certainly indicate significant growth, they're well within

formation. But the smart-edge, dumb-core paradigm gets you only so far, and there's a valid role for the network, as exemplified by the Comcast experience. As Lars observed, "It's not all about the edges."

---

**The network is required to provide neutral information about path conditions in a timely manner, while applications and transport protocols choose how to act on that information. But the smart-edge, dumb-core paradigm gets you only so far, and there's a valid role for the network.**

---

projected increases in gross network capacity and so no cause for concern.

Danny concluded by observing that the Internet appears to be at an inflection point as it transitions from a focus on connectivity to a focus on content.

### Sharing Means Caring

The final panellist was Lars Eggert, principal scientist at Nokia Research Center and Transport Area director at the IETF, who briefly introduced the IETF activities related to bandwidth management, or capacity sharing. Lars observed that the Internet is all about capacity sharing. The connectionless, best-effort, end-to-end nature of the Internet enabled it to scale and resulted in the tremendous innovation that we all now take for granted. Sharing Internet resources as the Internet does requires congestion control mechanisms at the transport layer and requires applications to be "social" in their behaviour toward each other. "Sharing means caring," as Lars explained.

The architectural principles of the Internet mean that in general, the responsibility is split between the applications and the network. The network is required to provide neutral information about path conditions in a timely manner, while applications and transport protocols choose how to act on that in-

The IETF toolbox includes TCP and TCP-friendly congestion control that allows hosts to determine their transmission rate according to path conditions based upon observed round-trip time and packet loss. Extensions and optimizations include Explicit Congestion Notification (ECN) and Active Queue Management (AQM). However, as Lars observed, "Mechanisms like ECN and AQM were developed in the 1990s, when core speeds were around 45 Mbps. Now we have those speeds in the access network. Stuff that we did back then for the core should be revisited to see what we could use in the access network."

A new IETF WG (Low Extra Delay Background Transport, or ledbat) is standardizing a congestion control algorithm to allow hosts to transmit bulk data without substantially affecting the delay seen by other users and applications. Another new WG (Multipath TCP, or mptcp) is endeavouring to extend TCP so as to enable one connection to transmit data along multiple paths between the same two end systems. This effectively pools the capacity and reliability of multiple paths into a single resource and enables traffic to quickly move away from congested paths.

The Application Layer Traffic Optimization, or alto, WG is focused on

### Bandwidth Bandwagon Panellists

Leslie Daigle, Internet Society  
(moderator)

Kenjiro Cho, Internet Initiative Japan

Lars Eggert, Nokia

Danny McPherson, Arbor Networks

Richard Woundy, Comcast

improving peer-to-peer application performance while simultaneously aligning peer-to-peer traffic better with ISP constraints. Providing peer-to-peer applications with network, topology, and other information should enable them to make better-than-random peer selection, thereby improving performance for the application and alignment with ISP preferences.

A new BoF meeting at IETF 76 was Congestion Exposure (conex), which targets exposing the expected congestion along an Internet path. This would be a new capability and could allow even greater freedom over how capacity is shared than we have today. “This is a very powerful mechanism that provides an information exchange between the network core and the edges that wasn’t there before, and it has lots of potential uses,” said Lars. Such a capability could be used for a variety of purposes, such as congestion policing, accountability, service-level agreements, and traffic engineering.

Finally, Lars drew attention to another BoF meeting taking place during IETF 76 on recommendations for home gateways: homegate, which is intended to collect requirements from disparate RFCs and provide an overview for implementers of home gateway devices. The goal is to improve the network experience for an end user using a home gateway to access the Internet.

There are already many tools to share Internet capacity fairly, effectively, and

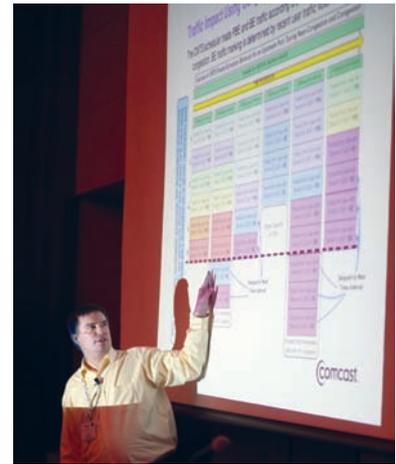
efficiently, and the IETF is designing new and better tools where needed. Lars concluded by noting that a lot could be gained by more consistently and appropriately using the tools we already have.

### A Balanced Approach to the Impact of Broadband

In discussion after the panel, both Lars Eggert and Richard Woundy observed that it is the impact of broadband on the network that has largely exposed a lot of these issues for congestion management. “I’m very glad that this discussion has picked up over the last few years with the rise of broadband,” said Lars. “We need mechanisms to handle the [new broadband] speeds safely.” Higher access speeds make it possible for individual end users to have a significant impact on the network.

Richard emphasized that it would be a mistake to conclude from all of this that ISPs want to stop investing. The concern is, rather, to ensure ISPs are able to deliver a good customer experience for all, even when traffic increases in unexpected ways (doubles overnight, for example). “That’s the kind of situation where congestion management makes sense, but it needs to be followed up with capacity upgrades,” Richard said. “You don’t do one without the other; otherwise, you’re just letting your service fall apart.”

*Details of the event, a set of slides, audio, and a transcript are available from the ISOC Web site <http://isoc.org/bandwidth/>.*



Panelist Richard Woundy at IETF 76

Photo/Internet Society



Bob Hinden asks a question during the openic portion of the IETF 76 panel discussion

Photo/Internet Society



IETF 76 participant and ISOC Board member Bert Wijnen attends the panel discussion

Photo/Internet Society

#### IETF Toolbox

alto: <http://www.ietf.org/dyn/wg/charter/alto-charter.html>

conex: <http://www.ietf.org/proceedings/09nov/agenda/conex.txt>

homegate: <http://www.ietf.org/proceedings/09nov/minutes/homegate.htm>

ledbat: <http://www.ietf.org/dyn/wg/charter/ledbat-charter.html>

mptcp: <http://www.ietf.org/dyn/wg/charter/mptcp-charter.html>

TCP road map: <http://www.ietf.org/rfc/rfc4614.txt>

## Google: IPv6 Is Easy, Inexpensive

By Carolyn Marsan

*IETF 76 proves a fitting backdrop for presentation of the first annual Itojun Service Award, which recognizes extraordinary dedication toward the development and deployment of IPv6.*

IPv6 is not rocket science. That's the message that the Google engineers who are the first winners of the Itojun Service Award for outstanding contributions to the development and deployment of IPv6 want to send to the IETF community. Lorenzo Colitti and Erik Kline were presented with the Itojun Service Award at the IETF meeting in Hiroshima, Japan. The two engineers have been leading Google's IPv6 development efforts for two years.

"We're up to a handful of people working on IPv6 almost 100 percent of the time," said Erik, an IPv6 software engineer at Google.

So far, Google supports IPv6 in its Search, Alerts, Docs, Finance, Gmail, Health, iGoogle, News, Reader, Picasa, Maps, Wave, Chrome, and Android products. Google is working on IPv6 for YouTube and Google Voice.

Lorenzo and Erik said the main lesson they've learned from Google's IPv6 development efforts is that it isn't very hard or very expensive to add support for IPv6 to existing Web services and applications. "As a content provider, you can get an IPv6 service up and running without changing all of the back-end stuff," Erik said. "You need to audit where IPv4 addresses are stored and

used, but you don't actually have to have deep and 100-percent-pure IPv6 throughout all the stacks. You can deploy IPv6 only where it makes sense."

The two Google engineers recommend taking a dual-stack approach to IPv6 development and to mirror IPv6 services as closely as possible to existing IPv4 services. "From the networking point of view, you want to use the existing infrastructure," Lorenzo said. "You want to dual stack everything you can, and design IPv6 as closely as possible to the existing IPv4 infrastructure."

Google is already seeing some benefits from its IPv6 development efforts, particularly in simpler and potentially lower-cost network management. "We can talk directly to the new LTE handsets and a bunch of IPv6 set-top boxes,"

Erik said. "We'll be able to talk to them directly, as opposed to only seeing them behind application proxies or NAT [network address translation] devices."

"We actually have had a couple of IPv6-only networks access Google over IPv6," Lorenzo said. "It's a more direct path, and it's better connectivity."

Lorenzo and Erik said they have not experienced any performance problems with IPv6 and that IPv6 did not cause the widespread outages that were erroneously attributed to it. "The outage that was blamed on IPv6 in May did not, in fact, affect IPv6 and was not due to IPv6," Lorenzo said.

Erik added that "IPv6 was the one thing that was up and running."

Lorenzo and Erik haven't discovered a significant drawback with IPv6 except for the time and effort that it takes to deploy. "It takes time if you have to wait for a vendor box to be fixed. If you find something that's broken, you have to try to work around the issues. But it's certainly not rocket science," Lorenzo said.

IPv6 deployment isn't expensive, either. "It costs less than you think it would," Lorenzo said. "You don't have to spend much money on it if it's part of your upgrade process."

The Itojun Service Award honours the memory of Dr. Jun-ichiro "itojun" Hagino, who passed away in 2007 at the age of 37. The award, established by the friends of itojun and administered by the Internet Society, recognizes and commemorates the extraordinary dedication exercised by itojun over the course of IPv6 development.

The Itojun Service Award focuses on pragmatic contributions to the development and deployment of IPv6 in the spirit of serving the Internet. The annual award includes a presentation crystal, a USD 3,000 honorarium, and a travel grant. 



Erik Kline and Lorenzo Colitti (from left) are presented the first Itojun Service Award by Jun-ichiro "itojun" Hagino's brother, his mother Kumiko Hagino, and Jun Murai

Photo/Internet Society

## ISOC Fellows Enjoy In-Person IETF Experience

By Carolyn Marsan

Meeting other network engineers who are facing and solving similar problems—that’s the main benefit that technologists from developing countries said they’re receiving as a result of the Internet Society’s fellowship programme. ISOC awarded fellowships to 12 technologists from developing countries to attend the IETF meeting in Hiroshima, Japan, in November 2009.

Participants from Argentina, Brazil, Chile, Colombia, Costa Rica, Fiji, India, Mauritius, and Pakistan were able to experience the Internet standards process in person. The ISOC Fellowship Programme receives significant financial support from leading Internet companies, including Afilias, Google, Intel, Microsoft, and the Nominet Trust.

While the IETF conducts most of its standards development work via e-mail, ISOC fellows said they valued the experience of face-to-face meetings with their counterparts from around the world.

Hassan Zaheer, manager of IP core at Worldcall Telecom Ltd. in Pakistan,

said the most useful aspect of attending the IETF meeting in Hiroshima was being able to talk freely to industry veterans with whom he had exchanged e-mails in the past. “The informal atmosphere is the most benefitting and at the same time very interesting,” said Hassan, who’s been working in IP network operations for a decade. “It can’t be possible to write down all of [the knowledge I gained] from corridor meetings and discussions at social gatherings.”

Hassan discovered that Comcast is facing the same problem of increasing demand for bandwidth from its broadband customers that Worldcall Telecom is seeing. He gained insight from Comcast that will help as Worldcall Telecom deploys DOCSIS 3.0, a cable television industry standard. “Comcast has the same problems as we are facing here, and I found good examples of how they resolved it using DOCSIS Quality of Service, which I can implement here also,” he said.

Hassan said it’s an honour to be an ISOC fellow. “For someone like me, who has been working on network technologies for so long and communicating with peers and involved in technology evolution, it’s the best opportunity to meet them in person and exchange views,” he said.

The free-flowing exchange of ideas and information at IETF meetings is what Gargi Bag, an Indian doctoral student at Ajou University in South Korea, said she valued most from her ISOC fellowship. Gargi’s area of research is the provision of mobility support for low-



Internet Society fellows and mentors at IETF 76 in Hiroshima, Japan

Photo/Internet Society

power, low-bandwidth devices that use the IETF’s 6LowPAN protocol. “The meeting provided me an excellent opportunity to meet people in both the 6lowpan and netlmm [Network-Based Localized Mobility Management] working groups and attend their sessions,” she said. “This helped me [resolve] the technical issues and gave me a clear direction to proceed with my ongoing research.”

Gargi said she found the other IETF meeting participants to be approachable and helpful. “IETF is attended by experts from every field related to networks, and it is relatively easy to approach anyone for technical discussions,” she said.

Gargi added that she has shared with her colleagues much of what she learned about the hierarchy of the IETF and how the standards body works. “I gave a talk to my fellow lab mates in the lab seminar related to this,” she said. Hugo Salgado, who’s with NIC Chile, said he was able to gain practical information about how best to deploy Domain Name System Security Extensions (DNSSEC) from early adopters in the IETF community. “We’re preparing to launch DNSSEC for the .CL zone,” Hugo said. “With DNSSEC, we’re following all the current and future work, clarifications, and operational best practices.”

Hugo added that he also found it interesting to hear the thoughts of IETF leaders about the issues of network neutrality and how best to deal with congestion control. “In Chile, there’s an ongoing discussion about network neu-

### IETF 76 Fellows and Mentors

Zartash Afzal Uzmi (Pakistan)

Mentor: Daniel King

Gargi Bag (Korea, Republic of)

Mentor: Sri Gundavelli

Fernando Gont (Argentina)

Ali Tufail (Korea, Republic of)

Mentor: Hamid Mukhtar

Muhammad Yousaf (Pakistan)

Mentor: Fred Baker

Hassan Zaheer (Pakistan)

Mentor: Mat Ford

### Returning Fellows

João Marcelo Ceron (Brazil)

Sandra L. Céspedes (Colombia)

Subramanian Moonesamy (Mauritius)

Terry Rupeni (Fiji)

Hugo Salgado (Chile)

Carlos Alberto Watson Carazo (Costa Rica)

Continued on next page

*ISOC Fellows, continued*

trality, where the government presented a law that the parliament is analyzing,” Hugo said. “This is a topic that is being discussed in many working groups.”

The Hiroshima meeting was Hugo’s fourth IETF meeting and his second as an ISOC fellow. “For my company, the IETF meeting is a great source of advice with the technologies that we use,” Hugo said. “Other people’s DNSSEC experiences are the things we need right now to launch this technology.”

Subramanian Moonesamy, a consultant for Eland Systems in Mauritius, said attending the IETF meetings makes it easier to resolve issues that flare up on the working group e-mail lists. Subramanian’s work involves implementing the protocols developed at the IETF. “Attending the IETF meeting allows me to have face-to-face discussions with the authors of the specifications and raise any issues I may have,” Subramanian said. “An IETF meeting offers an informal venue to resolve differences of opinion. That works out well in an IETF meeting as people are generally open and friendly.”

Subramanian attended both IETF 76 in Hiroshima, Japan, and IETF 70 in Vancouver, Canada, as part of the ISOC Fellowship Programme. “I can participate actively in the IETF because of the ISOC Fellowship to the IETF meeting, and I have improved my technical knowledge,” Subramanian said. “They offer a glimpse of the Internet of tomorrow from a technological perspective.”

Subramanian said the recipients aren’t the only ones to benefit from the ISOC Fellowship Programme. “The ISOC Fellowship Programme provides people from developing countries with the means to contribute to the IETF and make their voices heard,” Subramanian said. “It is also a benefit to the IETF to have more participants from developing countries.” 

## RFC Editor in Transition: Past, Present and Future

**By Leslie Daigle**

In April 2009, the RFC Editor published RFC 5540: 40 Years of RFCs, which summarized the RFC Series’ publication history. The series has been the technical publication series for Internet technology since long before there was an Internet Engineering Task Force. While the RFC Series is the publication vehicle for the IETF, it has been, and remains, scoped more broadly than that. This is captured in RFC 4844: The RFC Series and RFC Editor:

The RFC Series is the archival series dedicated to documenting Internet technical specifications, including general contributions from the Internet research and engineering community as well as standards documents.

For the past three of the four decades of the series’ history, the RFC Editor work has been carried out at the University of Southern California Information Sciences Institute (USC/ISI). At the time of this writing, as 2009 draws to a close, the RFC Editor role is facing another evolutionary step: The work involved in managing the overall series is being split up to recognize the different components of the editing, production, and archiving activities and to lay the groundwork to ensure its continued success, as outlined in RFC 5620: RFC Editor Model (Version 1).

At the IETF 76 plenary, USC/ISI and the role it has played in supporting the RFC Editor over the past 30 years were given special recognition. Some members of the team will move from USC/ISI to the RFC Editor’s new home, where they will continue their work. We took the opportunity to sit down and talk with current and future RFC Editor staff and advisory board members, including current RFC Editor staff members Bob Braden, Sandy Ginoza, and Alice Hagens, as well as Bob Hinden, who is a member of the RFC Editor advisory board.

### The people behind the RFC Editor

Jon Postel was the first RFC Editor,

starting it in 1969 as an activity to keep track of RFC Series documents. Bob Braden, who was then part of the ARPANET research programme, told how he got started with the RFC Series: “I wrote my first RFC in the early 1970s, when it was somewhere around RFC 100. I was at that point manager of programming for the Computing Center at UCLA, and ARPA wanted to connect it to ARPANET as a resource.” This was all pre-TCP/IP, and Bob’s staff had to implement file transfer and telnet. At the same time, Jon was a graduate student at the University of California, Los Angeles, and Bob worked with him as a colleague. It was before Jon got his Ph.D. and moved to SRI in 1973–74. In 1980, Jon moved to USC/ISI, taking the RFC editorship with him. Joyce Reynolds went to work for Jon at USC/ISI. She did much of the actual editing and became an important part of making the RFC Editor. Jon was responsible for quality control, running the operation, and generally being the series editor. When Jon died suddenly in 1998, Bob, who joined USC/ISI in 1986, and Joyce both felt a keen sense of loss. “Jon was a very remarkable guy in many ways,” Bob said. “We knew how much the RFC Series meant to Jon, and we volunteered to carry it on.”

Sandy Ginoza joined USC/ISI to work on the RFC Editor activity in 1999, just after Jon passed away. Alice Hagens came on board in 2005, taking on more of the computer-oriented aspects of the work.

## RFC Series

Although we tend to reference and read individual RFC documents, it is important to understand that there is significant value in the collection of published RFCs as a series. On the importance of the RFC Series, Bob Hinden said, “This community is IETF focused, but to the larger world not centred around the IETF, it’s really the RFCs that are how you build the Internet. One of the things that made the Internet possible was the RFC Series: that you could build things and deploy things without coming to IETF meetings was valuable.” Bob went on to outline his own experiences, such as meeting engineers in Taipei, for whom it was the first time they had ever met anyone who had written an RFC. Even the notion of going to an IETF meeting was in another dimension. “The RFC Series is what enables people to build products, networks, and the Internet,” he said.

And it is quite an active series. Currently, some 300 documents (10,000 pages) are published every year, and while it might be interesting to review the material to detect trends or arcs of work in the Internet technical community, that type of activity is beyond the current scope of the RFC Editor. Focusing on consistency of the series, Bob Braden wondered, “Will we eventually have good enough statistics from the errata system to gauge our error rate?”

The intent of the RFC Series is to serve the broader Internet community; it is not just for or by the IETF. Sandy’s perspective on the value of the Independent Stream of RFCs is that “it offers an alternate view than what happens in the IETF and what working groups have decided to take on as part of their chartered activities. It’s good to document that work was done, results were generated, lessons learned, etc. ‘We tried it; don’t do it this way.’ We often get asked why it’s called RFC when we’re not really requesting comments anymore, but that is the genesis, and the Independent

Stream keeps some of that alive.”

Bob Braden offered his own perspective on the Independent Stream. “Historically, the RFC Series is supposed to be larger than the IETF, and while Jon was alive, the editor did whatever he thought he ought to do; the community didn’t question it much.” However, in the absence of Jon as an authority figure, the community began to ask questions and build its own set of beliefs, eventually coming to believe that RFCs were only for the IETF. That was resolved with RFC 4846, which explained that there is a separate set of independent submissions that do not come through the IETF. “It’s not a big stream, not a lot of documents, but it is important philosophically,” Bob added. “The Internet community is bigger than the IETF.”

The RFC Series is, nevertheless, entwined with the IETF and its activities. For instance, the discussion of (IETF) intellectual property rights (IPR) has led to an impasse in assigning boilerplate to RFCs that allow the continued publication of the Independent Stream documents. That issue is being worked on and resolved, but it offers an example of some of the complexities—and frustrations—that can arise as part of the RFC Editor process. “The current situation—that the independent submissions cannot be published because we don’t know what the boilerplate is—is just terrible,” said Bob Braden.

Bob Hinden, who has been tracking the IPR work from the IETF side, agreed and elaborated on some important lessons learned: “The IETF created a process in the IPR working group that was focused on trying to provide a solution to what they perceived as a problem.



Alice Hagens, Bob Braden, and Sandy Ginoza are recognized at IETF 76 for their work with the RFC Editor

Photo/Internet Society

But they lost sight of the complexity and cost of implementing that solution compared with the actual risk of something bad happening. We have learned a lot about doing this in the future. This isn’t like a protocol spec where you fix a bug in the finite state machine. This has a real effect on people doing stuff. When you ask for legal opinions you get the answer about how to solve the problem, but that’s not the end of the process. You need to balance the cost of solving the problem with the risk of what you’re trying to avoid. Lawyers are supposed to give you the lowest-risk answer. You need to follow through with questions about likelihood and consequences. This is all great hindsight, and I hope we can apply it in the future.” Hinden also said he believes the current impasse could have been avoided if the new procedure had specified that it go into effect when appropriate supporting conditions were met, instead of on a specific flag day, such as the date of publication of the RFC.

The impacts of entwining the RFC Series and the IETF go both ways. For example, the RFC Series recognizes three levels of standards documents: Proposed, Draft, and Full. The expectation, documented in the IETF standards process, is that standards-track specifications should be published as Proposed and then advanced to Draft and Full as

*Continued on next page*

*RFC Editor in Transition, continued*

the specification gets tested commercially and acknowledged as appropriately mature to move to the next stage. In reality, as observed at the IETF 76 Wednesday night plenary, many of the important specifications that form the basis of the operating Internet are still published only as Proposed Standard. Bob Braden explained the history of the standards-track RFC maturity system this way: “Labels were invented whole cloth by the original IAB [Internet Architecture Board], who were a bunch of academics. At that point the Internet had not been commercialized—there were no commercial pressures—so we imagined that it made sense to step through progressions in a theoretical world. In the real world, companies are putting out products. There is no financial incentive for people to spend time advancing documents. Plus, the IETF is so large and there are so many working groups that we try to dispatch them as fast as we can; there is no one around to advance a document.” There have been, and will continue to be, proposals for moving important, current standards (such as BGP) forward in maturity or for collapsing the maturity scale and labeling system.

On the fun side of the RFC Series, there remains a tradition of “April 1st” RFCs. “That people want to participate in that is cool,” said Sandy. “And we get to see the runners-up and the really-not-so-good ideas!”

Alice agreed, adding that “there are high standards for straight-faced satire.”

**RFC Editor**

Traditionally, the RFC Editor has not only populated the series with new (approved) documents but also kept all the threads together in the RFC Series. Describing the role’s origins, Bob Braden pointed out that “originally, Jon was prince of his kingdom. As RFC Editor, he was honorary member of the IAB, informally called the protocol czar. He

used the RFC Editor position to actively prevent bad ideas from getting pushed. Jon imposed a consistency of style on the document series. You pick up RFC 1001 and compare it with 2001, and they look very similar.” Jon believed, and the RFC Editor continues to believe today, that consistency was a worthwhile attribute, promoting stability in the series.

---

**“One of the things that made the Internet possible was the RFC Series: that you could build things and deploy things without coming to IETF meetings was valuable.” — Bob Hinden**

---

Reflecting back, Bob Braden said, “In discussions over the last five years, people have expressed the view that we don’t need an RFC Editor just take an Internet-Draft and publish it. Drives me crazy. The implication is that it doesn’t matter whether it is good English, correctly referenced, consistent, etc. I can’t stand that view.” One of the arguments for such an approach to IETF document publishing is that editing can inadvertently alter, and thereby introduce errors to, text. But the RFC Editors understand that.

Alice said changes to text can be problematic, “partly because of the technical content and partly because it is a group process. It’s agreed-upon text. The idea is how precious the text is and how a slight change can make a large difference.”

Sandy agreed, adding that “for as many changes that get pushed back upon, there are many that make it through the process: for as many people as look at the document before it gets to us, there are things that escape them; there is often missing text, missing words.” According to Alice, with working group documents, people are often focused on getting the technical ideas right, but nobody has sat down and read the text from beginning to end. In addition, there are many in the community who are not native English speakers. It all comes back to the consistency and

professionalism of the output of the series.

**RFC Editing Process**

As the RFC Series has grown, achieving consistency has required the creation and refining of processes. “When Joyce and I took over,” said Bob Braden, “we built the Web site and regularized a

lot of things, and the community began to ask, ‘Why do you do it that way?’” In response, the editors started publicizing the style manuals they used. Joyce and Bob generated a lot of rules that have become institutionalized.

Of course, there is continuing evolution. Bob Braden noted that the addition of errata was his idea, although “it has turned out to be a much, much bigger deal than ever imagined, as is often the case,” he said, laughing. “Now we’re talking about adding image files to solve the problem of incorporating graphics in an ASCII RFC. John Klensin and I generated a plausible solution for that, and we hope to get it installed soon.”

It is important to note that there are some edits the RFC Editor will not make. According to Sandy, the RFC Editor tries to ensure consistency of terminology and to make recommendations that improve consistency within a document both in a technical sense and within the series. “We don’t change the active/passive voice,” she said. “We might suggest it, but we are concerned that it would affect the author’s intent.” Being conservative is key. Sandy said she was surprised by how “simple grammatical changes can have a serious technical impact; placement of a comma can make a big difference in how people read the document and what they implement.”

Working with authors is an important part of making the editing process suc-

cessful. Innovations such as having the RFC Editor help desk at IETF meetings and making the AUTH48 (final check of the RFC Editor's edits) more of an interactive dialogue have helped build community and create awareness of how to build a better document that conveys the meaning as intended. "It is extremely useful to get discussions started earlier, which lessens problems during AUTH48," said Alice. She added that it has also been useful to have face time with the developers of community-created tools, such as xml-2rfc and the ABNF checker, which have been instrumental in improving RFC production. Office hours, building relationships, and face time "all help make it about working together," said Sandy.

Looking forward, Sandy said she would like to see the RFC editing process (and series) "grow and continue to be more consistent, with better community relations and more transparency so authors can look at our site and better understand the process, instead of thinking their document has gone into a black box."

### On the Verge of Major Change

As this article is written, the RFC world is on the brink of major structural change. Following IAB-led community discussion, there is a new model for recognizing the components of activity that make up the RFC activities. ISI is handing off the RFC Editor activity, which will be taken up by separate organizations working together. The status of those structural changes (and their implementation) is noted in the sidebar (see box, this page), provided by IAB chair Olaf Kolkman. In October 2009, AMS was awarded two-year contracts to manage the RFC Production Centre and the RFC Publisher.

Sandy will be joining Association Management Solutions (AMS) as RFC Production Centre director and Alice will be joining as senior editor and information technology development project

manager. To the question of whether the current RFC advisory board will carry forward in the current format or will change, Bob Braden answered, "The current board serves two functions: It provides a supply of experienced people who review independent submissions, but it also gives the RFC Editor advice on policy issues. Some members of the advisory board are very strong members of the IETF in terms of policy advice. In forming the board, I tended to identify a subset of people within the IETF who have long IETF and publishing experience. In the new world there will be an RSAG [RFC Series Advisory Group], which will take over the policy discussions that are currently being conducted by the editorial board. In practice it will be the same people, at least for a while, but with separate duties. That separation is useful."

In considering the change of organizations, Sandy said the biggest thing in moving to AMS is that it is a more service-oriented environment. "In the new model," she said, "it is important that the ISE [Independent Submissions Editor] and RSE [RFC Series Editor] be respected individuals who are granted some of the independence the RFC Editor had at ISI."

Alice added that the biggest shift will be the institutional memory of Bob Braden, but, as she explained, "there were 10 years of overlap, so it's not going into a black hole. And much of it is written down in the document series. I'm confident that the continuity of the series won't be lost by the move to AMS."

Bob Hinden offered another perspective. "I think one of the positive things that has come out of the new model that has gotten lost is this: A lot of people in the IETF didn't understand where the series had come from, or why the IETF chose to use it," he said. "It is the formalization that there are different streams that have different rules. Before, this was confused with the IETF

standards process. Going forward we'll have the opportunity to use the RFC Series for other relevant Internet publication streams that have not been part of IETF. Now we have a framework that would allow that."

Although it is on the verge of major changes, the RFC Series and RFC Editor function are clearly continuing what has been a long process of constant evolution and change. This is just a new chapter in the series' history. 

*This article is composed of interviews conducted by Leslie Daigle and Lucy Lynch, and notes compiled by Mat Ford.*

### Reorganizing the RFC Editor

The RFC production centre and RFC publisher have been selected after an IAOC-led procurement process. AMS takes over the responsibility for these functions as of January 1, 2010. The two other functions that the RFC Editor model (RFC 5620) defines are the RFC Series Editor and the Independent Submissions Editor.

The priority of the IAB is to appoint an RFC Series Editor. After a first acquisition round, it was decided that during the first couple of months, the focus would be on managing the transition from Bob Braden's team at ISI to the multicomponent role as defined by the RFC Editor model. As such, the responsibilities and profile of the person performing that transition are likely to differ from the responsibilities and profile of the person performing the role after the transition.

At the time of this writing, the advisory committee that assists the IAB is working on a recommendation to the IAB for the RSE function. After that job is finished, the committee will continue by working on a recommendation for the ISE position.

In the meantime, Bob has been found willing and able to provide background and assistance for the transitional RSE and to keep the wheels in motion while that RSE has not yet started.

# Congestion Exposure: We're All in This Together

By Philip Eardley

One of the fundamental characteristics of the Internet's architecture is that capacity is shared on a packet-by-packet basis (there are no circuits). Today, several mechanisms are used to achieve such capacity sharing, including TCP, deep-packet inspection (DPI), and volume caps. The mechanisms for achieving capacity sharing are partly cooperative (such as between two users' TCP algorithms) and partly competitive (such as between TCP and DPI or between one user's 10 TCP flows and another's single TCP flow). This is part of what is called the tussle in cyberspace.<sup>1</sup>

Capacity sharing and congestion are two sides of the same coin. A (good) transport protocol needs to fill the bottleneck link, or else the network will be underutilized. That means a certain amount of congestion is a good thing. Thus, sharing out capacity is also about sharing out congestion. Note also that the amount of congestion a sender suffers is equal to the amount it causes.

All of this suggests that it is sensible to try to make capacity management more cooperative—among and between users, content providers, Internet service providers (ISPs), and carriers. We believe doing so will lead to more-effective use of bottleneck links and encourage deployment of more capacity where required.

Congestion exposure seeks to promote such cooperative capacity management. It would constitute a new capability for the Internet—one in which any node would be able to see the rest-of-path congestion; that is, the congestion between the node and the destination.

IETF 76 in Hiroshima, Japan, saw a well-attended BoF session, where it was decided that the IETF should work on congestion exposure. A new working group (WG) called conex is expected to be formed by IETF 77.

## How Congestion Exposure Works

Currently, information about congestion is visible only at the transport layer in the end systems (the information is

hidden from the network layer). We would like to see the congestion information in the header of each IP datagram, which would make it visible to all nodes in the network. This will require a new mechanism that enables senders to inform the network of the level of congestion they expect their packets to encounter.

Such a mechanism means having the sender include in the IP header an indication of the current congestion across the path to the destination. A router suffering congestion adds Explicit Congestion Notification (ECN) marks to packets (as defined in RFC 3168) or else drops them. (The latter course of action is less desirable, but proper discussion of this point is beyond the scope of this article.) The destination then reports back to the sender the total congestion, which completes the loop (Figure 1).

The whole-path-congestion information, which is added by the sender, does not get modified en route. So, by subtracting the congestion-so-far (the information carried by the ECN marks), any node can infer the rest-of-path congestion.

One obvious requirement for true functionality is accurate information. For example, can the sender lie about the amount of conges-

tion it is causing in order to gain an advantage? If so, one potential mechanism for preventing this type of cheating might involve having the final network on the end-to-end path check that the sender's declaration of the total congestion matches the actual congestion experienced. If there is a discrepancy, then the final network may impose a punishment, such as dropped packets.

## How to Use Congestion Exposure

First, let us consider how an ISP might use our new congestion exposure metric (we will discuss the benefits to end users shortly). Today, most ISPs use techniques like DPI, volume capping, and FairShare (see The Bandwidth Bandwagon on page 1 for an explanation of the FairShare system as deployed by Comcast) to deal with bandwidth shortages, especially in the busy hour. Essentially, they are trying to improve the experience of the majority of users at the expense of a limited group of high-bandwidth users and bandwidth-intensive applications. For instance, the ISP may use DPI to bias against applications it considers low value (perhaps peer to peer) in order to improve the experience for other users and applications.

Congestion exposure helps the operator by improving the granularity of the information available to inform bandwidth-management mechanisms. With congestion exposure, the operator sees the current congestion, which is an indication of the actual stress on capacity, as opposed to just a count of

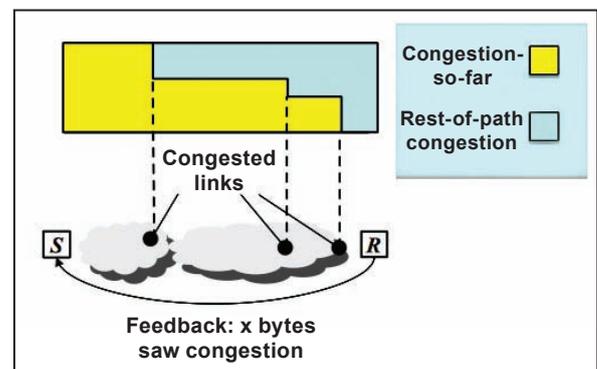


Figure 1

the volume over the course of a month (or the busy hour), which is an extremely crude metric. The operator also sees the whole-path congestion rather than just conditions within its own network.

Contracts between ISPs and their users could be modified to take account of the congestion exposure information, such as by attaching a congestion allowance clause. The allowance could be in the form of a token bucket, wherein a user is allowed to cause congestion (on the whole path) up to a particular rate and burst size, but anything higher would result in consequences, such as dropped traffic or traffic that gets forwarded at a lower priority. All contracts could look the same as today (at least to nongEEK users) and include a tier of options, such as basic, medium, and advanced.

The key difference between this new congestion exposure paradigm and what is happening today is that in the new paradigm, an end host's operating system could potentially optimize the user's experience by favouring some of its applications. The operating system deduces when the user's congestion allowance is endangered, and it then balances congestion between its applications. Perhaps during a period of heavy congestion, the user's videoconferencing would continue at full rate while a file download would be paused. Similarly, if the download happened to be more important, then that could be favoured instead.

Applied in this fashion, congestion exposure would encourage the use of scavenger transports (similar in aim to LEDBAT; see <http://www.ietf.org/dyn/wg/charter/ledbat-charter.html>) that preserve the user's congestion allowance for more-important applications. The extra benefit of congestion exposure is that it enables fully flexible yet simple differentiation of QoS (quality of service) under the control of the end user's applications and operating system.

By enabling greater end-system control and freedom over transmissions, congestion exposure makes it possible for operators to increase their network's efficiency. In particular, the combined utility of an operator's users will increase. It is exactly this kind of cooperation between users and networks that lies at the heart of our hopes for congestion exposure.

Interestingly, congestion exposure is application and protocol agnostic. Deciding which applications consume the user's congestion allowance is under the control of the user's particular preferences; in other words, the network doesn't care, so DPI is no longer needed. Thus, regulators requiring non-discriminatory traffic management should also be happy.

### Other Benefits of Congestion Exposure

There are other potential benefits of congestion exposure, for both users and networks. For example, when there is little congestion, the application can run at a much faster rate than TCP would achieve. When there is congestion, a sender is likely to favour short transfers over large ones (because they use up less congestion allowance); for example, a sender might favour Web browsing over peer to peer. As Figure 2 shows, the result is that the short transfer now completes much more quickly, while the larger one takes barely any longer.

We also believe that congestion exposure can improve an operator's incentive to invest in new capacity. One of the stumbling blocks to investment that we face today is that a few users tend to grab nearly all of the extra bandwidth while the cost is spread out over all users through fees. With congestion exposure, an operator is motivated to invest because the benefit is more evenly spread out (or targets those who want to pay for

it). An operator may also be able to identify which of its links have the greatest incipient demand and hence determine where to focus that investment.

Other potential benefits include new tools that exploit the congestion information for DoS mitigation, traffic engineering, and internetwork service-level agreements.

### The Future of Congestion Exposure at the IETF

The proposed conex WG will concentrate on the specification of how rest-of-path congestion information is carried in IP packets. This does, however, require standardizing a change to IP, which is not an insignificant step!

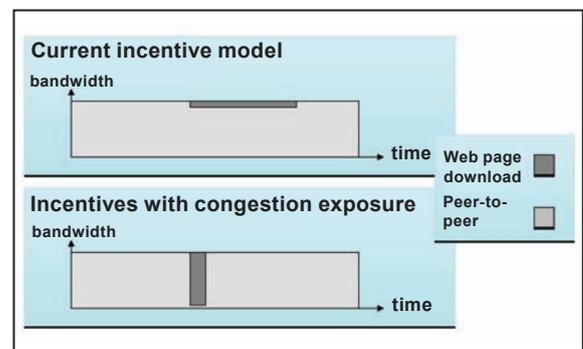


Figure 2

Conex will also work on how to transport the whole-path-congestion information from the destination to the sender, as well as how to prevent parties from being less than truthful about congestion. The WG will encourage experiments to determine which use cases are most useful and hence worthy of deeper consideration.

It is hoped that the work of the proposed conex WG will lead to better cooperation between users and networks—and so achieve better capacity sharing on the Internet. 

For more information, see <http://trac.tools.ietf.org/area/tsw/trac/wiki/re-ECN>.

### Reference

1. D. Clark, J. Wroclawski, K. Sollins, R. Braden. "Tussle in Cyberspace: Defining Tomorrow's Internet"; IEEE/ACM Transactions on Networking (TON), Volume 13, Issue 3 (June 2005), <http://portal.acm.org/citation.cfm?id=1074047.1074049>.

# The Kantara Initiative for Online Identity: A One-Year Progress Report

By J. Trent Adams and Eve Maler

Founded in April 2009, the Kantara Initiative was conceived as an open, global organization with the mission of promoting interoperability and technology harmonization across the myriad identity solutions available and under development. With the proliferation of single-protocol solutions being pursued, the founders of the Kantara Initiative set out to promote the deployment of heterogeneous protocols, standards, and solutions for vendors and end users within the entire network identity ecosystem.

As can be expected, it was a huge undertaking to corral enough industry and community support for this idea to move onto the world stage. Key to the success in setting it up was the contribution of time, effort, and intellectual capital from the Liberty Alliance. In 2008, the Liberty Alliance board of trustees foresaw the need for a new organization that would support all of the major industry solutions and community standards, promoting component-level interoperability while also addressing the related business, legal, and regulatory issues. Thus, the Liberty Alliance began the effort of reaching out to other organizations to form Kantara.

Among the organizations agreeing to found Kantara with this mission were the Internet Society, the Information Card Foundation, the DataPortability Project, and XDI.org. This extended set of organizations added to the existing industry leaders within Liberty, such as AOL, British Telecommunications, CA, Intel, Oracle, and Sun Microsystems. Together, they began working on structuring an inclusive environment conducive to multiple points of view and to collaboration at a global scale. The result is an organization with no membership fees to participate and with a truly transparent operating structure.

Since its inception, the Kantara Initiative has fostered a robust environment in which a wide variety of identity ecosystem challenges are being tackled. Among the active work streams are:

- Promoting component (i.e., protocol-level) interoperability
- Ensuring global interoperability within a vertical market but also across markets (e.g., finance to health care)
- Developing business and policy best practices related to end-user engagement (e.g., contractual intellectual property rights between entities, and privacy issues)
- Advancing government and regulatory compliance; helping guide legislation that honours the privacy of end users and protects user-managed access to users' data
- Fostering identity assurance programmes to support business and government trust requirements
- Promoting recognition of end-user usability needs in developing industry solutions

These threads can be seen in the various Kantara working groups. While all operate under one governance umbrella—and communication and coordination among groups are encouraged through joint membership on a Leadership Council—each working group has a distinct mandate and set of deliverables.

In order to support such a diverse set of activities, the Kantara Initiative took an unusual approach regarding the protection of intellectual property contributed to the groups in service of the Kantara charter. This approach (1) requires each group, at inception, to

select the intellectual property rights (IPR) regime that best fits its goals and the eventual standards organization to which specifications may be submitted and (2) requires participants to agree to the regime before taking their places at the table. Such flexibility allows each group a quick and easy way to get up and running with the IPR option that suits the group's members while helping protect developers of Kantara-incubated specifications.

Of the 18 groups chartered since April, some of the most active are:

- *E-Government*: Facilitates collaboration and discussion among Kantara groups with an interest in e-government identity management applications and services. This group acts as a forum to discuss best practices by government organizations on national, regional, and municipal levels and offers government-subject-matter expertise in the development of Kantara Initiative policy recommendations and specifications.
- *Healthcare Identity Assurance*: Designs, implements, and tests reference applications for secure access to health information. One example use case under consideration is for consumers to be able to access their health records with a standardized log-in system. Another is a way for emergency workers to access critical health information during emergencies or natural disasters.
- *Identity Assurance*: Fosters the adoption of trusted online identity services, identifying and resolving specific obstacles to their market and commercial acceptance. This group is actively working toward the development of a global standard framework necessary to support trusted identity service providers.
- *Information Sharing*: Identifies and documents use cases that illustrate the benefits and challenges of user-driven information sharing. By

focusing on the benefits and addressing the obstacles, the group specifies the policy and technology solutions that are required to enable a smooth and effective information flow.

- **Privacy and Public Policy:** Focuses on the interplay between privacy, technology, and policy, and aims to ensure that Kantara contributes to better privacy outcomes for users, data custodians, and other stakeholders. This group engages with a diverse range of privacy stakeholders, understanding their different perspectives, translating and mediating between them as necessary, and documenting privacy-related principles and good practices applicable to a broad range of prevalent technology platforms.
- **User-Managed Access:** Develops specifications that let an individual control the authorization of data sharing and service access made between online services on the individual's behalf, and to facilitate interoperable implementations. The group expedites the process of collaborating with different communities on a draft solution that meets their shared goals that cross community boundaries.

Rather than setting up another standards body, the Kantara Initiative focuses on incubation of ideas and concepts. If specifications emerge from the groups, they are then submitted to other standards-setting organizations

Rather than setting up another standards body, the Kantara Initiative focuses on incubation of ideas and concepts. If specifications emerge from the groups, they are then submitted to other standards-setting organizations for adoption and operational maintenance.

for adoption and operational maintenance. Each chartered group that anticipates producing specifications selects the standards body to which it expects to contribute its work when it is fleshed out. For example, some groups are targeting the W3C, others are looking to OASIS, and some have their sights on the IETF.

A prime example of this process is the User-Managed Access (UMA) working group. The work incubated in UMA illustrates how the Kantara Initiative fosters innovation, interoperability, and community. This group is developing specifications that empower an individual to control the authorization of data sharing and service access made between online services on an individual's behalf, thereby allowing for permissioned data sharing even if the user's entire set of data is hosted on many different servers.

To foster the adoption and building of a modular solution, the group is profiling and extending OAuth-related specifications already under development at the IETF (the ultimate destination identified for UMA specifications)

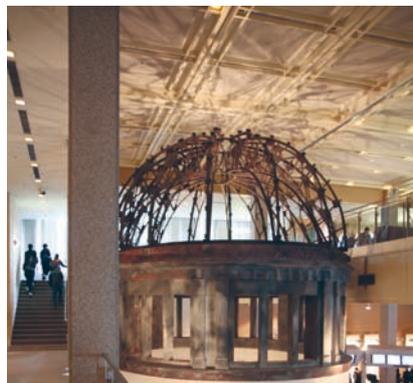
while attracting participation from a wide variety of stakeholders who might normally have difficulty participating in such an effort were a membership fee required. One of the group's goals is to facilitate multiple interoperable implementations, and it has worked with Kantara leadership to develop a bounty programme for attracting development interest in a UMA protocol validator.

With work like this under its umbrella, the Kantara Initiative has moved from concept to fully functioning reality with broad global support across all sectors of the identity ecosystem. And with a growing membership—including additions such as Neustar, PayPal, NTT, Danish National IT, Deutsche Telekom, and the government of Canada—the organization is healthy and beginning to hit its stride. Participation is key to the success of the mission, and with stakeholders easily able to sit at the table, the organization will surely realize its full potential.

To learn more and join the discussion, visit <http://kantarainitiative.org>.



The A-Bomb Dome, which survived the A-bomb blast, is part of Hiroshima Peace Memorial Park in Hiroshima, Japan.



Model of the Hiroshima Peace Memorial at the Hiroshima Peace Museum.



Scale model of Hiroshima city flattened after the detonation of the A-bomb. The red ball depicts the explosion point.

## IRTF Update

By Aaron Falk

Four research groups (RGs) met at IETF 76 in Hiroshima, Japan: Host Identity Payload; Scalable, Adaptive Multicast; Delay Tolerant Networking; and Routing. Most of the 13 RGs that make up the Internet Research Task Force (IRTF) are active.

Six IRTF drafts are awaiting publication by the RFC Editor. The publication process has been in suspension for several months pending revision of the IETF Trust License Provisions. The current expectation is that publication will resume in January 2010. One of the documents in the RFC Editor queue is draft-irtf-rfcs-05.txt, which defines the IRTF RFC stream. The draft was removed from the publication queue so that the rights language can be made consistent with that of the Independent Stream.

From time to time the IRTF considers proposals for new RGs. Several folks who have expressed interest in an RG on virtual networks have been meeting informally for some time. A Bar BoF on the topic was held at IETF 76 for the purpose of reaching agreement on a draft charter. Some interest has also been expressed in establishing an RG with a charter related to the Internet of Things.

The Routing RG met with the Internet Architecture Board at IETF 76 to discuss plans for making a recommendation to the IETF based on the group's work over the past two years in evaluating proposals for a future routing and addressing architecture. The work was motivated by concerns over forwarding table growth in the default-free zone. There has been a great deal of energy in the RG, and while progress has been made in defining the nature of the problem and gaining an understanding of the classes of solutions available, it will be a challenge for the group to converge on recommendations.

The End-to-End RG, which is chaired by Craig Partridge and Karen Sollins, has decided to close in January 2010 after 26 years. The RG has served as a focal point for several important concepts in Internet design, including slow start and improved round-trip time estimation, Random Early Drop, Integrated and Differentiated Services, Weighted Fair Queuing, PAWS (Protection Against Wrapped Sequence Numbers), and Transaction TCP. The end2end-interest mailing list will continue operation.

At each IETF, brief overviews of selected IRTF RGs are presented to help familiarize the IETF community with the research topics that are under discussion. At the IETF 76 plenary, short presentations described the Anti-Spam RG and the Scalable, Adaptive Multicast RG.

The Anti-Spam research group, chaired by John Levine, looks at open problems in topics related to combating unsolicited e-mail. The original hope was that the RG would create initiatives around which standardization could occur. Those initiatives have not yet been created. The group has an open membership, which includes representatives from industry and academia as well as independent participants. Combating spam has become a major industry, one that has led to the formation of an IETF working group (DKIM), trade groups (MAAWG and ESPC), and conferences (CEAS and presentations to Usenix). The RG has produced a document that provides details on Domain Name System blacklists and white lists (the document is awaiting publication), and it is working on a draft that covers management practices for blacklists. Another



Aaron Falk, IRTF Chair

---

Combating spam has become a major industry, one that has led to the formation of an IETF working group (DKIM), trade groups (MAAWG and ESPC), and conferences (CEAS and presentations to Usenix).

---



ISOC Fellows and mentors meet over lunch at IETF 76



IETF 76 take in the Bandwidth Bandwagon panel discussion



Attendees enjoy some down time at IETF 76 in Hiroshima, Japan. Photos/Internet Society.

work in progress for the RG covers taxonomies on antispam and spamming techniques (see <http://wiki.asrg.sp.am>). The RG serves as a good sounding board for folks who have antispam ideas.

The Scalable, Adaptive Multicast RG (SAMRG), chaired by John Buford and Thomas Schmidt, is investigating the development of unified approaches to multicast that take advantage of link-layer multicast, IP multicast, and application multicast when they are available. There are several active researcher/developer communities in the SAMRG, including application layer multicast over peer-to-peer networks; an experimental protocol called XCAST, which is optimized for medium-scale voice conferencing; native (link-layer) multicast; and applications that use multicast. Some of the RG's work is focused on extending the Automatic IP Multicast without Explicit Tunnels (AMT) protocol as an overlay that can tie together multicast-enabled clouds. Other work focuses on defining a common application programming interface and name space that would enable applications to take advantage of multicast when it is available. In the future, the RG plans to work on extending common simulation tools to hybrid multicast mechanisms as well as on developing a wide-area hybrid multicast test bed.

#### 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/DocProtoActions0503.shtml>.

# IETF Meeting Calendar

## IETF 77

21–26 March 2010  
Location: Anaheim, CA, USA

## IETF 79

7–12 November 2010  
Host: Tsinghua University  
Location: Beijing, China

## IETF 78

25–30 July 2010  
Host: SIDN  
Location: Maastricht, The Netherlands

## IETF 80

27 March–1 April 2011  
Host: TBD  
Location: Europe (provisional)

Register now for

## IETF 77

21–26 March 2010

Anaheim, California, USA

<https://www.ietf.org/registration/ietf77/ietfreg.py>

Early bird registration: USD 635 through 1700 PDT March 12, 2010

Regular registration: USD 785

Full-time students: USD 150 with on-site proof of ID

One-day pass: USD 200 (see Web site for details)

Special thanks to



for hosting IETF 76

The ISOC Fellowship to the IETF is sponsored by



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



## IETF Journal

**IETF 76**  
**Volume 5, Issue 3**  
**January 2010**

Published three times  
a year by the  
Internet Society

Galerie Jean-Malbuisson 15  
1204 Geneva  
Switzerland

Guest Editor  
Matthew Ford

Associate Editor  
Wendy Rickard

Contributing Writer  
Carolyn Marsan

Editorial and Design  
The Rickard Group, Inc.

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

E-mail

[ietfjournal@isoc.org](mailto: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

