Volume 6, Issue 1 • June 2010

the IETF® Journal

*A report from IETF 77, March 2010, Anaheim, California, USA. Published by the Internet Society in cooperation with the Internet Engineering Task Force**

Inside this issue

Authentication, P2P Spotlighted at IETF 77

From the Editor's Desk, by Mat Ford

Reducing the number of domain-specific username/password combinations we all have to carry around these days is one of the potential upsides of standardized federated authentication solutions. Leif Johansson of NORDUnet provides us with a review of developments and some of the exciting recent discussion that took place during IETF 77 in Anaheim, California, in his article, "It's the F-Word" (this page).

Photo/Mikhail Strizhov

Anaheim, California, site of IETF 77

Another area of growing interest within the Internet Engineering Task Force (IETF) is that of peer-to-peer technologies (P2P). And in this issue, we get a good overview of the work to date and future directions for P2P in the IETF (see page 16).

IPv6 deployment is a regular topic for discussion at IETF, and the IETF 77 meeting was no exception. ISOC organized a well-attended panel session, adjacent to the IETF meeting, that exposed the growing momentum behind IPv6 deployment (see page 9). In addition, there was output from the 3GPP/IETF workshop on IPv6 transition that took place in March 2010 (see page 10).

Reflecting on the successes and failures of an organization's working life is an important part of developing and maturing as an organization. A new initiative to catalogue the outcomes of IETF work offers just such an opportunity (see page 18), and the case history of Uniform Resource Names leads to an important conclusion about Internet development and the role of the IETF (see page 11).

Also in this issue are our regular columns from the chairs of the IETF, the Internet Architecture Board, and the Internet Research Task Force; coverage of the hot topics discussed during the plenary meetings; and an opportunity to get to know the ISOC Fellows to IETF 77 from around the world.

As always, we are hugely grateful to all of our contributors, and we welcome comments as well as suggestions for contributions to future issues. Readers can send email to ietfjournal@isoc.org.

It's the F-Word

By Leif Johansson

Everything is federated these days. In some cases, particularly when it is waved at a problem, the F-word is not well defined. In other cases there are clearly defined semantics for the word federation. Federated authentication is one of those cases.

The Enterprise Authentication Model

In the late 1990s, the purpose of authentication was to establish a security context using technologies

Continued on page 4

Message from the IETF Chair

By Russ Housley


Russ Housley, IETF Chair

With more than 1,350 individuals from 48 countries in attendance, IETF 77 provided evidence that the commitment to the work being done by the IETF is as strong as ever. The meeting, which was held in March 2010 in Anaheim, California, ended with significant progress being made in a number of areas.

Unfortunately, we were unable to secure a host for IETF 77. As a result, some of the amenities that attendees had come to expect, such as T-shirts and the Tuesday evening social event, were not provided.

Fortunately, sponsorship for the meeting was strong. We received contributions from Booz Allen Hamilton, Comcast, IPSO Alliance, National Cable & Telecommunications Association, Time Warner Cable, and the U.S. National Institute of Standards and Technology. I would like to offer my heartfelt appreciation to those organizations for their roles in making IETF 77 a successful meeting.

On behalf of the entire IETF community, I would like to recognize the contributions and dedication of the outgoing IAB, IESG, and IAOC members. They have provided a real service for the Internet community.

Since IETF 76, 11 new working groups (WGs) have been chartered, and 6 WGs were closed, resulting in a total of 125 chartered WGs. Between the two meetings, the WGs and their individual contributors produced 521 new Internet-Drafts and updated 1,181 existing Internet-Drafts, some more than once. The Internet Engineering Steering Group (IESG) approved 146 Internet-Drafts for publication as RFCs, and the RFC Editor published 124 new RFCs.

The Nominations Committee completed its work prior to IETF 77, and the membership of the Internet Architecture Board (IAB), IESG, and IETF Administrative Oversight Committee (IAOC) transitioned in April. On behalf of the entire IETF community, I would like to recognize the contributions and dedication of the outgoing IAB, IESG, and IAOC members. They have provided a real service for the Internet community.

In particular, the following dedicated individuals deserve everyone's thanks and appreciation: Gonzalo Camarillo (2008–10), Stuart Cheshire (2008–10), Gregory Lebovitz (2008–10), Andrew Malis (2008–10), and David Oran (2006–10) of the IAB; Ross Callon (RTG AD: 2006–10), Lisa Dusseault (APP AD: 2006–10), Pasi Eronen (SEC AD: 2008–10), Cullen Jennings (RAI AD: 2006–10), and Magnus Westerlund (TSV AD: 2006–10) of the IESG; and Fred Baker (2007–10) of the IAOC.

I challenge the incoming IAB, IESG, and IAOC members to provide comparable levels of service and dedication.

The RFC Editor has implemented the changes documented in RFC 5741. Please note that the upper left corner of the title page indicates the source of the RFC. I'm also happy to report that all of our RFCs now proudly read Internet Engineering Task Force (IETF) on the title page.

The RFC Editor has provided status pages for every RFC. Pick your favourite RFC and check it out at http://www.rfc-editor.org/info/rfcXXXX.

IETF 78 will be held in Maastricht, Netherlands, 25–30 July 2010, hosted by SIDN. As always, scheduling information for upcoming IETF meetings can be found at http://www.ietf.org/meeting/upcoming.html

I look forward to seeing you in July.

Words from the IAB Chair

By Olaf Kolkman


Olaf Kolkman, IAB Chair

Every year we would go through the same ritual. In late winter, sometimes in early spring, my mother would ask me to sort through my toys. I was allowed to keep most of them, but some had to go. It was every year, shortly before my birthday.

Not an easy task for a boy in primary school.

Keeping the Legos and the fischertechnik were easy choices. But what to get rid of? The Tin Car and the plastic soldiers I used to simulate an infantry? The Tin Car and small plastic model plane were in use nearly every week when I played airport. I hardly ever played with the toys individually; however, when I used them in pairs or triplets, the hours would dissolve as I busily relived a recent episode of the Thunderbirds.

Gaining fresh perspectives while cherishing previous ones is a good way to evolve—not only for technology but also for organizations.

Every so often, true magic occurred. I would combine Lego, fischertechnik, planes, soldiers, a wooden garage, and a train (one of those wooden ones they now sell at IKEA) to form a gigantic city. Me and the boys from the neighbourhood would spend half of our vacation playing with the city, which soon would be left to deteriorate when we had to attend school again. I recently touched base with one of my childhood friends on Facebook, and we mutually agreed that the inspiration for our technology and engineering careers had its roots in those spring breaks.

That spring break city must have been in the back of my mind as I sorted through the toys. I didn't want to get rid of the plastic barn or the gas station model; they were seminal pieces when constructing the fantasy world. But my mother was strict, stuff had to go, the amount of storage space in our house had an absolute limit, and there would not be room for birthday presents if I didn't act.

So I acted. I disposed of toys, and when my birthday came, I got new toys. Toys I could use to build other cities, cities that would still inspire. Since they were built from different pieces, they would look different, but their essence would be the same; they were always recognizable. I remember the year I had to throw out almost half of the toys, and when I did, they were replaced with new ones on my birthday. I was childishly sad when it happened but also eager to build a new city during spring break.

I want to avoid the risk of carrying the metaphor much further. Some folks may misunderstand and think that I compare people to toys. Far from that. However, gaining fresh perspectives while cherishing previous ones is a good way to evolve—not only for technology but also for organizations.

This year the NomCom replaced Gonzalo Camarillo, Stuart Cheshire, Gregory Lebovitz, Andrew Malis, and Dave Oran with Bernard Aboba, Ross Callon, Spencer Dawkins, Andrei Robachevsky, and Hannes Tschofenig. I would like to thank the folks who left. It was a pleasure to work with them.

The Internet Architecture Board will hold its retreat in June 2010. I hope we will build a city that provides us with a new perspective.

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

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/DocProtoActions0601.shtml

It's the F-Word, continued

such as GSS-API (Generic Security Service Application Program Interface), SASL (simple authentication and security layer), or TLS (transport layer security) between two endpoints of a communications channel together with identifiers that represented users. Back then authentication involved two parties: the client and the server.

In this model, which I'll refer to as enterprise authentication model, the user identity and the user identifier are one and the same. Using a directory service—often one based on LDAP (Lightweight Directory Access Protocol)—the server uses the user identifier to look up additional information about the user. For example, an IMAP (Internet message access protocol) server might store information about user mailboxes in the directory keyed by the user identifier that is used in authentication with the IMAP server. Simply put, the IMAP server can pull information about the user from a directory by using the user identifier as a lookup key.

This model still works well within a single organization and with a single directory. The success of products such as Microsoft Active Directory provides ample evidence that within the enterprise community, this model is alive and well today.

The limitations of the enterprise authentication model arise when organizations need to deploy shared (often business-to-business) applications or when they attempt to merge enterprise directories as a result of corporate mergers or acquisitions. Arguably the failure of the enterprise model to handle cross-organizational authentication deployments has been the result of the failure of LDAP to provide a scalable, vendor-neutral way to authorize secure access to protected data. The risks involved with allowing external access to critical enterprise directory data is typically viewed as unacceptable.

The Rise of SAML

As relationships among enterprises, service providers, and users became more interdependent, the challenge of obtaining lookup permissions from outside of an organization grew. In this new world, in order for a service in organization A to be able to accept users from organization B, the service in organization A needs permission to look up users in organization B's directory. Even if the number of organization A's and services is small, the problem of controlling access to organization B's directories quickly becomes unmanageable.

Initially the solutions seemed to focus on Web applications (for the most part, they still do, though that is starting to change). Toward the end of the 1990s, the need to manage centralized authentication for Web applications drove the development of the so-called enterprise Web single sign-on solutions.

Out of those solutions a number of open-source and commercial options evolved, many of which relied on the management of HTTP cookies. Those solutions were also intraorganizational in nature and, as the need to connect Web SSO (single sign-on) products across organizational boundaries grew, there arose a need for a standardized protocol for communication authentication information between applications.

In January 2001, the OASIS Security Services Technical Committee (SSTC) convened to begin work on what became SAML (security assertion markup language), a technology that has become one of the cornerstones of federated authentication.

The assertion can be thought of as a sealed envelope containing a statement to the effect that the user has successfully authenticated herself as well as the properties of the user the IdP wants to make known to the service provider.

SAML can support several use-cases but the most commonly deployed pattern is called Browser Web SSO. This particular profile of SAML involves three actors: the user, the identity provider, and the service provider. SAML uses XML-based messages and relies on public key cryptography (although not necessarily on public key infrastructure [X.509] or PKIX) to sign and encrypt those messages. In a typical authentication flow the user presents credentials to the identity provider (IdP), thereby proving her identity to the IdP. The IdP then gives the user a SAML message called an assertion, digitally signed with the key of the IdP and optionally encrypted with the key of the service provider. The assertion can be thought of as a sealed envelope containing a statement to the effect that the user has successfully authenticated herself as well as the properties of the user the IdP wants to make known to the service provider. If the service provider trusts the public key of the IdP, the service provider is able to consume the assertion. The assertion often contains an identifier of the user and—more important—additional attributes associated with the user, relieving the service provider from having to conduct additional lookups in directories.

In the SAML federation model, identity information—identifiers and additional attributes—is pushed from the IdP to the SP. Conversely, enterprise authentication operates according to a pull model. This might seem a small difference but it fundamentally changes the way applications (service providers) consume identities.

At this point, the astute reader will undoubtedly ask: "But what about key management?" Indeed, key management is the core of the matter. Some federations still rely on traditional PKIX-style hierarchical PKI for key management. Others, including many large-scale SAML federations, rely on an alternative key-management method. In this method, collections of keys are collectively signed, resulting in an object that behaves like a combination of a PKI certificate and CRL (certificate revocation list). Incidentally, this bag-of-keys model has been considered by the KARP (Keying and Authentication for Routing Protocols) working group as the basis for routing protocol key management (albeit in that case using symmetric keys).

Common to all approaches to key management is a federation that consists of those identity providers and service providers that share trust in a set of keys. Such a trust framework is called a ring of trust. The deployment of OpenID typically relies on a single global ring of trust that encompasses all OpenID IdPs, and SPs. In fact, it is entirely possible that the success of OpenID is due in large part to the absence of a requirement on key management. Conversely, SAML federations often do require key-and-trust management, which constitutes a major part of the work involved in running a SAML federation.

SAML has seen large-scale deployments in the research and education (R & E) sector worldwide. The Trans-European Research and Networking Association (TERENA) conducts an unofficial census of R & E federations through its REFEDS activity, which places the number of R & E federation users at somewhere around 10–12 million per month in the leading countries in Western and Central Europe and in the United States alone. These figures will undoubtedly grow rapidly in the next few years.

Examining Alternative Technologies for Building Identity Federations

Identity federations have been successfully built using other technologies. One of the largest such deployment is the eduroam federation, which uses Radius and EAP (extensible authentication protocol) to provide access to wireless networks (using 802.1xZ) on hundreds of university campus networks across the globe. The success of eduroam demonstrates that the driving force behind federated authentication is almost always the need to share resources across organizational boundaries.

The IETF has been mostly absent from this field but that may soon be changing. IETF 77, which was held in Anaheim, California, in March 2010, saw its first Moonshot project "bar BoF"[1] session. The aptly named project, which is partly funded by JANET (http://www.ja.net/) and by the GEANT3 project (http://www.geant.net/), aims high but the potential benefits justify the effort. It brings together a wide range of IETF standards including EAP, GSS-API (generic security service application program interface), and Radius together with SAML to construct a federation framework that may provide many current IETF protocols, including SSHv2, NFSv4, and IMAP, access to federated identity. If successful, it would extend identity federation beyond Web-only applications, but it would also provide a general trust-framework for the Internet.

Other work related to federated authentication that might end up being done in the IETF address alternative approaches that have been proposed to bridge Web-centric identity (such as SAML or OpenID) and SASL.

The driving force behind a lot of these efforts is the need to federate messaging, calendaring, as well as virtual worlds protocols, which are some of the more important examples of applications where the browser isn't the obvious choice of a client. The needs of the mobile market and its focus on apps may well turn out to be a boon for federated identity. While HTTP is often used as a protocol layer, typically through RESTful[2] service calls, the client is not always a browser in the traditional sense.

The driving force behind a lot of these efforts is the need to federate messaging, calendaring, as well as virtual worlds protocols, which are some of the more important examples of applications where the browser isn't the obvious choice of a client.

Some of the efforts described here have one foot in the IETF and one foot in other standards-development organizations. Some of the work is in even more loosely organized groups of volunteers. There is a huge potential for the IETF to become (and stay) involved in this process but it will require a certain amount of agility on the part of the IETF to keep on top of the changing landscape.

Notes

1. An informal birds-of-a-feather (BoF) session
2. See http://www.ietf.org/mail-archive/web/ietf-announce/current/msg07086.html

NomCom, Day Passes Top IETF Plenary Agenda

By Carolyn Duffy Marsan

Issues surrounding the nomination of new leaders and the availability of day passes at meetings were two of the hot-button issues discussed at the IETF Plenary session in Anaheim, California.

Mary Barnes, chair of the IETF's NomCom for 2009–10, gave a presentation about the many new appointments made to the Internet Engineering Steering Group and the Internet Architecture Board (IAB) this year. She urged the IETF community to get more involved in the NomCom process. "Approximately 10 percent of the IETF community participated," Mary said, asking plenary attendees to consider volunteering to be NomCom members or to contribute nominations or feedback on candidates.

Mary added that she'd like to see a more diverse set of candidates, including representatives from different regions of the world, from service providers as well as vendors, and of both genders. "We have multiple nominees and sitting members with the same affiliation," she said, pointing out that two area directors as well as the IETF chair receive security-related funding from the U.S. government. She said it's been a challenge to find nominees with certain kinds of expertise, particularly in the transport area.

Mary told attendees that lobbying for a particular candidate tends to backfire. "Lobbying for a specific nominee doesn't work," she said, adding that the NomCom must keep information confidential throughout the process. "The outcome and decisions for NomCom09 were not compromised by this activity, but lobbying and leaks have the potential to severely damage the process," she warned.

NomCom issues also were mentioned at the open-mic session. "Lobbying really indicates that the person doesn't understand how this process works and the environment here," said one IETF participant.

Another issue that generated discussion during the open-mic session concerned day passes. IETF Administrative Oversight Committee chair Bob Hinden said the IETF sold 124 day passes for the Anaheim meeting, adding that total attendance at the meeting was 1,234. "We were expecting [day pass buyers] to be people who had not been to an IETF meeting before," Bob said. "We were expecting to see a large number of additional attendees in Hiroshima, but we were surprised to see the numbers here. We actually sold more day passes here than in Hiroshima."

At the open-mic session, several IETF participants expressed concern about the selling of day passes at the Anaheim meeting, pointing out that this practice prevents newcomers from getting integrated into the organization.

Several attendees at the Open Authentication Protocol (OAUTH) working group meeting used day passes. "They had a very limited perspective of how we negotiate with each other in terms of the documents," said a member of this working group. "Many flew off with a negative impression of the group. I'm very concerned that they won't come back, actually."

Photo/Peter Löthberg

Outgoing IESG members (from left) Ross Callon, Lisa Dusseault, and Pasi Eronen

Photo/Peter Löthberg

IETF 77 participants relax between sessions

Another problem with day passes is that they prevent working groups from accomplishing enough business during the weeklong IETF meeting because so much work is done during informal hallway conversations. "One person who used a day pass was a chair of a working group. Another was a prolific writer of Internet-Drafts," said an IETF participant. "They were all in a financial position to pay for the week."

Bob said he'd look at the data covering who bought day passes and that a decision would be made about whether to allow day passes in the future.

IAB Plenary Addresses Internet Consolidation

The trend toward the consolidation of traffic and aggregation of privacy information in the hands of a few Internet infrastructure players was the topic of the IAB Plenary session in Anaheim.

Craig Labovitz, chief scientist at Arbor Networks, disclosed the latest data from his two-year study of the Internet traffic carried by 110 Internet service providers (ISPs). Arbor Networks and University of Michigan

Continued on page 8

IETF 77 At–A–Glance

Photo/Peter Löthberg

Registered attendees: 1350
Number of countries: 48
New WGs: 11
WGs closed: 6
WG currently chartered: 125
New Internet-Drafts: 521
Updated Internet-Drafts: 1181
IETF Last Calls: 128
Internet-Drafts approved for publication: 146

RFC Editor Actions (November 2009–March 2010)
113 RFCs published of which
- 53 Standards Track
- 5 BCP
- 46 Informational
- 9 Experimental

RFC Online: 30+ documents were put online
RFC repository rsync: "everything–ftp" now available
110 Internet-Drafts submitted for publication
- 74 submitted by the IETF WGs
- 27 submitted by IETF individuals
- 9 submitted by IRTF, IAB, and independent submissions combined

IANA Actions (November 2009–March 2010)
1400+ IETF-related requests processed
- 719 Private Enterprise Numbers
- 65 Port Numbers
- 47 TRIP ITAD Numbers
- 81 Media-type Requests

In addition, IANA
- Reviewed 92 I–Ds in Last Call
- Reviewed 116 I–Ds in IESG Evaluation
- Reviewed 107 I–Ds prior to becoming RFC and 59 contained actions for IANA

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

APP	rydeirde	Registry Data Escrow/Internet Registration Escrow
RAI	e2md	E. 164 to Metadata
TSV	conex	Congestion Exposure
GEN	wgdtspec	Review of Datatracker Specifications to Support
APP	rydeirde	Registry Data Escrow/Internet Registration Escrow

NomCom, continued

researchers monitored 14 terabits per second of traffic—approximately 25 percent of all interdomain traffic on the Internet.

What researchers found is that a massive build-out of data centres by leading Internet content providers such as Google and Comcast has changed the topology of the Internet. Traffic no longer flows from national backbone operators to regional access providers, to local access providers, to customers. Now, traffic is carried by large content providers and content delivery networks—dubbed "hypergiants"—which pass it to Internet exchange points or directly to consumers.

The research shows that instead of having ISPs be the top 10 carriers of Internet traffic, Google is now third and Comcast is now sixth. Google alone represents 6 to 10 percent of the Internet's interdomain traffic, Craig said. "Companies like Google are delivering more traffic than global transit carriers are," he said. "What we're seeing quickly evolve is a much flatter, much more densely interconnected Internet. There are significant routing, traffic, security, and economic implications." Craig said that that trend is causing new commercial models to evolve such as paid content and paid peering.

Another shift in Internet traffic that researchers noted is that the fastest-growing applications are video (up 67 percent), secure shell (up 47 percent), the virtual private network (up 36 percent), games (up 29 percent), and Web (up 25 percent). Applications that are declining most include peer to peer, news, and file transfer.

"There's a growing volume of Internet traffic that uses Port 80," Craig said. "We're seeing a rapid concentration of application traffic over an ever-smaller number of ports." Overall, Internet traffic is growing at 45 percent per year, which Craig called a significant but manageable growth rate. He said IPv6 traffic represented only 0.4 percent of Internet traffic as of last year.

Craig summed up the implications of his research for the IETF community, noting that he's seeing the "slow death of end to end" due to network address translation, firewalls, and siloed ecosystems.

In a separate presentation, privacy researcher Balachander Krishnamurthy from AT&T Labs–Research spoke of a similar consolidation trend whereby an increasing amount of personal information about Internet users is being gathered and stored by a smaller number of Internet companies. Balachander conducted a five-year study of 1,200 of the most popular consumer websites to see the cookies, Java scripts, and other mechanisms for gathering personal information about visitors that is sent to hidden third parties. Those third-party sites include ad networks, analytics companies, and content delivery networks.

Balachander found that the top 10 authoritative Domain Name System servers (ADNSs) connected to 78.5 percent of the visible information-gathering nodes on popular websites. Those top ADNSs were operated by ad-serving and traffic-measuring companies such as Doubleclick, Google, Yahoo! and Omniture.

Those top 10 domains have grown from 40 percent to nearly 80 percent of all hidden nodes over the past five years. "This situation is grimmer in the face of acquisitions," Balachander said, pointing out that Google purchased Doubleclick in 2007 and Adobe purchased Omniture in 2009. "In September 2009, the Google family reached over 70 percent, which is the highest by far among all third parties."

Balachander said that through such tools as the InPrivate Filtering feature of Internet Explorer 8.0, it may be possible for users to improve their filtering of third-party sites that gather privacy data.

Nomcom Voting Members
- Scott Brim
- Dave 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 Levkowetz, Tools advisor
- Jon Peterson, IAB liaison
- Tim Polk, IESG liaison
- Henk Uijterwaal, IAOC liaison
- Bert Wijnen, ISOC BoT liaison

ISOC Panel Notes Rise in IPv6-Related Activity

By Carolyn Duffy Marsan

Momentum surrounding IPv6 is picking up, and IETF participants should be ready for it to snowball soon, according to an Internet Society panel held in Anaheim, California, during the IETF meeting.

Leslie Daigle, chief Internet technology officer at ISOC, said she saw an increase in IPv6-related activity during 2009. She pointed out that Japan published its IPv6 action plan last year, while the U.S. government required IPv6 in its acquisition regulations. Australia moved up—to 2012—the deadline for having its whole government transitioned to IPv6.

Leslie also said that such Internet service providers (ISPs) as Hurricane Electric, Verizon, and Comcast were stepping up their efforts to deploy IPv6. "This is anecdotal evidence, but it's also a very different picture than we saw a year ago in terms of ISPs stepping up and announcing plans that they are deploying IPv6," Leslie added. "There are also some demonstrating real, live, successful networks."

Leslie pointed out that Google, Netflix, and YouTube were among the content providers making services available on IPv6. "When it starts to snowball, you should expect that those customers will be looking at your services and wanting to access them over IPv6," she said. "Increasingly, there are customers out there."

She noted that China Mobile added 88 million new subscribers in 2008 and was expecting similar growth in 2009. "Pv6 addresses are the only option for networks of this scale," she said.

Leslie urged content providers, service providers, and application developers to prepare to reach customers through IPv6. "There's certainly motion on IPv6," she said. "There's some sense that if not now, at some point in the foreseeable future this is going to be in a snowball effect."

Geoff Huston, chief scientist at APNIC and a longtime IETF participant, said he's been trying to measure IPv6 deployment. He researched three sets of data: Border Gateway Protocol table entries, DNS queries, and dual stack Web server access. He said the number of routing table entries for IPv6 has grown from 1,000 to 3,000 from 2008 to 2010. "This is good news," he said. "In terms of routing, IPv6 is growing faster than we thought." But, he pointed out, IPv6 still represents less than 1 percent of IPv4 routing table entries, which top 300,000.

Photo/Internet Society

ISOC's IPv6 panellists Geoff Huston (left) and Jason Livingood

Geoff said it was hard to quantify IPv6 activity by looking at DNS data, but by studying Web server ratios, he estimated that IPv6 represents 1 percent of Internet traffic today. "Use of IPv6 has increased over the past four years to hit 1 percent of traffic," he said, adding that "the number of folks doing 6to4 tunneling as a percent of IPv6 traffic is decreasing rapidly, while the number of folks doing Teredo is really low."

Geoff said the measurements of IPv6 deployment are problematic because of the rate at which IPv4 addresses are being consumed. "If you really wish as an industry to avoid some of the more dramatic problems that might come up, we have to do some work on IPv6," he said.

Jason Livingood, executive director of Internet Systems Engineering at Comcast, said customer response to the ISP's announcement of IPv6 trials this year has been very strong. Comcast is testing three IPv6 transition mechanisms developed by the IETF: 6RD, dual stack lite, and native dual stack over cable and fiber. "The response has been great," Jason said. "We've been very, very pleasantly surprised. We had 5,500 volunteers sign up in a 9- or 10-day period."

Panelist David Temkin, network engineering manager at Netflix, said he was surprised at how easy it has been

Photo/Internet Society

Audience members at a panel organized by the Internet Society on IPv6

Continuedn on next page

3GPP/IETF Workshop on IPv6 Transition in 3GPP Networks

By Dan Wing

With the expectation of significant growth among mobile networks, of increased market share in IP-enabled phones, and of severe limitations in the IPv4 address space, it would be difficult to overstate the growing need for IP addressability. Hence, service providers, which have been talking about IPv6 for years, are starting to move in the direction of both IPv4+IPv6 and IPv6-only mobile networks. Toward that goal, 3GPP (3rd-Generation Partnership Project, the standards development organization tasked with developing standards for third-generation networks) and the IETF held a joint workshop on IPv6 in cellular networks in San Francisco in March 2010. The workshop was sponsored by China Mobile.

The two-day workshop featured more than 40 presentations covering everything from problems to solutions, to successes with regard to deployment of IPv6 on cellular networks. Many of the operators in attendance revealed they are conserving public IPv4 address space by using RFC 1918 space internally and operating large-scale network-address-translation (NAT)44 devices. However, due to the size of those operators' internal networks, it's likely that RFC 1918 space will not be large enough to accommodate the traffic without splitting the network into overlapping segments—a solution that complicates the delivery of services provided by the carrier itself and that raises concern about operating NAT44 devices in perpetuity.

Most operators plan to provide both IPv4 (using RFC 1918 space) and IPv6 for handsets, and to NAT the IPv4 traffic. Some concern was expressed that with most Internet content being IPv4, merely enabling IPv6 on handsets is not enough to move traffic to IPv6. However, Google was cited as an example of an IPv6 content provider. By placing IPv6 on handsets, IPv6 content is accessed directly (due to RFC 3484 preference rules), thus avoiding the operator's NAT44. One operator at the workshop announced plans to deploy IPv6-only handsets that will use NAT64 (and DNS64) to access IPv4 content. Doing so will mean that all traffic from those handsets will be IPv6 over that operator's network. Some of the newer cellular services, such as automated machine-to-machine communication, are expected to function more smoothly over IPv6.

The meeting concluded with participants reiterating that tools already exist to facilitate IPv6 deployment and that, for the most part, new work is unnecessary. Additional work is required in the BEHAVE WG on stateful NAT64 (which provides a mechanism for an IPv6-only handset to share an IPv4 address), on per-interface NAT44, and in the SOFTWIRE WG, where consideration will be given to gateway-initiated dual stack lite.

Further discussion on IPv6 in cellular networks can be found on the 3gv6 mailing list at https://www.ietf.org/mailman/listinfo/3gv6.

ISOC Panel,continued from page 9

to deploy IPv6. "We rely on a CDN [content delivery network] for the bulk of our movie streaming. We host our own website and most of the content that goes behind that. Both the internal integration of our website and our corporate network and the external integration with Limelight for an IPv6 CDN was very straightforward," David said.

Magnus Westerlund, a researcher at Ericsson Research, said he is seeing cellular operators in Europe waking up to the reality that IPv6 is imminent, although few are deploying it. "Everybody is waiting," he said. "It could happen anytime soon."

Photo/Internet Society

IPv6 panel audience members

Photo/Internet Society

Dave Temkin speaking at IPv6 panel at IETF 77

Photo/Internet Society

IPv6 panellist Jason Livingood

Photo/Internet Society

The Curious History of Uniform Resource Names

By Leslie Daigle

Sometimes it's hard to judge whether an engineering effort has been successful or not. It can take years for an idea to catch on, to go from being the butt of jokes to becoming an international imperative (IPv6). Uniform Resource Names (URNs), which are part of the Uniform Resource Identifier (URI) family, are conceptually at least as old as IPv6. While not figuring in international directives for deployment, they—and the technology engineered to resolve them—are still going concerns.

The curious thing is that since the URN working group (WG) concluded in 2002, these two aspects (the URN itself and its resolution system) have had almost completely independent histories. Does the world have a universally supported, persistent resource-naming infrastructure for Internet applications, as envisaged at the outset of the URN WG in 1996? No, not by any stretch of the imagination. Were the six-plus years of IETF engineering effort therefore wasted? No, not that either. Rather, the URN work has contributed certain important building blocks that have been used and reused in efforts that have followed.

Does the world have a universally supported, persistent resource-naming infrastructure for Internet applications, as envisaged at the outset of the URN WG in 1996? No, not by any stretch of the imagination.

In the Beginning: The Intent of URNs

Along with considerable information about the expected and intended state of the Internet Information Infrastructure Architecture, RFC 1737: Functional Requirements for Uniform Resource Names (http://www.ietf.org/rfc/rfc1737.txt), published in December 1994, outlines the requirements of Uniform Resource Names. Key to all of this is the intended purpose of a URN, which is identified as providing "a globally unique, persistent identifier used for recognition, for access to characteristics of the resource, or for access to the resource itself." And, further, "A URN identifies a resource or unit of information. It may identify, for example, intellectual content, a particular presentation of intellectual content, or whatever a name assignment authority determines is a distinctly namable entity. A URL identifies the location or a container for an instance of a resource identified by a URN. The resource identified by a URN may reside in one or more locations at any given time, may move, or may not be available at all."

The list of functional requirements for URNs is not that extensive, but it is constraining. That URNs are to be persistent, as well as global in scope and uniqueness, is pretty clear from the stated purpose. Additionally, the document stipulates that URNs are to be scalable (assignable to anything conceivably available on the network for hundreds of years) while supporting legacy naming systems, allowing independent assignment of identifiers by autonomous "naming authorities," and still allowing "resolution" of URNs—that is, translation from the URN into one or more URLs.

RFC 1737 notes that it does not address the question of requirements for resolution, thereby leaving that question open. That open question was the source of many heated discussions within the URN WG over the years, with some proponents fiercely demanding "sub-second resolution!" as an imperative, while others wanted to first ensure distributed and resilient services.

A Technology Structure to Support the Intent

The URN WG was established in 1996, after proponents of several specific URN proposals had come to a high-level understanding of how to support a diversity of potential applications and needs for URNs while maintaining a generalized standard and support infrastructure. The key was to allow for independence in resolution systems while binding identifiers together under a single generic, URI-consistent syntax with a discovery system for the resolution services.

Put simply, and as captured in RFC 2141: URN Syntax (http://www.ietf.org/rfc/rfc2141.txt), URN identifier syntax is

"urn:" <NID> ":" <NSS>

where <NID> is a namespace identifier (to distinguish between different schemes of persistent identifiers, with different authorities, etc.) and where <NSS> is the namespace-specific string. Within certain important constraints to synchronize with URI syntax, the namespace-specific string can be structured in whatever way the authority for the namespace wishes. This allows for either structured or unstructured namespaces as well as either human-readable or machine-oriented identifiers. And each namespace is completely independent of the next: each is free to reuse the same strings to different purposes.

To support that simple-yet-flexible identifier system, some level of discovery

Continued on next page

Curious History,continued

system was needed in order to be able to find the relevant final resolution services for the different namespaces. From RFC 2276: Architectural Principles of Uniform Resource Name Resolution (http://www.ietf.org/rfc/rfc2276.txt), the architectural principles for URN resolution were based on two assumptions: "In general, we must assume that almost any piece of the supporting infrastructure of URN resolution will evolve. In order to deal with both the mobility and evolution assumptions that derive from the assumption of longevity, we must assume that users and their applications can remain independent of these mutating details of the supporting infrastructure. The second assumption is that naming and resolution authorities may delegate some of their authority or responsibility; in both cases, the delegation of such authority is the only known method of allowing for the kind of scaling expected. It is important to note that a significant feature of this work is the potential to separate name assignment, the job of labelling a resource with a URN, from name resolution, the job of discovering the resource given the URN. In both cases, we expect multitiered delegation."

At the time, there was only one place to look for support of such multitiered delegation: the Domain Name System (DNS). So the URN WG developed a resolution discovery system, rooted at URN.ARPA and defining Naming Authority Pointer (NAPTR) DNS resource records to suit in 2000, which was originally defined in RFC 2915: The Naming Authority Pointer (NAPTR) DNS Resource Record (http://www.ietf.org/rfc/rfc2915.txt) and was updated by RFC 3403: Dynamic Delegation Discovery System (DDDS): Part Three: The Domain Name System (DNS) Database (https://tools.ietf.org/html/rfc3403).

Simply put, the discovery system works by starting with <NID>.URN.ARPA and using the content of the retrieved NAPTR resource records to identify subsequent steps in discovering where and how to resolve a particular URN. Practically speaking, this means that the authority for resolving ISBN-based URNs could rest with an international ISBN body, while new namespaces for computing activities could be built to partition resolution responsibilities among several subauthorities that are not traditional publishers at all. Equally important, this distribution of the underlying resolution authority could change over time, because the discovery system provides dynamic rules for directing requests to appropriate authorities.

The authority for resolving ISBN-based URNs could rest with an international ISBN body, while new namespaces for computing activities could be built to partition resolution responsibilities among several subauthorities that are not traditional publishers at all.

The overall principles are simple—explained in less than 500 words!—but the system also provided for a great deal more complexity and power to address the many side issues that have been raised during all the years that identifiers for the Internet have been discussed.

Initial Steps—and Immediate Divergence of Intent and Structure

Over the next few years, a number of URN namespaces were registered with the Internet Assigned Numbers Authority (IANA), per the registration process outlined in RFC 2611: URN Namespace Definition Mechanisms (http://www.ietf.org/rfc/rfc2611.txt) (and updated by RFC 3406: Uniform Resource Names [URN] Namespace Definition Mechanisms, http://tools.ietf.org/html/rfc3406, in 2002). The key purposes of a formal registration process for URNs included ensuring that some conscious effort was put into securing a piece of real estate in the URN namespace identifier space, including a review of the principles of URNs (persistence, global uniqueness) and an indication of how URNs in the namespace are meant to be resolved. Notably, none of these registered namespaces elected to use the global resolution discovery service based in the DNS. However, almost immediately, another application was found for this dynamic, distributed approach to resolution in the work of mapping E.164 (telephone) numbers to Internet telephony resources, in the ENUM WG.

RFC 2916: E.164 number and DNS (http://tools.ietf.org/html/rfc2916) described this first non-URN application using NAPTR DNS resource records. Notably, the ENUM work was not defining a URN namespace; in other words, it was not attempting to describe a use of E.164 telephone numbers as if they were persistent identifiers of an Internet resource. Rather, the feature that ENUM wanted to leverage was the ability to put the control of the association of the number to a set of available services into the hands of the number holder. Subscribers ought to be able to control where their related Session Initiation Protocol (SIP) service terminates, for example. At the same time, the E.164 number space is hierarchically managed, so delegations are made (and managed) at higher levels in the telephone number tree than the simple end telephone number.

Generalization

The ENUM use of NAPTR records brought to light that at least one of the URN assumptions had more general applicability; in other words, it could be assumed that identifier and resolution systems for many applications might feature delegation of some authority and/or responsibility for mapping identifiers to resources, with an expectation of multitiered delegation.

With that in mind, the relevant specifications were updated and refined to produce a more generic definition of the dynamic delegation discovery system, or DDDS, that supported URNs:

- RFC 3401: Dynamic Delegation Discovery System (DDDS) Part One: The Comprehensive DDDS (http://www.ietf.org/rfc/rfc3401.txt)
- RFC 3402: Dynamic Delegation Discovery System (DDDS) Part Two: The Algorithm (http://www.ietf.org/rfc/rfc3402.txt)
- RFC 3403: Dynamic Delegation Discovery System (DDDS) Part Three: The Domain Name System (DNS) Database (http://www.ietf.org/rfc/rfc3403.txt)
- RFC 3404: Dynamic Delegation Discovery System (DDDS) Part Four: The Uniform Resource Identifiers (URI) (http://www.ietf.org/rfc/rfc3404.txt)
- RFC 3405: Dynamic Delegation Discovery System (DDDS) Part Five: URI.ARPA Assignment Procedures (http://www.ietf.org/rfc/rfc3405.txt)

RFC 3401 is the umbrella document providing the road map for this comprehensive set of specifications. RFC 3402 provides the definition of the generic principles of the DDDS approach, independent of any application or implementation. RFC 3403 ties this general architectural specification to the already extant specification—the DDDS as

implemented in DNS, using NAPTR records. RFC 3404 and RFC 3405 were intended to make clear how other applications could make use of the DDDS approach and NAPTR records. Those last two documents were the completing pieces to bring the URN (and URI) resolution approaches in line.

With these documents in hand, it was now possible for any new application to make use of this sort of dynamic delegation discovery system using the DNS. A number of applications seemed to need it, such as SIP. However, the set of RFCs was evidently daunting; in some cases, applications considered defining (and deploying) entirely new DNS resource records instead. DDDS can appear overpowering in its generalized state.

Powerful is good, but when you're looking for a good tool to deal with nails, you don't really want to be directed to the machine shop to carve a well-balanced handle and then forge a solid steel head (no matter how good the instructions are for constructing your own hammer).

It became apparent that a number of potential uses for DDDS/NAPTR were not being realized—in part because of that evident complexity. At the same time, there were some similarities between those potential uses; perhaps they represented a class of DDDS application.

To deal with that and to hopefully provide a more accessible tool for application specifications seeking to provide some (DNS-based) discovery of application services and protocols, a generic, simplified DDDS application was specified. For application services that use it, starting from any unique key mapped into a domain name, the S-NAPTR ("straightforward NAPTR") DDDS application (RFC 3958: Domain-Based Application Service Location Using SRV RRs and the Dynamic Delegation Discovery Service [DDDS], http://www.ietf.org/rfc/rfc3958.txt) defines how to find a domain's preferred server for a given application and protocol. This gives domains a flexible and dynamic approach to homing services. It gives application specification writers a complete hammer. S-NAPTR uses only the core features of the NAPTR-based DDDS. It fit many needs but fell short of serving a few more. U-NAPTR (RFC 4848: Domain-Based Application Service Location Using URIs and the Dynamic Delegation Discovery Service [DDDS], http://www.ietf.org/rfc/rfc4848.txt) extends S-NAPTR to allow for returning fully formed URIs at the end of the dynamic delegation process of DNS lookups.

Powerful is good, but when you're looking for a good tool to deal with nails, you don't really want to be directed to the machine shop to carve a well-balanced handle and then forge a solid steel head (no matter how good the instructions are for constructing your own hammer).

Continued on page 16

ISOC Fellows Gain Skills, Network at IETF 77

Since 2006, the Internet Society (ISOC) Fellowship to the IETF programme has provided a critical link between the IETF and Internet technology professionals in developing regions. The fellowship, which operates under the aegis of ISOC's Next Generation Leaders Programme, has since provided dozens of opportunities for men and women to meet face-to-face with the network engineers who are grappling with Internet-related problems and issues.

In March 2010, six technologists travelled to Anaheim, California, as first-time fellows and three others attended as returning fellows. Here's what some of them are doing—and what they are saying about their experience at IETF 77.

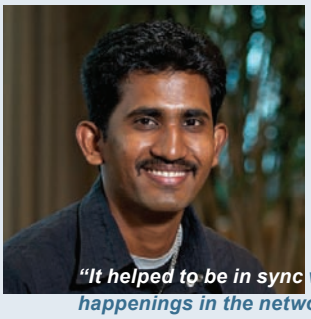

"It helped to be in sync with the happenings in the networking community apart from approaches to getting my drafts moved up the ladder to become RFCs."
— Palanivelan Appanasamy (India)

Educated at the University of Madras, Chennai, Palanivelan is currently employed by Cisco Systems in Bangalore, India, where he is a senior member of the next-gen router test team and is responsible for technical inputs as well as for providing guidance for the team. What did Palanivelan appreciate most about IETF 77? "The opportunity to share, learn, and discuss ideas with the best in the business."

"I enjoy participating in discussions with people as well as the opportunity to contribute in some of the areas important to the African continent, such as IPv6, congestions, and DNS security."
— Jean-Robert Hountomey (Benin)

Born and raised in Cotonou, the largest city and the economic capital of Benin, Jean-Robert is now in Minnesota, where he is studying computer forensics and working as a consultant on network architecture design and security. His primary interests within the IETF are IPv6, the Anti-Spam Research Group, DNS, and DNS Security Extensions. The best part of IETF 77 for him, Jean-Robert said, was "meeting friendly people" and being able to talk with the people who make the Internet run.

"Attending an IETF meeting is always a very good opportunity to build relationships and [conduct discussions] with key people involved actively in standardization work."
— Afaf El Maayati, Returning Fellow (Rabat, Morocco)

Born in Marrakech, Morocco, Afaf now lives in the capital city of Rabat, where she works at the National Telecommunications Regulatory Agency—or ANRT (the organization delegated by the Internet Corporation for Assigned Names and Numbers as the administrator of the .ma ccTLD). At ANRT, Afaf serves as a project manager responsible for auditing the registration and management of .ma, for studying and verifying the accreditation requests of .ma registrars, and for mediation of conflicts concerning .ma domain names, among other activities. Within the IETF, she is especially interested in Domain Name System Operations and DNS Extensions.

"I think the meeting opened up the IETF to me. I have a better understanding of how it works and a greater desire to participate."
— Kondwani Masiye (Blantyre, Malawi)

Born and raised in Blantyre, Malawi, Kondwani was educated at the University of Malawi's Chancellor College and is now employed by Malawi Telecommunications Limited. His work involves deploying new IP and data systems for Malawi Telecoms' service provider or carrier network. He works primarily in project management; network designing and engineering; and implementation of various systems as well as router, switch, and server configuration. "I also do a bit of network management, and I design custom solutions for our customers," he wrote via email. Down the road, Kondwani would like to develop his expertise as a network architect because he's interested in a number of areas within the IETF, including IPv6, multiprotocol label switching, virtual private networks, DNS, and Border Gateway Protocol. What would he consider the most gratifying aspect of attending IETF 77? "I liked the interaction with people in my field, who are extremely brilliant, and the openness of the participants."

Photos/Internet Society

"The ISOC plenary session on IPv6 provided an opportunity to look at the exhaustion of IPv4 [addresses] and why the move to IPv6 should be taken seriously."
— Sakaio Manao (Suva, Fiji)

Originally from Nukulaelae, Tuvalu, Sakaio is now based in Suva, Fiji, where he works at the SOPAC (Pacific Islands Applied Geoscience Commission) Secretariat managing the SOPAC local- and wide-area networks while studying to become a Microsoft Certified IT Professional. Of particular interest to Sakaio are IPv6 operations and management. IETF 77 gave him the opportunity to be "exposed to a high level of technical opinions, discussions, and expertise."

"The IETF experience was great exposure for me. I had the opportunity to interact with the RFC contributors and share my views, particularly the developing-world perspective, with them."
— Mohibul Hasib Mahmud, Returning Fellow (Dhaka, Bangladesh)

A native of Dhaka, Bangladesh, Mohibul has a bachelor's degree from the National University of Bangladesh and an M.B.A. from the Asian University of Bangladesh. He is currently pursuing ISOC's Next Generation Leaders eLearning Programme, an online course made available at Diplo Foundation. Mohibul currently works at BRAC BDMail Network Ltd, a Bangladesh-based Internet service provider, where he manages the core network team, conducts network planning, and maintains upstream links. "My interests are in network routing, security, quality of service, and network management," he wrote via email. In the future, he would like to learn more about Internet technology and policy as well as the "Internet ecosystem" and to contribute to the Internet standards development process. What does he enjoy most about IETF meetings? "I enjoy the interaction among the participants as well as the networking opportunities. IETF [meetings] provide larger bandwidth for the participants for interaction, which is not possible in the mailing list."

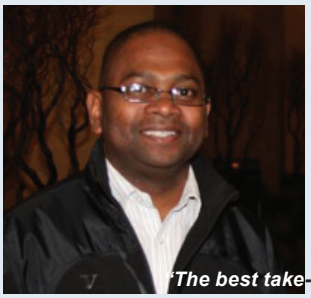

"The best take-home [from IETF 77] was positive contribution to some of the sessions as well as having hands-on experience with IETF issues, establishing new research collaborative networks, and making friends."
— Idris A. Rai (United Republic of Tanzania)

Born in Zanzibar and educated in Turkey (B.Sc. and M.Sc. in electrical and electronics engineering) and France (Ph.D. in networks and computer science), Idris now works at Makerere University in Kampala, Uganda, where he is responsible for teaching communications networks programmes, supervising graduate students at the master's and doctoral levels, and managing academic programmes and student affairs for the Faculty of Computing and Information Technology. He also conducts research in network protocols, overlay networks, and mobile computing. In the short term, Idris is hoping to secure a full professorship by 2011 and perhaps a top management position before 2015. In the long term, he hopes to create and launch businesses in the area of communications networks. Mainly, he says, he hopes to "honestly contribute in order to make a huge difference, wherever I am."

IETF 77
First-time Fellows
Palanivelan Appanasamy (India)
Mentor: Keyur Patel
Jean-Robert Hountomey (Benin)
Mentor: Alain Aina
Sakaio Manoa (Fiji)
Mentor: Phil Roberts
Kondwani Masiye (Malawi)
Mentor: Joel Jaeggli
Idris A. Rai (Tanzania, United Republic of)
Mentor: Fred Baker
Gustavo Ramos (Brazil)
Mentor: Joao Damas

Returning Fellows
Afaf El Maayati (Morocco)
Mohibul Hasib Mahmud (Bangladesh)
Dessalegn Mequanint Yehuala (Ethiopia)
Noah Sematimba (Uganda)

ISOC Fellows and Returning Fellows at IETF 77 in Anaheim, California

Curious History,continued from page 13

Today's Reality

According to the IANA registry, there are 40 formal URN namespaces registered today. (See http://www.iana.org/assignments/urn-namespaces/urn-namespaces.xhtml.) The namespaces range from identifiers for IETF protocol resources to the Digital Video Broadcasting Project, to 3GPP, to ISBN. Very diverse communities of interest have rallied around the URN concept of a persistent, unique global identifier and established a namespace for their purposes. None of these use the formally established resolution mechanism (DDDS). However, approximately 25 RFCs reference the DDDS/NAPTR RFC for uses from ENUM to SIP.

The ECRIT WG emergency services discovery work (LoST) uses U-NAPTR, and the DIME WG is considering S-NAPTR for its application.

Finally, discussions are under way to review and revise the URN syntax document, currently under way on the reestablished urn@ietf.org mailing list.

So, both URNs as a concept and the underlying technology are alive and well, if not exactly living together in married bliss.

The Big Takeaway?

While not classically successful, the URN work produced output that clearly has had value for high-impact derivative works. The work started with a vision for persistent identifiers, with a dedicated group of IETF workers interested in the problem space, and with a perceived market need for those identifiers. That single driving market never actually appeared. Reviewing the 40 registered formal URN namespaces, one would find it hard to detect a single unifying theme between them that could have led to a single resolution system that would have supported all of them.

So, the building-blocks approach has been the most successful: Successive waves of IETF engineers have picked up on the blocks that fit their own needs and reused them. And so, we continue to build the overall Internet infrastructure not by single unifying services that solve yesterday's problems today but, rather, by building blocks that support continued, organic evolution.

The Peer-to-Peer Invasion

By Enrico Marocco and David Bryan

It was soon dubbed "the basement meeting." In March 2005, during IETF 62 in Minneapolis, a group of some of the most active participants in what is now the RAI (Real-time Applications and Infrastructure) area got together for an informal Bar BoF (birds-of-a-feather) meeting. The goal was to share thoughts and ideas about making SIP (session initiation protocol) user agents work without centralized elements—such as SIP proxies and SIP registrars—and be capable of reproducing their functionality with distributed algorithms run by the endpoints themselves. The benefit would be a significant reduction in capital and operational costs, allowing virtually everyone to set up and run their own SIP service. The topic had been discussed informally for quite some time, and the recent launch of Skype with similar technology had proved it was possible.

The basement meeting was not the first time that a group of IETF participants gathered together to discuss standardizing peer-to-peer (P2P) protocols in the IETF. It was, however, the beginning of a trend leading to the creation of several IETF working groups (WGs) dealing with P2P technologies.

P2PSIP: The IETF Takes Up P2P

Both the Internet industry and the academic world have shown interest in the topic, but drawing the line between research and engineering issues and narrowing the scope of the WG required a great deal of energy and much discussion. It took almost two years for the proponents of the new work to agree on a charter and get a group approved. Finally, in 2007, between IETF 67 and IETF 68, the P2PSIP WG was created, chartered to specify a protocol for building a distributed overlay to provide SIP registration and routing functionality. Four years and 10 meetings later, the base protocol specification, the RELOAD distributed hash table, is ready for WG Last Call and the first P2P protocol Standard Tracks RFC is about to see the light of day.

The P2P Infrastructure Workshop: LEDBAT and ALTO Are Born

The following year, other P2P-related topics made it to the IETF. In the P2P Infrastructure (P2PI) workshop organized by the RAI area directors, Internet service providers, vendors, and applications developers gathered together at Massachusetts Institute of Technology in Boston to discuss

the issues created for Internet infrastructures by P2P traffic. Two areas of improvement were identified.

The first observation participants made at the P2PI meeting was that traffic generated by file-sharing applications interferes with delay-sensitive traffic, degrading the user experience of other applications on the same network. To address that problem, the LEDBAT WG was chartered to work on an experimental congestion control algorithm that yields to TCP. The new algorithm would be intended for use in bulk-transfer applications, including P2P file-sharing applications. Today, significant progress has been made in specifying the LEDBAT algorithm, backed by an implementation integrated in the most popular BitTorrent client and by significant simulation work carried out in universities and research labs.

The second point of concern noted was that P2P applications establish overlays on top of the Internet infrastructure with little or no knowledge of the underlying network topology. That lack of information leads peers to make suboptimal choices. For example, it is common for a peer who needs to obtain information from another peer to choose randomly, possibly picking one located on the other side of the planet. Such random selection may ignore many peers who are topologically closer and that therefore likely would provide better performance. Like LEDBAT, the ALTO working group was created following the P2PI workshop, chartered to produce a protocol enabling network operators, Internet service providers, and others with network topology information to share it with P2P applications, thereby enabling those applications to improve the peer selection process. Sharing the information is expected to benefit both end users and network operators, resulting in better application performance and more-rational use of network resources.

Recent Developments: PPSP and DECADE

More recently, the Peer-to-Peer Streaming Protocol (PPSP) WG was formed, driven by the success of applications like PPLive and Spotify, which use P2P methods to distribute real-time content. The primary goal of the group—officially approved in Anaheim, California, during IETF 77—is to specify protocols that enable peers to synchronize and exchange media chunks of live or time-shifted content. Drawing on the widely deployed P2P applications in the wild, the PPSP architecture envisions two different kinds of nodes: peers and trackers. Trackers act as facilitators, maintaining and distributing information about the shared content and about which peers store which portions. Peers are the nodes that, with the support of trackers, find the location of the content they want, establish connections among themselves, and actually store, send, and receive the content. The design of a new transfer protocol between the peers has been ruled out of scope—meaning, the group will likely reuse a protocol such as RTP (Real-time Transport Protocol) or HTTP.

Finally, the DECADE WG, approved right after IETF 77, is the latest WG chartered to work on P2P issues. The upstream network links of home users are often relatively small; and bandwidth-greedy applications such as file sharing and real-time streaming can quickly saturate the link when uploading content. To address that problem, DECADE's advocates propose a mechanism whereby users can store content they are sharing in the network—ahead of the last-mile bottleneck. The DECADE group seeks to specify a protocol for fine-grained control of such network storage. The protocol is intended to be integrated in P2P applications, thereby enabling them to store, retrieve, and manage data remotely as well as to define policies for sharing such data with other peers and to control the utilization of resources like bandwidth and number of connections.

P2PRG: IRTF Efforts on P2P

While many aspects of P2P technology are mature enough for standardization efforts in the IETF, there are still many areas of active research in P2P. The Peer-to-Peer Research Group (P2PRG) seeks to provide a direct channel for sharing the latest research advances in the area with the IETF engineering community.

The first observation participants made at the P2PI meeting was that traffic generated by file-sharing applications interferes with delay-sensitive traffic, degrading the user experience of other applications on the same network.

While this group had been dormant for some time—taking three years off following IETF 65—the P2PRG had a well-attended session during IETF 74. Since then, the group has become very active. The members meet regularly, and they recently published an informational RFC addressing security in P2P systems for real-time communications. Topics recently discussed include research issues related to P2P traffic optimization, real-time content distribution, and P2P group management.

P2P and the IETF

While recent studies show a slight reduction in the fraction of network traffic attributable to P2P, these applications have become extremely popular, and they are being used for a variety of purposes. The maturity of the

Continued on next page

Peer-to-Peer, continued

applications is demonstrated by the fact that in some cases—for example, the BitTorrent protocol—several independent interoperable implementations coexist in the wild. This is generally a clear sign that the technology is mature and ready for standardization. Many other applications are widely deployed among significant user communities, but those users are forced to choose between multiple, noninteroperable applications, which limit opportunity to share information and which often result in duplicated development and distribution of content. As shown by the active P2P

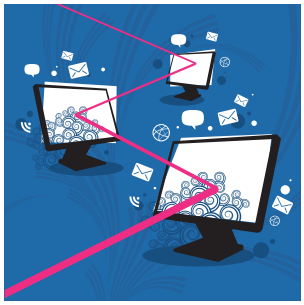

participation at IETF, many of these communities have a strong desire for interoperable standards for these applications and a willingness to contribute ideas and work to achieve that goal.

P2P is an increasingly important network technology that is rapidly becoming pervasive. Not every application would benefit from a P2P approach; however for those that would, the IETF is uniquely positioned to specify P2P protocols that harmonize with the rest of the Internet protocols. Armed already with the insight, the experience, and a vast reservoir of collective wisdom on how the Internet works, the IETF community is best able to evaluate the complex protocol design choices that P2P protocols present. There is a vibrant and growing community within the IETF that is actively seeking to apply P2P to real-world engineering problems on the Internet.

IETF Outcomes: An interview with Dave Crocker

In February 2010, IETF chair Russ Housley announced the launch of a new wiki dealing with "IETF Outcomes." The wiki, which can be found on the IETF tools site at http://trac.tools.ietf.org/misc/outcomes/, features technologies and services that were developed within the IETF and that represent notable successes and failures. It is the result of a collaborative effort by IETF participants—who are invited to use it to provide feedback about the utility of IETF work—and it is a mechanism for facilitating public understanding of IETF work and its impact.

Photo/Internet Society

Dave Crocker at IETF 77 in Anaheim, California

The *IETF Journal* took the opportunity of IETF 77 in Anaheim, California, to meet with Dave Crocker—a driving force behind the creation of the wiki—to chat about the motivations that gave rise to its development and about expectations for its future.

> Wikis possess a classic grassroots quality: they are developed by the community, they are transparent, and they permit resolving disagreement through open debate.

IETF Journal: What motivated the creation of the wiki?

Dave Crocker: For many of us, the usual measure of success is the publication of an RFC, but we're in the communications business, and in the real world this involves closing the loop with feedback. The need for assessment was clear; the question was how to do it. I focussed on finding a way to help the community develop an internal sense of accountability. Wikis possess a classic grassroots quality: they are developed by the community, they are transparent, and they permit resolving disagreement through open debate. To get this started, I talked with a few people over the space of about a month. It began as a simple table, but a wiki became the obvious choice once the need arose to support continuing change, provided by the community. In classic Internet terms, it scales better. After the initial group exercise stabilized, I approached the IETF management. As a grassroots, ongoing exercise, the status of the wiki is inherently informal, which nicely matches its placement in the tools.ietf.org portion of the IETF website. There's a mailing list to go with it, which is there to discuss issues with the wiki in general. We're slowly but surely seeing people taking the initiative to contribute to it.

IETF Journal: How do you measure the success of a standard in the marketplace? Is it always subjective?

DC: There's some text in the wiki that describes ratings, but in general it's pretty subjective. One of the columns of the wiki table is called usage, which might seem an odd term, but ultimately, the reason we make stuff is so it gets used. What does it mean to get used? I don't think that having software implement a spec qualifies it as a success. I think having somebody use that software makes it a success. The difference is very important. The IETF is driven largely by an industry that produces things, not by an industry that uses those things. The rating system is only a five-point scale, from complete failure to massive success. When you're doing survey research, that's as many points as you want for a casual audience. If there's a lot of debate about the rating for a given standard, then we don't know enough to rate it.

It's not only usage that matters; it's also the extent to which a piece of work prompted derivative works. It turns out you can have something that's a complete failure but that triggers derivative work of importance. An example of that is PEM [Privacy Enhanced Mail], which generated a lot of useful outcomes, even though the protocol itself was a complete failure. (For a more detailed discussion on this general point, see Leslie Daigle's article on URNs on page 11.)

IETF Journal: What's the incentive for somebody to update the wiki with information about a standard that has failed but that may have involved considerable effort to create?

DC: That's a very interesting question. Frequently, people are brutally honest about their own work. The IETF environment encourages that level of honesty. People don't beat themselves up in public very often, but within the community there is respect for learning what didn't work and then using that information. Not surprisingly there is a normal tendency for people to point out others' failures. So if somebody wants to hurt somebody else's feelings by putting an entry in the wiki, the only relevant question is: Is the criticism accurate? I am aware of the concerns about the possible social and political downsides. I myself have had concerns about the wiki format, that it could create competition between areas. But what's bad about that?

Engineers tend to project acceptance of their wonderful ideas, but the market doesn't work that way. In the earlier days of the IETF, market pull was a consideration when chartering new work.

One subtlety that has developed as a result relates to some IETF technical efforts that had a number of false starts, such as DNSSEC [Domain Name System Security Extensions]. A number of long-time DNS experts worked on this topic because it's so important. As a result, we've developed multiple entries to try to capture multiple phases of work and differing outcomes.

IETF Journal: Do we need a methodology that is applicable to other standards development organizations?

DC: This methodology for producing outcomes ratings of IETF work is so simple that I'd expect it could be applied to any group; whether groups want to or not is their choice. But note that as a grassroots tool, it does not require the blessing of the organization. It could be interesting to try to generalize to the W3C [World Wide Web Consortium].

IETF Journal: What do you think the IETF might learn from the development of this evaluation tool?

DC: Given that this is done with subjective, coarse-grained data, I hope that all we learn are subjective and coarse-grained things. We may see that some areas of work have better track records than do others. The most interesting thing I hope we'll learn is some sense of which approaches to doing work tend to be successful and which approaches tend not to be successful. That's ambitious to hope for, and it requires a lot of effort and thinking, but it would be pretty nice if we could get there.

I think debates over what is the right way to assess a particular effort are very useful because getting clarity about what succeeds and what doesn't will help the next time. Just getting people to worry about the long term could be the biggest benefit. Engineers tend to project acceptance of their wonderful ideas, but the market doesn't work that way. In the earlier days of the IETF, market pull was a consideration when chartering new work. Now we measure only whether there are people interested in working on something, so we end up with things being worked on for a long time that don't always get used. I hope the wiki inspires people to think first about who's going to use what they're interested in creating.

The discussion at the IETF plenary—where, among other things, debate concerned workload on the IESG—made me think that while computer networking is about sharing limited resources, we also need to do this for ourselves. Improving quality-control mechanisms can help the IETF leadership decide how to ease its workload. I hope the wiki can be a part of that.

IETF Journal: Thanks for your time, Dave.

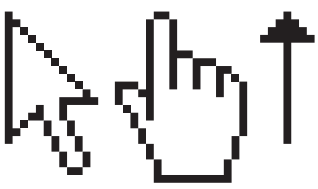

IRTF Update

By Aaron Falk

Aaron Falk, IRTF Chair

At each IETF meeting, the Internet Research Task Force (IRTF) chair presents a short status report. This article summarizes the report made to the IETF 77 plenary.

Seven IRTF research groups (RGs) met at IETF 77, including Delay Tolerant Networking RG (DTNRG); Internet Congestion Control RG; Host Identity Payload RG; Peer2Peer RG; Routing RG; Scalable, Adaptive Multicast RG; and Virtual Networks RG. The Internet Architecture Board reviewed the scope and progress of the Scalable, Adaptive Multicast RG. Ten of the 13 RGs are meeting, have active mail lists, or both. The three quiescent groups are Mobility Optimizations RG, Network Management RG, and Public Key Infrastructure-Next Generation RG.

An RFC series for the IRTF was created in 2009 by way of RFC 5743. However, some necessary changes in copyright policy and other boilerplate prevented publication of RFCs in the series for many months. The logjam was removed in March 2010, and eight IRTF RFCs have since been published, including documents from five different RGs. Seven additional drafts, mostly from the DTNRG, are in review and should be submitted to the RFC Editor soon.

A new RG on virtual networks (VNRG)—chaired by Martin Stiemerling of NEC and Joe Touch of the University of Southern California Information Sciences Institute (USC/ISI)—was chartered this spring. Virtual networks are appearing in test beds, data centres, the GRID, and cloud services as a way of providing flexible resource allocation and management. However, the approaches used in the global Internet, as part of the test beds and within business and organizational enterprises, are quite different. One question the group is examining is how to identify and bind processes and virtual machines to virtual networks. The VNRG also hopes to establish a common framework and terminology for virtual networks.

Virtual networks are appearing in test beds, data centres, the GRID, and cloud services as a way of providing flexible resource allocation and management.

The IRTF sponsored a tutorial on NetFPGAs in conjunction with IETF 77. The NetFPGA platform enables researchers and instructors to build high-speed, hardware-accelerated networking systems. The platform can be used in the classroom to teach students how to build Ethernet switches and Internet Protocol routers by using hardware rather than software. Researchers can use the platform to prototype advanced services for next-generation networks. More information about NetFPGA can be found at http://netfpga.org.

The End-to-End RG closed after 26 years. Among the many significant contributions the group made to the IETF are slow start and improved round-trip time estimation, Random Early Drop, Integrated and Differentiated Services, Weighted Fair Queuing, PAWS, and Transaction TCP. While the End2end RG was a closed group, it maintained an active and open mailing list. The list will continue as an independent service to the community at USC/ISI's Postel Center. More information can be found at http://mailman.postel.org/mailman/listinfo/end2end-interest.

We are trying out some new ideas for improving the IRTF. A new mailing list, irtf-discuss@irtf.org, has been created to encourage community input on proposed research groups. (See https://www.irtf.org/mailman/listinfo/

Photo/Peter Löthberg

ISOC's Karen O'Donoghue (left) and Lucy Lynch take a break during IETF 77

Photo/Peter Löthberg

IETF 77 attendees enjoy a break

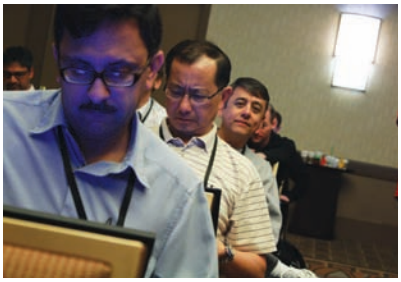

Photo/Peter Löthberg

IETF 77 participants attend the opening plenary

Photo/Peter Löthberg

SIDN, which will be hosting IETF 78, prepares for the next meeting while at IETF 77

irtf-discuss.) Another proposal being considered is the creation of a regular open IRTF meeting (similar to an IETF area meeting) as a venue for research topic proposals and other related discussions that do not necessarily fit within the context of an RG meeting. Finally, we are adding dots to IRTF RG chair badges.

An informal Bar BoF (birds of a feather) was held at IETF 77 on the research issues in the broad area of Internet of Things. There are many ways of viewing this topic, but one way is to look at two classes of use cases: The first is silicon cockroaches, which are small, ubiquitous objects, such as embedded sensors, RFIDs, asset tracking systems, and biomedical devices, and the second is machine-to-machine systems, such as cyberphysical systems, actuators, building networks, energy systems, and automotive systems/networks. Some of those systems have interesting characteristics that influence how the devices interact with the network, such as:

- Order(s) of magnitude bigger than the Internet, in number of endpoints
- No computers or humans at endpoint
- Inherently mobile, disconnected, unattended

Given those characteristics, many possible research topics were identified, such as security, privacy, authentication, naming, authority (by people and by devices), discovery, management, maintenance, policy, preferences, presence (of people and of devices), location, capabilities, services, information model, and coordination. One group has gone off to try to sketch out a charter for an RG.

Finally, a few additional proposals are cooking for new RGs on social informatics and Internet protocols; on economics, law, and policy; and on privacy in the cloud. Look for updates on those topics as they mature.

The Internet Research Task Force promotes research of importance to the evolution of the future Internet by creating focused, long-term, and small Research Groups working on topics related to Internet protocols, applications, architecture and technology. See http://www.irtf.org.

IETF 77 Plenary Snapshot

The IETF Trust adopted the new Trust Legal Provisions (TLP 4.0) breaking a logjam of almost a year in the publication of independent submissions. Blue sheets are used as physical proof of attendance at IETF meetings.

Outgoing IESG/IAB/IAOC Members

Fred Baker
Ross Callon
Gonzalo Camarillo
Stuart Cheshire
Lisa Dusseault
Pasi Eronen
Cullen Jennings
Gregory Lebovitz
Andrew Malis
David Oran
Magnus Westerlund

Appointments

Incoming IAOC Member

- Eric Burger

Incoming IAB Members

- Bernard Aboba
- Ross Callon
- Spencer Dawkins
- Andrei Robachevsky
- Hannes Tschofenig

Incoming IESG Members

- Stewart Bryant
- Gonzalo Camarillo
- David Harrington
- Peter Saint-Andre
- Sean Turner

IAOC Report

2009 Financial Summary

- Entered 2009 with serious concern about economy and its effect on our finances
  — Planned for downturn with contingency budget
- ISOC provided $150K stimulus funds that the IAOC used to lower registration fee
- ISOC set aside $350K contingency fund to cover shortfall if needed

2010 Budget

- IAOC adopted $5.3M budget
- Registration Fee: $635 (steady for 3 years)
- Budget includes IETF tools investment of $575K
  —Data Tracker extension for Author and WG
  —Secretariat tools in python–Django
  —RFC Services
  —Program Management

RFC Editor Restructuring

- New Model
  —RFC Production Centre
  —RFC Publisher
  —RFC Series Editor
  —Independent Submissions Editor

Production Centre and Publisher Status

- RFC Production Centre
  —Transition from ISI to AMS completed
- RFC Publisher
  Transition from ISI to AMS completed
- 109 RFCs published since 1 January 2010

Contract Cycles

- Secretariat
  —Contract extended thru 2010
- One year Extension or RFP in 2011
- RFC Editor Services Contracts
  —Contract with ISI extended through 30 June 2010 to assist with transition

Photo/Peter Löthberg

Photo/Peter Löthberg

  —New contracts in 2010 for Production Centre and Publisher
- 6-year terms with 2-year reviews
- Independent Submissions Editor
  —5-Year term with review at 2-year intervals with extension opions
- Transitional RSE
  —Contract through IETF80
- NOC IDIQ Contracts with VeriLan and Swisscom thru 2011
  —Extension or new RFP in 2011 for next period IETF 77

IETF Tools Contract Status

- Database Project Manager appointed
  —Henrik Levkowetz under contract
- Three Master Service Agreements executed for IDIQ Python--Django development
  —First task order to IOLA for IESG Data Tracker award soon
- Working Group and Author Requirement specification development
  —Ed Juskevicius under contract
- AMS developing Secretariat Tools in python--Django
- Additional Task Orders and RFPs forthcoming

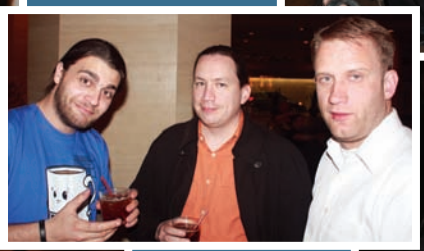

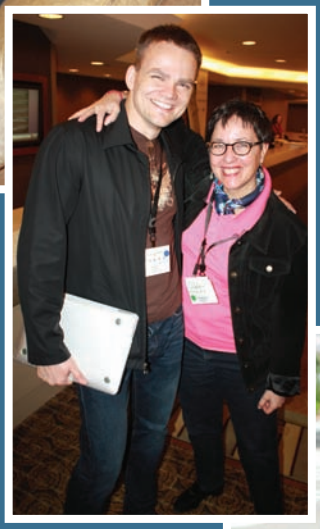

Photos/Peter Löthberg

IETF Meeting Calendar

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

IETF 79
7–12 November 2010
Host: Tsinghua University
Location: Beijing, China

IETF 80
27 March–1 April 2011
Host: TBD
Location: Prague, Czech Republic

IETF 81
24-29 July 2011
Host: TBD
Location: Quebec City, Canada

Register now for

IETF 78

25–30 July 2010

Maastricht, The Netherlands

https://www.ietf.org/registration/ietf78/ietfreg.py

Early bird registration: USD 635 if paid in full prior to 1700 PDT July 16, 2010
Regular registration: USD 785 USD if paid after early bird cutoff date and time
Full-time students: USD 150 with on-site proof of ID
One-day pass: USD 350 (see website for details)

The ISOC Fellowship to the IETF is sponsored by

Afilias Global Registry Services, Google, intel, Microsoft, nominet trust

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

Afilias Global Registry Services, ARIN, Cisco, comcast, Ripe NCC, WIDE Project

IETF® Journal
IETF 77
Volume 6, Issue 1
June 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.

Photos property of the Internet Society unless otherwise noted

Email
ietfjournal@isoc.org
Find us on the Web at
http://ietfjournal.isoc.org

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

Internet Society