



A report from IETF 81, July 2011, Quebec City, Canada. Published by the Internet Society in cooperation with the Internet Engineering Task Force*

Inside this Issue

From the Editor's Desk	1
DANE: Taking TLS Authentication to the Next Level Using DNSSEC.....	1
Message from the IETF Chair	2
Words from the IAB Chair	3
ISOC Panel Addresses Regulation, Innovation, and the Internet	8
IAB Plenary Tackles IPv6, Privacy Issue	10
New Technology Demo: PCP.....	13
Internet Society Fellows Are Welcome Additions to IETF Standards Work.....	15
IETF Ornithology: Recent Sightings	16
IETF At-A-Glance	18
IRTF Update	19
Calendar	20

From the Editor's Desk

By Mat Ford

The recent and widely publicized intrusion into the DigiNotar certificate authority infrastructure amply highlighted the weaknesses of the systems widely in use today to secure online communications. The advent of DNSSEC deployment raises the intriguing possibility of using the DNS as a secure repository for certificates in the future. In our cover article, Richard Barnes offers a detailed overview of the DANE working group's efforts to make this possibility a technical reality (see below).

Attendees of the IETF 81 meeting in Quebec City will have noticed the invitations to visit a demonstration of the Port Control Protocol in the terminal room. For a detailed explanation of the purposes of this new IETF technology and the demo that was presented, see page 13.

In our IETF Ornithology column, we again provide an overview of the proceedings and outcomes of the Birds-of-a Feather (BoF) meetings that took place during IETF 81. These are frequently some of the most interesting and accessible meetings for the general observer as participants seek to explain the background to and motivations for new work topics (see page 16). Also in this issue are our regular columns from the IETF and IAB chairs, highlights from the ISOC panel on Internet evolution, coverage of hot topics discussed during plenary meetings, and an opportunity to meet the ISOC Fellows to IETF 81.

I'm pleased to announce new subscription options to enable receipt of the *IETF Journal* in advance of the next IETF meeting. Have the latest edition sent to you as soon as it is available in hardcopy or via email by subscribing at: <https://www.isoc.org/apps/events/ietf-subscription.php>.

As always, we are hugely grateful to all our contributors. Please send comments and suggestions for future issues to ietfjournal@isoc.org.

DANE: Taking TLS Authentication to the Next Level Using DNSSEC

By Richard L. Barnes

Abstract

Authentication of Domain Name System (DNS) names for Transport-Layer Security (TLS) endpoints is a core security challenge in many Internet protocols, most famously Hypertext Transfer Protocol (HTTP). Today, the cryptographic bindings that underlie TLS authentication are asserted in Public Key Infrastructure for X.509 (PKIX) certificates issued by third-party certification authorities (CAs). The DNS-based Authentication of Named Entities (DANE) working group is developing protocols that allow certificates to be bound to DNS names using Domain Name System Security Extensions (DNSSEC). These protocols will enable additional assurances for the traditional, PKIX-based model, as well as enabling domain holders to assert certificates for themselves, without reference to third-party certificate authorities. With this increased responsibility, however,

Continued on page 4



* The articles published in the IETF Journal are not intended to reflect the opinions or the position of the IETF or the Internet Society. See <http://www.ietf.org>.

Message from the IETF Chair

By Russ Housley

The IETF participants are energetic! The work of the IETF is highly relevant!

IETF 81 was held in Quebec City, Canada. It was a successful meeting, attended by 1,057 people from 46 countries. This first-time meeting in Quebec City was wonderfully hosted by Research In Motion (RIM)—convention center facilities were very comfortable and Tuesday evening's social event at the Musée de la Civilisation was well attended. Comcast and .ca were our sponsors, and Bell and Telus provided network connectivity. Thanks to all for your support.

Many working groups made significant progress at IETF 81, and it was a genuine pleasure to see so many talented people engaged and collaborating.

Since IETF 80, five working groups (WGs) have been chartered and five have closed—our count remains steady at 121 WGs. Between meetings, the WGs and their individual contributors produced 553 new Internet-Drafts and updated 1,138 existing Internet-Drafts, some more than once. The Internet Engineering Steering Group (IESG) approved 100 Internet-Drafts for publication as RFCs. The RFC Editor published 149 new RFCs.

The BEHAVE Working Group has essentially finished its work on mechanisms to help transition from IPv4 to IPv6. I look forward to the day when the vast majority of Internet traffic is using IPv6, and these mechanisms are removed from the Internet. The HOMENET Working Group is an important step in this direction. This new working group is working on specifications for IPv6 for residential networks.

The IETF continues to improve its tools. The Datatracker provides a great deal of visibility into the processing of the documents in the IETF stream. The Datatracker was recently extended to include visibility into actions within Working Groups. Over the next few months, it will be further extended to provide visibility into the processing of documents in the IRTF, IAB, and Independent Submission streams.

IETF 82 will take place in Taipei, Taiwan, 13–18 November 2011, and will be hosted by the Taiwan Network Information Center (TWNIC). Scheduling information for the upcoming IETF meetings can always be found at <http://www.ietf.org/meetings/meetings.html>. I look forward to seeing you there. 



Russ Housley, IETF Chair

The IETF continues to improve its tools. The Datatracker provides a great deal of visibility into the processing of the documents in the IETF stream. The Datatracker was recently extended to include visibility into actions within Working Groups. Over the next few months, it will be further extended to provide visibility into the processing of documents in the IRTF, IAB, and Independent Submission streams.

The mission of the Internet Engineering Task Force is to make the Internet work better by producing high-quality and relevant technical documents that influence the way people design, use, and manage the Internet. See <http://www.ietf.org>.

Recent IESG Document and Protocol Actions

A full list of recent IESG Document and Protocol Actions can be found at <https://datatracker.ietf.org/iesg/ann/new/>

Words from the IAB Chair

By **Bernard Aboba**

The IETF 81 Technical Plenary included a report on World IPv6 Day organized by Leslie Daigle, as well as a series of presentations on the Web Privacy Tustle, organized by Alissa Cooper. World IPv6 Day presentations included reports from Facebook (Donn Lee), Google (Lorenzo Colitti), Yahoo! (Igor Gashinsky), Telefonica (Carlos Ralli Ucendo), and Cisco (Mark Townsley). For the session on the Web Privacy Tustle, Jens Grossklags, Fred Carter, Andy Zeigler, and Alissa Cooper provided their thoughts.¹



Bernard Aboba, IAB Chair

The IAB held its annual retreat 12–13 May 2011, at VeriSign in Sterling, Virginia. During the retreat, the IAB reviewed the Programme and Initiative framework that it put in place during last year's retreat, and agreed to commit to the framework as an organizational tool going forward.

Programmes are long-term activities that are scoped and managed by the IAB and expected to continue over multiple IAB terms. These include IANA Evolution (led by Olaf Kolkman), Internationalization (led by Dave Thaler), ITU-T Coordination (led by Andrei Robachevsky), Liaison Oversight (led by Spencer Dawkins), Privacy (led by Alissa Cooper), and the RFC Editor (RSOC) (led by Joel Halpern).

Initiatives are short-term activities that can be completed in one tenure, usually resulting in an RFC, statement, or presentation. These include IPv6 for IAB Business (led by Bernard Aboba), IP Evolution (led by Danny McPherson), HTTP/Web Evolution (led by Jon Peterson), and DNS (led by Jon Peterson).

Minutes of the IAB retreat have been posted² and since the retreat, a description of each programme and initiative, as well as the membership list, has been made available on the IAB website.³

One of the topics discussed at the IAB retreat was interactions with governments. In the months since the retreat, those interactions have continued. On 28 July, the IAB responded to the “Internet Assigned Numbers Authority Further Notice of Inquiry” from the National Telecommunications and Information Administration.⁴

On 16 September, the IAB sent a letter relating to global interoperability in emergency services to the European Commission.⁵

Another topic discussed at the IAB retreat was liaison management. In response to “Some IESG Thoughts on Liaisons,”⁶ the IAB drafted a reply,⁷ and began acting on the recommendations. In August, Scott Mansfield was appointed as IETF Liaison to ITU-T for MPLS (replacing Stewart Bryant), and John Drake was appointed as IETF Liaison to ITU-T SG-15 for the optical control plane (replacing Adrian Farrel). Previously, on 30 June, the IAB announced the appointment of Eliot Lear as the IETF Liaison Manager to the ITU-T (replacing Patrik Falström).

The IAB extends its thanks to Adrian, Stewart, and Patrik for their service to the community.

On 29 June, the IAB announced a Call for Comment on “The RFC Editor Model (Version 2).” Joel Halpern is revising the document to incorporate feedback from the community. The search for an RFC Series Editor, first announced in July,⁸ continues.

The IAB has issued draft reports on the privacy workshop hosted at MIT by the IAB, W3C, ISOC, and MIT CSAIL in December 2010,⁹ as well as the smart object workshop hosted on 25 March, 2011, in Prague.¹⁰

The views and positions documented in the workshop reports are those of the workshop participants and do not necessarily represent the views of the IETF, W3C, IAB, IESG, or ISOC.


Continued on next page

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

Words from the IAB Chair, continued from page 3

On 7 June 2011, the IAB met with the W3C Technical Architecture Group to discuss common architectural issues in security and privacy, as well as ongoing work within the IETF and W3C on real time communications and location.¹¹

On 27 June, the IAB responded to ARIN's request for guidance regarding draft policy ARIN-2011-5.¹²

The IAB appointed Ole Jacobsen as IETF representative to the 2012 ICANN Nomcom, and reappointed Thomas Narten as IETF liaison the ICANN Board of Directors. 

References

1. Technical plenary materials, <http://www.ietf.org/proceedings/81/technical-plenary.html>
2. IAB Retreat minutes, <http://www.iab.org/documents/minutes/minutes-2011/iab-minutes-2011-05-12/>
3. IAB Activities, <http://www.iab.org/activities/>
4. IAB response to the IANA FNOI, <http://www.iab.org/2011/07/28/iab-responds-to-internet-assigned-numbers-authority-iana-functions-further-notice-of-inquiry/>
5. IAB letter to the European Commission, <http://www.iab.org/2011/09/16/iab-sends-letter-to-the-european-commission-on-global-interoperability-in-emergency-services/>
6. "Some IESG Thoughts on Liaisons," <http://trac.tools.ietf.org/group/iesg/trac/wiki/LiaisonThoughts>
7. IAB Response to "IESG Thoughts on Liaisons," <http://www.iab.org/documents/correspondence-reports-documents/2011-2/iab-response-to-some-iesg-thoughts-on-liaisons/>
8. RFC Series Editor Search, <http://www.iab.org/2011/07/14/rfc-series-editor-search/>
9. Privacy Workshop Report, <http://tools.ietf.org/html/draft-iab-privacy-workshop>
10. Smart Objects Workshop Report, <http://tools.ietf.org/html/draft-iab-smart-object-workshop>
11. IAB Meets with W3C TAG, <http://www.iab.org/2011/06/07/iab-meets-with-the-w3c-technical-architecture-group-tag/>
12. IAB Response to ARIN, <http://www.iab.org/2011/06/27/iab-responds-to-arin-request-for-guidance-regarding-draft-policy-arin-2011-5/>

DANE: Taking TLS Authentication to the Next Level Using DNSSEC, continued from page 1

DNS operators will play a more critical role in securing applications. So while DANE holds the promise of more direct authentication, it will also create some new security challenges.

Introduction

TLS is used as the basis for security features in many modern Internet application service protocols in order to provide secure client-server connections (RFC 5246). It underlies secure HTTP and secure email (RFC 2818, RFC 2595, RFC 3207), and provides hop-by-hop security in real-time multimedia and instant-messaging protocols (RFC 3261, RFC 6120). In all of these applications, the service that the user ultimately wants to connect to is identified by a DNS domain name (RFC 1034). A user might enter "https://example.com/" into a web browser or send an email to "alice@example.com". One of the main purposes of using TLS in these cases is thus to authenticate the server as a legitimate representative of the domain name, in other words, to assure the user that the entity on the other end of the connection actually represents "example.com". Note that this applies to DTLS as well, since it uses the same handshake as TLS.

Today, a server asserts its right to represent a domain by presenting a PKIX digital certificate (RFC 5280). A client then has to evaluate the certificate to decide whether or not the certificate is sufficient to establish the server's identity. This decision is usually based on two criteria:

1. whether the certificate contains the desired domain name (RFC 6125), and
2. whether the certificate is issued under a trusted certification authority (in the PKIX parlance, a "trusted CA" or "trust anchor").

In order to apply this second criterion, the client must choose some set of trust

anchors. In current systems, the choice of trust anchors is mostly out of the user's hands. For example, all modern browsers and operating systems come with default trust anchor lists. Users can add to this list, using their browsers' "Accept this certificate?" dialogs, although it can be very difficult to remove trust anchors from the default list.¹

The result of the current, manual model is that trust anchors end up having very broad authority. With no way to discover which trust anchor should be vouching for a particular set of domain names, current systems allow any trust anchor to issue certificates for any domain name. The risks of having broadly trusted CAs have recently become clear, as attackers were able to break into two small CAs and create fraudulent certificates for Google and Facebook, among others.^{2,3} The importance of default trust lists has also driven up the cost to application providers of deploying TLS, leading some to rely on self-signed certificates, which provide no authentication at all (without some out-of-band negotiation).

The IETF DANE working group⁴ was chartered to develop ways to use DNSSEC to improve TLS authentication of domain names. With the advent of DNSSEC, it is possible for clients to obtain authenticated data directly from zone operators. In the context of TLS in particular, DNSSEC should allow clients to securely ask the operator of a domain about which certificates they should accept as credentials for that domain. The domain operator might specify constraints on certificate validation or even supply a new trust anchor. As with most new technologies, there will be a few challenges in migration; most notably, DNS operators will play a far larger role in the security of Internet applications. But if these challenges can be overcome, then DANE should increase both the



Quebec City, site of IETF 81

security of applications and the ability of domain operators to secure the services they offer.

DANE Records

If the goal of DANE is to allow domain operators to make statements about how clients should judge TLS certificates for their domains, then what sorts of statements should DANE allow them to make? The DANE use cases document (draft-ietf-dane-use-cases) lays out three major types of statements:

1. CA Constraints: The client should only accept certificates issued under a specific CA.
2. Service Certificate Constraints: The client should only accept a specific certificate.
3. Trust Anchor Assertion: The client should use a domain-provided trust anchor to validate certificates for that domain.

All three of the above statements can be viewed as constraining the scope of trust anchors. The first two types limit the scope of existing trust anchors; the third provides the client with a new trust anchor (still within a limited scope). More on these in a moment.

The current draft DANE protocol defines a DNS Resource Record type TLSA for describing TLS associations—statements about what

certificates are associated to a domain (draft-ietf-dane-protocol). Each TLSA record has three basic fields:

- Usage: Which type of statement this record is making
- Selector/Matching: How a TLS certificate chain should be matched against this record (e.g., by exact match, by public key, or by SHA-1 digest)
- Certificate for Association: The actual data against which the TLS certificate chain should be matched

These records are stored under the target domain with a prefix that indicates the transport and port number for the TLS server. For example, if Alice runs a secure web service at example.com, and wants to tell clients to only accept certificates from Charlie's CA, she could provision a TLSA record under `_443._tcp.example.com` with the following contents:

- Usage: CA constraint
- Selector/Matching: SHA-1 digest
- Certificate for Association: SHA-1 digest of Charlie's certificate

When client Bob wants to connect to "https://example.com", he can find these TLSA records and apply Alice's constraints when he validates the server's certificate.

Adding Constraints to PKIX

The major objective of the CA constraints and service certificate constraints is to guard against mis-issue of certificates. A certificate is mis-issued when a CA issues a certificate to an entity that does not represent the domain name in the certificate. Mis-issue can come about in many ways, including via malicious CAs, compromised CAs (as in the Comodo and DigiNotar example above), or CAs that are simply misled as to the attacker's identity through fraud or other means. Today, mis-issue can be difficult to detect, as there is not a standard way for clients to figure out which CAs are supposed to be issuing certificates for a domain. When an attacker issued false certificates for the Google Gmail service under the DigiNotar CA, it was noticed only because a vigilant user posted to a Gmail help forum.⁵

By contrast, domain operators know exactly which CAs they have requested certificates from, and, of course, which specific certificates they have received. With DANE, the domain operator can securely convey this information to the client. For example, to guard against the DigiNotar attack, Google could have provisioned a TLSA record expressing a CA constraint with their real CA (which is not DigiNotar) or a certificate constraint with their actual certificate. Then DANE-aware clients would have been able to immediately see that the DigiNotar certificates were improperly issued and possibly indicative of a man-in-the-middle attack.

Empowering Domain Operators

According to data from the EFF SSL Observatory, which scans the whole IPv4 address space for HTTPS servers and collects their certificates, approximately 48 percent of all HTTPS servers present self-signed certificates.⁶ A number of other servers present

Continued on next page

DANE: Taking TLS Authentication to the Next Level Using DNSSEC, continued from page 5

certificates issued under CAs that are not in the major default trust anchor lists. For example, the United States Air Force web portal⁷ uses a certificate issued under a Department of Defense CA that is not trusted by Firefox. In the current environment, most clients cannot authenticate these servers at all; they have to rely on users manually accepting certificates, hopefully with some out-of-band information. As a result, these servers and their users are highly vulnerable to man-in-the-middle attacks against their supposedly secure sessions.

DANE trust anchor assertions enable the operators of a domain to advertise a new trust anchor, under which certificates for that domain are issued. Using these records, clients can dynamically discover what trust anchors they should accept for a given domain instead of relying on a static list provided by a browser or operating system.

It may seem odd to talk about a domain supplying a client with trust anchors, since trust anchor provisioning is typically a very sensitive activity. If an attacker is able to install a trust anchor

into a victim's trust anchor store, then the attacker can masquerade under any name by issuing certificates under that name. In fact, the PKIX working group has defined an entire protocol for managing trust anchors (RFC 5934).

DANE ensures that trust-anchor provisioning is secure by applying scoping, and verifying its scoping using DNSSEC. DANE trust anchor assertions are scoped to a particular domain name, so even if an attacker can introduce a false trust anchor, he can only use it to spoof a single name. Furthermore, trust anchor assertions must be DNSSEC signed, meaning clients can verify that the entity providing the trust anchor represents the domain in question. Ultimately, the client still has to have a list of trust anchors configured, but DNSSEC trust anchors instead of PKIX trust anchors. Of course, in principle, a client only needs one trust anchor for DNSSEC—the root zone trust anchor. Since control of the DNS root doesn't change often, it makes sense that it be statically configured!

The ability of a domain operator to explicitly indicate a trust anchor for a

domain is obviously very powerful. It may be tempting to ask whether this is really the only use case that DANE needs, i.e., whether the constraint cases above are needed at all. The answer is that the constraint cases are useful as a way to fold in PKIX validation with external CAs, in addition to domain-asserted trust anchors. Most obviously, this feature is useful in transition, when not all clients are DANE-aware. But even in the longer term, it's possible that CAs will provide added value over DANE. For example, although DANE is designed to bind certificates to domain names, CAs can vouch for bindings of certificates to other things, including the legal identity and physical location attested to in Extended Validation certificates.⁸

Transition Challenges

As described above, DANE offers valuable new security properties for TLS authentication. But, as with most IETF technologies—especially security technologies—there are both challenges to overcome and potential pitfalls.

The most significant constraint for DANE deployment is DNSSEC deployment. On the server side, this is not a huge issue considering DNSSEC support is spreading fairly rapidly. On the client side, things may be more difficult. While there are DNS libraries with robust DNSSEC support, many of the major DNS APIs that are used by applications do not provide information about the DNSSEC status of the results returned. In order to implement DANE, application developers may have to refactor their DNS support, in addition to querying for new record types. If more sites come to rely on DANE, this could also draw increasing attention to the various types of intermediaries that cause DNSSEC breakage (e.g., home gateways that improperly set DNS flags).

Adding DNSSEC to the TLS connection process can also add significant



latency to the TLS connection process. In addition to completing the TLS handshake and certificate validation, the client has to wait for several DNS round-trips and then validate the chain of DNSSEC signatures. These combined delays can add up to multiple seconds of latency in connection establishment. Especially for real-time protocols such as HTTPS, SIP, or XMPP, such delay is clearly undesirable.

The primary mechanism proposed to mitigate these delays is to have the server pre-fetch all of the relevant DNSSEC records (i.e., DS, DNSKEY, and RRSIG records chaining back to the root). Then the server can provide a serialized version of the DNSSEC records in the TLS handshake, saving the client the latency of the required DNS queries. The details of this mechanism, however, are still being worked out among the DANE, TLS, and PKIX working groups (draft-agl-dane-serializechain). A prototype version is now available in the Google Chrome web browser.⁹

Security Considerations

From a security perspective, the major impact of DANE is the new role that DNS operators will play in securing Internet applications. DNSSEC has meant that DNS operators have more security functions; DANE deployment will give them an explicit impact on application security, acting as arbiters of who can authenticate under a given name in TLS. Particularly if services make use of trust anchor assertions, DNS operators will play an analogous role to the one CAs play today—meaning, a compromise in a DNS operator will enable an attacker to masquerade as a victim domain (albeit for a more limited set of domains due to DANE's constraints on names). In this way, DNS operators are likely to inherit many of the security troubles that CAs experience today and will need to strengthen their security posture accordingly.


Another more subtle risk arises from the fact that the operator of a DNS zone is not always the same as the entity that is authorized to control the contents of the zone, which we can call the domain holder. We used the phrase domain operator above, since DNSSEC only protects DNS information between the operator's name server and the client—it doesn't say that what's provisioned in the name server is authorized by the domain holder. When a domain is operated by a third party, the third party is a point of vulnerability between the client and the holder of the domain. If the domain operator provides false DANE information through malice or compromise, then a client will not be able to distinguish it from genuine DANE information. To some extent, this risk is not new; many current CAs authenticate requests for domain certificates based on information under the control of the domain operator, domain operators can already influence the credentialing process. With DANE, however, the vulnerability is much easier to exploit as the DNS operator needn't trick a third party. This vulnerability is also fundamental to protocols that rely on DNSSEC for security. The implications for DANE are discussed in detail in the DANE use cases document (draft-ietf-dane-use-cases). The main mitigation is simply increased care on the part of domain holders to ensure that domain operators are not behaving badly.

Conclusions

For many years now, Internet applications have relied on assertions by third-party certification authorities to ensure that a server holding a particular private key was authorized to represent a domain. The promise of DANE is a more direct interaction between clients and the domains they interact with, secured by DNSSEC. In the short run, DANE can be deployed as an adjunct to the current system of certificates and authorities, adding constraints to

better protect domains. In the long run, DANE will enable domain operators to vouch for their own names.

The transition and security issues that face DANE are largely the growing pains of DNSSEC. It's not that DANE is causing these problems itself; rather, the problems arise because DANE is the first real usage of DNSSEC that is expected to be widely deployed. So while it may be difficult to mitigate some of the security issues raised by DANE and to enable more robust DNSSEC support in applications and gateways, these changes will ultimately make it simpler for applications to use DNSSEC for other purposes.

The DANE working group is making consistent progress on its deliverables and there already exist some prototype deployment tools. Their use cases document has been approved by the IESG (draft-ietf-dane-use-cases), and the document defining the TLSA record type is maturing (draft-ietf-dane-protocol). On the client side, a variant of DANE has been implemented in Google Chrome; on the server side, there are prototype tools available to generate DANE records and to generate DNSSEC-stapled certificates based on DANE records.^{10,11} 

References

1. <http://arstechnica.com/apple/news/2011/09/safari-users-still-susceptible-to-attacks-using-fake-diginotar-certs.ars>
2. <http://blogs.comodo.com/it-security/data-security/the-recent-ra-compromise/>
3. http://www.theregister.co.uk/2011/09/06/diginotar_audit_damning_fail/
4. <http://tools.ietf.org/wg/dane>
5. <http://www.google.co.uk/support/forum/p/gmail/thread?tid=2da6158b094b225a&hl=en>
6. <http://www.eff.org/observatory>
7. <https://www.my.af.mil/>
8. http://cabforum.org/Guidelines_v1_2.pdf
9. <http://www.imperialviolet.org/2011/06/16/dnssecchrome.html>
10. <https://dane.xelerance.com/>
11. <http://www.imperialviolet.org/2011/06/16/dnssecchrome.html>

ISOC Panel Addresses Regulation, Innovation, and the Internet

By Carolyn Duffy Marsan

What will drive the Internet's evolution in the future: market forces or government regulation? This was the topic debated by a panel of experts at an Internet Society-sponsored luncheon held in Quebec City in conjunction with the IETF 81 meeting.

Panelists, including experts from across the IETF community, predicted a future of increasing regulation given how the Internet has become critical infrastructure for both government and industry.

Leslie Daigle, chief Internet technology officer at ISOC, explained that three key aspects of the Internet need to be retained, regardless of how the network and its underlying protocols evolve: permissionless innovation, open access, and collaboration.

The idea that anyone can innovate on the Internet "is really a fundamental and important aspect of the Internet

experience and under the hood that we need to preserve," said Daigle. "That comes to play in ensuring the technology and protocols that we have are ... building blocks. They are not full systems or closed systems that are so tied to one purpose that they can't be used in any other fashion."

Daigle added that, unlike broadcast television, it's important that the Internet remain an open medium that fosters collaboration—both for the common good and for self interest.

"The network is no longer a playground," she added. "This creates an environment where, perhaps, the notion

"The network is evolving whether we like it or not, whether we plan it or not. This is not the network it was a few years back. It's changing, and it's changing in a very unplanned and interesting way. But we're really losing the ability to reason about what the concatenation of functionality along the path is."

—Mark Handley
Professor of Networked Systems
University College London

of the evolution of the Internet is different than it was a decade ago."

Mark Handley, professor of networked systems at University College London, talked about some of the technical challenges facing Internet evolution. As an example, he pointed to two protocols that he has worked on—SIP and Multipath TCP—which evolved in ways he never would have predicted when they were being designed.

"We built a lot of flexibility into

SIP—perhaps more flexibility than we should have," said Handley. "But it has survived the process, and it has been distorted into being used in lots of different ways that we never thought of."

Handley's research shows that middleboxes run by access networks and content delivery networks are poorly understood from an operational perspective and create significant technical challenges to designing new protocols or adding new functionality.



Panel speaker Geoff Huston

"The network is evolving whether we like it or not, whether we plan it or not," said Handley. "This is not the network it was a few years back. It's changing, and it's changing in a very unplanned and interesting way. But we're really losing the ability to reason about what the concatenation of functionality along the path is."

Handley warned that protocol designers who propose doing more functions over UDP or HTTP to get around middleboxes on TCP will merely cause a situation where middleboxes reappear one layer higher on the network stack.

"We can't evolve the Internet in any kind of reasonably controlled manner, nor even reason about how it works, unless we understand the motivation behind these middleboxes, why they are there, and whether we can then retroactively try to fit them into some kind of architecture," he said.

Bernard Aboba, principal architect with the Lync Division at Microsoft, predicted more government regulation of the Internet because the Internet is putting tax-generating industries out of business.

“The Internet has fairly rapidly transitioned from a plucky challenger to the incumbent, and the former incumbent is on the verge of obsolescence or bankruptcy,” said Aboba, pointing to the recent failure of bookstore chains and other publishers. “Plain old telephone service may become a relic by the end of the decade. There are entire industries that are potentially being obsoleted by the Internet. This has effects within our communities and on a global basis, which are quite considerable.”

For example, state and local governments are seeing their tax revenues decline as more sales occur on e-commerce sites. As telephone revenues dwindle, so do the taxes collected on telephone bills that cover emergency services for the disabled.

“The pace of Internet change is so rapid in many cases that what we have is the need to rethink entire areas of regulation within a very short period of time,” said Aboba.

One issue is the looming obsolescence of plain old telephone service, which the U.S. Federal Communications Commission reports could occur as soon as 2018. According to Aboba, preparing for the telephone system to be turned off in seven years would be a “monumental” engineering and regulatory task.



Audience member Lee Howard

“The argument over how [the Internet] is to be regulated and taxed is just beginning,” said Aboba. “Ultimately we will have to rethink what rules apply and see what rules make sense on the Internet because it is inherently not a local thing, it is a global thing.”

Geoff Huston, chief scientist at APNIC, heralded the Internet as a poster child of the power of open, unregulated markets. He said it was because of its unfettered nature that the Internet was able to transform technology, business and human life during the last decade. However, he warned that this era of openness is ending.

“The Internet has been an absolute triumph of open markets. But markets fail, markets distort, monopolies form, [and] cartels form. How do you prevent that in something like the Internet? And, indeed, what’s happening in the Internet is that volume economics are creating massive monopolies,” he said. “What’s happening is the market is ossifying. Innovation doesn’t work very well when you are trying to service two billion users.”

Huston points out two issues that will hamper Internet innovation in the future: 1) the sheer size of the Internet and the fact that it is the public telecommunications system of the world, and 2) the idea that network neutrality is going by the wayside in favor of device manufacturers approving content. This creates a scenario where the incumbents can set the terms and conditions of the challenges, something that hasn’t happened in the Internet before.

“All of a sudden we’re recreating what we found so disastrous in the 1960s with



Audience member and IETF Fellow Khoudia Gueye Sy

the telephone companies. We’re recreating massive amounts of ubiquitous control through carriage,” said Huston, adding that free markets created this problem and that it’s up to regulators to fix it.

Huston predicted that if regulators fail to protect the Internet’s openness, it will result in only a handful of major companies offering products and services to Internet users for many years to come. Then the demand for innovation will build up over a decade or two until there is revolutionary change.

“True salvation ... is going to come from the regulatory sector, but we’re asking an awful lot, perhaps too much,” said Huston, “because we’re asking for that very delicate, light touch that keeps the incumbents to the level where innovation is still possible... where your bright idea actually has the ability to redefine tomorrow’s business.”

Huston predicted that if regulators fail to protect the Internet’s openness, it will result in only a handful of major companies offering products and services to Internet users for many years to come. Then the demand for innovation will build up over a decade or two until there is revolutionary change.

IAB Plenary Tackles IPv6, Privacy Issues

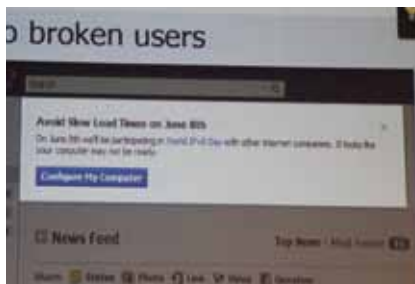
By Carolyn Duffy Marsan

The IAB's plenary in Quebec City featured a recap of World IPv6 Day as well as a discussion of privacy considerations on the Web.

World IPv6 Day Panel

World IPv6 Day, held 8 June 2011, was a tremendous success in terms of encouraging Web sites and content delivery networks to conduct a 24-hour trial of IPv6, said Leslie Daigle, chief Internet technology officer for the Internet Society.

"Facebook, Google, Yahoo, and more than 1,000 other Web sites turned on IPv6 on their front door," Daigle said, adding that the goal of the Internet



Preparation for World IPv6 Day on Facebook.

Society-organized event was to motivate ISPs, hardware makers, operating system vendors, and other Web companies to prepare their services for IPv6 as well as to understand the issues related to IPv6 transition. "It was a good day for the Internet coming together."

Daigle said two-thirds of the participating Web sites left IPv6 on after the event was over.

"There was no large-scale breakage. The DDoS fears did not pan out. Overall it was a success," Daigle said. "We moved the needle on IPv6 deployment."

According to Lorenzo Colitti, network engineer for Google, the most important thing that Google did for World IPv6 Day was to warn users that they may suffer from IPv6 brokenness and offer them information about how to fix the problem.

As a result of these warnings, he continued, World IPv6 Day was business as usual for Google.

"We got 65 percent more traffic on World IPv6 Day," Colitti said. "And the day after, we kept most of the traffic... Some services left IPv6 on for everyone, including YouTube and Mobile Gtalk."

One reason that Google users suffered from less IPv6 brokenness than expected is that Google implemented a technique called "Fast Fallback" in its Chrome browser. This feature allows users with broken IPv6 connectivity to automatically revert to IPv4. It is also available in Chrome and Firefox 7.

"We saw an 80 percent to 90 percent reduction in dual-stack brokenness," Colitti said. "If all browsers behaved like this, we would publish our Quad A records (instead of using whitelisting). The browser versions with Fast Fallback were 99.995 percent as reliable as IPv4..."

"There was no large-scale breakage. The DDoS fears did not pan out. Overall it was a success. We moved the needle on IPv6 deployment."

—Leslie Daigle
Chief Internet Technology Officer
Internet Society

Apple is adding this robustness in OS X Lion. All we need is IE to follow suit."

Google, which serves 60 percent of the IPv6 Internet, said that the bulk of IPv6 adoption globally is in two networks: France's Free and Japan's KDDI. KDDI, for example, distributed IPv6 to 15 percent of its user base in the five weeks prior to World IPv6 Day.

"One ISP by itself made a significant difference in the adoption of IPv6 in Japan," Colitti added.

Donn Lee, a network engineer at Facebook, said the social media site saw more than 1 million IPv6 users on World IPv6 Day, which represented 0.20 percent of its users. The number of users suffering from dual-stack IPv6 brokenness was 0.02 percent, which was down from measurements of 0.03 percent prior to the event.

"We sent a message to broken users that said on June 8 we'll be participating in World IPv6 Day and it looks like your computer may not be ready," Lee said, adding that IPv6 brokenness "seems to be declining after World IPv6 Day."

The preparation work helped Facebook to experience no technical difficulties on World IPv6 Day. In fact, Facebook permanently dual stacked its developer's site after the event ended.

Facebook was pleasantly surprised that all of its Facebook Connect Widgets—served up by 2.5 million Web sites—worked on World IPv6 Day in dual-stack mode.

"We also found with moving to IPv6 that there is nothing to fear," Lee said.

"It works. It's ready to deploy. It wasn't that hard to do."

Igor Gashinsky, a principal architect with Yahoo, said World IPv6 Day not only helped encourage IPv6 adoption, it also helped shrink the number of users suffering from IPv6 brokenness from 0.078 percent to 0.022 percent in less than two years.

“Over 30 different Yahoo markets participated... All of the markets are served in 10 different data centers,” Gashinsky said. “Our initial plan was to complete IPv6 in all of these sites sometime in 2010, but we had problems with just about every vendor’s implementation in their switching gear. As a result, we have seven IPv6 proxy locations.”

After extensive preparations to its infrastructure along with outreach to its users about the issue of IPv6 brokenness, Yahoo had a successful World IPv6 Day experience.

“We served over 2.2 million users via IPv6. We saw over 1 million visits to Yahoo’s IPv6 help pages. At the peak, 0.229 percent of our users were via IPv6,” Gashinsky said. “Over 1.2 million IPv6 users were from France, which was almost double the U.S.”

Still, Gashinsky pointed out that IPv6 is not widely deployed enough for Yahoo to support IPv6 permanently.

“That was a lot of work for 0.229 percent of our users,” he said. “We need more IPv6 access... Can we break single digits, please, and then we can talk about leaving it on?”



IESG member Ron Bonica



The IAB plenary panel address a filled-to-capacity room.

Telefonica enabled Web sites in Spain and Brazil for World IPv6 Day. Network engineer Carlos Ralli said its user brokenness stayed below 0.04 percent.

“We saw no related call center activity. This is good news,” Ralli added.

Cisco experienced the highest percentage of IPv6 users on World IPv6 Day, with 1.11 percent of its traffic coming to www.cisco.com via IPv6. Even better, the company saw zero tech support calls related to World IPv6 Day. As a result of this positive experience, Cisco left one of its Web sites—www.scansafe.com—up in dual-stack mode when the event ended.

“Cisco customers and users are more interested in IPv6 than the broader population of users,” said Cisco Fellow Mark Townsley, who pointed out that 2.26 percent of the company’s logged-in users on World IPv6 Day had IPv6 capabilities.

Privacy Panel

Jens Grossklags, assistant professor at Pennsylvania State University’s College of Information Sciences and Technology, kicked off the privacy panel with a discussion of several experiments that he has conducted on how privacy concerns impact consumer behavior.

Grossklags has discovered that whether people were identified as privacy fundamentalists, profiling concerned, identity concerned or unconcerned about privacy, they gave away more information about themselves online than they had anticipated.

“Across all scenarios, the degree of information revelation is higher than you would expect from a rational consumer,” he said, adding that privacy is about “hard decision making over time, with actions now having consequences later.”

In one experiment, Grossklags discovered that user attitudes were affected when consumers received before and after notifications with warnings about the risks of downloading software. However, when he paid people to download unknown executables, they were willing to forego their privacy concerns.

“People who should have known better participated once the price was right,” he said. “Seventy percent of the participants knew it would be dangerous to download unknown programmes, but all of them did it anyway.”

His takeaway is that users are subject to immediate gratification, and that the users with more protection such as

Continued on next page

IAB Plenary Tackles IPv6, Privacy Issues, continued from page 11

“Privacy risks exist in most technologies, even ones that might appear to have little risk.”

—Andy Zeigler
Programme Manager, Internet Explorer Team
Microsoft

antivirus software are more likely to take risks, such as downloading executables, than those who are unprotected.

Fred Carter, senior policy and technology advisor for the Office of the Information and Privacy Commissioner for Ontario, gave a regulator’s perspective on information privacy issues. He said the key issues around information privacy online were minimizing the use, sharing, and collection of personally identifiable information (PII), thereby enhancing data security and engaging individuals in managing their own PII.

Carter discussed the concept of “Privacy by Design,” which has been adopted by privacy regulators worldwide. Privacy by Design has seven foundational principals, which include having privacy as a default setting, offering full functionality along with privacy, and providing visibility and transparency about information gathering.

Privacy by Design “is gaining ground as a high-level normative framework,” Carter says. “Work is still needed on how to operationalize it and apply it to information infrastructures, networked [systems], and related engineering standards.”

Carter urged IETF participants to consider the Privacy by Design framework in protocol design. “The next stage is to have people like yourselves apply it to particular cases and teach us what the best practices are,” he added.

Andy Zeigler, a programme manager with Microsoft’s Internet Explorer team, gave IETF participants background on privacy-related issues that have cropped

up for browser makers with such technologies as CSS 2.1 and geolocation.

“Privacy risks exist in most technologies, even ones that might appear to have little risk,” Zeigler said, pointing out that its best to take privacy into consideration when authoring specifications. “Privacy risks can be very difficult to fix after a spec is implemented.”

Zeigler discussed how users are surreptitiously tracked as they browse the Web and that there are many benefits to this tracking, including personalization. “The problem is the ownership and control of the information,” he said.


Zeigler pointed out that Microsoft was an early adopter of privacy controls with its support of the W3C’s P3P standards. But P3P proved to be too complex for users, too simple for nuanced business relationships, and was not being implemented by many Web sites.

Today, Microsoft is supporting the idea of Tracking Protection Lists, which block tracking content, in IE9. This version of its Web browser also supports the “Do Not Track” header for HTTP.

To wrap up the privacy panel, Alissa Cooper described the IAB’s Privacy Programme, which aims to develop privacy thinking within the technical standards community. Cooper outlined an approach to protocol design that would involve systemic privacy threat modeling similar to how security considerations are taken into account during the standards development process.

Cooper outlined several challenges for IETF participants as they scope privacy threats, including the diversity of user privacy preferences, a lack of incentives for supporting privacy features, and the fact that common practices or laws might dictate the emphasis for online privacy.

“One of the big questions that we are trying to grapple with is how do we become more systemic at building threat models for privacy,” Cooper said. “That begs the question of how to decide what threats are in scope and what threats are not in scope.”

Cooper asked participants to review a document (draft-morris-privacy-considerations-03) and comment. 



Quebec City streetscape painting by Carolyn M. Maynard

New Technology Demo: PCP

By Tina Tsou, Christian Jacquenet, and Xiaohong Deng

Several organizations collaborated to carry out the Port Control Protocol (PCP) demonstrations during the IETF Quebec City meeting. PCP is a simple, flexible, lightweight protocol that is being designed by the IETF to address some of the issues raised by the forthcoming IPv6 transition period where access to some legacy IPv4 content requires the control of firewall or network address translator capabilities for the dynamic allocation of transport-layer port numbers.

The demonstration was inspired by the IETF mantra of “rough consensus and running code” to expose IETF technologies to real operator requirements and scenarios, and to develop the technology innovation and provide feedback to the IETF community.

One demonstration consisted of a Universal Plug-n-Play (UPnP) to PCP interworking function (<http://data-tracker.ietf.org/doc/draft-bpw-pcp-upnp-igd-interworking>) that has been implemented in a Customer Premises Equipment (CPE) device provided by

France Telecom Orange while the PCP server was implemented in Huawei’s NE40E router that also supports Dual-Stack Lite (DS-Lite) and Carrier Grade NAT (CGN) capabilities.

The Internet Gateway Device (IGD) machinery is used between computer and UPnP/IGD to request the allocation of a port number to the CPE so that a pinhole can be created accordingly to allow access to the content requested by the terminal from the Internet. The demo has proven the capability of PCP after only a few months of IETF specification effort. This demo provided the first experimental implementation of draft-ietf-pcp-base-12 and draft-bpw-pcp-upnp-igd-interworking-02. It was a pioneering effort and helped to inspire thought and discussion leading to a better deployment of this technology.

The first demonstration scenario comprised two computers, both with BitTorrent file-sharing software installed and both connected to two different CPEs, so that the terminals could display in real time the difference between BitTorrent clients with and without UPnP-PCP interworking functions supported for file-exchange purposes. The BitTorrent software of one computer either randomly selects a port number or uses the port number specified by the user to listen on. The computer then uses UPnP/IGD to interact with the CPE. With a pinhole assigned by the CGN, access to the BitTorrent client from the Internet is made possible. Compared to a second computer without a pinhole assigned on the CGN, the first BitTorrent client often had a faster download speed because when clients can be accessed remotely they are able to see more peers, thereby improving file-sharing performance.

In the second demonstration scenario, the CPE requested several

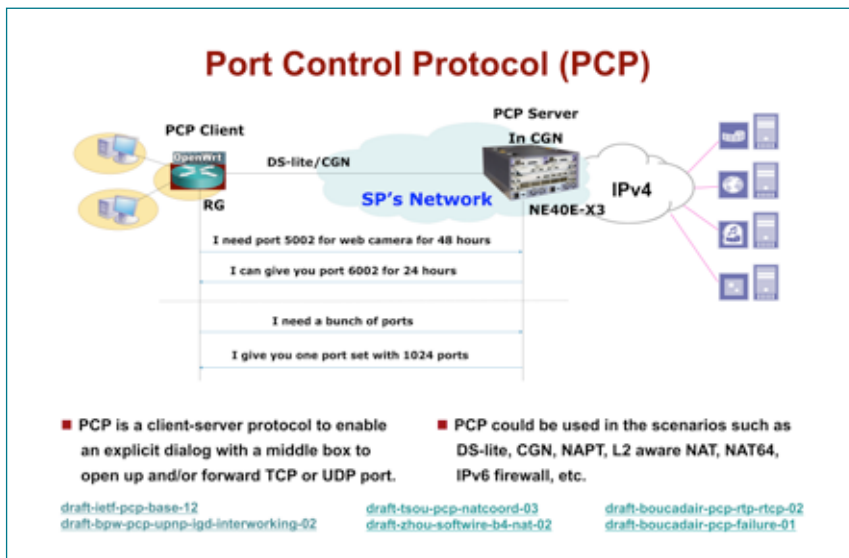


Figure 1: Functional Picture

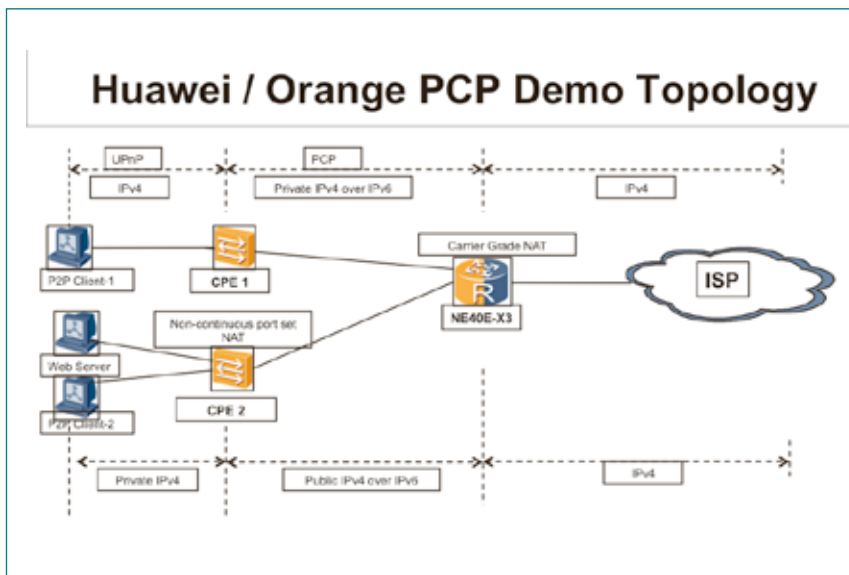


Figure 2: Demo Topology Picture

Continued on next page

New Technology Demo: PCP, continued

sets of noncontiguous ports (utilizing draft-tsou-pcp-natcoord-03 and draft-zhou-softwire-b4-nat-02). Upon receipt of the corresponding PCP request sent by the client, the PCP server requested the CGN to assign port forwarding and to bypass NAT on the requested port ranges. The NAT function was performed on the CPE from this point forward, thus reducing the NAT processing requirement on the CGN router.

In the third scenario, users were given the ability to form the PCP request from a specific web portal—this reflects a context in which the customer is offered the capability to explicitly request the port number(s) needed to ensure that their content, maintained in their premises, can be accessed from the Internet. Once the port number request has been fulfilled, the PCP Client embedded in the CPE then forwards the corresponding PCP Request message to the PCP server. Different options within the PCP packet can be set dynamically.

A further demonstration, provided by China Telecom, showcased Lightweight 4over6 (<http://tools.ietf.org/html/draft-cui-softwire-b4-translated-ds-lite-01>). This is an IPv4/IPv6 transitional solution developed by China Telecom, which uses the PCP protocol to dynamically allocate port-restricted addresses to subscribers. Through this approach, users in the demo room can simultaneously access IPv4 and



Xiaohong Deng, France Telecom (left), talks Port Control Protocol with Tobias Gondrom.



Port Control Protocol Demo staff John Wang (left) and Susan Hares (right), Huawei Technologies Canada.

This demo provided the first experimental implementation of draft-ietf-pcp-base-12 and draft-bpw-pcp-upnp-igd-interworking-02. It was a pioneering effort and helped to inspire thought and discussion leading to a better deployment of this technology.

IPv6 services over an IPv6-only access network. In order to reduce the volume of state that must be maintained in the network, the Lightweight 4over6 approach maintains the NAT capabilities on the client. The core network need only maintain per-subscriber state instead of per-session state, thus the workload of the core network can be reduced significantly.

Internet Systems Consortium demonstrated a PCP server controlling a software-based DS-Lite CGN running on a commodity Linux netbook. An OpenWrt-based CPE provided the UPnP-IGD and NAT-PMP Interworking Functions.

An unmodified BitTorrent client was used to demonstrate the interworking function, while a modified BitTorrent client demonstrated an application-based PCP client communicating directly with the PCP server to request port mappings.


The PCP-enabled BitTorrent client was also used to perform basic interoperability testing with the Huawei PCP server, the first time such interoperability testing has been attempted.

Demonstration team personnel included:

- **France Telecom Orange:**
Christian Jacquenet, Xiaohong Deng, Mohamed Boucadair, Gu Daqing, Wang Lan
- **Huawei Technologies Canada:**
Susan Hares, Tina Tsou, Thomas Zhang, Cathy Zhou, John Wang, Victor Marin, Bill Weng, David Gao, Gary Jan
- **ISC:** Francis Dupont, Paul Selkirk
- **China Telecom:** Chongfeng Xie, Qiong Sun



John Wang, Huawei Technologies Canada (right), shares his Port Control Protocol expertise.



The Port Control Protocol Demo staff.

Internet Society Fellows to the IETF Increase Participation on a Global Scale

Six members of an Internet Society Fellowship programme took part in the 81st meeting of the Internet Engineering Task Force (IETF). The Internet Society Fellows to the IETF programme, which operates under the aegis of the Internet Society's Internet Leadership programme, began in 2006 and is today an established Internet Society activity. Through a competitive process, the Internet Society selects Fellows from a talented pool of applicants from around the world. Fellows infuse IETF meetings with their diverse experience and local expertise about how the Internet works in their communities.

"Since the inception of the Fellows to the IETF programme, the Internet Society has selected and supported 57 engineers from more than 25 developing and emerging economies," said Toral Cowieson, senior director of Internet leadership at the Internet Society. "Having these

engineers engaged in this critical standards work helps ensure representation of a broad range of viewpoints and the ongoing development of globally relevant and effective Internet standards."

As in prior years, each first-time Fellow was assigned a mentor to assist them in networking with others with similar Internet technology interests and to advance specific standards work. These Fellows will continue to participate in the IETF's standards development process and will bring back to their local communities the knowledge and expertise they gained from their IETF experience.

"We are also very excited about the caliber and diversity of Fellows that will be joining us at IETF 82 in Taipei, Taiwan," continued Cowieson. "Of the 13 new and returning Fellows joining us, we'll have representation from ten countries, all of whom will bring their unique perspective to the meeting."

Fellowship Alumni from IETF 81



Caciano Machado (Brazil)
Network Operator
Federal University of Rio Grande do Sul



Amir Qayyum (Pakistan)
Professor and Head of Electronic
Engineering Department
M. A. Jinnah University, Islamabad



Jeronimo Bezerra (Brazil)
Returning Fellow
Network Analyst, Point of Presence of
National Network of Research; Master's
student, Federal University of Bahia



Richard Mikisa (Uganda)
Manager IP Networks
Orange Uganda Limited



Nestor Tiglao (Philippines)
PhD student, Department of Electrical and
Computer Engineering INESC-ID/INOV
Technical University of Lisbon



Khoudia Gueye Sy (Senegal)
Returning Fellow
Head of Network and Internet Services
University Cheikh Anta Diop of Dakar

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

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

IETF Ornithology: Recent Sightings

Compiled by Mat Ford

Getting new work started in the IETF usually requires a Birds-of-a-Feather (BoF) meeting to discuss goals for the work and to help assess the level of interest in and support for new work. In this article, we'll review the BoFs that took place during the IETF, their intentions and outcomes. If you're inspired to arrange a BoF meeting, please be sure to read RFC 5434, Considerations for Having a Successful Birds-of-a-Feather (BoF) Session. Full descriptions of the BoFs that were proposed in the run-up to the IETF 81 meeting can be found on the wiki, here: <http://trac.tools.ietf.org/bof/trac/wiki/WikiStart>.

REPUTE—Reputation Services



Swainson's Thrush
(*Catharus ustulatus*)

Description: This was a Working Group-forming BoF meeting to gauge momentum and interest in working on protocols and other specifications for providing a general-purpose reputation framework. In the open Internet, making a meaningful choice about the handling of content requires an assessment of its safety or trustworthiness. This is based on a trust metric for the owner (identity) of an identifier associated with the content, to distinguish (likely) good actors from bad actors. The generic term for such information is *reputation*.

This working group would develop mechanisms for reputation reporting by independent services. One mechanism would be for a basic assessment of trustworthiness. Another would provide a range of attribute/value data that is used as input to such an assessment.

Proceedings: <http://www.ietf.org/proceedings/81/repute.html>

Outcome: Discussion was rather wide-ranging during the meeting but focused down to working specifically on reputation services for email. Several concerns were expressed about the complexity of this

work if not very narrowly focused, and the fact that prior work in this space hasn't gained traction historically. Nonetheless, people are interested in doing the work and there were a lot of people interested in seeing the output, so chartering discussions for a Working Group will continue.

MULTRANS—Multicast Transition

Description: This BoF meeting considered the question of how to devise the means whereby existing multicast mechanisms can operate successfully when signaling and content must traverse one or more IP version boundaries.

Proceedings: <http://www.ietf.org/proceedings/81/multrans.html>

Outcome: The BoF identified the concrete scenarios that were of most importance, namely IPv4 sources with both IPv4 and IPv6 receivers to support IPTV applications. The discussion concerning whether or not these applications are required to work interdomain did not conclude. Several gaps in understanding were identified and the existing list of requirements needs further work and input. The meeting did not reach the point of considering the question of whether a Working Group could be formed, so a second BoF meeting during IETF 82 seems likely.

In the open Internet, making a meaningful choice about the handling of content requires an assessment of its safety or trustworthiness. This is based on a trust metric for the owner (identity) of an identifier associated with the content, to distinguish (likely) good actors from bad actors. The generic term for such information is 'reputation'.

CICM—Common Interface to Cryptographic Modules

Description: The Common Interface to Cryptographic Modules (CICM) (pronounced *kick-em*) defines an abstract API for the security services provided by cryptographic modules developed by multiple vendors. The API is intended to support high assurance cryptography, security domain separation, and enhanced module, key, and channel management capabilities that are vendor neutral. The purpose of a CICM Working Group would be to publish an API for high assurance cryptographic devices and to provide guidance for any new submissions related to high assurance cryptos.

Proceedings: <http://www.ietf.org/proceedings/81/cicm.html>

Outcomes: It wasn't clear from the presented materials how this work could fit in the IETF, and the question of whether it would be a better fit for the IRTF was raised. Significant charter reworking is required and more detailed elaboration of how the proposed work could impact on existing IETF protocols.

It wasn't clear from the presented materials how this work could fit in the IETF, and the question of whether it would be a better fit for the IRTF was raised. Significant charter reworking is required and more detailed elaboration of how the proposed work could impact on existing IETF protocols.

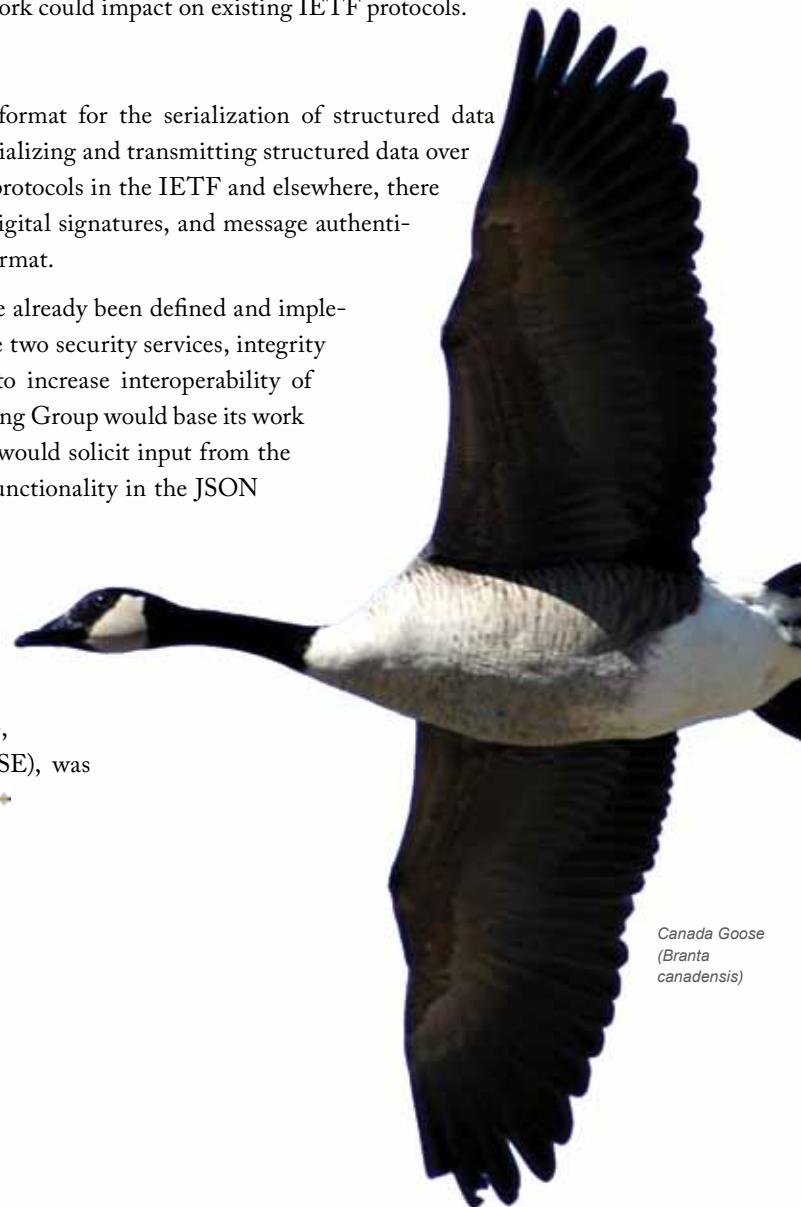
WOES—Web Object Encryption and Signing

Description: Javascript Object Notation (JSON) is a text format for the serialization of structured data described in RFC 4627. The JSON format is often used for serializing and transmitting structured data over a network connection. With the increased usage of JSON in protocols in the IETF and elsewhere, there is now a desire to offer security services such as encryption, digital signatures, and message authentication codes (MACs) for data that is being carried in JSON format.

Different proposals for providing such security services have already been defined and implemented. This proposed Working Group's task is to standardize two security services, integrity protection (signature and MAC) and encryption, in order to increase interoperability of security features between protocols that use JSON. The Working Group would base its work on well-known message security primitives (e.g., CMS), and would solicit input from the rest of the IETF Security Area to be sure that the security functionality in the JSON format is correct.

Proceedings: <http://www.ietf.org/proceedings/81/woes.html>

Outcome: This BoF meeting went very well. There was clear consensus about the work plan and very little controversy on the points raised. A draft charter for the WG, to be called Javascript Object Signing and Encryption (JOSE), was submitted to the community for review on 30 August. 🍌



Canada Goose
(*Branta canadensis*)



IETF 81 At-A-Glance

Registered attendees: 1057

Newcomers: 133

Number of countries: 46

IETF Activity since IETF 80 (March–June 2011)

New WGs: 5

WGs closed: 5

WG currently chartered: 121

New Internet-Drafts: 553

- 183 updated
- 43 updated more than once

Updated Internet-Drafts: 1138

IETF Last Calls: 128

Internet-Drafts approved for publication: 100

RFCs published: 149

- 63 Standards Track and 5 BCP
- 66 Informational and 15 Experimental

IANA Activity since IETF 80 (March–June 2011)

Processed 1532 IETF-related requests, including:

- 776 private enterprise number requests

- 45 port number requests
- 54 TRIP ITAD number requests
- 32 media-type requests

Reviewed 137 I-Ds in Last Call and reviewed 130 I-Ds in IESG Evaluation

Reviewed 132 I-Ds prior to becoming RFCs and 74 of them contained actions for IANA

Cumulative percentage average for IETF-related requests: 95%

Protocol registries conversation to XML: 78% complete

RFC Editor Activity (March–June 2011)


Digital signatures for RFCs

- Discussed on rfc-interest list (subject: [rfc-i] Signing RFCs)
- As of 19 July 2011: RFC Editor will sign RFCs in same manner as I-Ds (as defined in RFC 5485)

Conclusion of FYI RFC sub-series

- Draft-iesg-rfc1150bis-01.txt approved
- Obsoletes RFC 1150
- Moves RFC 1150 to Historic status

Be the First Host on Your LAN to Receive the *IETF Journal*!



Receive the latest edition of the *IETF Journal* as soon as it is available—in hardcopy or via email. Subscribe today at:

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

IRTF Update

By Lars Eggert

During IETF 81 in Quebec City, Canada, seven research groups held meetings:



- Delay-Tolerant Networking RG
- Peer-to-Peer RG
- Network Management RG
- IP Mobility Optimizations RG
- Host Identity Protocol RG
- Scalable Adaptive Multicast RG
- Virtual Networks RG

In addition, a first IRTF Open Meeting was held. The purpose of an Open Meeting is to allow interested IETF attendees to get a quick overview of all current IRTF activities and to discuss topics that are of IRTF-wide relevance, such as proposals to form new research groups. It is expected that IRTF Open Meetings will be held regularly during IETF meeting weeks.

The IRTF Open Meeting is also where the awardees of the Applied Networking Research Prize (ANRP) give their invited talks. The ANRP is awarded for recent results in applied networking research that are relevant for transitioning into shipping Internet products and related standardization efforts, and is supported by the Internet Society in coordination with the IRTF. For IETF 81, two ANRPs were awarded. The awardees were Mattia Rossi for his research into reducing BGP traffic, and Beichuan Zhang for his research into “green” traffic engineering. The ANRP selection process for IETF 82 is currently underway, with decisions expected at the end of September 2011. See <http://irtf.org/anrp> for details.


Several of the IRTF’s current research groups are nearing the end of their chartered work. These include the Traffic Modeling RG, the Anti-Spam RG, the Host Identity Protocol RG, and the Routing RG. These groups are discussing whether to take on new work or to close. Researchers in the relevant areas are encouraged to suggest topics for potential future collaboration.

On the IRTF RFC Stream, eight new RFCs were published since IETF 80. Six of those RFCs came out of the Delay-Tolerant Networking RG, which has now published the bulk of the experimental DTN specifications and supporting documents. The Routing RG and the IP Mobility Optimizations RG each published one RFC.

Please join the IRTF discussion list to stay informed: <http://www.irtf.org/mailman/listinfo/irtf-discuss>. 



Applied Network Research Prize winner Mattia Rossi.



Applied Network Research Prize winner Beichuan Zhang.

IETF Meeting Calendar

IETF 82

13–18 November 2011
Host: Taiwan Network Information
Center (TWNIC)
Location: Taipei, TW

IETF 83

25–30 March 2012
Host: TBD
Location: Paris, FR

IETF 84

29 July–3 August 2012
Host: Google
Location: Vancouver, BC, CA

IETF 85

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

For more information about past and upcoming

IETF Meetings

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

Special thanks to



for hosting IETF 81

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



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



IETF® Journal

IETF 81
Volume 7, Issue 2
October 2011

Published three times
a year by the
Internet Society

Galerie Jean-Malbuisson 15
1204 Geneva
Switzerland

Editor
Mat Ford

Associate Editors
Megan Kruse
Wendy Rickard

Contributing Writer
Carolyn Marsan

Editorial and Design
The Rickard Group, Inc.

Editorial Board
Bernard Aboba
Leslie Daigle
Mat Ford
Russ Housley
Lucy Lynch
Wendy Rickard
Greg Wood

Email
ietfjournal@isoc.org

Find us on the Web
<http://ietfjournal.isoc.org>

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

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

