the IETF® Journal

Volume 7, Issue 1 • July 2011

*A report from IETF 80, March 2011, Prague, Czech Republic. Published by the Internet Society in cooperation with the Internet Engineering Task Force**

Inside this issue

Making the Internet Work Better: Bufferbloat and Secure Routing Infrastructure at IETF 80

From the Editor's Desk, by Matthew Ford

The IETF has been working on securing the routing infrastructure of the Internet for many years. The Secure Inter-Domain Routing (sidr) working group has recently been rechartered and it is, therefore, timely to review the status of the current efforts in the IETF on this topic. Contributing authors Geoff Huston and Randy Bush do just that in our cover article, Securing BGP and SIDR (this page).

Prague, Czech Republic, site of IETF 80

A new feature in this issue of the *IETF Journal* is our IETF Ornithology column (page 16), which provides an overview of the proceedings and outcomes of the Birds-of-a Feather (BoF) meetings that took place during IETF 80. These are frequently some of the most interesting and accessible meetings for a general observer as participants seek to explain the background to and motivations for new work topics.

Jim Gettys has been on the campaign trail since late last year trying to raise awareness of the issues caused by excessively large buffers in network hardware and software. His presentation to the Transport Area Open Meeting (http://www.ietf.org/proceedings/80/slides/tsvarea-1.pdf) during IETF 80 generated a lot of discussion, and we present an article from him on the topic in this issue (page 14). You can follow progress on this topic at http://www.bufferbloat.net/.

Also in this issue are our regular columns from the IETF and Internet Architecture Board chairs, highlights from the Internet Society panel on metrics for IPv6 deployment, coverage of the hot-topics discussed during the plenary meetings, and an opportunity to get to know the Internet Society Fellowship to IETF 80 Fellows from around the world.

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

Securing BGP and SIDR

By Geoff Huston and Randy Bush

For many years the Internet's fundamental elements—names and addresses—were the source of basic structural vulnerabilities in the network. With the increasing momentum behind the deployment of Domain Name System Security Extensions (DNSSEC) there is some cause for optimism that we have the elements of securing the name space now in hand, but what about addresses and routing? In this article we look at current efforts within the IETF to secure the use of addresses within the routing infrastructure of the Internet, and the current work of the Secure Inter-Domain Routing (SIDR) working group.

Continued on page 8

Message from the IETF Chair

By Russ Housley


Russ Housley, IETF Chair

The work of the IETF remains relevant and energetic!

IETF 80 was the second IETF meeting held in Prague, Czech Republic (the first was IETF 68 in March 2007) and, once again, CZ.NIC did a wonderful job as host. It was a very successful meeting, attended by 1,196 people from 49 different countries. The facilities at the Hilton were comfortable and the unique social event on Tuesday evening at the Municipal House was well attended. On behalf of the IETF community, I would like to express our appreciation to Comcast and Trilogy for their generous sponsorship as well as to Dial Telecom, which provided network connectivity. Thank you to all for your support.

Many working groups (WGs) made significant progress at IETF 80. It was a genuine pleasure to see so many talented people engaged and collaborating.

On 3 February 2011, IANA assigned the last five IPv4 address blocks to the Regional Internet Registries (RIRs). At a very nice ceremony in Miami, Florida, USA, each RIR received one of the final address blocks.

Since IETF 79, 11 new WGs have been chartered and 14 were closed for a total of 121 WGs. Between the meetings, the WGs and their individual contributors produced 608 new Internet-Drafts and updated 1,237 existing Internet-Drafts, some more than once. The Internet Engineering Steering Group approved 164 Internet-Drafts for publication as RFCs and the RFC Editor published 104 new RFCs.

As you probably already know, on 3 February 2011, IANA assigned the last five IPv4 address blocks to the Regional Internet Registries (RIRs). At a very nice ceremony in Miami, Florida, USA, each RIR received one of the final address blocks. We have all known that the IPv4 address space would run out this year. In preparation, the IETF developed IPv6, which has long been ready for deployment; the IPv4 run out simply increases the urgency for IPv6 deployment. The explosive growth of the Internet can only continue with the bigger address space offered by IPv6.

I wish to emphasize that the depletion of the IANA IPv4 address pool is not a crisis. The Internet has not been significantly different since the beginning of February; there have not been any notable changes caused by the empty IANA IPv4 address pool. While this is not a crisis, there is a need for action to ensure that the Internet can continue to grow. The transition to IPv6 requires the attention of many actors. Our parents, spouses, and children will be largely unaware of the transition from IPv4 to IPv6—for them, this milestone may be insignificant—but making the transition will ensure their continued amazement at the endless possibilities offered by the growing Internet. Thanks for all that you have done to help make the transition to IPv6 possible. I hope you will continue working to make it a reality.

IETF 81 will take place in Quebec City, Canada, on 24–29 July 2011 and will be hosted by RIM. 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.

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 http://www.isoc.org/ietfjournal/DocProtoActions0701.shtml

Words from the IAB Chair

By Bernard Aboba


Bernard Aboba, IAB Chair

At the March 2011 IETF meeting, the Internet Architecture Board (IAB) selected new leadership for both the IAB and the Internet Research Task Force (IRTF). Lars Eggert was named chair of the IRTF, succeeding Aaron Falk. I was named IAB chair, succeeding Olaf Kolkman. The IAB would like to express its gratitude to both Aaron and Olaf for their years of dedicated service to the community.

The IETF 80 Technical Plenary focused on the evolution of browser functionality and some of the implications for the future of applications. Jon Peterson organized the session, which included presentations by IAB alumni Harald Alvestrand, Leslie Daigle, and Jonathan Rosenberg as well as Henry Thompson from the W3C TAG.

Since IETF 79, the IAB has been involved in two workshops. The first was the Internet Privacy Workshop, which was held in December 2010 at the Massachusetts Institute of Technology (MIT) Computer Science and Artificial Intelligence Laboratory. Sponsored by the IAB, the Internet Society, MIT, and W3C, the workshop focused on how to ensure that the architecture and technology of the Internet, including the Web, are developed in a way that respects user privacy. The issue of privacy was covered from a number of perspectives, including legal/regulatory and technical. Technical discussions addressed all levels of the stack, including network layer privacy (Tor) and issues associated with privacy on the Web. The workshop minutes have been posted. Alissa Cooper will be editing the workshop report.

The second workshop, Interconnecting Smart Objects with the Internet, was held in Prague, Czech Republic, in March 2011. Participants explored design considerations for protocols running on constrained devices, such as energy efficiency. Hannes Tschofenig is working on the minutes of that workshop and will be editing the workshop report.

The IAB has both appointed the RFC Series Oversight Committee (RSOC) and issued an RFC Editor Model (Version 2) draft. The draft, which is under discussion on the RFC Interest list (rfc-interest@rfc-editor.org) as well as within the IAB, reflects a year of experience with RFC 5620 (RFC Editor Model Version 1). That document is being edited by Joel Halpern.

With respect to the evolution of the IANA function, the IAB provided comments in response to the National Telecommunications and Information Association's Request for Comments on the IANA Functions. It also published RFC 6220: Defining the Role and Function of IETF Protocol Parameter Registry Operators. Kudos to the RFC Editor team for handling the expedited publication, despite the compressed time frame.

As I write this, the IAB has just completed its annual retreat, held 12–13 May 2011 at Verisign in Sterling, Virginia, USA. Danny McPherson handled the logistics. The major focus of the retreat was the IAB's work plan for the next 12 months and the refinement of the programme structure laid out by the IAB a year ago. In the next edition of the *IETF Journal*, we will talk about the retreat in more detail.

References

- Privacy Workshop position papers, slides and minutes: http://www.iab.org/about/workshops/privacy/
- Smart Objects Workshop agenda, position papers, and presentation slides: http://www.iab.org/about/workshops/smartobjects/
- RSOC appointment: http://www.ietf.org/mail-archive/web/ietf-announce/current/msg08652.html
- The RFC Editor Model (Version 2): http://tools.ietf.org/html/draft-iab-rfc-editor-model-v2
- RFC-Interest list: https://www.rfc-editor.org/mailman/listinfo/rfc-interest
- The NTIA "Request for Comments on the IANA Functions": http://www.ntia.doc.gov/frnotices/2011/fr_ianafunctionsnoi_02252011.pdf

Continued on page 4

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.

ISOC Panel Debates Metrics for IPv6 Progress

By Carolyn Duffy Marsan

What metrics should the Internet engineering community use to track deployment of IPv6? That was the question debated by a panel of experts at an Internet Society event held on 29 March in conjunction with the IETF meeting in Prague, Czech Republic.

The panel, entitled "IPv6: How will we know we've arrived?", considered various measurements of IPv6 deployment, including traffic statistics, hardware capability, support from existing Web services, and the creation of new applications.

Leslie Daigle, chief Internet technology officer at the Internet Society, said a statistic such as having 20 percent of the Internet's traffic using IPv6 by the end of 2011 is beyond a stretch goal. That's because CZ.NIC, Arbor Networks, and others estimate that IPv6 currently represents around 0.1 percent or 0.2 percent of all Internet traffic.

"We're looking for a tipping point,'" Leslie said. "We're looking for the point where there is enough IPv6 in the world that it gets the attention of business decision makers who need to be signing the checks and the work plans to actually deploy IPv6. [We need recognition] that this is not an experiment anymore but is, in fact, business viable."

Lee Howard, director of network technology at Time Warner Cable, said the U.S.-based ISP turned up its first IPv6 customers last year and would conduct a residential IPv6 trial in spring 2011. He said the driving force for Time Warner Cable to adopt IPv6 is to avoid the costs associated with carrier-grade network address translation (NAT).

Instead of measuring IPv6 traffic, Lee recommends measuring IPv6-based connections. He hopes to have as many network connections as possible traverse IPv6 rather than carrier-grade NAT, so he is monitoring the progress of connection-oriented applications such as online gaming to IPv6.

"We talk about iTunes or Google Maps as opening many, many sessions at once as being key applications for us," Lee says. "Where IPv6 will be most needed…is in applications that open many sessions at once. Those are the kind of places I am looking for IPv6 and where I want to measure success."

Carlos Ralli, an IPv6 expert at Telefónica, said IPv6 progress can be tracked by watching the activities of various parts of the Internet community, including network operators, network equipment vendors and software developers

"We have no doubt that the main driver is IPv4 exhaustion," Carlos says. "But we believe once we will be there, there will be people exploiting the end-to-end communications" with new applications.

Vesna Manojlovic, a technical course developer with RIPE NCC, said the European Internet registry has devised two measurements to track IPv6 deployment.

The first measurement, dubbed IPv6 RIPEness, is a rating system that awards network operators stars when they do the following: receive IPv6 address space, make their IPv6 addresses visible in RIPE NCC's routing information service, have a route6 object in the RIPE Whois database, and enable reverse DNS. Vesna said network operators that pass all of these tests receive four stars on a list maintained by RIPE NCC.

"This is not a measurement of traffic; this is just capability," Vesna said. "These are the prerequisites for pushing the bits around. If they don't have this, probably they are not going to have traffic either."

Interrnet Society panellists, from left, Ondrej Filip, Vesna Manojlovic, Carlos Ralli, and Lee Howard

Manojlovic said 40 percent of all network operators in RIPE NCC's region have achieved all four stars on the IPv6 RIPEness measurement. "Last year, it was about 30 percent. Now it's 40 percent. So by the end of the year, 50 percent is quite realistic," she added.

RIPE NCC splits out its IPv6 RIPEness data on a per-country basis, too. "This is very useful for regulators and governments because when they compare themselves to their neighbours and see that their neighbours are doing

Words from the IAB Chair, continued

- The IAB response to NTIA: http://www.ntia.doc.gov/comments/110207099-1099-01/attachments/IAB%2DIANA%2DNOI%2Edoc
- Other responses: http://www.ntia.doc.gov/comments/110207099-1099-01/
- IAB Programs and Initiatives: http://www.ietf.org/mail-archive/web/ietf-announce/current/msg07731.html
- IETF 80 proceedings: https://datatracker.ietf.org/meeting/80/materials.html

better than they are, they want to catch up," she added.

RIPE NCC also tracks another metric on a global basis, which is the number of IPv6-enabled Autonomous System Numbers (ASNs). This data is available by region and by country.

"The global average is 9 percent. The Czech Republic has 15 percent of ASNs IPv6-enabled, and Holland has 35 percent," Vesna says. "I cannot set any goal for that [measurement], but it could be something interesting to look at in the future."

CZ.NIC, which operates the .cz top-level domain, has found that 7 percent of the domains in the Czech Republic have Quad A records—the type of DNS records used for name resolution with IPv6—and 20 percent of domains have Quad A records for name servers.

"That sounds pro-mising," says Ondrej Filip, CEO at CZ.NIC. "But on the other hand, if you look at the exchange point, and look at the ratio of IPv6 and IPv6 traffic, the IPv6 traffic is just 0.1 percent of the total traffic at the exchange point. We believe we are quite good on the side of the content providers, but we still need to work a little bit on the side of the ISPs."

Ondrej said another sign of IPv6 progress in the Czech Republic is that the country's largest ISP, Telefónica, and its most popular Web site, Seznam.cz, have agreed to offer IPv6 this year.

Leslie summed up the panel discussion by pointing out that the group did not agree on one single metric for measuring progress in IPv6 deployment.

"Of all the things we discussed today, the one that was not mentioned…is turning IPv4 off because that's so far in the future and it's not anywhere near our midterm goals," she said, adding that "the key is to make sure that the Internet keeps working" during this transition period from IPv4 to IPv6.

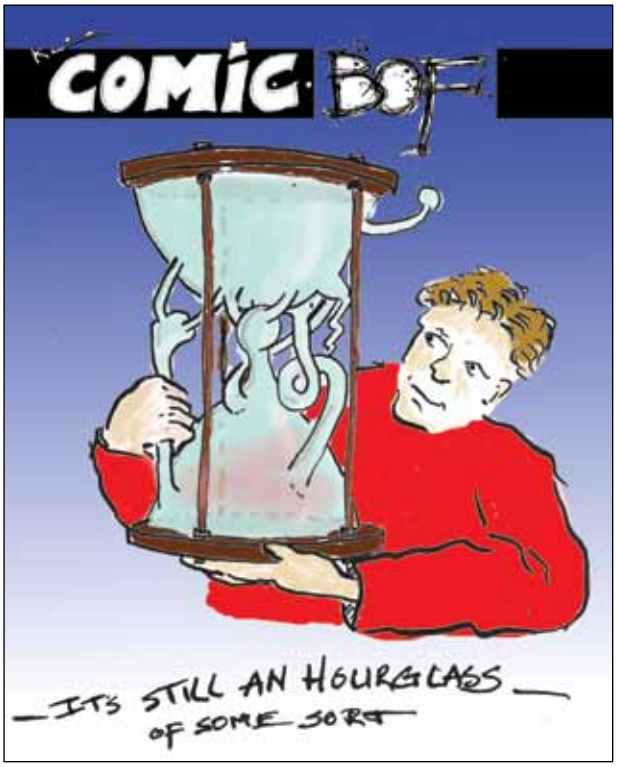

Technical Plenary Tackles Role of IETF in Application Protocols

By Carolyn Duffy Marsan

What role—if any–should the IETF play in the development of application protocols? That was the hot-button issue debated by expert panelists at the Internet Architecture Board's technical plenary session on March 28 in Prague, Czech Republic.

The panellists noted that today's Internet application developers tend to favour an open source approach, such as publishing their own application programming interfaces (APIs), rather than participating in standards bodies such as the IETF. The reason developers are doing this is because it dramatically shortens their time to market versus going through the traditional standards development process.

"These guys have achieved really fast scale for interprovider messaging without standards," says Jonathan Rosenberg, chief technology strategist at Skype. "It's interesting because it's all about a new model for delivery of apps to users…. Nowhere during this process did anybody need to show up at an IETF meeting and ask for a standard…. All the intermediaries are cut out and the dependencies are gone, and that's what has brought life and innovation to these apps."

Several panelists promoted the idea of the IETF getting involved in new standards-development work that could underpin these APIs and provide basic, interoperable functionality to next-generation browsers.

"The consequence of the browser as a delivery platform is that there are an enormous amount of APIs required," said Henry Thompson, who serves as liaison between the IETF and the W3C. "With so many new APIs…what happens to the One Web for All goal?"

Thompson kicked off the panel with an overview of activities at the W3C's Technical Architecture Group related to future Web architecture. Thompson said that both the IETF and the W3C were

Continued on next page

Technical Plenary, continued

facing some of the same difficulties in engaging participants in application-oriented work, and he recommended that they cooperate more in the future.

"The fact that IETF and the W3C have proceeded in relatively amicable parallel tracks for many years without regular exchange of people is, in hindsight, unfortunate," Henry said. "Having a little more first person interaction between the two groups is good for all of us."

Henry said W3C is in the process of reconsidering its Web architecture documents in light of the fact that the Web is no longer a collection of documents, but instead has evolved into a collection of documents, data and applications.

"Tensions have arrived because HTTP, HTML and browsers were not designed to deliver applications, they were designed to share documents," Henry said. "Our new Open Web Platform is a platform for innovation, consolidation and cost efficiency…. These new Web architecture documents are pretty much squarely in W3C territory, but many of the more recent concerns have drawn us into IAB territory. Clearly, it's time to give some thought to demarcation."

Harald Alvestrand, a former IETF chair, urged the group's participants to continue to be involved in the development of application-oriented standards with an overriding goal of making the Internet work better. He promoted the idea of Real Time Communications Web (RTC Web), a proposed working group that would create a set of specifications that would be useful for the interoperability of future browsers.

"We should have a uniform interface in the browser so you can send media from one browser to another without all this plug-in stuff," Alvestrand said. "You should be able to have compatibility between browsers. If I use Firefox, and you use Opera, it doesn't matter; we are compatible. But in order to actually work, it has to be matched with uniform APIs inside the browsers so the downloadable applications can run anywhere naturally."

Alvestrand said that having uniform APIs will help address the fact that end users have different kinds of platforms—with varying screen sizes, microphones and cameras – and that applications need to detect those differences and behave accordingly. "This makes interfaces complicated unless we are really good at designing them," he said.

Having a standard real-time communications platform for browsers would foster innovation because it is so expensive to build a proprietary one, Alvestrand argued.

"If the interfaces are standard and universally deployed, anyone can write a video-using application. It's cheap," he said. "If you have standard APIs and functionality availability, you can just try something, put it out there, and see if anyone uses it."

In summation, Harald said that the Internet's new applet paradigm requires more standards, not fewer standards.

"The IETF needs to lose its fear of the type of protocol called an API," he said. "The IETF needs to take responsibility for making sure the whole thing works up to a level where it can actually be accessed. The IETF needs to work together with other organizations so the

> "It's interesting because it's all about a new model for delivery of apps to users…. Nowhere during this process did anybody need to show up at an IETF meeting and ask for a standard…. All the intermediaries are cut out and the dependencies are gone, and that's what has brought life and innovation to these apps."
>
> — Jonathan Rosenberg

Photo/Peter Löthberg

Outgoing IAOC and IESG members on stage during the Plenary

appropriate wisdom…is made available to the right people at the right time."

Jonathan, one of the lead authors of the IETF's Session Initiation Protocol (SIP), discussed the successes and failures of this VOIP protocol as an example of how the group's standards are getting adopted—or not adopted—in the application space. He pointed out that the IETF began work on SIP in 1999 and has published more than 100 technical specifications related to SIP in the intervening nine years.

"By anybody's metrics, [SIP] has been a heavily successful protocol," Rosenberg says. "It has been implemented in hundreds, if not thousands, of products. It is making people money, and it is making people happy…The protocol met the needs of an industry at the time."

However, Jonathan points out that SIP has failed at providing functionality beyond that offered by the Public Switched Telephone Network, doing little to deliver new features. He attributes this failure to the economics of telecommunications carriers and how it takes them years, if not decades, to introduce new functionality.

"If you're a service provider that wants to roll something out, you have to get enough community interest to get something massively marketable enough to get it through, which means it's probably pretty vanilla technology," he said. "Something weird and unusual is not going to make through the curve. That means it's pretty difficult to roll out and deploy innovative new stuff."

Jonathan noted that if end users wanted innovative features with interoperability, then the market would meet that need. Instead, service providers are offering simple, commodity features on slower timelines, and they're putting more energy into developing proprietary features, which are quicker to develop and offer differentiation.

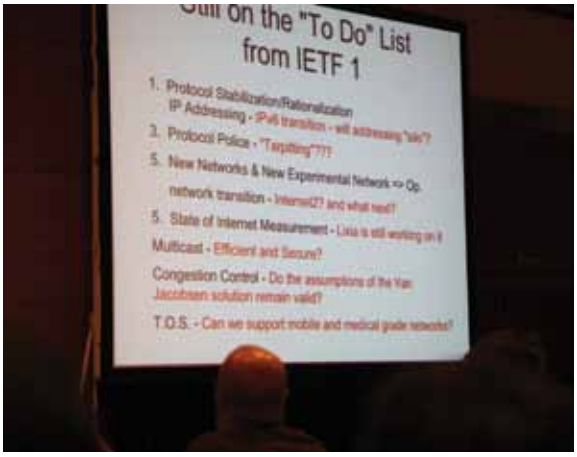

Photo/Peter Löthberg

Outstanding action items from IETF 1!

One trend "is the popularity of service providers publishing REST APIs as a new form of inter-domain interoperability also without standardization," Rosenberg says. "The model there is a service provider…writes some code that sits on a server somewhere, they write some code that gets sent to a client through an app store or client. They deploy it, and then they publish it. Then they are done."

Jonathan added that these service providers "are innovative, they have been moving fast, and they have dramatically shortened the time to market. These guys produce software and deploy services in weeks or months, as opposed to years or decades following the telecom innovation cycle."

In the past, service providers waited until standards were set before deploying new applications. Today, "the need for having inter-provider standards is gone," Rosenberg concluded. "Standards are moving from being first, to being last, if ever."

Leslie Daigle, chief Internet technology officer at the Internet Society, wrapped up the panel with a recommendation that the IETF focus on creating standards for the building blocks that developers can use to create innovative applications and services. As an example, she pointed to HTTP, which is a protocol for linking and sharing content that is serving as a key building block for many of today's proprietary applications.

"We have a bad track record of predicting the future of the Internet," Leslie said. "That is why we specify building blocks and not buildings."


Dave Crocker at the mic at the Internet Society panel at IETF 80 in Prague, Czech Republic


Photo/Peter Löthberg

IAOC members on stage during the IETF 80 plenary


Photo/Peter Löthberg

IETF 80 participants at the open-mic session duing the plenary

Securing BGP and SIDR, continued from page 1

We will look at the approach taken by the SIDR Working Group and examine the architecture and mechanisms that have been adopted as part of this study. This work was undertaken in three stages: the first concentrated on the mechanisms to support attestations relating to addresses and their use; the second looked at how to secure origination of routing announcements; and the third looked at how to secure the transitive part of BGP route propagation.

Supporting Attestations about Addresses through the Resource Public Key Infrastructure (RPKI)

Prior work in the area of securing the Internet's routing system has focused on the operation of the Border Gateway Protocol (BGP) in an effort to secure the operation of the protocol, and validate, as far as is possible, the contents of BGP Update messages. Some notable contributions in more than a decade of study include Secure-BGP (S-BGP) [sBGP], Secure Origin BGP (soBGP) [soBGP], Pretty Secure BGP (psBGP) [psBGP], IRR [IRR], and the use of an AS RR in the DNS, signed by DNSSEC [DNS].

The common factor in this prior work was that they all required, as a primary input, a means of validating basic assertions relating to origination of a route into the inter-domain routing system: that the IP address block and the AS numbers being used are valid and that the parties using these IP addresses and AS numbers in the context of routing advertisement are properly authorized to so do.

The approach adopted by SIDR for the way in which trust is formalized in the routing environment is through the use of Resource Certificates. These certificates are X.509 certificates that conform to the PKIX profile [PKIX]. They also contain an extension field that lists a collection of IP resources (IPv4 addresses, IPv6 addresses and AS Numbers) [RFC3779]. These certificates attest that the certificate's issuer has granted to the certificate's subject a unique "right-of-use" for the associated set of IP resources, by virtue of a resource allocation action. This concept mirrors the resource allocation framework of the IANA (Internet Assigned Numbers Authority), the regional Internet registries (RIRs), operators and others, and the certificate provides a means for a third-party (relying party) to formally validate assertions related to resource allocations [sidr-arch].

The hierarchy of the RPKI is based on the administrative resource allocation hierarchy, where resources are distributed from the IANA to the RIRs, to Local Internet Registries (LIRs) and end users. The RPKI mirrors this allocation hierarchy with certificates that match current resource allocations

The hierarchy of the RPKI is based on the administrative resource allocation hierarchy, where resources are distributed from the IANA to the RIRs, to Local Internet Registries (LIRs) and end users. The RPKI mirrors this allocation hierarchy with certificates that match current resource allocations (Figure 1).

The Certification Authorities (CAs) in this RPKI correspond to entities that have been allocated resources. Those entities are able to sign authorities and attestations, and to do so they use specific purpose End Entity (EE) certificates. This additional level of indirection allows the entity to customize each issued authority for specific subsets of number resources that are administered by this entity. Through the use of single-use EE certificates, the issuer can control the validity of the signed authority through the ability to revoke the EE certificate used to sign the authority. As is often the case, a level of indirection comes in handy.

Signed attestations relating to addresses and their use in routing are generated by selecting a subset of resources that will be the subject of the attestation, by generating an EE certificate that lists these resources, and by specifying validity dates in the EE certificate that correspond to the validity dates of

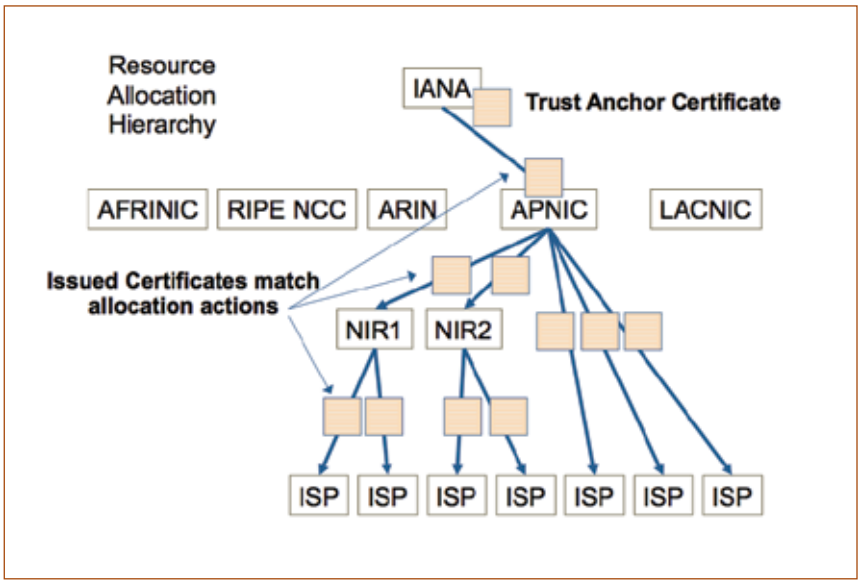

Figure 1: Hierarchy of the RPKI

the authority. The authority is published in the entity's RPKI repository publication point. The RPKI makes conventional use of Certificate Revocation Lists (CRLs) to revoke certificates that have not expired, but which are no longer valid. Every CA in the RPKI regularly issues a CRL according to the CA's declared CRL update cycle. A CA certificate may be revoked by an issuing authority for a number of reasons, including key rollover, the reduction in the resource set associated with the certificate's subject, or termination of the resource allocation. To invalidate an object that can be verified by a given EE certificate, the CA that issued the EE certificate can revoke the corresponding EE certificate.

The RPKI uses a distributed publication framework, wherein each CA publishes its products (including EE certificates, CRLs and signed objects) at a location of its choosing. The set of all such repositories form a complete information space, and it is fundamental to the model of securing BGP in the public Internet that the entire RPKI information space is available to every Relying Party (RP). It is the role of each RP to maintain a local cache of the entire distributed repository collection by regularly synchronizing each element in the local cache against the original repository publication point. To assist RPs in the synchronization task, the each RPKI publication point uses a manifest. A manifest is a signed object that lists the names (and hash values) of all the objects published at that publication point. It is used to assist RPs to ensure that they have managed to synchronize against a complete copy of the material published at the CA's publication point.

The utility of the RPKI lies in its ability to validate digitally signed information and, therefore, give relying parties some confidence in the validity of signed attestations about addresses and their use. The particular utility of the RPKI is not as means of validation of attestations of an individual's identity or their role, but as a means of validating their authority to use IP address resources. While it is possible to digitally sign any digital object, it has been suggested that the RPKI system uses a very small number of standard signed objects that have particular meaning in the context of routing security.

The utility of the RPKI lies in its ability to validate digitally signed information and, therefore, give relying parties some confidence in the validity of signed attestations about addresses and their use.

Securing Route Origination

The approach adopted by SIDR to secure origination of routing information is one that uses a particular signed authority, a Route Origination Authorization (ROA) [ROA]. An ROA is an authority created by a prefix holder that authorizes an AS to originate one or more specific route advertisements into the inter-domain routing system. An ROA is a digital object formatted according to the Cryptographic Message Syntax specification (CMS) [RFC 3852] that contains a list of address prefixes and one AS number. The AS is the specific AS being authorized to originate route advertisements for one or more of the address prefixes in the ROA. The CMS object also includes the EE resource certificate for the key used to verify the ROA. The IP Address extension in this EE certificate must encompass the IP address prefixes listed in the ROA's contents.

The ROA conveys a simple authority. It does not convey any further routing policy information, nor does it convey whether or not the AS holder has even consented to actually announce the prefix(es) into the routing system. The associated EE certificate is used to control the validity of the ROA and the CMS wrapper is used to bind the ROA and the EE certificate within a single signed structure in a secure fashion.

There is one special ROA, one that authorizes AS 0 to originate a route. This is a "negative" authority, used to indicate that no AS has authority to originate a route for the address prefix(es) listed in the ROA.

If the entire routing system were to be populated with ROA's, then identification of an invalid route advertisement would be directly related to detection of an invalid ROA or a missing ROA. However in a more likely scenario of partial use of ROA's (such as when only some legitimate route originations are authorized in a ROA), the absence of an ROA cannot be interpreted simply as an unauthorized use of an address prefix. This leads to the use of a tri-state validation process for routes. If a given route matches exactly the information contained in an ROA whose EE certificate can be validated in the RPKI (a "valid" ROA) then the route can be regarded as a "valid" origination. Where the address prefix matches that in a valid ROA, but the origination AS does not match the AS number in the ROA, and there are no other valid ROAs that explicitly validate the announcing AS, then the route can be considered to be "invalid". Also, where the address prefix is more specific than that of a valid ROA, and there are no other valid ROAs that match the prefix, then the route can also be considered "invalid". Where the prefix in a route is not described in any ROA and is not a more specific prefix of any ROA, then the route has an "unknown" validation outcome. These three potential

Continued on next page

Securing the BGP and SIDR, continued

outcomes can be considered a set of relative local preferences. Routes whose origin can be considered "valid" are generally proposed to be preferred over routes that are unknown, which, in turn, can generally be preferred over routes that are considered invalid. However, such relative preferences are a matter to be determined by local routing policy. Local policies may choose to adopt a stricter policy and, for example, discard routes with an invalid validation outcome [sidr-roa-validation].

The way in which ROAs are used to validate the origin of routes in BGP differs from many previous proposals for securing BGP. In this framework the ROAs are published in the RPKI distributed repository framework. Each RP can use the locally cached collection of valid ROAs to create a validation filter collection, with each element of the set containing an Address, prefix size constraints and an originating AS. It is this filter set, rather than the ROAs themselves that are fed to the local routers [sidr-rpki-router]. (An example of the way in which ROAs can be used to detect prefix hijack attempts is shown in Figure 2).

The model of injecting validation of origination into the BGP domain is an example of a highly modular and piecemeal deployment. There are no changes to the BGP protocol for this origin validation part of the secure routing framework.

The process of securing origination starts with the address holder, who generates local keys and requests certification of their address space from the entity from whom their addresses were allocated or assigned. With this CA resource certificate, the address holder is then in a position to generate an EE certificate and a ROA that assigns an authority for a nominated AS to advertise a route for an address prefix drawn from their address holdings. The one condition here is that if an address holder issues a ROA for an address prefix providing an authority for one AS to originate a route for this prefix, then the address holder is required to issue ROAs for all the AS's that have been similarly authorized to originate a route for this address prefix. The address holder publishes this ROA in their publication point in the distributed RPKI repository structure.

Relying Parties can configure a locally managed cache of the distributed RPKI repository and collect the set of valid ROAs [rcynic]. They can then, via the dedicated RPKI cache-to-router protocol [rpki-rtr], maintain, on a set of "client" routers the set of address prefix/originating AS authorities that are described in valid ROAs. This information can be used by the BGP-speaking router as an input to the local route decision process.

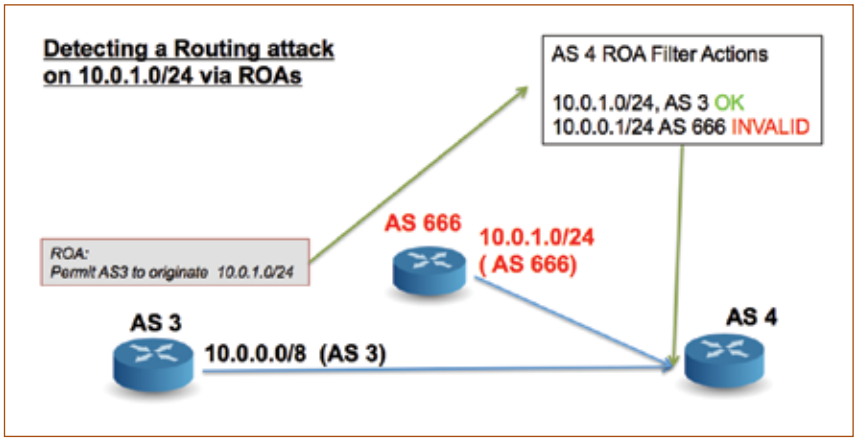

Figure 2: Use of ROAs to detect unauthorized Route Origination

This model of operation supports incremental deployment, wherein individual address holders may issue ROAs to authorized routing advertisements independent of the actions of other address holders. Also, ASs may deploy local validation of route origination independently of the actions of other ASs. And given that there are no changes to the operation of BGP, then there are no complex interdependencies that hinder piecemeal incremental deployment of this particular aspect of securing routing.

Securing Route Propagation—BGPsec

Origin validation as described earlier does not provide cryptographic assurance that the origin AS in a received BGP route was indeed the originating AS of this route. A malicious BGP speaker can synthesize a route as if it came from the authorized AS. Thus, it is very useful in detecting accidental misannouncements, but origination validation does little to prevent malicious routing attacks from a determined attacker.

In looking at the operation of the BGP protocol, some parts of the protocol interaction are strictly local between two BGP-speaking peers, such as advising a peer of local attributes. Another part of the BGP protocol is a "chained" interaction, in which each AS adds information to the protocol object. This attribute of a BGP update, the AS Path, is not only useful to detect and prevent routing loops, it is also used in the BGP best path selection algorithm.

A related routing security question concerns the validity of this "chained" information, namely the AS Path information contained in a route. Within the operation of the BGP protocol, each

AS that propagates an update to its AS neighbours is required to add its AS number to the AS Path sequence. The inference is that at any stage in the propagation of a route through the inter-domain routing system, the AS Path represents a viable AS transit sequence from the local AS to the AS originating the route. This AS Path attribute of a route is used for loop detection. Locally, the AS Path may also be used as an input to a local route policy process, using the length of the AS Path as route metric.

Attacks on the AS path can be used to subvert the routing environment. A malicious BGP speaker may manipulate the AS Path to prevent an AS from accepting a route by adding its AS number to the AS Path, or it may attempt to make a particular route more likely to be selected by a remote AS by stripping out AS's from the AS Path. Accordingly, it is important to equip a secure BGP framework with the ability to validate the authenticity of the AS Path presented in a BGP update [kapela/pilosov].

In attempting to validate an AS path there are a number of potential validation questions. The first and weakest question is, Are all ASs in the AS Path valid ASs? A slightly stronger validation question is, Do all the AS pairs in the AS Path represent valid AS adjacencies (where both ASs in the pairwise association are willing to attest to their mutual adjacency in BGP)? A even stronger question is, Do the sequence of ASs in the AS Path represent the actual propagation path of the BGP route object? This last question forms the basis for the SIDR activity in defining an AS Path validation framework, BGPsec. This is an attempt to assure a BGP speaker that the operation of the BGP protocol is operating correctly and that the content of a BGP update correctly represents the inter-AS propagation path of the update from the point of origination to the receiver of the route. This is not the same as a policy validation tool and it does not necessarily assure the receiver of the route that this update conforms to the routing policies of neighbouring BGP speakers. This route also does not necessarily reflect the policy intent of the originator of the route. The BGPsec framework proposed for securing the AS Path also makes use of a local RPKI cache, but it includes an additional element of certification. The additional element of the security credentials used here is an extension to the certification of AS numbers with a set of operational keys and their associated certificates used for signing update messages on eBGP routers in the AS. These "router certificates" can sign BGP update attributes in the routing infrastructure, and the signature can be interpreted as being a signature made "in the name of" an AS number.

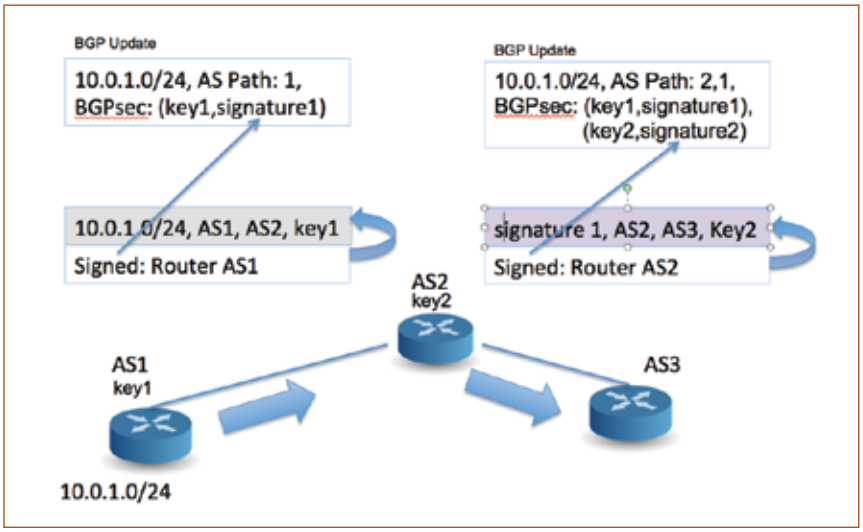

Figure 3: BGPsec AS Path Protection

In the BGPsec framework, eBGP speaking routers within the AS have the ability to "sign" a BGP update before sending it. In this case, the added signature "covers" the signature of the received BGP update, the local AS number, the AS number to which the update is being sent, as well as a hash of the public key part of the router's key pair used to sign route updates. The couplet of the public key hash and the signature itself is added to the BGP protocol update as a BGPsec update attribute. As the update traverses a sequence of transit ASes each eBGP speaker at the egress of each AS adds its own public key hash and digital signature to the BGPsec attribute sequence (Figure 3).

This interlocking of signatures allows a receiver of a BGP update to use the interlocking chain of digital signatures to validate (for each AS in the AS Path) that the corresponding signature was correctly generated "in the name of" that AS in the AS path, and that the next AS in the path matches the next AS in the signed material. The "forward signing" that includes the AS to which the update is being sent prevents a man-in-the-middle attack of the form of taking a legitimate outbound route announcement destined for one neighbour AS and redirecting it to another AS. But this signing of the AS Path is not quite enough to secure the route update, as the AS Path needs to be coupled to the actual address prefix by the originator of the route. The route originator needs to sign across not only the local AS and the AS to whom the route update is being sent, but also the address prefix and the expiry time of the route. This allows the path to be "bound" to the prefix and prevents a man-in-the-middle splicing a signed path or signed path fragment against a different prefix.

Continued on next page

Securing the BGP and SIDR, continued

If the signatures that "span" the AS Path in the BGP update can all be validated, then the receiver of the BGP update can validate, in a cryptographic sense, the currency of the routing update. It can also validate that the route update was propagated across the inter-AS routing space in a manner that is faithfully represented in the AS Path of the route.

The expiry time of the EE certificates used in conjunction of signed route updates introduces a new behaviour into BGPsec. In the context of BGP, an announced route remains current until it is explicitly withdrawn or until the peer session that announced the route goes down. This property of BGP introduces the possibility of "ghost route" attacks in BGP, wherein a BGP speaker fails to propagate a withdrawal in order to divert the consequent misdirected traffic from its peers. In BGPsec, all route advertisements are given an expiry time by the originator of the route that corresponds to the notAfter time of the EE certificate used to sign the protocol update, after which the route is to be considered invalid. The implication is that an originator of a route is required to re-advertise the route, and refresh the implicit expiry timer of the associated digital signature at regular intervals.

This approach to route update validation is not quite the "light-touch" of origination validation. In this case the mechanism requires the use of a new BGP attribute and negotiation of a new BGP capability between eBGP peers. In turn, this means that the model of incremental deployment is one that is more "viral" than truly piecemeal. By "viral" we mean that this is a model of incremental deployment in which direct eBGP peers of a BGPsec-speaking AS will be able speak BGPsec between themselves in a meaningful way. In turn these adjacent AS's can offer to speak BGPsec with their eBGP peers, and so on. This does not imply that BGPsec deployment must necessarily start from a single AS, but it does imply that communities of interconnected AS's all speaking BGPsec will be able to provide assurance via BGPsec on those routes originated and propagated within that community of interconnected ASs. It also implies that the greatest level of benefit to adopters of secure BGP will be realized by ASs that adopt BGPsec as a connected community of ASs.

There are other changes to the behaviour of BGP that are implied by this mechanism. BGP conventionally permits "update packing" where a number of address prefixes can be placed in a single update message if they share a common collection of attributes, including the AS Path. At this stage it appears that such update packing would not be supported in secure BGP, and each update in secure BGP would refer to a single prefix. Obviously this would have some impact on the level of BGP traffic, but early experiments suggest not at an unreasonable cost.

BGP conventionally permits "update packing" where a number of address prefixes can be placed in a single update message if they share a common collection of attributes, including the AS Path. At this stage it appears that such update packing would not be supported in secure BGP, and each update in secure BGP would refer to a single prefix.

There are further impacts on BGP that have not been fully quantified in studies to date. The addition of a compound attribute of a signature and a public key identifier for every AS in the AS Path has size implications on the amount of local storage a secure BGP speaker will need to store these additional per-prefix per-peer attributes. It has has broader implications if used in conjunction with current proposals for multi-path BGP where multiple paths, in addition to the "best" path are propagated to eBGP peers. Also, the computational load of validation of signatures in secure BGP is significantly higher in terms of the number of cryptographic operations that are required to validate a BGP update.

However, BGPsec is not intended to "tunnel" across those parts of the inter-domain routing space that do not support BGPsec capabilities. When an update leaves a BGPsec realm, the BGPsec signature attributes of the route are stripped out, so the storage overheads of BGPsec are not seen by other BGP speakers. Similarly, the periodic updates that result from the expiry timer should not propagate beyond the BGPsec realm. If the boundary is prepared to perform BGP update packing to non-BGPsec peers then even the unpacked update overhead is not carried outside of the BGPsec realm.

It is also noted that the "full" load of BGPsec would only necessarily be carried by "transit" ASs; that is, those ASs that propagate routes on behalf of other ASs. Historically we see some 15 percent of ASs are "transit" ASes, while all other ASes behave as "stub" ASes that only originate routes and do not appear to transit routes for others. Such stub ASes can support a "light weight" simplex version of BGPsec that can either point default a default route to its upstream AS provider, or trust its upstream ASs to perform BGPsec validation. In this case the stub AS needs to provide BGPsec signed originated routes to its upstream ASs, but no more.

Conclusion

The work on the specification of the RPKI itself and the specification of origin validation is nearing a point of logical completion of the first phase of standardization within the IETF, and the working draft documents are being passed from the working group into the review process leading to their publication as proposed standard RFCs. The RIRs are in the process of launching their RPKI services based on these specifications and the initial deployment of working code has been made by a number of parties, who are also working on integration of origination validation in BGP implementations.

The work on securing the AS Path is at an earlier phase in the development process and the initial design material is being considered by the SIDR Working Group. It is expected to take a similar path of further review and refinement in light of developing experience and study of the proposed approach.

There are further impacts on BGP that have not been fully quantified in studies to date. The addition of a compound attribute of a signature and a public key identifier for every AS in the AS Path has size implications on the amount of local storage a secure BGP speaker will need to store these additional per-prefix per-peer attributes.

The RPKI has been designed as a robust and simple framework. As far as possible, existing standards, technologies, and processes have been exploited, reflecting the conservatism of the routing community and the difficulty in securing rapid, widespread adoption of novel technologies.

Acknowledgements

The work described here is the outcome of the efforts of many individuals who have contributed to securing BGP over a period that now spans two decades, and certainly too many to ensure that all the contributors are recognized here. Instead, the authors would like to acknowledge their work and trust that the mechanisms described here are a faithful representation of the cumulative sum of their various contributions.

References

[sBGP] S. Kent, C. Lynn, and K. Seo, “Secure Border Gateway Protocol (S-BGP),” *IEEE Journal on Selected Areas in Communications,* vol. 18, no. 4, pp 582-592, April 2000.

[soBGP] R. White, “Securing BGP through secure origin BGP,” *Internet Protocol Journal,* vol. 6, no. 3, September 2003.

[psBGP] P. van Oorschot, T. Wan and E. Kranakis, “On Interdomain Routing Security and Pretty Secure BGP (psBGP),” *ACM Transactions on Information and System Security,* vol. 10, no. 3, July 2007.

[IRR] G. Goodell, W. Aiello, T. Griffin, J. Ioannidis and P. McDaniel, “Working Around BGP: An Incremental Approach to Improving Security and Accuracy of Interdomain Routing,” *Proc. of Internet Society Symposium on Network and Distributed System Security* (NDSS, 2003), February 2003.

[DNS] T. Bates, R. Bush, T. Li and Y. Rekhter, “DNS-based NLRI origin AS verification in BGP,” Internet Draft, July 1998.

[PKIX] D. Cooper et al., “Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile,” Request for Comment RFC5280, May 2008.

[RFC3779] C. Lynn, S. Kent and K. Seo, “X.509 Extensions for IP Addresses and AS Identifiers,” Request for Comment RFC3779, June 2004.

[sidr-arch] M. Lepinski, S. Kent, “An Infrastructure to Support Secure Internet Routing,” work in progress (Internet Draft), February 2008.

[sidr-cert-profile] G. Huston, G. Michaelson, R. Loomans, “A Profile for X.509 PKIX Resource Certificates,” work in progress (Internet Draft), September 2008.

[ROA] M. Lepinski, S. Kent, D. Kong, “A Profile for Route Origin Authorizations (ROAs),” work in progress (Internet Draft), July 2008.

[RFC3852] R. Housley, “Cryptographic Message Syntax (CMS),” Request for Comment RFC3852, July 2004.

[sidr-rpki-router] R. Bush, R. Austein, “The RPKI/Router Protocol”, work in progress (Internet Draft), March 2011.

[kapela/pilosov] http://wwwwww.wired.com/threatlevel/2008/08/revealed-the-in/, August 2008.

Bufferbloat: Dark Buffers in the Internet

By Jim Gettys

We have conflated "speed" with "bandwidth." As Stuart Cheshire wrote in "It's the Latency, Stupid" (http://rescomp.stanford.edu/~cheshire/rants/Latency.html), "Making more bandwidth is easy. Once you have bad latency, you're stuck with it." Bufferbloat is the existence of excessively large (bloated) buffers in systems, particularly network communication systems.

Bufferbloat is now (almost?) everywhere.

Today's routers, switches, gateways, broadband gear, and so on have bloated buffer sizes to where we often measure latency in seconds, rather than microseconds or milliseconds.

Telephone standards for maximum desirable latencies are in the 150–200 ms range, and human perception for some latency is as low as 10 ms. You can never get that time back. Any unnecessary latency beyond the minimum imposed by the speed of light is too much.

Although some buffering is required to smooth bursts in communications systems, we've lost sight of fundamentals: packet loss is (currently) the only way to signal congestion in the network, and congestion-avoiding protocols such as TCP rely on timely congestion notification to regulate their transmission speeds.

What happens when we put large or truly bloated buffers into our systems, in a misguided attempt to avoid all packet loss, or when we aim to eke out almost unmeasurable increases in performance on an artificial benchmark, or just because the buffer memory doesn't cost us anything and happens to be there?

Really bad things happen (see http://gettys.wordpress.com/2010/12/06/whose-house-is-of-glasse-must-not-throw-stones-at-another/), as John Nagle's cogent explanation, RFC 970 (from 1985!), describes:

A datagram network with infinite storage, first-infirst- out queuing, and a finite packet lifetime will, under overload, drop all packets.

Some of the buffers we now observe in the Internet are effectively infinite in size. More is not necessarily better. More is often worse.

Not all packet loss is evil: some packet loss can be essential for correct operation. But once your bloated buffers fill, there's no timely congestion notification by packet loss or explicit congestion notification (ECN), and eliding notification has destroyed the congestion avoidance servo loop in transport protocols. Only the buffers on either side of the bottleneck (lowest bandwidth) link fill, and if those buffers are not managed, they can and do fill completely, inducing much higher packet loss than that you attempted to avoid. Other buffers in the path, remaining nearly empty, remain dark and undetectable.

Bufferbloat induces painful latencies for you and all others who share your network path. Any application that saturates a link with bloated buffers can induce bufferbloat pain: uploading videos to YouTube, emailing messages with large images attached, backing up large files or file systems, downloading large files, such as ISO images, a Linux distribution image, or a movie via Bittorrent, watching Netflix, and even visiting certain kinds of webpages can all fill these buffers.

Any semblance of interactivity of your network is gone; any hope for good teleconferencing or voice over IP, or fragging your opponent before they frag you, is lost. Even Web browsing becomes painful, and applications often fail entirely. Wonder no more why your network connections are so poor. This is why the "Internet is slow today" refrain is so common.

With modern TCP stacks (almost everything except Windows XP), even a single TCP transfer can induce bufferbloat suffering. The problem is not limited to TCP; UDP-based protocols are equally capable of filling bloated buffers. But you never see these buffers until they start to fill, and you can observe them only indirectly; they are "dark buffers" — like dark matter in the universe.

Last year the ICSI Netalyzr group proved that our broadband Internet technologies—cable, DSL, and FIOS alike—suffer badly from bufferbloat. And bufferbloat isn't confined to these technologies, but has also infected our home routers, 3G networks, and even our operating systems (see http://conferences.sigcomm.org/imc/2010/papers/p246.pdf).

For example, Linux typically has at least two major contributors to bufferbloat—the network stack's transmit queue and the ring buffers in the device driver—but several more places exist where buffers can hide. These buffers are often hundreds of packets in size on modern hardware, and we can find such large transmit rings on similar hardware on other operating systems. They might lurk in line cards in network gear, in modems, or elsewhere —the hiding places are endless.

Bufferbloat also infects many of our corporate and ISP networks. I believed

we had solved congestion problems for Internet routers with the development of active queue management (AQM) algorithms such as random early detection (RED; see www. icir.org/floyd/papers/early.pdf), but the cottage industry of more than 100 papers on tuning RED proves this belief incorrect. Classic RED 93 can't solve our wireless problems. Although failing to use AQM when possible might be misguided, all those RED tuning papers help us understand why some network operators (both corporate and public) do not trust RED and are reluctant to enable it.

Won't adding more bandwidth help? Usually not. Buffering has been growing, frequently faster than bandwidth, over generations of often upward-compatible technologies. Plug a current device into a previous generation link, and your buffers become insanely large, even in the rare case those buffers were static sized "correctly." They will then be sized for maximum theoretical bandwidth over maximum latency paths, often much larger than you will ever experience. Yet the actual bandwidth available varies, often by orders of magnitude. This demonstrates that a single static answer seldom exists regarding the correct buffer size in any system.

Adding bandwidth can even make your suffering worse: for example, if you have more broadband bandwidth than 802.11 bandwidth, the bottleneck shifts to that hop, where your laptop and home router bufferbloat is often even worse than in the broadband link. Now those "dark buffers" cause your pain.

We must systematically stamp out bufferbloat wherever it occurs in our systems by managing buffers at all times wherever they appear. We can mitigate the worst bufferbloat by eliminating the grossly mis-sized static buffers, but actual solutions require serious work, further research, and the use of some form of AQM, in its most general sense.

We're also rapidly destroying TCP slow start, with independent changes by both Web server and Web browser. I even fear for the Internet's stability.

We have a large mess on our hands that spans hardware, software, firmware, operating systems, home routers, broadband, 3G, 802.11, and just about everywhere I have looked. We're all in this bloat we built together, and had better work together to row to shore quickly. Dark storm clouds surround us. Only together will we shine a light on all the dark buffers hidden in the Internet. I've been drawing together experts across all these problems at bufferbloat.net—please help out.

Internet Society Fellows Bring First-Hand Experience of Standards Processes to Developing Countries

Four information technology professionals from Asia, Europe, and South America attended their first IETF meeting in March 2011 as part of the Internet Society's Fellowship to the IETF Programme. The programme, which operates under the aegis of the Internet Society's Next Generation Leaders Programme, enables Internet technologists from developing regions to participate more fully in the IETF's standards work by facilitating their attendance at an IETF meeting.

IETF 80 First-time Fellows

Roman Arcea (Moldova)
Mentor: Jean-Michel Combes
Paventhan Arumugam (India)
Mentor: Marcelo Bagnulo Braun
Luis Balbinot (Brazil)
Mentor: Juergen Quittek
Suhaidi Hassan (Malaysia)
Mentor: Fernando Gont

IETF 80 Returning Fellows

Zartash Afzal Uzmi (Pakistan)
Palanivelan Appanasamy (India)
Azael Fernandez Alcantara (Mexico)
Fernando Gont (Argentina)
Idris A. Rai (Tanzania, United Republic of)

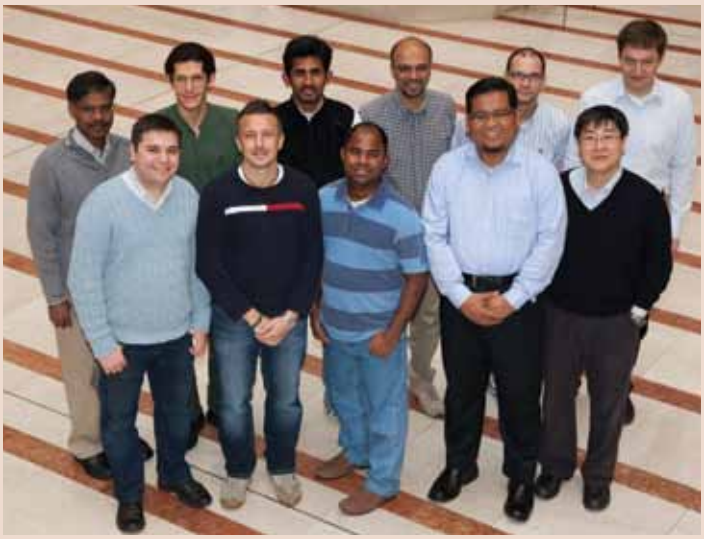

Internet Society Fellows and Returning Fellows at IETF 80 in Prague, Czech Republic

IETF Ornithology: Recent Sightings

Compiled by Mat Ford, with inputs from Harald Alvestrand, Francois Le Faucheur, Alexey Melnikov, Brian Rosen and Tim Chown.

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 new regular feature of the *IETF Journal*, we will review the BoFs that took place during the most recent IETF meetings and summarize their intentions and outcomes.

> BoF meetings have a very different tone than do [working group] WG meetings. The purpose of a BoF is to make sure that a good charter with good milestones can be created and that there are enough people willing to do the work needed in order to create standards. Some BoFs have Internet-Drafts already in process, whereas others start from scratch.
>
> An advantage of having a draft before the BoF is to help focus the discussion. On the other hand, having a draft might tend to limit what the other folks in the BoF want to do in the charter. It's important to remember that most BoFs are held in order to get support for an eventual working group, not to get support for a particular document.
>
> Many BoFs don't turn into WGs for a variety of reasons.
>
> *The Tao of IETF (http://www.ietf.org/tao.html)*

If you are 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 took place during IETF 80 in Prague, the Czech Republic, can be found on the wiki at http://trac.tools.ietf.org/bof/trac/wiki/WikiStart.

RTCWEB - Real Time Communication on the World Wide Web

Description: Many implementations have been made that use a Web browser to support direct, interactive communications, including voice, video, collaboration, and gaming. In these implementations, the Web server acts as the signalling path between these applications, using locally significant identifiers to set up the association. Up until now, such applications have typically required the installation of plugins or nonstandard browser extensions. There is a desire to standardize this functionality so that these types of applications can be run in any compatible browser and allow for high-quality real-time communications experiences within the browser.

Minutes: http://www.ietf.org/proceedings/80/minutes/rtcweb.txt

Outcome: This effort, spearheaded by Google, but with coinitiative takers including Cisco, Ericsson, Mozilla, and Skype, aims to establish standards for enabling browsers to send audio and video

between each other without the need for media intermediaries, which would make interactive video a commonplace rather than a specialized undertaking.

The meeting was very well attended (more than 250 attendees) and showed both strong support for letting this go forward and quite good consensus that we had achieved "roughly the right level" in the proposed charter. The Friday work session laid out some of the challenges ahead, including thorny issues such as what happens when intellectual property rights claims meet "mandatory to implement."

The work will go forward in cooperation with the W3C.

CDNI—Content Distribution Network Interconnection

Description: There is an emerging requirement for interconnecting content delivery networks (CDNs) so they can interoperate as an open content delivery infrastructure for the end-to-end delivery of content from Content Service Providers (CSPs) to end users. This BoF is an opportunity to discuss the proposed development of IETF standards to facilitate such CDN interconnection. These standards might include protocols for the following.

There is an emerging requirement for interconnecting content delivery networks (CDNs) so they can interoperate as an open content delivery infrastructure for the end-to-end delivery of content from Content Service Providers (CSPs) to end users. This BoF is an opportunity to discuss the proposed development of IETF standards to facilitate such CDN interconnection.

Exchange of metadata between CDNs

Exchange of transaction logs and monitoring information

Exchange of request-routing information

Exchange of policies & capabilities

Content management/flushing

Minutes: http://www.ietf.org/proceedings/80/minutes/cdni.htm

Outcome: Approximately 120 people attended the CDNI BoF. The first presentations established the multiple-use cases that network service providers have for CDN interconnection (such as footprint extension, off-net delivery, offload and fail-over, multivendor, and over the top content providers) and introduced the key missing touch-points needed across CDNs.

The CDN interconnection experiments presented next confirmed the feasibility of content delivery across two CDNs while bringing to light the limitations resulting from the lack of a standard approach. A poll of the audience revealed a strong agreement that there was a real problem that needs solving.

The next few presentations articulated the overall CDNI architecture as well as the functional role and requirements for each of the four inter-CDN interfaces that would be in the scope of the CDNI WG (CDNI Request-Routing, CDNI Distribution Metadata Exchange, CDNI Logging, and CDNI Control). The resulting discussion concluded that the CDNI WG would need to document the threat analysis, including discussion of the CDNI trust model and privacy issues.

The next discussion centred on the question of whether CDNI required new protocol or protocol extensions. The key outcome is that the CDNI work is expected to involve definition of new schemas (such as the definition of a Web services-style of interface), but no new schema languages and no new protocol or protocol extensions.

A show of hands indicated consensus that the problem was well understood and that the IETF was the right place to solve it. After review of the draft charter and some real-time tweaks, there was consensus in the room that the draft charter identifies the right set of deliverables and that a WG

Continued on next page

IETF Ornithology, continued

should be created with that charter. Many individuals volunteered to be editors or reviewers of the CDNI deliverables.

The charter has been refined and the the IESG recently announced formation of the CDNI working group in the Transport Area. The CDNI WG will hold its first meeting during IETF 81.

Plasma–Policy Augmented S/MIME

Description: Current S/MIME mechanisms provide cryptographic access to a message based on the identity of the recipient at the time of transmission. Any additional access-control enforcement depends on the configuration of the recipient's email client. Several Internet-Drafts have been submitted that establish a more robust access-control mechanism, where cryptographic access to a message is only granted after the access check.

Current S/MIME mechanisms provide cryptographic access to a message based on the identity of the recipient at the time of transmission. Any additional access-control enforcement depends on the configuration of the recipient's email client. Several Internet-Drafts have been submitted that establish a more robust access-control mechanism, where cryptographic access to a message is only granted after the access check.

This proposed WG would develop a framework for enforcing a more robust access-control mechanism, based on existing CMS, S/MIME, and SAML-based policy-enforcement standards. The WG will also develop any necessary extensions to these base protocols specific to this problem space.

Minutes: http://www.ietf.org/proceedings/80/minutes/plasma.txt

Outcomes: Despite some original concerns that the work is only of interest to a limited group of people, the meeting was well attended and there was active participation in the discussion of related use cases. Some interest in using a ABFAB-capable solution together with the policy server was expressed.

One area director questioned why sending secure email can't be done using only Web protocols (without SMTP at all), eliminating the need for S/MIME. Proponents answered that users still like to use email for many tasks, so building upon/fixing existing secure email is a desired goal.

At the end of the meeting several participants (in addition to the BoF proponents) expressed their desire to work on something in this space. Several people were also interested in use of the proposed architecture for non-email use cases (such as with XMPP or for website access controls). As there was no strong consensus that a WG should be formed, the proposed charter was not discussed. This will be done on the mailing list at an appropriate time.

Some questions were raised about whether the proposal presented was the best design and whether something else (which can turn out to be more generic) can be used. The discussion on the mailing list is ongoing. It is currently unclear whether this proposal will go on to form a WG.

PAWS–Protocol to Access White Space Database

Description: As nations struggle to provide radio spectrum for use with wireless Internet bandwidth, there are problems with incumbents in a given band. The incumbents may use the spectrum inefficiently, especially geographically; for instance, there may be large regions where a particular band is not used at all and others where use is only in a portion of the band.

Recently, techniques have been developed that attempt to share spectrum between incumbents and new users. The generic term for this is "white space." For example, in over-the-air TV bands, spectrum is divided into channels. In any one area it is possible that not all channels have TV transmitters in range. There is a desire among many regulators to make this prime spectrum available for

Internet access and other uses, as long as the new use does not interfere with existing TV band use.

The U.S. Federal Communications Commission has freed up its digital TV spectrum for unlicensed use and has crafted rules and regulations that require compliance. The available spectrum and associated channels for use vary on a regional basis. In the United States, as in many other countries, there are other incumbents besides the television broadcasters. Spectrum and channel availability is dynamic while spectrum use requires verification of availability prior to use (and periodically on an ongoing basis). The U.S. regulator has decided that all new users of the TV band white space ("TV Band Device" or TVBD) must query a database with the location of the TVBD and receive from the database a list of available channels that it may use.

Multiple databases containing the information about available channels for use at a given location are expected to exist. A device is required to query the database for available channels and associated information. There are several scenarios that the U.S. regulation permits, which include a simple tower/client arrangement where the tower queries the database on behalf of itself and its client TVBDs and ad-hoc mobile networks where at least one TVBD in the ad-hoc network has another path to the Internet that can query the database.

Minutes: http://www.ietf.org/proceedings/80/minutes/paws.txt

Outcome: The BoF went very well. After everyone was brought up to the same level on what White Space is, how it works, and what the group wanted an IETF WG to do, participants engaged in a number of lively discussions on privacy and security issues and how the group may be able to achieve its goals of supporting database queries that are independent of physical layer, spectrum band, or nation. Several suggestions were made for charter language improvement in these areas.

A number of people in the room indicated that they would contribute or review contributions. It is common at this stage to have little opposition to creating a WG but only a relatively small number of people willing to do work. This BoF attracted roughly two to three times the usual number of people indicating they would be active participants.

The IAB and IESG have been discussing the charter. It does look like the group will have approval to form a Working Group, albeit with some further changes to the scope and, specifically, in the security and privacy text on the proposed charter. An interesting aspect of this work is that the participants did not suggest a specific area, as the work does not seem to fit neatly into any of the area definitions particularly well but, rather, it touches on a number of them. In the end, it seems like we it be in the Apps area.

Renum–Site Renumbering

Description: As outlined in RFC 5887, renumbering, especially for medium to large sites and networks, is viewed as an expensive, painful, and error-prone process, often avoided by network managers as much as possible. Some would argue that the very design of IP addressing and routing makes automated renumbering intrinsically impossible. In fact, managers have an economic incentive

to avoid renumbering, with many resorting to private addressing and Network Address Translation (NAT). Consequently, mechanisms for managing the scaling problems of wide-area (BGP4) routing that require site renumbering are often dismissed as unacceptable. Even so, renumbering is sometimes unavoidable.

The task of the proposed RENUM Working Group is to (1) identify specific renumbering problems in the context of site-wide renumbering and (2) develop point solutions and system solutions to address those problems (or, if appropriate, to stimulate such development in other Working Groups). The principal target will be solutions for IPv6, but solutions that apply equally to IPv4 may be considered.

Minutes: http://www.ietf.org/proceedings/80/minutes/renum.txt

Outcomes: While there was an encouraging level of interest expressed in the BoF in terms of people willing to work on this topic, it was also recognized (quite rightly) that renumbering is a huge area and that trying to solve all renumbering issues in one go is unrealistic. The discussions in the BoF raised a number of ways that the work could be more focused. Therefore, since the meeting, the BoF has been working on scoping the potential WG charter to allow achievable, but useful, goals to be defined.

There is significant concern over the potential increase in the size of backbone routing tables should large numbers of end sites adopt IPv6 PI addressing. Thus, focusing the work on (managed) IPv6 enterprise networks and (unmanaged) SOHO networks is one avenue to scope the work more realistically. As IPv6 becomes more widely deployed, it is expected that we would have developed appropriate new operational and protocol elements so that site renumbering will be seen as a more routine event.

A first step for a new WG could thus be undertaking scenario-based descriptions, including documentation of current capabilities and best practices. These texts can then contribute to a gap analysis that also draws on existing published work in RFC 4192 and RFC 5887. In parallel, a review of current IP address-management practices would be performed to determine if a more appropriate model for supporting site renumbering might be devised. This initial programme of work would not lead immediately to new protocols or practices, but a subsequent WG recharter would identify how and where such work could be done.

Implicitly, such a focus would mean a 'RENUM6' WG would not cover renumbering avoidance methods, ISP renumbering (except the ISP's role in site renumbering), or IPv4 renumbering. he IPv6 Site Renumbering (6renum) working group has been formed in the Operations and Management Area and the full charter is at http://datatracker.ietf.org/wg/6renum/charter/.

IETF 80 At–A–Glance

Registered attendees: 1196

Newcomers: 173

Number of countries: 49

New WGs: 11

WGs closed: 14

WG currently chartered: 121

New Internet-Drafts: 608

- 191 updated
- 54 updated more than once

Updated Internet-Drafts: 1237

IETF Last Calls: 151

Internet-Drafts approved for publication: 164

The Internet Society's Leslie Daigle moderates the Internet Society panel discssion at IETF 80

RFC Editor Actions (November 2010–February 2011)

RFCs published: 74 (about 1940 pages)

I-Ds submitted for publications

- 81 IETF WGs
- 32 IETF individuals
- 9 IRTF, IAB, and independent combined

Zartash Afzal Uzmi, ISOC Fellowship to the IETF Fellow

Marcelo Bagnulo Braun, ISOC Fellowship to the IETF Fellow Mentor

IESG Observation

RFC Ediror Average Processing Time

- Edit: 3.5 weeks
- RFC Edit: 1.9 weeks
- Auth48: 2.0 weeks

IANA Activity since IETF 79 (November 2010-February 2011)

Processed 1468 IETF-related requests, including:

680 private enterprise number requests

- 79 port number requests
- 47 TRIP ITAD number requests
- 15 language subtag requests
- 46 media-type requests

Reviewed 136 I-Ds in Lat Call and reviewed 134 I-Ds in IESG Evaluation

Reviewed 117 I-Ds prior to becoming RFCs and 57 of them contained actions for IANA

Suhaidi Hassan, ISOC Fellowship to the IETF Fellow

IRTF Update

By Lars Eggert

On behalf of the Internet Reseach Task Force (IRTF), I am pleased to report that of the 24 nominations received for the first award of the Applied Networking Research Prize (ANRP), two submissions were awarded the Applied Networking Research Prize. Each submission was reviewed by two to four members of the selection committee, according to a diverse set of criteria, including scientific excellence and substance, timeliness, relevance, and potential impact on the Internet.

The awardees are Mattia Rossi for his research into reducing BGP traffic[1] and Beichuan Zhang[2] for his research into "green" traffic engineering.

Both researchers have been invited to present their findings at the IRTF Open Meeting, to be held during IETF 81, 24–29 July 2011, in Quebec City, Canada.

Future Calls for Nominations

A call for nominations will be issued this summer targeting IETF 82 in Taipei, Taiwan, in November 2011. For 2012, we are considering switching to a yearly nomination/award cycle, with a call for nominations in late 2011 that will cover all three IETF meetings planned for 2012.

Those interested in receiving future calls for ANRP nominations should subscribe to the IRTF-Announce mailing list. You are also encouraged to join the Internet Society (http://www.Internet-Society.org/join) to stay informed of other networking research initiatives.

About the ANRP

The ANRP is awarded for recent results in applied networking research that are relevant for transitioning into shipping Internet products and related standardization efforts. Researchers with relevant, recently published results are encouraged to apply for this prize, which offers the opportunity to present and discuss their work with the engineers, network operators, policymakers, and scientists who participate in the IETF and the IRTF. Third-party nominations for this prize are also encouraged. The goal of the ANRP is to recognize the best new ideas in networking and bring them to the IETF and IRTF, especially in cases where they would not otherwise see much exposure or discussion.

The Applied Networking Research Prize (ANRP) consists of the following:

- a cash prize of USD 500
- an invitation to speak at the IRTF Open Meeting
- a travel grant to attend the week-long IETF meeting that covers airfare, hotel, registration, and stipend
- recognition at the IETF plenary
- an invitation to related social activities
- the potential for additional travel grants to future IETF meetings, based on community feedback

The ANRP will be awarded three times per year, in conjunction with the three annual IETF meetings.

Applicants must nominate a peer-reviewed, recently-published, original journal, conference or workshop paper. Both self nominations (nominating one's own paper) and third-party nominations

1. Geoff Huston, Mattia Rossi, and Grenville Armitage. A Technique for Reducing BGP Update Announcements through Path Exploration Damping. *IEEE Journal on Selected Areas in Communications (JSAC),* Vol. 28, No. 8, pp. 1271–1286, October 2010
2. Mingui Zhang, Cheng Yi, Bin Liu, and Beichuan Zhang. GreenTE: Power-Aware Traffic Engineering. *Proc. IEEE International Conference on Network Protocols (ICNP),* pp. 21–30, October 2010

(nominating someone else's paper) are encouraged. The nominee must be one of the main authors of the nominated paper.

The nominated paper should provide a scientific foundation for possible future IETF engineering work or IRTF experimentation, analyze the behavior of Internet protocols in operational deployments or realistic testbeds, make an important contribution to the understanding of Internet scalability, performance, reliability, security or capability, or otherwise be of relevance to ongoing or future IETF or IRTF activities.

Applicants must briefly describe how the nominated paper relates to these goals, and are encouraged to describe how presentation of these research results will foster their transition into new IETF engineering or IRTF experimentation, or otherwise seed new activities that will have an impact on the real-world Internet.

The goal of the ANRP is to foster the transitioning of research results into real-world benefits for the Internet. Therefore, applicants must indicate that they (or the nominee, in case of third-party applications) are available to attend the respective IETF meeting in person and in its entirety.

Nominations for the ANRP are not considered to be contributions to the IETF. However, the invited talks at the IRTF Open Meeting are considered to be contributions to the IETF, and the IETF "Note Well" does apply to them.

Selection Process

A small selection committee composed of individuals who are knowledgeable about the IRTF, IETF, and the broader networking research community will evaluate the submissions against the selection criteria. The goal is to select one or two submissions for the ANRP during each application period. All applicants will be notified by email.

The current ANRP selection committee includes:

- Mark Allman, ICIR
- Lou Berger, LabN
- Ross Callon, Juniper
- Lars Eggert, Nokia Research Center
- Olivier Festor, INRIA
- Mat Ford, the Internet Society
- Andrei Gurtov, HIIT
- Al Morton, AT&T Laboratories
- Bruce Nordman, LBL
- Jörg Ott, Aalto University
- Stefano Previdi, Cisco
- Martin Stiemerling, NEC Laboratories

The ANRP is supported by the Internet Society as part of its Internet Development and Evolution Awards Programme, in coordination with the IRTF.

IETF Meeting Calendar

IETF 81
24-29 July 2011
Host: Research in Motion (RIM)
Location: Quebec City, CA

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

For more information about past and upcoming

IETF Meetings

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

Special thanks to

for hosting IETF 80

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 80
Volume 7, Issue 1
July 2011

Published three times a year by the Internet Society

Galerie Jean-Malbuisson 15
1204 Geneva
Switzerland

Editor
Mat Ford

Associate Editor
Wendy Rickard

Contributing Writer
Carolyn Marsan

Editorial and Design
The Rickard Group, Inc.

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