IETF Journal

Volume 4, Issue 3
February 2009

*A report from IETF 73, November 2008, Minneapolis, Minnesota. Published by the Internet Society in cooperation with the Internet Engineering Task Force**

Inside this issue

The Past Meets the Present at IETF 73

From the Editor's Desk, by Mirjam Kühne

Since it was first published some 30 years ago, the IP model has emerged as one of the most influential technological developments of our time. As it turns out, its evolution is also one of the more interesting stories in the history of the Internet. At IETF 73, Internet Architecture Board member Dave Thaler gave a well-received presentation called Evolution of the IP Model, which was based on an Internet-Draft published last November. An article based on the presentation, which Dave adapted specifically for the *IETF Journal*, appears on this page.

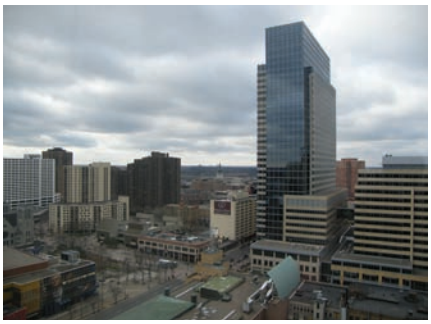

Photo by Peter Löthberg

View of Minneapolis, site of IETF 73

Jon Postel was remembered throughout the week, most notably at a private dinner in Minneapolis commemorating the 10-year passing of Jon and honouring this year's winner of the Internet Society's Jonathan B. Postel Service Award (see page 14). Jon's memory also was deeply felt during the Plenary, particularly when Dave, as part of his presentation, invoked Jon's famous credo: Be conservative in what you send; be liberal in what you receive.

This edition of the *IETF Journal* features several other articles of note. Fred Baker offers a new perspective on the transition from IPv4 to IPv6 (see page 16). Kevin Chege of the Kenya Education Network and Mat Ford give us an inside look at the impact of bandwidth-intensive applications on low-bandwidth regions of the global network (see page 18). Leslie Daigle discusses Internet security and stability in her piece about unwanted traffic on the Internet (see page 19). And Geoff Huston returns with an in-depth look at the role that resource certification may play in interdomain routing (see page 21).

We hope you enjoy this issue!

Evolution of the IP Model

By Dave Thaler

In the technical plenary, the Internet Architecture Board (IAB) presented its work on the Evolution of the IP (Internet Protocol) model. For purposes of this work, the IP model refers to the service model exposed by the IP layer to upper-layer protocols and applications (figure 1, page 7). That is, the IP model can be viewed either as a set of behaviours that can be relied on by higher layers or as a set of expectations that higher layers have around IP. In this sense, it is similar to a loosely defined contract that has evolved over time.

A Short History Lesson

In the beginning, IP was first published in 1978 as an Internet Experiment Note (IEN). At the time, IENs were a separate series from RFCs but later merged into the RFC series. After several updates as IENs, IP became RFC 760 in 1980; and finally, the one we cite today, RFC 791, was published in 1981. There was considerable evolution in IP during those three years.

Continued on page 7

Message from the IETF Chair

By Russ Housley


Russ Housley, IETF Chair

I am pleased to report that IETF 73, which was held in Minneapolis in November 2008, was a highly successful meeting. While the total number of attendees was down (937), the number of countries represented was up (52). Many people attribute the reduced attendance to the global economic downturn. That may be true, yet the work of the IETF remains relevant, and the people who came were enthusiastic. Many working groups made significant progress in Minneapolis.

Google was the meeting host and certainly made everyone feel welcome. The social event was well attended, and everyone had a fun, game-filled evening. The site network was subcontracted to VeriLAN Networks, whose staff, working with a group of dedicated volunteers, provided a very sound network.

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

Since IETF 72, two new WGs were chartered and five WGs were closed. We have about 115 chartered WGs. Between the meetings, the WGs and their individual contributors produced 389 new Internet-Drafts and updated 887 Internet-Drafts, some of them more than once. The Internet Engineering Steering Group approved 75 Internet-Drafts for publication as RFCs. The RFC Editor published 97 new RFCs.

During IETF 73, one of the hot topics during the several sessions and many hallway discussions was IPv4 and IPv6 coexistence. The discussion of requirements for NAT–PT (network address translation–protocol translation) continues from the previous meeting. Throughout the week, an IPv6-only network was available for those who were interested in experiencing the Internet without IPv4.

Using WebEx, several WGs conducted an experiment aimed at accommodating remote participants. In one WG session, a presentation was made by a participant in another location. In the plenary, WebEx was used in addition to the usual audio streaming to enable remote participants to follow the presentations. Enabling fruitful remote participation is one way the IETF will ensure that important work gets accomplished despite the potential for reduced meeting attendance brought on by the global economic downturn.

I wish to extend a special thank-you to the authors of *Beautiful Security,* a new book being published by O'Reilly Media. All contributing authors are donating all royalties to the IETF. This contribution is greatly appreciated.

I look forward to IETF 74 in San Francisco on 22-27 March 2009 and IETF 75 in Stockholm on 26-31 July 2009. Scheduling information regarding the next IETF meetings may always be found via http://www.ietf.org/meetings/meetings.html. I look forward to seeing you there.

New BoF Meetings

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

Applications Area
oauth: Open Web Authentication

Internet Area
multimob: Multicast Mobility

Correction

The article, IPv6 Deployment: Lessons from the Trenches, which appeared in the last issue of the *IETF Journal* (Vol. 4, No. 2, October 2008), was inadvertently published without credit being given to the author. The author is Gregory M. Lebovitz


Olaf Kolkman, IAB Chair

IETF 73 Facts and Figures

Registered attendees 937
Countries 52
New WGs 2
Closed WGs 5
WGs Chartered 115
New Internet-Drafts 389
Updated Internet-Drafts 887
IETF Last Calls 98
Approvals 75
(July–October 2008)
97 RFCs published of which
- 59 Standards Tracks
- 6 BCP

102 Internet–Drafts submitted for publication
- 43 waiting for submission of a missing reference
- 78 submitted by IETF

IANA Actions (July–Oct. 2008)
Processed 1,466 IETF-related requests of which:
- 840 Private Enterprise Numbers
- 97 Port Numbers
- 44 TRIP ITAD Numbers
- 30 media-type requests

Reviewed 69 I–Ds in Last Call
Reviewed 81 I–Ds prior to becoming RFC
- 65 contained IANA actions

Words from the IAB Chair

By Olaf Kolkman

"The Internet Architecture Board (IAB) has a number of responsibilities, one of which is to maintain the relationships between the IETF and external organizations. RFC 2850 describes the process in the following manner:

> *The IAB acts as representative of the interests of the IETF and the Internet Society in technical liaison relationships with other organizations concerned with standards and other technical and organizational issues relevant to the world-wide Internet."*

Not completely coincidentally, these are the same words I used when I opened this column in the October 2007 issue of the *IETF Journal*. I was recently reminded of that role when Sha Zukan, undersecretary of the United Nations, invited the IETF and the IAB, through the Internet Society (ISOC), to provide an annual performance report on the steps the organization has taken toward "enhanced cooperation" on public policy issues pertaining to the Internet.[1]

Enhanced cooperation is a term that was coined during the 2005 World Summit on the Information Society. It is a fairly political term that can be seen as an attempt to shift the focus from control over the Internet to discussions about the roles of policy makers, governments, and other stakeholders. It is clear that the IETF has a role in promoting enhanced cooperation, particularly since we are one of the stakeholders in what is commonly referred to as the Internet's multistakeholder model. The IAB cooperates with ISOC to explain, clarify, and improve that model; in other words, multiple stakeholders cooperate to take their responsibility in managing and maintaining their part of the Internet's technical and policy environment. The stakeholders involved include various SDOs (standards-development organizations), such as the IETF, as well as the RIRs (Regional Internet Registries), ICANN (the Internet Corporation for Assigned Names and Numbers), user communities, and governmental and intergovernmental organizations.

With the multistakeholder model in mind, the IAB report[2] highlighted the open and international nature of the IETF and its relationship to other organizations. It emphasized our commitment to the open development and evolution of the Internet protocol suite, and it underscores ISOC's vision of an Internet that benefits all people throughout the world.

The role of the IETF within the multistakeholder model is a serious one. It takes real effort to participate responsibly in the various initiatives that are related to the multistakeholder process. Fortunately, ISOC is assisting us by handling the public policy and governance issues that concern the IETF. They do so by tracking developments, such as those within the Internet Governance Forum, and by flagging issues on which the IAB needs to take action on behalf of the IETF.

That process allows us to focus on technical issues, such as assumptions about the evolution of the IP model, one of the IAB's technical work items,[3] on which, I think it is fair to say, we had a successful technical plenary at this past meeting.

1. https://wiki.tools.isoc.org/@api/deki/files/73/=UNrequest20080312.pdf
2. http://www.iab.org/documents/correspondence/2008-11-26-IETF-response-UN-enhanced-cooperation.pdf
3. http://tools.ietf.org/html/draft-iab-ip-model-evolution

Plenary Report

By Wendy Rickard

Note: This is not a complete report of the plenary sessions; rather, it is a summary of the highlights of the discussions. All IETF 73 presentations can be found at http://www.ietf.org/meetings/past.meetings.html.

In a departure from its usual agenda, IETF 73 merged the administrative and technical plenaries into one session. Internet Architecture Board (IAB) chair Olaf Kolkman described the change as serving two purposes: One was an attempt to "try and open up more time in the agenda" and the other was to mitigate some of the venue logistics associated with rearranging rooms for the plenaries.

Following a few opening comments, Olaf introduced Chris DiBona of Google, which served as the host of IETF 73. "We're proud to be hosting this event," said Chris. "I use that term carefully and deliberately. We're proud because of the important work you do." Chris described the IETF as a place where open source and open standards are enhanced. "The work you do," he said, "is what keeps the Internet free."

The technical plenary was turned over to IAB member Dave Thaler, who gave a presentation on the history and evolution of the IP model. (See page 1 for an article by Dave based on his talk.) The presentation, which generated overwhelmingly positive feedback among attendees, elaborated on Dave's recently submitted Evolution of the IP Model Internet-Draft (draft-iab-ip-model-evolution-01.txt). "A couple of years ago,"

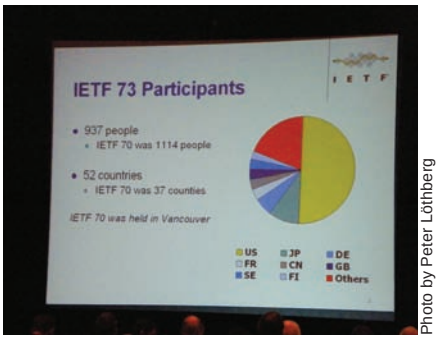

Photo by Peter Löthberg

said Dave, "Lixia [Zhang] once commented that while we talk about these assumptions, maybe someone should write down what this service model is. That is partly what got this started." In addition to his detailed discussion of the IP service model and how it has evolved over time, Dave described how the Internet-Draft documents the properties of the IP layer as they are seen by upper-layer protocols and applications. The document also looks at properties that, if changed, could cause problems, and it provides much-needed guidance for protocol designers and implementers.

"This work is well done," said Dave Crocker during the question-and-answer portion of the technical plenary. "It is reflective, integrative, and practical. By turning out a paper with lessons learned, we are reminded that there are indeed lessons we need to learn."

John Klensin agreed, emphasizing the importance of taking a pragmatic and sober approach to the creation of a historical record. "Too often, institutional memory can get lost and be replaced by mythology," he said.

Drawing heavily on the wisdom of the late Jon Postel, Dave Thaler encouraged IETF participants to frame their understanding of the evolution and future development of the IP model around Jon's oft-repeated quote: "Be liberal in what effects you accept, and conservative in what effects you cause."

Administrative Updates

As part of the administrative portion of the IETF 73 plenary, IETF chair Russ Housley announced that the "code sprint" on the Saturday before the meeting "was a big success." Deployed during the week were Datatracker 2.09 and then version 2.12. Russ pointed out that the interface "has a new look and feel" and that even more would be happening in the coming weeks. He publicly thanked all of those involved, including Glen Barney, Lars Eggert, Pasi Eronen, Bill Fenner, Jelte Jansen, Tero Kivinen, Henrik Levkowetz, Alexey Melnikov, Chris Newman, Robert Sparks, and Magnus Westerlund. Russ also announced that all of the royalties of the soon-to-be published book *Beautiful Security* by John Viega, et al, are being donated to the IETF.

Photo by Peter Löthberg

Google's Chris DiBona (left) accepts a plaque from IETF chair Russ Housley

IAOC Report

An IAOC operations report from the IETF Administrative Oversight Committee (IAOC), which was put together by IAOC chair Jonne Soininen and IETF administrative director Ray Pelletier, reviewed the IETF's financial position for the past year and announced that the organization is expected to break even in 2008. Unfortunately, registration revenues have been underperforming over the past few meetings, due mainly to a reduction in registrations. Jonne and Ray noted that registration revenue for IETF 73 was down sharply, approximately USD 115,000 below budget.

The revenue forecast for 2009 is ex-

"This work is well done," said Dave Crocker during the question-and-answer portion of the technical plenary. "It is reflective, integrative, and practical. By turning out a paper with lessons learned, we are reminded that there are indeed lessons we need to learn."

pected to be USD 3.6 million. Meeting attendance is expected to decline, and registration fees are expected to increase from USD 635 to USD 675. Interim meetings could add approximately USD 328,000 in registration revenue. Ray and Jonne anticipate expenses in 2009 to reach USD 5.4 million, due in large part to the impact of several extraordinary expenses, such as USD 261,000 in interim meeting expenses and a USD 210,000 increase in IT infrastructure enhancement costs.

Given the uncertainties of the global economy, the IAOC has proposed contingency plans for 2009. Assuming that meeting attendance falls by 20 to 30 percent, expenses can be cut in some areas, such as food and beverages; some equipment costs and support travel; and credit card fees. It was pointed out, however, that a number of other expenses couldn't be reduced, such as expenses for the Information Sciences Institute contract for RFC Editor services, expenses for the Association Management Solutions contract for secretariat services, the Network Operations Centre (NOC) expenses for meetings, and the IETF Trust expenses.

The discussion over dealing with budget shortfalls generated a number of comments by attendees. Dave Crocker expressed the need to increase opportunities for remote meetings, as opposed to relying so heavily on the larger IETF meetings. Russ responded that while there is a lot of work going into a remote meeting scheduled for the Real-Time Applications and Infrastructure area in January 2009, in the Internet area the interim meeting was creating a significant burden for the organizers. "We want to experiment with shifting this burden away from the working-group chairs to the secretariat," he said.

Meeting hosts are being sought for all meetings beginning in 2010. Companies interested in hosting an IETF meeting should contact Drew Dvorshak at dvorshak@isoc.org.

Photo by Peter Löthberg

IETF 73 participants meeting in the hotel lobby

IETF Trust

Ed Juskevicius delivered a report on the IETF Trust, beginning with an announcement of legal provisions pertaining to IETF documents. Work on a new policy, as requested by the community in RFC 5377, has been completed, and the effective date of the new policy was set for 10 November 2008 (to coincide with the publication of RFC 5377 and RFC 5378). The newly published Legal Provision policy can be found at http://trustee.ietf.org/license-info/.

New boilerplate language, which is now required for all new submissions to the IETF, was announced, and a transition plan for including the text in new submissions was approved. According to the policy, new submissions can use either the old or the new copyright text through 16 December 2008. After 01h00 UTC on that day, only submissions with the new boilerplate language will be accepted.

The IETF Trust has started work on the updating of applicable document generating and verifying tools and templates for the new boilerplate text. The new text is as follows:

Copyright (c) [insert year] IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (http://trustee.ietf.org/license-info) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Additional boilerplate text for some documents includes the following:

Submission Compliance for all Internet-Drafts

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Derivative Works and Publication Limitations

(See Section 6.c of "Legal Provisions" policy.)

Two different sets of text are possible for use on *some* working group documents (but *never* on any standards-track document.

A brief discussion on the need to police the IETF logo and trademark followed, with some concern expressed by Ed over a handful of cases where the IETF logo is being used in social networking venues. Ed pointed out that it's necessary to notify the offenders or risk losing the trademark.

Continued on next page

Plenary, continued from page 5

While there was general agreement among attendees over the need to police the IETF logo, a few issues were raised in response to the discussion about copyrights. During the question-and-answer portion of the plenary, Sam Hartman questioned the validity of the copyright process. "What happens if I am updating an old document, with the old license, or if I change jobs, and due to an employment contract, I don't own the license?" he asked. According to Sam, 99 percent of the old documents were developed under the previous intellectual-property-rights policy, which, he said, granted fewer rights. "I, as the author of an update, can't give you rights that weren't provided in the old document," he said.

Margaret Wasserman agreed with Sam that the new policy might need to be reviewed. "Let's say, for example, that at some point someone wants to make a new type of IPv6 address," she said. "And suppose we can't find Bob Hinden. How can we give his rights to the IETF Trust for the expanded license when all he gave before were the more restrictive rights?" After some debate over whether the potential glitches in the new policy presented a problem for the IETF process, Jonne Soininen responded that he would look into the issue.

NomCom Report

Nominations Committee (NomCom) chair Joel Halpern said he was pleased with the response to the recent call for volunteers. The IETF NomCom makes appointments to fill open slots on the IAOC, the IAB, and the IESG (Internet Engineering Steering Group). Ten people are selected randomly from a pool of volunteers and there has been a critical need over the past year for volunteers. (For more information, see http://www.ietf.org/nomcom/.)

By November 2008, 99 qualified volunteers had expressed an interest in serving. Of those, 10 members were selected and all are serving. The list of members can be found at https://wiki.tools.ietf.org/group/nomcom/08/.

Currently, the NomCom is collecting feedback on issues related to areas, collective bodies, processes, and individuals, and it invites community feedback and input on any other areas community members count as important. To date, feedback has been solicited from a number of individuals. Joel expressed interest in hearing from folks the committee might not think to ask. Feedback can be sent to nomcom08@ietf.org.

Currently, the NomCom is collecting feedback on issues related to areas, collective bodies, processes, and individuals, and it invites community feedback and input on any other areas community members count as important.

NOC Report

Noah Weis reported that VeriLAN staff and volunteers staffed the NOC. On Sunday, the guest room wireless network was successfully switched on, and all public areas were wireless. Special thanks were extended to the University of Minnesota, including Frank DiGravina, David Farmer, Tim Peiffer, and Dan Westacott, as well as to Cisco, Infobox, and Juniper for the network hardware.

In a question-and-answer session aimed at the IESG, Pete Resnick raised the issue of being better prepared for IETF meetings. Pete expressed concern that too many of the WG meetings are focusing on presentations rather than on interactive discussions. "Presentations are not the same as discussions," Pete said. "The presentations need to be sent in advance so that the meeting time is devoted to the work."

"Pete makes an excellent point," said Internet Research Task Force chair Aaron Falk, who said it is the responsibility of the WG chairs to make the best use of face time. "Area directors can help," he said. "They could contact WG chairs in advance to see how they're doing."

Not all agreed that this approach was prudent or even practical. Charles Perkins expressed some disapproval. "I don't understand the emphasis on getting things done early," he said. "The fact is, people are deadline driven. And they are working like crazy."

Comic BoF

Photo by Peter Löthberg

IETF 73 participants enjoyed a night out at Gameworks in Minneapolis

Evolution of the IP Model, continued from page 1

The evolution didn't stop there, however, and the IP model has continued to change over the years to meet new demands. Some of those changes were intentional. Some were because we found deficiencies. Others were the result of new capabilities. Often, the changes were a consequence of trying to do something else.

By 1989, there was already some confusion concerning the IP model. RFC 1122 was written in an effort to clear up some of that confusion as well as to extend the service model. There are plenty of other RFCs that offered advice on various specific aspects of the IP model, and as a result, to gain an understanding of the IP model, one needed to search many RFCs.

Another RFC appeared in 2004, which is probably the one that is closest in spirit to this work. That one—RFC 3819—offered advice for link-layer protocol designers on how to minimize the impact on layers above the link layer. Hence, it dealt with the service model at the bottom of IP, whereas the present work deals with the service model at the top of IP.

Through it all, many applications and higher-layer protocols have been built on top of IP. Besides the things that were actually documented in those RFCs, they made various assumptions about IP. Those assumptions today are not listed in one place. They're not necessarily that well-known. They're not necessarily thought about when changes are being made. And, as we'll see, increasingly they're not even true.

The goals of the IAB work were first, to collect these assumptions (or, increasingly, myths) in one place—or at least provide references to the various other places that already have them—and second, to document to what extent those assumptions are true and to what extent they are not. Beyond that, we were interested in providing some guidance for the community.

The collected assumptions were previously presented to various subsets of this community. For example, much of the information about the assumptions was presented in the Internet area meeting at IETF 72 in Dublin, and another subset was presented in the EXPLISP BoF. The IAB solicited input from the community—or at least those subsets of the community—so we could go off and come up with guidance. We have now done that between Dublin and Minneapolis. The focus in the technical plenary was thus to discuss the IAB guidance as a working session.

The goals of the IAB work were first, to collect these assumptions (or, increasingly, myths) in one place—or at least provide references to the various other places that already have them—and second, to document to what extent those assumptions are true and to what extent they are not.

Assumptions

The basic IP service model described in RFC 791 indicated that senders are able to send to an address without signalling a priori. Receivers can listen on some address they've already obtained—without signalling a priori. Packets can be of variable size, and there's no guarantee of reliability, ordering, or lack of duplication. That's the model that we held up as the great IP service model.

That left a lot unstated, however. RFC 1122 added some clarification—

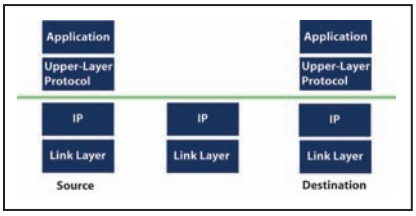

Figure 1: The model exposed by IP to higher-layer protocols and applications

for example, with respect to defining the notion of strong-host (or end-system) versus weak-host models, both of which were allowed and supported on different platforms. On one hand, in the strong-host model, when a host sends a packet from a particular source address, it has to send the packet out on an interface that corresponds to that source address. Similarly, if a packet comes in to a particular destination address, it has to arrive on an interface corresponding to that destination address, or the host will drop it. In the weak-host model, on the other hand, the host can send and receive packets on any interface. Routers, for example, follow the weak-host model when forwarding is enabled.

So while RFC 1122 added some clarification, we had two different behaviours and so ended up with two different variants of the IP model. Since different platforms implemented different behaviours, applications could not rely on a single behaviour.

Even with RFC 1122 and other RFCs, many other assumptions made by upper-layer protocols and applications are not well documented. Such assumptions or myths generally fall into four categories: assumptions about IP connectivity, assumptions about addressing, assumptions about upper-layer protocol extensibility, and assumptions about security.

When talking about claims that may or may not be true, Snopes.com has a nice model that includes a claim, examples, and a status of true or false or partially true. We will use this same model

Continued on next page

Evolution of the IP Model, continued from page 7

below. Note that all of the assumptions we will talk about are at best partially true, but that hasn't stopped applications from making the assumption anyway.

Assumptions about IP Connectivity

The document covers a number of connectivity-related assumptions, three of which we will mention here by way of example.

Claim: Reachability is symmetric, or "If I can reach you, you can reach me." Some examples of upper-layer protocols and applications that do this include request-response protocols; in other words, if a request can reach you right now, then the response can come back to me. That's a fairly small time window. Then there are much broader assumptions, such as, "If I can reach you today, then you can reach me tomorrow."

There are lots of reasons that this is not entirely true. For example, we have NATs, firewalls, one-way media such as satellite links, and even wireless situations such as 802.11 ad hoc mode, whereby if my radio is stronger than yours, I can get a packet to you, but you can't get one back to me. There have also been some efforts to try to make this claim be more true. For example, RFC 3077 was one effort to help restore this for some of these cases. Today, request and response paradigms usually work, but not callbacks over a much longer time frame.

Claim: Reachability is transitive, or "If I can reach you and you can reach her, then I can reach her." The same sorts of things (NATs, firewalls, packet radio technologies, etc.) interfere with this assumption too, so today you not only have lack of symmetry, but you also have lack of transitivity.

Claim: The latency of the first packet that you send to a destination is typical of what you'll see after that. This assumption is commonly made when picking from a set of candidate servers or addresses. Many applications and protocols send a packet to each of them and then use the one that responds first, assuming that it's the best one to talk to. A number of things interfere with that assumption today. First, the first packet may have additional latency due to a routing-related cache miss in an intermediary (e.g., ARP or flow-based routers). While resolving the next hop, the packet may be queued in the meantime. When comparing two different destinations, if one of them needs to do a resolution and one doesn't, there can be a difference from what you might expect. Second, there are a number of protocols that have the notion of path switching. For example, we see this in Mobile IPv6, Protocol Independent Multicast–Sparse Mode (PIM-SM), the Multicast Source Discovery Protocol (MSDP), and various Routing Research Group proposals, wherein packets initially follow one path and then quickly switch to a more efficient path. That means the first burst may have a much longer latency than subsequent packets do. If you have another destination that is already switched, you might unduly think that one is closer, but it might not be. So, as a result, if you make that assumption, you can end up with, in some cases, highly suboptimal choices, and that can result in longer paths, lower throughput, and a higher load on the Internet.

In terms of IAB guidance around IP connectivity-related assumptions, we first observed that the reasons they are no longer true—or at least much less true—tend to fall into two categories: either they are effects of something done at the link-layer independent of IP or they are effects of things done specifically by network-layer technologies.

Usually, the link-layer effects are not intentionally trying to break IP. Designers don't set out trying to invent a link type that is difficult to run IP over, but it's when defining IP over them that we inadvertently create problems. RFC 3819 gives good advice to link-layer designers about what they can do. The other piece of guidance that the IAB added for those of us who define IP over various link types is that we try to recognize the gaps mentioned in the document and compensate for them as much as is practical. Examples of where the IETF has actually made attempts at doing this is RFC 3077, which attempts to compensate for unidirectional links such as satellite links, and RFC 2491, which attempted to compensate for non-broadcast-capable links. A notable gap today is in the area of compensating for lack of transitivity, such as with IP over 802.11 ad hoc mode.

As for network-layer technologies that

> As for network-layer technologies that interfere with these assumptions, in the IETF we like the notion of reachability. We say that everyone should be able to talk to everyone. That's true, of course, until we start getting stuff we don't like getting.

Photo by Peter Löthberg

IETF 73 plenary attendees line up at the mic.

interfere with these assumptions, in the IETF we like the notion of reachability. We say that everyone should be able to talk to everyone. That's true, of course, until we start getting stuff we don't like getting. Most of us then realize we don't actually want to be reachable by everyone; we want to be reached only by the good guys. The notion of restricting reachability to only some portion of those who might want to communicate with us is already a part of the current IP model. An example is with IPsec (Internet Protocol Security), which is now a core part of IP itself. The point is that blocking communication to or from unauthorized parties is legitimate.

When reachability is affected for reasons beyond simply restricting access to authorized parties only, the IETF should attempt to proactively avoid such hindrances for new technologies—or solve them for existing technologies (e.g., 802.11 ad hoc). Referring back to figure 1, the IP model is what is exposed to the transport layer and above on the source and destination. One approach designers sometimes use to avoid some of the effects (e.g., nontransitivity) of odd link types is to hide the link from hosts that run upper-layer protocols and applications and use such links only between routers.

We give the following principle, with wording inspired by another principle from the late Jon Postel: "When defining a protocol, be liberal in what effects you accept from lower layers, and conservative in what effects you cause to upper layers." Using the general principle, being liberal and conservative, we have roughly two categories of work. Perhaps half of the work in the IETF is around the network layer and below, and half is around the transport layer and above. Many of us are actually in both camps. Work at the higher layer should avoid making the assumptions when practical—and at least consider them in the writing of requirements and applicability statements. Work at lower layers should avoid making the assumptions be less true when practical and similarly document any remaining effects on the assumptions made by upper layers so that other designers and administrators are aware of the impact.

Note the use of the word practical earlier. To illustrate what we mean by this, let's look at a specific example. IAB RFC 4903 on multilink subnet issues talks about non-broadcast multiaccess (NBMA) links. An example of an NBMA link is 6to4, which is not intended to go across a particular network but across the Internet. It doesn't support multicast, but we often don't have multicast deployed across the public Internet anyway. So 6to4 is an example wherein trying to add multicast over it may not be practical. It would add complexity that would not be needed until you actually have wide-area multicast that would be needed across the same environment.

We give the following principle, with wording inspired by another principle from the late Jon Postel: "When defining a protocol, be liberal in what effects you accept from lower layers, and conservative in what effects you cause to upper layers."

Assumptions about IP Addressing

Claim: Addresses are stable over long periods of time. Once upon a time, that was mostly true, at least until we started having things like the Dynamic Host Configuration Protocol (DHCP) and hosts that move around. In common application programming interfaces (APIs), such as the sockets API that many of us use, applications call a name resolution API, such as gethostbyname or getaddrinfo, and then connect to one or more resolved addresses. When a name is resolved in DNS, the requester gets back a time to live (TTL) together with the addresses, but this TTL is not present in the API, and so, applications may cache the answers for longer than indicated by the DNS, resulting in eventual communication with the wrong entity.

Photo by Peter Löthberg

Dave Thaler speaking at the technical plenary at IETF 73

We also see some efforts that are intentionally, or as a side effect, trying to restore this assumption to some effect. Proxy Mobile IPv6 (RFC 5213), for example, tries to restore it for some level of mobility within a network. Protocols such as Mobile IP and the Host Identity Protocol (HIP) try to provide stable addresses to some extent by adding an additional stable address that an application can use. Hence, if applications that make this assumption use the stable address, then they work better.

Claim: A host has only one address and one interface. Unfortunately, there exist many applications that resolve a name to a set of addresses and then simply pick the first one and use it. We saw this in a lot of applications when we started porting them to use IPv6, for example. Other applications use an address to identify

Continued on next page

Evolution of the IP Model, continued from page 9

a user or a machine and get confused if multiple users or multiple machines use the same address or if the same user or machine uses multiple addresses. Another common problem is that there are many DHCP options for per-machine information, whereas DHCP options are obtained over a particular interface from a particular network, and often, it's not mentioned how the host then converts per-interface information (such as a set of DNS servers) to machine-wide information when it gets multiple answers from different interfaces or networks.

Photo by Peter Löthberg

IETF 73 plenary audience

So, of course, this assumption is much less true today. Many hosts have multiple interfaces, and hosts have both IPv4 and IPv6 addresses even if they have only one interface. The use of virtual private networks (VPNs) is also fairly common and results in multiple interfaces. To some extent, protocols like Mobile IP and HIP are trying to restore this by adding another "address" that applications that make this assumption can use and be isolated from the use of multiple other addresses the host may have.

Claim: An "address" used by an application is the same as the "address" used for routing. What some call an ID/locator split is an example of when this is not true. Many applications have assumptions, however incorrect they may be, about the relationship between proximity in the address space and proximity in the topology. That is, if you and I have similar addresses, you must be close to me, and hence you're a better peer to talk to than someone with an address that appears to be much less similar.

Similarly, some applications or protocols have a service select addresses to put in a referral to a client based on the client's address and how it relates to the potential candidates' addresses. This assumption is certainly not true with tunnelling to and from hosts, because applications see the address in the inner IP header, whereas routing uses the address in the outer IP header. Similarly, it is not true with IP/locator split schemes that split them at a host.

Again, the assumptions mentioned earlier are examples that serve to motivate the guidance that follows, and more assumptions are discussed in the draft.

> It's also worth mentioning that changes to any assumptions, not just assumptions about security itself, can impact security if some application or upper-layer protocol bases its behaviour on that assumption.

So, what does the IAB think about these? If we look back at architectural principles of the Internet, there's a good statement in RFC 1958: "In general, user applications should use names rather than addresses." If only that were true!

Today we have many APIs that unnecessarily expose addresses to applications, and many applications have to deal with the concept of an address only because they need to use it to open a transport connection instead of being able to connect by name, as Stuart Cheshire discussed in the plenary in Dublin. Today it's often an implementation issue, not a protocol issue, but there are also protocols defined to carry only IP addresses instead of carrying names when doing a referral or request for a later callback.

In general, anything that's already dependent on some naming system should try to avoid using addresses and use only names. As a side effect, this happens to ease the transition to IPv6 because applications that know nothing about IP addresses generally work without changes.

Assumptions about Upper-Layer Extensibility

Claim: New transport-layer protocols can work across the Internet. Figure 2 shows the hourglass from Steve Deering's presentation at the IETF 51 plenary back in 2001 regarding the "waist" of the Internet. Besides TCP and UDP, it shows "…". The IP model is not static, and neither are other layers. Some applications use raw sockets and make this assumption—or at least want this assumption to become more true. But today devices such as NATs and firewalls support only UDP and TCP or, even worse, support only HTTP. As a result, many applications and protocols (such as the whole Web Services architecture) today are built on top of HTTP instead of TCP or UDP, and we even see IP over HTTP, resulting in an architecture more like that shown in figure 3.

Claim: If one stream to a destination can get through, then so can others. For example, you have applications that open multiple connections to get better throughput, and you have applications

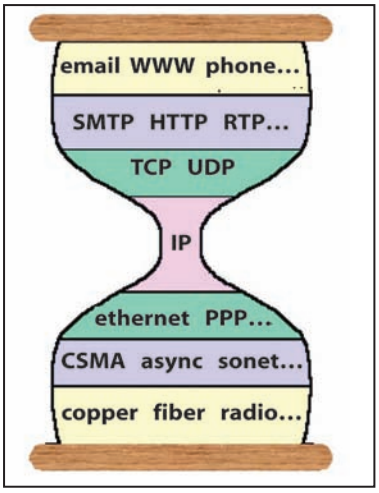

Figure 2: Steve Deering's hourglass showing the "waist" of the Internet

such as FTP that open separate connections for data and control. However, a number of factors may interfere with that assumption. Some firewalls may block specific ports, for example. Also, some middleboxes keep per-connection state and may run out of memory or ports when an application attempts multiple connections. This has come up in discussions of carrier-grade NATs, for example. Just because you can get one, doesn't necessarily mean you can get a hundred.

In considering these assumptions, we observe that RFC 791 doesn't actually describe what the requirements were, but there's a great paper by Dave Clark referenced in the document that does list the requirements that were discussed when IP was first designed. One such requirement was to support the widest possible range of applications by supporting a variety of types of service at the transport level.

The issues with this today arise either in the name of security—or as a side effect of something else, such as address shortage—and the same guidance applies as for the IP connectivity-related assumptions.

Assumptions about Security

In terms of security, the examples are well-known, including modifications to packets in transit, privacy, and forged source addresses. Fortunately, RFC 3552, which talks about how to write security considerations, already has excellent guidance for what people should do about these assumptions.

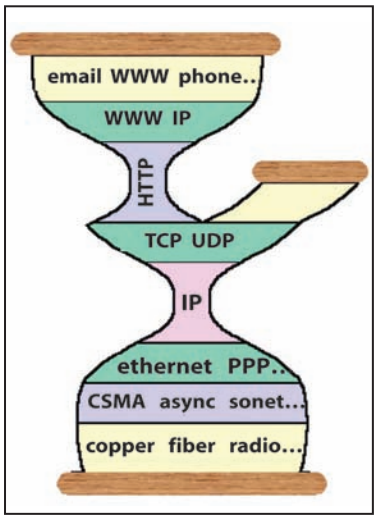

Figure 3: An updated hourglass showing an architecture often seen today

It's also worth mentioning that changes to any assumptions, not just assumptions about security itself, can impact security if some application or upper-layer protocol bases its behaviour on that assumption. For example, consider an application that binds to an IP address, running on a host with two interfaces: one on a "safer" network and one on some untrusted network that you don't want to do certain things on. If the application binds to an address on the good network, will it see only traffic that comes in across the good network? There are applications that assume this is the case. We saw earlier, however, that this is true for strong-host systems and false for weak-host systems. As a result, great care should be taken when making an assumption less true. Upper layers should also carefully consider the impact of basing security on any such assumption.

Of course, many assumption violations are actually done in the name of security even though they break some applications.

Conclusions

Unless you can enumerate all possible applications that are run, any changes to the assumptions listed in the document will probably break some applications. We've realized that it becomes harder and harder over time to evolve the IP model, because there are more and more things that might have assumptions built in. Still, the IP model is not static, and continuing to evolve it to meet new demands is important. Changes must be made with extreme care, however. Adding functionality that has no impact unless the upper layer asks for it is generally safe, but fewer entities will actually use it.

For those who make changes to the network layer or below, write down the effects on upper layers—for example, as part of requirements and applicability statements. For those who work on technologies at the transport layer or above, avoid these assumptions whenever practical, and if you do depend on any, write them down—for example, in requirements and applicability statements.


Photo by Peter Löthberg
Dave Thaler takes questions from the audience at IETF 73

ISOC Fellows Attend IETF 73

By Wendy Rickard

Photo by Kevin Craemer

Internet Society Fellowship to the IETF fellows, mentors, and ISOC staff attend IETF 73 in Minneapolis

Before IETF 73, Terry Rupeni had a relatively good understanding of the work being done by the IETF. What he couldn't quite grasp was the work flow. After a few days, the mild-mannered engineer from Fiji was clearly in the flow, and he was becoming much more comfortable with the temperature, both inside and outside the hotel.

As a newcomer to the IETF, Terry's trip to Minneapolis was made possible as part of the Internet Society (ISOC) Fellowship to the IETF programme. What did he think? "It's a little cold," he said of the city. The temperature inside was a different story. "It can get a little hot in there," Terry said, referring to both the working-group meetings and the plenary sessions. "Here people go right up to the mic," he said. "Where I come from, it's not part of our culture to speak up in this manner. In Fiji, you don't voice your opinion in public."

Understanding the IETF from the inside out can be a cultural shock for many of the engineers who come from less-developed regions. IETF fellow Jean Philemon Kissangou, who serves as technical manager for an Internet service provider in the Republic of the Congo and as director on the Board of the regional Internet registry AfriNIC, imagined that IETF meetings were "a complicated thing." Carlos Alberto Watson Carazon, who is from Costa Rica and who has attended LACNIC (Latin American and Caribbean Internet Addresses Registry) meetings, found the IETF experience uplifting. "The IETF respects opinions of people from other countries, and they respect other points of view," he said.

As the ISOC fellowship programme moves into its fourth year, the infusion of diverse cultures into IETF meetings is being felt both by fellows and by the IETF.

The Challenges Back Home

Being better prepared to tackle the challenges of managing or expanding access to the Internet in rural, remote, or developing regions is what motivates many of the fellows to apply to the fellowship programme. In Fiji, where Terry works as a network analyst at the University of the South Pacific, the Internet is still in its nascent stages. A government monopoly that ended nearly a year ago has given rise to four ISPs instead of one. The result has been increased competition and lower prices.

> The challenge for engineers like Terry lies in how to expand access beyond cities and towns to rural areas, where access is generally limited.

Since the monopoly ended, Internet use has taken off in Fiji, mainly in the area of personal use, and Terry is quick to comment on the impact that that kind of use has had on the culture. People are learning to use Facebook, and e-mail is becoming the norm. And, as in most places, the mainstreaming of the Internet in Fiji has fundamentally changed how people in the country communicate.

The challenge for engineers like Terry lies in how to expand access beyond cities and towns to rural areas, where access is generally limited. Since power supplies can be problematic, efforts are being made to explore solar energy as a way of powering networks, and the new ISPs are beginning to invest in wireless technologies.

As a whole, when it comes to the Internet, Terry said that Fijians are interested and engaged. "There are not too many technological challenges," he said. "We learn from developed countries. And more and more computers are coming from Malaysia."

Philemon said that in the Congo, while access continues to be a problem in his country, the challenges he faces these days are IPv6 related. "Our network is only IPv4," he said. "We need to adopt. I need to explain the benefits." As in Fiji, it's been difficult to fund and build an infrastructure that would connect the rural regions of the Congo. Access in those areas is made available mainly via satellite, though Philemon said the government is now dedicated to laying down fibre. And while the Internet in the cities is increasingly available at Internet cafés and at some private companies, in schools, teachers complain that they don't have access. "Students would like to use the Internet, but they can't," Philemon said. "Even within the government Internet use is limited."

In Carlos's corner of the world, Internet penetration is at roughly 24 percent. The problem, he said, is that in Costa Rica there are only two ISPs, and the government runs both of them. "With a government-controlled ISP," he said,

"getting connected can take up to six months." Getting connected through a cable company is easier but considerably more expensive. Even so, said Carlos, cost is not the main issue; the main issue is access. "Most folks can afford the Internet," he said. "They just can't easily get it."

Carlos believes strongly in the benefits the Internet offers, particularly in the areas of education, health care, and business development. He said he would like the government to do more to improve and expand access. "They should get behind open source," which, he said, is not just more affordable but higher quality.

The Polynesian island nation of Tuvalu is composed of four reef islands and five atolls. It is there that IETF 73 fellow Tenanoia Veronica Simona serves as an IT manager at the Tuvalu Telecommunication Corporation (TTC), one of two ISPs in the region (the other one is operated by the government). Currently, the Internet is available only on the main island of Funafuti, which is the capital of Tuvalu. According to Tenanoia, TTC provides WiFi connections primarily for government entities that are located outside the government building but also to some nongovernmental organizations. "As of last month" she wrote, "we have also started to provide ADSL connections to people at home." Next they are targeting schools in general, but primary schools in particular. One of two secondary schools is currently accessing the Internet from Vaitupu, one of the outer islands.

According to Tenanoia, the majority of the population isn't able to access the Internet yet, but she is hopeful that that will soon change. The challenge they face is their capacity to roll out the Internet to the outer islands, of which there are eight and all of which are separated by the Pacific Ocean. TTC is meeting that challenge by preparing to build a system that will prove capable enough to expand Internet access to the majority of the population in the outer islands. "We at TTC are trying to secure funds for this purpose," she wrote. "It is our priority."

Tenanoia said the Internet is an important enabler in her part of the world. She said it can cut the cost of telephone communications, and it allows students to conduct research more quickly and efficiently than they can now. And, like the other fellows, she said access to the Internet is critical to adequate health care. "Our medical personnel are now able to conduct research online for solutions from the main hospital in Tuvalu," she wrote.

Impressions from a Returning Fellow

"I was like an alien in Philadelphia," Mohibul Hasib Mahmud said in between meetings at IETF 73. In Minneapolis, the IETF fellow from Bangladesh hit the ground running.

Mohibul said the Internet is growing in his country, albeit slowly. He described it as an evolution in terms of technology development. Since 2006, though, it has been a revolution. That's when fibre was laid and Bangladesh was no longer dependent on satellite. Since then, usage has grown exponentially, in terms of both business and personal use.

Photo by Wendy Rickard

Fellow Jean Philemon Kissangou at IETF 73 in Minneapolis

Working with an ISP is always challenging, said Mohibul. As a technologist, he must constantly grow with the infrastructure. Over the past year or two, subscriptions have increased, and there is a growing need for more IP addresses, so they work with the regional Internet registries. They also work hard to create services that would meet clients' needs, especially now that there are at least a couple of hundred ISPs in his country.

According to Mohibul, regardless of the advances, penetration in his country is still too low and there's not much in the way of a telecom infrastructure. Internet access is available mainly in the cities, but even there, he said, there are problems, including a lack of computers. "Without computers," he said, "there isn't much of a way to take advantage of the Internet."

Even with the challenges, Mohibul is pleased with the impact the Internet has had on communications in his country. "Lots of people go overseas for jobs," he said. "Often, they are out of touch with family and friends for a long time. Telephoning is expensive. With the Internet, they are able to stay in touch."

IETF 73 Fellows and Mentors

Jean Philemon Kissangou (Congo)
Mentor: Alain Aina

Carlos Watson Carazo (Costa Rica)
Mentor: Roque Gagliano

Tenanoia Veronica Simona (Tuvalu)
Mentor: Fred Baker

Terry Rupeni (Fiji)
Mentor: David Farmer

IETF 73 Returning Fellows

Burmaa Baasansuren (Mongolia)

Mohibul Hasib Mahmud (Banglasdesh)

Veaceslav Sidorenco (Moldova)

Jonathan B. Postel Service Award Granted to EsLaRed

By Wendy Rickard

Photo by Kevin Craemer

Lois Postel, Postel Award winner Ermanno Pietrosemoli, and ISOC president Lynn St. Amour at the Postel Service Award ceremony in Minneapolis in November 2008.

Internet Society president Lynn St. Amour announced at the IETF 73 plenary that Fundación Escuela Latinoamericana de Redes (EsLaRed) had been granted the coveted Jonathan B. Postel Service Award. It is the first time in the 10-year history of the award that it has been given to an organization rather than to an individual.

The 2008 award commemorates the 10-year passing of Internet pioneer Jonathan B. Postel, who is best known for being editor of the RFC series and for administering the Internet Assigned Numbers Authority (see article, next page). A private dinner and award ceremony, which took place later in the week, was attended by Jon Postel's brother, Russ Postel, and Jon's mother, Lois Postel, who presented the award to EsLaRed president Ermanno Pietrosemoli.

Ermanno spoke of the announcement of the award as an emotional moment for both him and his organization. "We have been working for many years in the shadows, without much public fanfare or recognition," he said. "This is a very special occasion, and I am very deeply grateful to the Internet Society and to the IETF, which have been supporting us for the past 10 years." At the ceremony, Ermanno recalled having attended the INET meeting in Geneva in 1998, where he had the opportunity to meet and to get to know Jon, whom he described as a luminary. "I really feel very honoured to be somewhat humbly associated with his name," he added.

For the past 16 years, the little-known, Venezuela-based nonprofit EsLaRed has been training a new—and in some cases, the first—generation of Internet trainers and professionals, many of whom are forging Internet access in remote and under-served areas within and outside South America. EsLaRed's efforts to facilitate scientific and technical progress in Latin America and the Caribbean have been instrumental in forming what is today a vibrant and dynamic Internet community in the region.

In August 2008, EsLaRed participated in the effort to build a high-speed, 162-kilometre long wireless network in Malawi. The network is being used to enhance medical and educational applications at the University of Malawi. EsLaRed is also responsible for the design of a wireless data network in the Galápagos Islands.

What makes EsLaRed's approach to network training unique is its insistence on teaching the technology that is available. "Even if we could afford more-sophisticated technologies, what would be the point?" Ermanno asked during an interview in Minneapolis. "It's more important to have students master what they have and what they need to get the job done." In one case, students were taught how to build an antenna out of a can, and then they were instructed to take their antennas home and make them work. "Most of the time, people don't see technology as belonging to them," he said. "It's not part of their daily life. So for them, to build an antenna and realize that it works makes a big difference in their lives."

Seeing the difference that the Internet makes on people's lives has been a key motivator for both Ermanno and EsLaRed. "I'm very excited because we are in a part of the world that isn't a hot spot for technology," he said. "So we can see how the Internet actually changes lives."

Since 1992, EsLaRed, with support from its worldwide and regional sponsors, has conducted network training workshops nearly every year in locations throughout Latin America and the Caribbean, including Argentina, Brazil, the Dominican Republic, Ecuador, Mexico, and Peru. "We have been working at this for many years," said Ermanno, "and it has been hard at times."

At the ceremony, Lynn read a statement called Remembering Jon that was prepared by Vint Cerf . "Always a strong believer in the open and bottom-up style of the Internet," wrote Cerf, "Jon would . . . be pleased to see that the management of the Internet address space has become regionalized and that there are now five Regional Internet Registries cooperating on global policy and serving and adapting to regional needs as they evolve. He would be equally relieved to find that the loose collaboration of DNS root zone operators has withstood the test of time and the demands of a hugely larger Internet, showing that their commitment has served the Internet community well."

In Memory of Jon Postel

By Wendy Rickard

"If the Net does have a god, he is probably Jon Postel." — The Economist

Last October marked the 10-year anniversary of the passing of Internet engineer, standard-bearer, and icon Jon Postel. While Jon's work contributed in countless ways to the advancement and smooth functioning of the Internet, it was his role as RFC editor—a role he created and held from April 7, 1969, until his untimely death on October 16, 1998—and his work with the IANA (Internet Assigned Numbers Authority) that are of particular significance to the IETF community.

To those who knew him, Jon was a brilliant and astute engineer whose soft-spoken manner belied his dogged pursuit of excellence, a characteristic especially evident in his role as editor of the RFC document series. His feedback to authors was not confined to grammar and phrasing, for which he was a stickler; it included any potential inconsistencies, ambiguities, and duplications of effort. Jon's longtime colleague Joyce Reynolds, wrote in RFC 2555 that while operating systems and computers have changed over the years, "Jon's perseverance about the consistency of the RFC style and the quality of the documents remained true."

Equally impressive, especially in hindsight, was Jon's ability to anticipate the need for keeping track of the work. "Somehow, Jon knew, even 30 years ago, that it might be important to document what was done and why, to say nothing of trying to capture the debate for the benefit of future networkers wondering how we'd reached some of the conclusions we did (and probably shake their heads...)," wrote Vint Cerf, also in RFC 2555.

It wasn't always easy. According to Jake Feinler, "Jon often took merciless flak from those who wanted to continue discussing and implementing, or those whose ideas were left on the cutting-room floor. Somehow he always managed to get past these controversies with style and grace and move on."

Jon not only edited the RFC series; he also authored or coauthored more than 200 of them (see http://www.postel.org/postel.html). In September 1981, he wrote RFC 791, which described the Internet Protocol, as well as RFC 792 (the Internet Control Message Protocol) and RFC 793 (Transmission Control Protocol). In 1982, he authored RFC 821, which defined the Simple Mail Transfer Protocol.

According to Bob Braden, who worked with Jon at the Information Sciences Institute of the University of Southern California, there were many aspects of the IETF culture that matched Jon very well. "Dedication to making things that work, a never-ending attempt to keep protocols as simple and powerful as possible, and a slight counter-cultural tinge all characterized Jon," Bob wrote in a memo to the IETF after learning of Jon's passing. Jon's prescience, talent, and meticulous approach to his work are characteristics that stand out for many engineers and Internet developers from that time. "As far as I know," wrote Bob, "Jon had no model to follow when he wrote RFCs 791, 792, and 793, yet the result was a model that I personally have spent nearly 20 years studying and trying to emulate. Jon's contribution was not just the skill and grace of his editorial style; in writing these documents, Jon determined much of the detailed content, interpreting and elaborating the ideas of others to produce one seamless whole."

In January 1980, Jon wrote the statement that many say most accurately described his philosophy toward life: "In general, an implementation should be conservative in its sending behaviour, and liberal in its receiving behaviour."

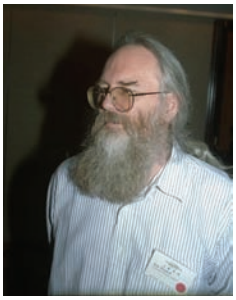

Photo by Peter Löthberg

Jon Postel

The comment, which appeared in RFC 760, was amended a year later by Jon in RFC 793 to the oft-repeated: "Be conservative in what you do; be liberal in what you accept from others."

On the 10-year anniversary of his passing, Vint, who wrote Jon's obituary and published it under the playful title of RFC 2468 (I Remember IANA), reminds us of the significance of that statement, not just for Jon, but also for the community at large. "[Although] he meant [it] in the context of detailed protocols, it also serves as a reminder that in a multistakeholder world, accommodation and understanding can go a long way towards reaching consensus or, failing that, at least toleration of choices that might not be at the top of everyone's list."

Vint has served as a prolific interpreter of Jon's legacy, describing him as the network's Boswell[1]. However, it was Jon's "devotion to quality and his remarkable mix of technical and editing skills that permeate many of the more monumental RFCs that dealt with what we now consider the TCP/IP standards," wrote Vint. "Many bad design decisions were reworked thanks to Jon's stubborn determination that we all get it 'right.' As the editor, he simply would not let something go out that didn't meet his personal quality filter. There were times when we moaned and complained, hollered and harangued, but in the end, most of the time, Jon was right and we knew it."

1. James Boswell was an 18th Century Scottish lawyer, diarist, and author whose name has become a term used to describe a constant companion and observer. See http://en.wikipedia.org/wiki/James_Boswell

IPv4/IPv6 Coexistence and Transition

By Fred Baker

The Internet faces a transition from its traditional IPv4 to IPv6, with a period of coexistence. Here is one technologist's view of the road ahead for the Internet Protocol and IP networks from the perspective of work happening in the IETF.

The time is rapidly approaching when the last of the IPv4 address space will be allocated. Even though options are being considered that would enable the trading of IPv4 address space as a commodity (which has important implications for routing) as well as for sliding an ISP interconnection layer underneath the IP protocol, networks requiring new address space in large amounts will deploy IPv6. Therefore, some form of coexistence and, eventually, a transition are inevitable. To that end, the IETF has explored several transition mechanisms, most of which are described in RFC 4213.

According to thefreedictionary.com, a transition is a "Passage from one form, state, style, or place to another." As such, a protocol transition from IPv4 to IPv6 requires two events:

- IPv6 must be turned on in routing, on application servers and services, and on the peer or client systems that use or participate in those services
- IPv4 must be turned off—at least in the network

Two serious questions arise from that observation.

- *How long a time period should be allowed between those events?* Does one turn IPv4 off and IPv6 on simultaneously, or does one turn on IPv6, allow a time interval to elapse, and turn IPv4 off later? If the latter, how long a time interval is rational?
- *How are IPv6 datagrams carried in an IPv4 network before IPv6 is turned on?* How are IPv4 datagrams carried in an IPv6 network after IPv4 is turned off? There are two broad categories of solutions: tunnelling solutions and translation solutions.

IPv4/IPv6 Coexistence

From the IETF's perspective, the optimal approach for existing networks is to focus not on *transition* but on coexistence. Turn on IPv6 now and start using it; turn IPv4 off at some point in the future when it is no longer a business requirement. Therefore, in the opinion of the IETF, network administrators should:

- Turn on IPv6 routing in their existing IPv4 networks.
- Contract IPv6 service with their upstream, peer, and downstream neighbours.
- Use the IPv6 protocol in addition to IPv4 in their applications and services both on server equipment and on their clients.

In doing so, network administrators will likely find software and hardware that are old or for some reason cannot be upgraded. They should schedule those upgrades as their budget allows.

The reason to support coexistence of this type should be obvious: If IPv6 isn't working or if another network does not yet support IPv6, the affected applications or services will remain available via IPv4.

Providing coexistence in network layer routing can be accomplished in any one of three ways:

- Enabling IPv6 on routers that carry IPv4
- Enabling IPv6 on other routers as a parallel network internal to the customer-perceived network
- Enabling IPv6 on a separate parallel network directly visible to neighbouring networks and customers

Building parallel networks is obviously far more expensive than turning on IPv6 on the existing equipment.

Overlay Networks

To date, before native IPv6 routing and applications were turned on, IPv6 has been in use via overlay networks that were built using tunnels or multiprotocol label switching. Initially, 6BONE and 6NET routed IPv6 through static tunnels. Dynamic approaches have since been developed, though, such as 6to4 (RFC 3056), Teredo (RFC 4380), and ISATAP (RFC 5214).

One solution that has been proposed for broadband access networks involves having the ISP continue to manage an IPv4 network, with IPv6 running as a tunnel overlay between the customer-provided-equipment (CPE) router and an ISP-identified tunnel endpoint. In this model the ISP does not offer native IPv6 service. As a transitional deployment step, the ISP indicates the IPv4 address of a tunnel endpoint to CPE routers when it configures them, and it includes an IPv6 prefix using DHCP-PD. This allows the ISP to traverse the parts of its network that are not yet ready to support native IPv6 forwarding.

The tunnelling of IPv4 through IPv6 is analogous to the tunnelling of IPv6 through IPv4, though only static tunnels are defined at this point.

Translation Technologies

Translation between IPv4 and IPv6 is not generally considered a viable long-term strategy, if only because it begs the question of the size of the address. If IPv6-to-IPv4 translation is sufficient to address the systems to which a user needs access, then one needs only to reallocate the existing IPv4 address space to solve the problem. Translation is, however, generally recognized as a necessity in certain cases to provide connectivity between IPv6-only and IPv4-only systems or networks. The issues that arise relate in part to path- MTU

(maximum-transmission-unit) detection, which is often problematic in IPv4 networks but is required in IPv6 networks. Other issues involve supporting applications that are not designed on a client/server architecture or that require a sophisticated firewall traversal mechanism.

The Stateless IP/ICMP (Internet Control Message Protocol) Translation Algorithm (SIIT) (RFC 2765) is implemented in a translating router. The router advertises one or more IPv4 prefixes (perhaps host addresses) in IPv4 routing and a prefix in IPv6 routing. By defined transforms, it translates between IPv4 and IPv6 and between ICMP and ICMPv6.

NAT-PT (RFC 2766) extends the SIIT concept with a DNS application layer gateway. The gateway replicates A records from the IPv4 network as AAAA records carrying SIIT-compliant addresses in the IPv6 domain and advertises A records for the IPv6 hosts with SIIT addresses in the IPv4 domain. Doing this statelessly, however, implies either host routing in the IPv6 network, which has scaling issues, or a small IPv6 domain—nominally a single LAN—attached to a much larger IPv4 domain, because the upper 96 bits of the address in the IPv6 domain are static.

There is one problem with SIIT and NAT-PT: they are designed to enable small IPv6 islands to operate within a general IPv4 network, and they do not scale well in a more general deployment. Hence, the IETF is at this point working on next-generation translation technologies intended to support more general deployment, based on operational experience with SIIT, NAT-PT, and CERNET-CNGI's (China Education and Research Network–China Next-Generation Internet's) IVI prototype. This is expected to help larger networks deploy new services using IPv6-only networks before they become able to get all of their existing users to turn on IPv6. As IPv6 becomes generally deployed, the need for translation disappears and one can expect the technology to disappear, overtaken by events.

No doubt, in the process of deployment, issues will arise, just as they do when deploying services using IPv4. Eventually, IPv6 deployment and coexistence implies a transition from IPv4 to IPv6; if there is reason to step into coexistence, eventually there will be users with only IPv6 service, and communicating with them will force other networks to follow.

Issues and Objections

Specific objections surround the business issues associated with IPv4-to-IPv6 transition and coexistence, including the costs of transition, the readiness of the protocol and its implementations, and the larger problems of routing and addressing.

Business Matters

If, as forecasts project, it becomes necessary to deploy IPv6 in order to obtain large amounts of address space inexpensively, then a company that fails to deploy IPv6 fails to offer connectivity to those new markets. Geoff Huston and others project that instead of wide-scale IPv6 deployment, a market for IPv4 address space will develop, with providers leasing or selling address space among themselves. In short, IPv6 connectivity is likely to be cheaper to deploy than IPv4 connectivity in the long term, and the revenue that connectivity brings is likely to be the same regardless of protocol. With IPv6, long-term profit potential is likely to be greater.

Operational and Capital Costs

As previously noted, one way to minimize business risk when deploying IPv6 is to deploy it in a network separate from the one running user services for IPv4. Running IPv4 and IPv6 at the same time will cost more than running either one alone.

Once IPv6 usage gets widely deployed in the network and IPv6 has been shown to be sufficient for the purpose, it would be wise to start removing IPv4 A and MX records in the DNS. This will enhance the use of IPv6 by taking IPv4 usage out of service, but it will do so without actually taking the service offline. If issues arise, restoring the DNS records restores IPv4 service. At some point, it should become clear that there is no IPv4 usage of the affected services, and IPv4 support is no longer a business requirement.

Protocol Implementation and Readiness

No doubt, in the process of deployment, issues will arise, just as they do when deploying services using IPv4. Eventually, IPv6 deployment and coexistence implies a transition from IPv4 to IPv6; if there is reason to step into coexistence, eventually there will be users with only IPv6 service, and communicating with them will force other networks to follow. In time, the business case for maintaining IPv4 connectivity will become questionable.

The Way Forward

As stated earlier, this author believes that IPv6 deployment will help address problems that the Internet is starting to experience and will experience in detail in a few years. From this author's perspective, the coexistence model is the least painful form of transition. It has monetary costs and other risks, but that is true of all transitions. This one also has a safety net built in, which more sudden approaches do not.

KENET: A Bandwidth Management Case Study

By Kevin Chege and Mat Ford

As part of ongoing efforts to better understand and respond to the variety of networking issues that can loosely be classed as stemming from bandwidth-intensive activities (see "The Internet and Bandwidth-Intensive Activities," *IETF Journal,* Volume 4 Issue 2), the Internet Society invited Kevin Chege of the Kenya Education Network (KENET) to attend IETF 73 in Minneapolis. Kevin has considerable personal experience dealing with the impact that bandwidth-intensive applications can have on a relatively low-bandwidth region of the global network. Attending the IETF meeting allowed Kevin to share his perspective with the engineers working in the newly formed ALTO (Application-Layer Traffic Optimization) and LEDBAT (Low Extra Delay Background Transport) working groups (WGs) and to learn from them how these activities might help him in his work.

Established in 1999, KENET connects educational institutions and research centres in Kenya with the goal of distributing knowledge throughout the country and making sure that the research and education communities have access to the Internet. Currently, there are 8 member institutions directly connected to the main node, which is in the capital city of Nairobi, and more than 40 additional member institutions that participate in the network by way of copper leased lines that are made available by Telkom Kenya, one of Kenya's main backbone providers. The main node consists of a 2 megabits-per-second uplink via a leased line, a 1 Mbps uplink via VSAT, and a 1.5 Mbps downlink via VSAT. The members' bandwidth usage typically ranges from 64 kilobits per second to 960 Kbps. Their leased lines terminate at KENET on E1 lines with a maximum 2 Mbps capacity. In some cases, members have their own VSAT downlinks due to the limitations of their leased line and use KENET only for uplink capacity.

To address some of the bandwidth limitations, KENET has begun working on a World Bank-sponsored project aimed at improving access for its member institutions by migrating from a copper-based infrastructure to fibre. This project is expected to reach completion by late 2009. Although the initial bandwidth purchase will be on VSAT, it is hoped that KENET will benefit from the price reduction and increased bandwidth availability that may be possible when undersea cables arrive in 2009. However, even when the cables become operational in mid-2009, it is expected that Internet bandwidth will be in short supply due to both financial pressures and limitations to the physical infrastructure.

The sort of problems and issues Kevin sees in his day-to-day work include WAN links from member institutions being saturated with P2P traffic, sometimes to the almost total exclusion of other traffic. This naturally stifles use of the Web for productive research by students and staff at the institutions served by KENET. Other important issues are related to (1) the lack of well-formulated IT policies that could help manage specific types of applications on the network, (2) the need for more detailed and widespread network monitoring, and (3) the need for better training of local network administrators, many of whom know what they want to achieve but are less clear on how to go about achieving it. This knowledge imbalance often leads to ineffective and, in some cases, counterproductive solutions being deployed.

One of the more successful strategies that has been adopted over time by KENET is regular training for staff on subjects like network management, security, and network monitoring using open-source tools. More-aggressive use of access-control lists and the widespread deployment of Web caches, spam, and antivirus filters have also helped mitigate many of the original issues. Detailed proposals have been developed to integrate a bandwidth-management and optimization tool for deployment at all member sites. This solution, which is yet to be developed, aims to provide a simplified interface to free, open-source tools, which will enable local network administrators to monitor and manage their local connectivity without needing highly specialised training beforehand. Having a standard solution for member network bandwidth management would also ease KENET's job of supporting the member networks remotely. It is hoped that such a solution could be applied to other upcoming National Research and Education Networks in the region, many of which could experience similar problems as they develop.

KENET and the member institutions want the flexibility to be able to push back on certain bandwidth-intensive applications at certain times of day, for example, when the network is under simultaneously high demand from researchers and staff.

KENET and the member institutions want the flexibility to be able to push back on certain bandwidth-intensive applications at certain times of day, for example, when the network is under simultaneously high demand from researchers and staff trying to do their research at the educational institutions. From this perspective, the sort of solu-

tion that the LEDBAT WG is working on should benefit KENET. The LEDBAT WG charter identifies a common scenario wherein applications experience large delays in the presence of P2P applications uploading over thin home uplinks; for KENET its entire WAN is peppered with thin uplinks. The potential benefits of a deployable solution in this space are therefore much greater than just the well-known use-case of an ADSL user simultaneously trying to run a P2P application and place a VoIP call. (It is less clear that the direction being pursued by the ALTO WG will be as useful for the KENET situation given that KENET's goal isn't so much keeping P2P traffic on-Net as it is ensuring P2P traffic gets out of the way whenever there are other, more-interactive applications trying to share the network resources. Because most content desired for download by P2P clients in Kenya is presumably off-Net, the sort of policy tools that an ALTO solution could provide may be ineffective in this instance.)

Speaking to this point directly at IETF 73, Kevin said he had observed a lot of engineers at the IETF making assumptions about the capacity of end-user connections based on their own, narrow domestic experiences. In reality, users in low-bandwidth networks have habits that are similar to those in high-bandwidth networks. They want to visit social networking sites, they want to download software and multimedia files, and they want to play online games, all of which may be bandwidth intensive and could effectively cripple the network for other users. The bandwidth-to-host ratio, poor-quality bandwidth and infrastructure, and, in most cases, the increased latency caused by the use of VSAT all make this a serious problem for network engineers in Kevin's home region. Bandwidth-management-related WGs at the IETF should ensure that their solutions consider and address the kinds of problems faced by KENET in its day-to-day network management role.

Revisiting Unwanted Traffic

By Leslie Daigle

In March 2006, the Internet Architecture Board (IAB) held an invitational workshop to look at the problem of unwanted traffic. The official workshop report was published as RFC 4948, and a more complete discussion of the implications appears in an article by Elwyn Davies that was published in the *IETF Journal* in December 2007 (Volume 3, Issue 3).

The workshop noted that the primary source of unwanted traffic comes from the so-called underground economy—that is, individuals who make use of the open systems of the Internet by leveraging hacked hosts and routing hardware to carry out activities for financial gain. Many of those activities, such as spam, stretch the bounds of civil behaviour; others are outright illegal, such as selling stolen credit card information.

Recognizing that the development cycle for new technologies was too long for specific development plans, the focus at the time was on mitigating strategies. It was never the plan to stop there, of course. Now, almost three years later, it is valuable to look back and see what pieces of development have occurred in the interim that will address some of the core issues of Internet security and stability.

Some of the vulnerabilities stated in RFC 4948 are the following:

"BGP route hijacking: in a survey conducted by Arbor Networks, route hijacking together with source address spoofing are listed as the two most critical vulnerabilities on the Internet. It has been observed that miscreants hijack bogon prefixes for spam message injections. Such hijacks do not affect normal packet delivery and thus have a low chance of being noticed."

"Everyone comes from Everywhere: in the earlier life of the Internet it had been possible to get some indication of the authenticity of traffic from a specific sender based for example on the Time To Live (TTL). The TTL would stay almost constant when traffic from a certain sender to a specific host entered an operators network, since the sender will 'always' set the TTL to the same value. If a change in the TTL value occurred without an accompanying change in the routing, one could draw the conclusion that this was potential unwanted traffic. However, since hosts have become mobile, they may be roaming within an operator's network and the resulting path changes may put more (or less) hops between the source and the destination. Thus, it is no longer possible to interpret a change in the TTL value, even if it occurs without any corresponding change in routing, as an indication that the traffic has been subverted."

"Packet source address spoofing: there has been speculation that attacks using spoofed source addresses are decreasing, due to the proliferation of botnets, which can be used to launch various attacks without using spoofed source addresses. It is certainly true that not all the attacks use spoofed addresses; however, many attacks, especially reflection attacks, do use spoofed source addresses."

Key areas of routing infrastructure security work are being pursued in the SAVI (Source Address Validation Improvement) and SIDR/RPSEC (Secure InterDomain Routing/Routing Protocol Security) working groups.

SAVI seeks to define a finer-grained mechanism for source IP address validation than ingress filtering. As described in its charter, "Partial solutions exist to prevent nodes from spoofing the IP source address of another node

Continued on next page

Revisiting Unwanted Traffic, continued from page 19

in the same IP link (e.g., the 'IP source guard'), but are proprietary. The purpose of the […] Working Group is to standardize mechanisms that prevent nodes attached to the same IP link from spoofing each other's IP addresses." The potential for mischief with such within-network ("on link") spoofing should not be discounted: compromised hosts could provide packets masquerading as critical infrastructure responses, such as spoofing gateway Address Resolution Protocol (ARP) packets and injecting false Dynamic Host Configuration Protocol (DHCP) responses, among others. Generally, having a finer granularity for validating source IP addresses and treating validation outcomes would be helpful in a number of situations. As part of network management policy options, depending on the situation, it might be desirable to block spoofed packets or merely log packets that appear to be spoofed. Therefore, the SAVI work is an important piece in the puzzle of preventing and mitigating unwanted traffic.

RPSEC was chartered to document the security requirements for routing systems and, in particular, to produce a document on Border Gateway Protocol (BGP) security requirements. Complementarily, the scope of work in the SIDR working group is to formulate an extensible architecture for an interdomain routing security framework and developing security mechanisms that fulfil requirements that have been agreed on by the RPSEC working group. The first order of business for SIDR is to develop an architecture and framework for a repository to allow formal validation of routing activities. This will take the form of an accessible database of formally verifiable descriptions of who has the right to use particular IP addresses or to announce routes for a given autonomous system. Deploying this will require cooperation among the Regional Internet Registries, network operators, and others. However, it is the necessary foundation for secure interdomain routing, source address verification (beyond individual network boundaries), and other important routing security mechanisms.

Key areas of routing infrastructure security work are being pursued in the SAVI (Source Address Validation Improvement) and SIDR/RPSEC (Secure InterDomain Routing/Routing Protocol Security) working groups.

These working groups are active, and clearly, their output will be useful in addressing some of the vulnerabilities identified by the IAB workshop. SAVI results could certainly help mitigate issues in address spoofing at all levels. Securing the routing infrastructure would make it considerably harder to inject bogus routes into the global routing fabric. It is hard to overestimate the importance of the eventual win from developing and deploying these technologies: route hijacking, both deliberate and accidental, has happened on a global scale, and the ramifications were felt up through "Layer 9" (global politics).

Another area of long-standing infrastructure security development is, of course, DNSSEC (DNS Security). While the technology has been available for quite a while, there is now some movement toward deployment. Some ccTLDs, such as .se, have deployed it. One gTLD, .org, has announced plans to deploy it. And there are discussions about whether, or how, to sign the root itself—a critical step toward ensuring the feasibility of a complete DNSSEC infrastructure and permitting authentication of DNS results. While DNSSEC itself does not prevent spam or phishing, it is a critical piece of infrastructure that provides the foundation for reliable, trustable infrastructure that will address those issues more directly.

While it is common to shrug and sigh at the inconvenience of unwanted traffic, it is an important area to address because it goes to the question of the Internet's evolution. The Internet has been developed, deployed, and built out based on a model of voluntary adherence to open standards and cooperative activity. That makes for an infrastructure that is immune to mandated change, which is both a feature and a challenge. The Internet is certainly no longer treated as the research network it was originally; it has become an integral part of the fabric of day-to-day lives, businesses, and civil organization in all developed—and many developing—countries.

To preserve the characteristics of operation that facilitate innovation at the edges, we need to demonstrate that the Internet can, in fact, evolve to meet the increased requirements: more security, moving from implicit trust in its operators to reasonable expectations of trust of the infrastructure itself.

Three years ago, the IAB workshop identified some key areas for concern. As noted earlier, the IETF's work has continued to develop more proactive, viable, infrastructure-securing technologies. The next, and perhaps most critical, step is to move toward global adoption and deployment of those technologies to address the network scourge that is unwanted traffic.

Photo by Peter Löthberg

Former IETF chairs Fred Baker (left) and Harald Tveit Alvestrand at IETF 73

Resource Certification

By Geoff Huston

Opinions vary as to what aspect of the Internet's infrastructure represents the greatest common vulnerability to the security and safety of Internet users, but it is generally regarded that attacks that are directed at the network's infrastructure are the most insidious, and in that case the choice is probably between the Domain Name System (DNS) and the inter-domain routing system.

The question of how to improve the robustness of these functions has been a long-standing topic of study. For the DNS it appears that there is convergence on DNSSEC (DNS Security) as the technical solution to securing DNS resolution operations, and the focus of attention in this space has shifted from technical behaviour to issues relating to operational deployment. It has been a long haul for DNSSEC, and to say there is an end in sight may well be premature at this stage, but there are definite signs of progress in this space. The same cannot be said of progress with securing routing—and particularly in securing inter-domain routing. Here there is still much to be done in order to achieve reasonable consensus on what technical measures to adopt, let alone the second step: the study of how such measures could be deployed across the Internet.

The IETF's approach to addressing the topic of securing inter-domain routing has followed a conventional IETF path. The first step has been to consider the nature of various vulnerabilities that exist within today's inter-domain routing system and then develop a set of requirements that should be addressed in any solution space—without necessarily defining what such a solution may be. Once the enumeration of requirements achieves a suitable level of consensus within the community, it will then be possible to commence work on standardising solutions. In the case of securing inter-domain routing, the first steps were undertaken in BoF sessions and in the subsequently formed Routing Protocol Security Requirements (RPSEC) working group. This work is almost complete, and apart from some definitive statement relating to a requirement for securing the autonomous system (AS) path attribute in BGP (Border Gateway Protocol), the set of requirements for securing inter-domain routing is now in a close-to-final state (draft-ietf-rpsec-bgpsecrec). The task of the Secure Inter-Domain Routing (SIDR) working group is to standardise technologies that can meet these requirements.

So where does resource certification come into the picture?

Public Key Cryptography

One commonly used security technology is public key cryptography. As long as a suitable amount of vague hand waving is used, the technique can be easily explained. The approach uses a pair of keys, A and B. Anything enciphered with key A can be deciphered only with key B and vice versa, and knowledge of the value of one key does not lead to discovery of the value of the other key. Key A is kept as a closely guarded secret, while key B is openly published. If I want to send you a message that only you can decipher and read, I should encrypt it using your public key. If I want to send you a message that only I could've sent (nonrepudiation), then I'll generate a digital signature of the message by using my private key. That way any attempts to alter the message will also be detectable.

This latter approach, of using keys to generate digital signatures of messages, lies at the heart of DNSSEC, because DNSSEC adds public keys and digital signatures to the DNS. A DNS query can generate a response that lists both the DNS answer and the digital signature of that answer. The DNS can also be queried to retrieve the public key that is used for signing all the components of that zone, so that the digital signature can be verified and the query agent can be assured that the response is a genuine one. But how can the key itself be verified? IN DNSSEC the hierarchical nature of the DNS itself is exploited by having each zone "parent" sign the keys of its delegated "children." So the zone key can be verified by retrieving the parent's signature across that zone key, and so on to the root of the DNS. As long as the query agent knows beforehand the value of the public key used to sign the root zone of the DNS and as long as DNSSEC is used universally, all DNS responses can be verified in DNSSEC.

The approach of using keys to generate digital signatures of messages lies at the heart of DNSSEC, because DNSSEC adds public keys and digital signatures to the DNS. But how can a key be itself verified?

While this approach works in the interlocked hierarchical structure of the DNS, when we turn our attention to securing the use of IP addresses and AS numbers in the context of inter-domain routing, then there is no comparable hierarchy to exploit. In such cases a common solution is to turn to digital certificates.

A digital certificate is a digitally signed public attestation by a certification authority that associates a subject's public key value with some attribute of the subject. A very typical application is in identity certification, where the certification authority is attestation that the

Continued on next page

Resource Certification, continued from page 21

holder of the private key whose matching public key is provided in the certificate has met the authority's certification criteria to be identified by a particular name. Digital certificates are useful because they are able to reduce the number of trust points in a security domain, so that each individual member of the domain does not have to validate identity and exchange public keys with every other member of the domain but can undertake a single transaction with a certification authority that is trusted by all the members of the domain. As long as every member of the domain carries the public key of the certification authority and can access all issued digital certificates, the members of the domain can verify each other's attestations and digital signatures.

Of course, digital certificates are used for far more than attestations of identity and can encompass the authority to perform specific tasks, undertake particular roles, or grant permissions and right-of-use authorities. It is that last-use case that is relevant to resource certification.

Resource Certificates

A resource certificate is a conventional X.509 certificate that conforms to the PKIX profile (RFC 5280) with one critical component—namely, a certificate extension that lists a collection of IP number resources (IPv4 addresses, IPv6 addresses, and AS numbers) (RFC 3779).

These certificates attest that by virtue of an associated resource allocation, the certificate's issuer has granted to the entity represented by the certificate's subject a unique right-of-use of the associated set of IP number resources listed in the certificate's extension. The unique right-of-use concept mirrors the resource allocation framework, where the certificate provides a means of third-party validation of assertions related to resource allocations (draft-ietf-sidr-arch).

By coupling the issuance of a certificate by a parent Certification Authority (CA) to the corresponding resource allocation, a test of a certificate's validity including the IP number resource extension can also be interpreted as validation of that resource allocation. Signing operations that descend from that certificate can therefore be held to be testable—under the corresponding hierarchy of allocation. In other words, if you received your address block from a particular Regional Internet Registry (RIR), then only that RIR can issue a resource certificate for you that includes your public key and the allocated number resources. Anything you sign using your private key can be verified via the RIR's issued certificate.

A resource certificate is a conventional X.509 certificate that conforms to the PKIX profile (RFC 5280) with one critical component—namely, a certificate extension that lists a collection of IP number resources (IPv4 addresses, IPv6 addresses, and AS numbers) (RFC 3779).

Unlike certificates that relate to attestations of identity, resource certificates are not necessarily long-lived. When an additional allocation action occurs, the associated resource certificate is reissued with an IP number resource extension that matches the new allocation state. In the case of a reduction in allocated resources, the previously issued certificates are explicitly revoked once the new certificate is issued. In other cases there is no explicit revocation of the older certificates.

The intention here is that any instrument signed by the subject's private key that relates to an assertion of resource control, whether it's a protocol message in a routing protocol or an administrative request to an ISP to route a prefix or an assertion of title over the right-of-use of a number resource, can be validated through the matching public key contained in the certificate and the IP number resource that are enumerated in this certificate. The resource certificate itself can be verified in the context of a resource certificate Public Key Infrastructure.

The Resource Certificate Public Key Infrastructure

The Resource Certificate Public Key Infrastructure (RPKI) describes the structure of the certification framework used by resource certificates. The intent of the RPKI is to construct a robust hierarchy of X.509 certificates that allows relying parties to validate assertions about IP addresses and AS numbers and their use.

The structure of the RPKI as it relates to public use of IP number resources is designed to precisely mirror the structure of the distribution of addresses and ASs in the Internet, so a brief description of this distribution structure is appropriate. The Internet Assigned Numbers Authority (IANA) manages the central pool of number resources. The IANA publishes a registry of all current allocations. The IANA does not make direct allocations of number resources to end users or Local Internet Registries (LIRs). Instead, it allocates blocks of number resources to the RIRs. The RIRs perform the next level of distribution: allocating number resources to LIRs, National Internet Registries (NIRs), and end users. NIRs perform allocations to LIRs and end users, and LIRs allocate resources to end users (figure 1).

The RPKI mirrors this allocation hierarchy. One interpretation of this model would see the IANA manage a root

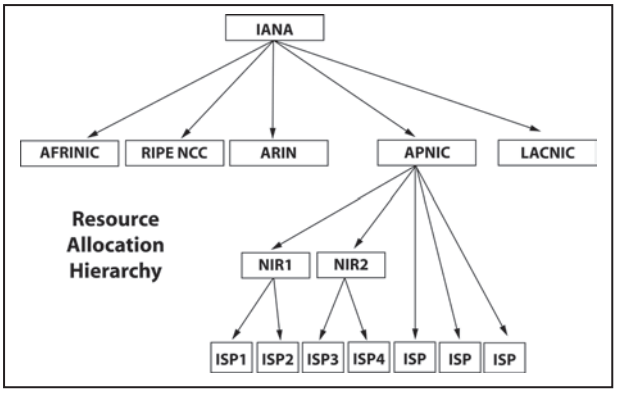

Figure 1. Address Distribution Hierarchy for the Internet

RPKI key, and using this key, the IANA would issue a self-signed root certificate and also issue subordinate certificates to each of the RIRs, describing in the resource extension to the certificate the complete set of number resources that have been allocated to that RIR at the time of issuance. The certificate would also hold the public key of the RIR and would be signed by the private key of the IANA. Each RIR would issue certificates that correspond to allocations made by that RIR, where the resource extension to those certificates lists all the allocated resources, and the certificate includes the public key of the recipient of the resource allocation, signed with the private key of the RIR. If the recipient of the resource allocation is an LIR or an NIR, then it too would issue resource certificates in a similar vein (figure 2).

The common constraint within this certificate structure is that an issued certificate must contain a resource extension that contains a subset of the resources that are described in the resource extension of the issuing authority's certificate. This corresponds to the allocation constraint that a registry cannot allocate resources that were not allocated to the registry in the first place. One implication of such constraint is that if any party holds resources allocated from two or more registries, then it will hold two or more resource certificates in order to describe the complete set of its resource holdings.

Validation of a certificate within this RPKI is similar to conventional certificate validation within any PKI—namely, establishing a chain of valid certificates that are linked by issuer and subject from a nominated trust anchor Certification Authority (CA) to the certificate in question. The only additional constraints in the RPKI are that every certificate in this validation path must be a valid resource certificate and that the IP number resources described in each certificate are a subset of the resources described in the issuing authority's certificate.

Within this RPKI, all resource certificates must have the IP addresses and AS resources present as well as marked as a critical extension. The contents of these extensions correspond exactly to the current state of IP address and AS number allocations from the issuer to the subject.

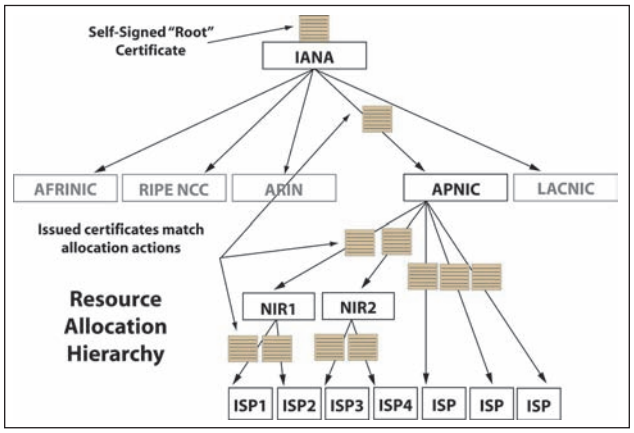

Figure 2. RPKI Resource Certificate Hierarchy

Any holder of a resource who is in a position to make further allocations of resources to other parties must be in a position to issue resource certificates that correspond to these allocations. Similarly, any holder who wishes to use the RPKI to digitally sign an attestation needs to be able to issue an End Entity (EE) certificate to perform the digital signing operation. For that reason, all issued certificates that correspond to allocations are certificates with the CA capability enabled, and each CA certificate is capable of issuing subordinate CA certificates that correspond to further sub-allocations and subordinate EE certificates that correspond to generation of digital signatures on attestations.

The RPKI makes conventional use of Certificate Revocation Lists (CRLs) to control the validity of issued certificates, and every CA certificate in the RPKI must issue a CRL according to the CA's nominated 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 the authority or attestation that was signed by a given EE certificate, the CA issuing authority that issued the EE certificate simply revokes the EE certificate.

Resource certificates are intended to be public documents, and all certificates and objects in the RPKI are published in openly accessible repositories. The set of all such repositories forms a complete information space, and it is fundamental to the model of securing the public Internet's inter-domain routing system that the entire RPKI information space be available. Other uses of the RPKI might permit use of subsets, such as the single chain from a given end-entity certificate to a Trust Anchor, but routing security is

Continued on next page

Resource Certification, continued from page 23

considered against all known publicly routable addresses and AS numbers, and so all known resource certification outcomes must be available. In other words, the RPKI's intended use in routing contexts is not a case where each relying party may make specific requests for RPKI objects in order to validate a single object, but one where each relying party will perform a regular sweep across the entire set of RPKI objects in order to ensure that the relying party has a complete picture of the RPKI information space. This aspect of the RPKI represents some interesting challenges, in that rather than having a single CA publish all the certificates produced in a security application at a single point, the RPKI permits the use of many publication points in a widely distributed fashion. Each CA is able to issue RPKI objects and publish them using a locally managed publication point. It is incumbent upon relying parties to synchronise a locally managed cache of the entire RPKI information space at regular and relatively frequent intervals. For that reason, the RPKI has introduced an additional mechanism in its publication framework—namely, the use of a manifest to enable relying parties to determine whether they have been able to retrieve the entire set of RPKI published objects from each RPKI repository publication point or whether there has been some attempt to disrupt the relying party's access to the entire RPKI information set. It also implies that the RPKI publication point access protocols should support the efficient function of a synchronisation comparison, so that a locally managed cache of the RPKI needs call for the uploading of only those objects that have been altered since the previous synchronisation operation.

Signed Attestations and Authorities

The underlying intent of digital certificates—and resource certificates in particular—is in terms of supporting a transitive trust relationship that allows a relying party to verify the authenticity of a signed artefact through verification of the signer's key using the PKI. So the obvious question is, What artefacts are useful to sign?

Much of the motivation for resource certificates has come from a desire to underpin efforts in securing aspects of inter-domain routing. This goes well beyond securing the individual point-to-point connection used between BGP speakers and refers to the matter of verifying the authenticity of the payload of the BGP exchange. The specific question that may be posed is, How can a BGP speaker validate the authenticity of the route object being presented to it?

The underlying intent of digital certificates—and resource certificates in particular—is in terms of supporting a transitive trust relationship that allows a relying party to verify the authenticity of a signed artefact through verification of the signer's key using the PKI. So the obvious question is, What artefacts are useful to sign?

The approach being studied by the SIDR working group is to use structured attestations, where, like the digital certificate itself, the attestation is structured in an ASN.1 digital object, and this object is signed using a signing formation that is itself a piece of structured ASN.1—namely, the Cryptographic Message Syntax (CMS) (RFC3852).

The first of these attestations relates to the ability to verify the authenticity of the "origination" of an inter-domain routing object. This refers to the address prefix and the originating AS, and the questions that this verification function is intended to answer are:

Is this a valid address prefix and AS number? Have these resources been allocated through the IP number resource allocation process?

Has the holder of the title of right-of-use" for the address prefix authorised the AS holder to originate a routing advertisement for this prefix?

Here an address holder is authorising a particular ISP to generate a route announcement for the address holder's particular address prefix. In this case, the prefix holder would generate an EE resource certificate with the IP number resource extension spanning the set of addresses that match the address prefixes that are the intended subject of the routing authority and would place validity dates in the EE certificate that correspond to the intended validity dates of the routing authority. The signed authority document would contain the Autonomous System number that is being authorised in this manner; a description of the range of prefixes that the prefix holder has authorised; and the EE certificate. The document would be signed by the EE certificate's private key by using a CMS signing structure. The resultant object is published in the RPKI distributed publication repository as a Routing Origin Authorization (ROA). A relying party can validate the ROA by checking that the digital signature in the ROA is correct, indicating that the authority document has not been tampered with in any way since it was signed, that the resources in the associated EE certificate encompass the prefixes specified in the document, and that the EE certificate itself is valid in the context of the RPKI by verifying that there is an issuer/subject chain of valid certificates that link one of the relying party's nominated Trust Anchors to the EE certificate.

The ROA itself is valid as long as the signing EE certificate is valid. To withdraw the authority prior to the expiration of the EE certificate, the ROA publisher can simply revoke the EE certificate. This leads to the concept of one-off-use EE certificates in the RPKI, where a key pair and a corresponding EE certificate are generated in order to sign a single attestation or authority. If the authority's lifetime is extended, the authority is reissued with a new EE certificate and with a new digital signature; and, as noted, the authority can be prematurely terminated through revocation of the EE certificate, so that at no stage is there a need to reuse the original signing private key. Once the private key has been used to sign this object, the key is destroyed, alleviating to some extent the total key management load.

In any security system, knowledge of what is authorised is helpful, but knowledge of what has not been authorised is perhaps even more helpful. For ROAs there is a situation analogous to DNSSEC, where DNSSEC is most effective from a client's perspective once the entire DNS space is DNSSEC signed. Where there are gaps in the DNSSEC signing chains, the client is left in an uncertain state regarding the verification outcomes of the unlinked DNS subhierarchies. The same could apply to ROAs, in that in an environment where not every originated route object has a published ROA, the absence of an ROA does not necessarily indicate an unauthorized route origination. If one of the objectives of this study is to define a framework that can unambiguously identify the unauthorized use of IP number resources in routing (route hijacks) even in a world where ROAs are used in a piecemeal fashion, then one possible refinement to the ROA model is the introduction of a comparable negative authority, the Bogon Origin Attestation (BOA).

In this case, the prefix holder generates a signed attestation, or BOA, in a manner similar to the ROA but does not provide any originating AS. Instead, the BOA refers to "all originating ASs" and has the semantic interpretation that any use in the routing space of this address prefix described in the BOA, or any more specific address prefix, should be regarded as unauthorized and the route should be discarded.

While this makes the detection of route hijacks more direct in a world of piecemeal use of ROAs, there is now the added complication of having both positive and negative authorities. The proposed resolution of this is to use a relative priority rule that ROAs take precedence over BOAs, so that if both a valid ROA and a valid BOA exist that describe the origination component of a route, then the route can be regarded as authorised.

It should be noted, however, that at this stage these concepts are works in progress, and are part of the SIDR working group's agenda of study, and the working group has not as yet reached any consensus position regarding the decision to advance these proposals onward along the Internet Standards Process.

Also on the near-term horizon, SIDR is examining approaches to secure the AS Path in BGP updates. The RPSEC (Routing Protocol Security Requirements) working group has explored two approaches in this space. One involves an incremental multiple-signature technique that allows a receiver of a BGP update to verify that the AS path described in the update is matched by a sequence of interlocking AS digital signatures using the RPKI. At the same time as an AS adds its own AS to the AS path prior to further eBGP propagation of the route update, the AS would digitally sign over an analogous sequence of AS signatures. This approach allows a receiver to perform a match of the AS sequence in the AS Path with the AS number sequence identified in the AS signature block. A match here would indicate that the BGP update has indeed been sequentially passed along the sequence identified by the AS Path. This approach was originally proposed in the sBGP design and has attracted some comment related to the computation overhead associated with the application and validation of these AS Path signature sequences. An alternative approach has been one that is described by RPSEC as being less rigorous and refers to a "feasibility" check that checks that each pair of ASs represented in the AS Path has an associated verifiable assertion of inter-AS adjacency that is digitally signed by both ASs.

It should also be noted that this activity of addressing aspects of improving the robustness of inter-domain routing has some previous context. In many parts of the Internet, some degree of routing integrity is managed through the use of Internet Routing Registries (IRRs) and the publication of routing policies through the use of Routing Policy Specification Language (RPSL) objects. While opinions vary as to the robustness of the security offered by the IRR approach, at the very

Continued on next page

In many parts of the Internet, some degree of routing integrity is managed through the use of Internet Routing Registries (IRRs) and the publication of routing policies through the use of Routing Policy Specification Language (RPSL) objects. While opinions vary as to the robustness of the security offered by the IRR approach, at the very least it can mitigate some weakness in the routing system through the use of a second check that can be used to filter the information that is being provided in a BGP feed.

Resource Certification, continued from page 25

least it can mitigate some weakness in the routing system through the use of a second check that can be used to filter the information that is being provided in a BGP feed. The weaknesses in the IRR system tend to relate to the consistency, completeness, and authenticity of the IRR data. In many cases, trust in the integrity of the data relies on the admission practices of the IRR itself, and individual data objects cannot be verified by clients of the IRR. One possible way to address this has been through the use of Routing Policy System Security (RPSS) measures, but the adoption of these measures has not been widespread, and the question still remains for the client that even if an IRR object was authenticated upon admission, it does not mean that when the object is subsequently used by an IRR client, the information reflects the current situation, and the information could well be invalid or not reflect the current policies of the IRR object's author.

One possible approach, being considered by the SIDR working group, is to implement the RPSS authentication models by using object signing in the context of the RPKI. For example, the RPSS assumption that routes should be announced only with the consent of the holder of the origin AS number of the announcement and with the consent of the holder of the address space implies in RPSS that both parties should authorise the entry of a route object into the IRR. Translating this into an analogous model by using the RPKI would require that a route object be signed with the digital signatures of both the AS holder and the address space holder, and a IRR client can verify this route object at the time of use by verifying both digital signatures. Either the address space holder or the AS holder can revoke its authorisation by revoking the EE certificate used to sign the route object, and the verification is independent of the particular IRR that has published the route object. It's also a possibility that the IRR itself can be folded into the RPKI distributed publication repository framework, as there is no particular requirement in such an environment for a disparate collection of IRRs with their own partial collections of routing policy information, although at this stage this is heading into the realm of more advanced speculation about the potential for application of resource certificates and digital signatures to RPSL and the IRR framework.

The various approaches to routing security studied so far offer a wide variety of outcomes…. The next step appears to include the task of establishing an appropriate balance between the overheads of operating the security framework and the extent to which efforts to disrupt the routing system can be successfully deflected by such measures.

Putting Resource Certificates into Context

Resource certificates and the associated RPKI represent a major part of any effort to construct a secure inter-domain routing framework. An RPKI, even partially populated with signed information, allows BGP speakers to make preferential selections to use routing information where the IP address block and the AS numbers being used are recognised as valid to use and that the parties using these IP addresses and AS numbers are properly authorised to so do. The RPKI can also be used to identify instances of unauthorized use of IP addresses and attempts to hijack routes.

However, the RPKI represents only one part of a larger framework of securing inter-domain routing, and the next step is that of applying the RPKI to the local BGP processing framework. There is also the need to move beyond validation of route origination and look at the associated issue of validation of the AS Path and potentially to consider the most challenging task: that of attempting to validate whether the initial forwarding decision associated with a route object actually represents the correct first hop along a usable forwarding path for packets to reach the network destination.

The issues here include not only a consideration of what can be secured and validated but also issues of scalability and efficiency in terms of deployment cost. The various approaches to routing security studied so far offer a wide variety of outcomes in terms of the amount of routing information that is validated, the level of trust that can be placed in a validation outcome, and the overheads of generating and validating digital signatures on routing information. The next step appears to include the task of establishing an appropriate balance between the overheads of operating the security framework and the extent to which efforts to disrupt the routing system can be successfully deflected by such measures.

The RPKI has been designed as a robust, simple framework. As far as possible, existing technologies and processes have been exploited, reflecting to some extent a level of conservatism on the part of the routing community and the difficulty in securing widespread acceptance of novel technologies.

References and Further Reading

The following documents provide further detail about the IETF work on resource certification. The Internet-Drafts listed here are still work in progress, and while they are reflective of the areas of activity of the SIDR working group, they do not necessarily represent finished work.

Internet-Drafts

Requirements

[draft-ietf-rpsec-bgpsecrec] BGP Security Requirements, B. Christian, T. Tauber, eds., work in progress, Internet-Draft, draft-ietf-rpsec-10.txt, November 2008. The report of the consensus outcomes of the RPSEC working group in enumerating the requirements for securing inter-domain routing. The outstanding topic in this report remains in the area of AS Path validation and the level of requirement associated with the two approaches described in the report.

Architecture

[draft-ietf-sidr-arch] An Infrastructure to Support Secure Internet Routing, M. Lepinski, S. Kent, work in progress, Internet-Draft, draft-ietf-sidr-arch-04.txt, November 2008. An overview of the RPKI approach, describing the RPKI, the distributed repository structure, and common operations.

Resource Certificates

[draft-ietf-sidr-res-certs] A Profile for X.509 PKIX Resource Certificates, G. Huston, G. Michaelson, R. Loomans, work in progress, Internet-Draft, draft-ietf-sidr-res-certs-15.txt, November 2008. The specification of the Resource Certificate.

RPKI Repository Structure

[draft-ietf-sidr-repos-struct] A Profile for Resource Certificate Repository Structure, G. Huston, G. Michaelson, R. Loomans, work in progress, Internet-Draft, draft-ietf-sidr-repos-struct-01.txt, October 2008. A description of the proposed distributed publication repository structure for the RPKI, including contents, access protocols, and object name conventions.

[draft-ietf-sidr-rpki-manifests] Manifests for the Resource Public Key Infrastructure, R. Austein et al., work in progress, Internet-Draft, draft-ietf-sidr-rpki-manifests-04.txt, October 2008. A specification for repository manifests. Manifests are signed constructs that describe all the objects currently loaded into a repository publication point and are used by relying parties as a means of ensuring that a local RPKI repository cache is correctly synchronised against the authoritative original publication point.

[draft-ietf-sidr-rescerts-provisioning] A Protocol for Provisioning Resource Certificates, G. Huston, R. Loomans, B. Ellacot, R. Austein, work in progress, Internet-Draft, draft-ietf-sidr-rescerts-provisioning-03.txt, August 2008. A proposed protocol for use between a subject and a certificate issuer to ensure that certificate requests, the IP number resource allocation state, and the issued certificate status are correctly synchronised. This extends the conventional certificate request model into a transaction protocol that also includes the ability to perform certificate revocation requests and status queries from the subject.

RPKI Signed Objects

[draft-ietf-sidr-roa-format] A Profile for Route Origin Authorizations (ROAs), M. Lepinski, S. Kent, D. Kong, work in progress, Internet-Draft, draft-ietf-sidr-roa-format-04.txt, November 2008. The specification of the syntax for signed ROAs.

[draft-ietf-sidr-bogons] A Profile for Bogon Origin Attestations (BOAs), G. Huston, T. Manderson, G. Michaelson, work in progress, Internet-Draft, draft-ietf-sidr-bogons-02.txt, October 2008. The specification of the syntax for signed BOAs.

[draft-ietf-sidr-roa-validation] Validation of Route Origination in BGP Using the Resource Certificate PKI, G. Huston, G. Michaelson, work in progress, Internet-Draft, draft-ietf-sidr-roa-validation-01.txt, October 2008. The specification of the semantics of ROAs and BOAs and the manner in which these objects may be interpreted in terms of the integration of these origination security credentials onto a BGP route selection process.

Certificate Policy and Practice Statements

[draft-ietf-sidr-cp] Certificate Policy (CP) for the Resource PKI (RPKI), K. Seo, R. Watro, D. Kong, S. Kent, work in progress, Internet-Draft, draft-ietf-sidr-cp-04.txt, November 2008. A description of the certificate policy that applies to all certificates issued within the RPKI framework.

Continued on next page

Resource Certification, continued from page 27

[draft-ietf-sidr-cps-irs] Template for an Internet Registry's (IR's) Certification Practice Statement (CPS) for the Resource PKI (RPKI), D. Kong, K. Seo, S. Kent, work in progress, Internet-Draft, draft-ietf-sidr-cps-irs-04.txt, November 2008. A template for the Practice Statement used by IRs to describe their operational practices in the issuance and management of resource certificates.

[draft-ietf-sidr-cps-isp] Template for an Internet Service Provider's Certification Practice Statement (CPS) for the Resource PKI (RPKI), D. Kong, K. Seo, S. Kent, work in progress, Internet-Draft, draft-ietf-sidr-cps-isp-03.txt, November 2008. A template for the practice statement used by ISPs to describe their operational practices in the issuance and management of resource certificates.

Individual Submissions

[draft-huston-sidr-aao-profile] A Profile for AS Adjacency Attestation Objects, G, Huston, G. Michaelson, work in progress, Internet-Draft, draft-huston-sidr-aao-profile-00.txt, September 2008. The specification of the syntax for a pairwise inter-AS routing adjacency attestation.

[draft-kisteleki-sidr-rpsl-sig] Securing RPSL Objects with RPKI Signatures, R. Kisteleki, J. Boumans, work in progress, Internet-Draft, draft-kisteleki-sidr-rpsl-sig-00.txt, October 2008. The specification of the addition of RPKI digital signatures to RPSL Objects in the context of an Internet Routing Registry.

[draft-manderson-sidr-fetch] RPKI Repository Retrieval Mechanism, T. Manderson, G. Michaelson, work in progress, Internet-Draft, draft-manderson-sidr-fetch-00, October 2008. A proposed mechanism to use the manifest as the basis of performing a synchronisation operation between a local RPKI cache and a source point.

RFCs

[RFC 5280] Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile, D. Cooper et al., RFC 5280, May 2008.

[RFC 3779] X.509 Extensions for IP Addresses and AS Identifiers, C. Lynn, S. Kent, K. Seo, RFC 3779, June 2004.

[RFC 3852] Cryptographic Message Syntax (CMS), R. Housley, RFC 3852, July 2004.

[RFC 2622] Routing Policy Specification Language (RPSL), C. Alaettinoglu et al., RFC 2622, June 1999.

[RFC 2725] Routing Policy System Security, C. Villamizar et al., RFC 2725, December 1999.

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/DocProtoActions0403.shtml.


Aaron Falk, IRTF Chair

IRTF Report

By Aaron Falk

Here are a few items on Internet Research Task Force (IRTF) developments since IETF 72:

- The composition of the Internet Research Steering Group (IRSG) has been made public and can be seen at http://www.irtf.org/chair.
- The document defining the IRTF publication stream has been finalized and is enqueued at the RFC Editor.
- Four IRTF RFCs have been published since IETF 72. Three are from the Delay-Tolerant Networking Research Group (dtnrg) on Licklider Transmission Protocol (LTP), and one is from the Network Management Research Group on Simple Network Management Protocol measurements.
- Draft-irtf-asrg-dnsbl (DNS Blacklists and Whitelists), which received substantial discussion on the IETF mailing list leading up to IETF 72, will be published as an IRTF RFC to document current practices. In response to the IETF list comments, it will include security considerations reflecting the IETF last-call comments.

In a departure from the usual IRTF reporting, the presentation at IETF 73 included introductions to two IRTF Research Groups: the Delay-Tolerant Networking Research Group (dtnrg) and the Peer-to-Peer Research Group (p2prg).

Delay-tolerant networking (DTN, sometimes called disruption-tolerant or disconnection-tolerant networking) is based on a model that makes no assumption that the sender and receiver are concurrently connected to the network. Data transfer is achieved using a multiparty communications model, wherein helpers provide a store-carry-forward mechanism for messages using extensible (Uniform Resource Identifier–based) naming. DTN supports multipath routing and caching to address connectivity disruptions. The dtnrg has been encouraging development, testing, and deployment of DTN protocols.

The exciting news from dtnrg is that the Bundle Protocol (BP) is now being tested in space. The code was uploaded to the United Kingdom's Disaster Monitoring Constellation satellite, which is operated by Surrey Satellite Technology Ltd. (SSTL). SSTL, Cisco Systems, and NASA Glenn conducted a test by downloading a 150-megabyte fragmented image of Earth by using the dtnrg's BP. In addition, NASA's Jet Propulsion Laboratory recently completed an experiment with DTN that involved using BP and LTP out to a distance of 15 million to 25 million kilometres. A report on results was presented at the dtnrg session.

The IRTF report at the IETF 73 plenary also included an introduction to the Peer-to-Peer Research Group. Peer-to-peer (P2P) networks exhibit a symmetric relationship between hosts. They are distributed, they scale to large numbers of nodes and users, they are autonomous, and they support anonymity. Well-known P2P systems include Usenet and BGP. BitTorrent and Skype are examples of new P2P systems.

The p2prg has been developing a new charter that will encourage (1) research on P2P/network traffic optimization (beyond ALTO (Application-Layer Traffic

Continued on next page

IRTF Update, continued from page 29

Optimization)); (2) security, privacy, anonymity, and trust; (3) improving interoperation between different P2P systems; (4) information storage, reliability, and retrieval in P2P systems; and (5) gaining a better understanding of P2P system performance and user behaviour "in the wild."

Additionally, five research groups met during the week of IETF 73. The dtnrg was summarized earlier, and notes from the other research group meetings follow.

Host Identity Protocol Research Group (hiprg)

The HIPRG met at IETF 73 and discussed two individual submissions. The first was a proposal to generalize HIP to include object-to-object (and not strictly host-to-host) communications. The second draft proposed to carry geolocation data explicitly in the HIP protocol. The meeting also received implementation updates from Boeing's HIP-based overlay deployment and the HIP for Linux project's recent work.

IP Mobility Optimizations Research Group (mobopts)

The meeting focussed on current documents, including IP Location Privacy, Multicast Mobility, and Media Independent Preauthentication. The latter two are completing research group last call, and the IP Location Privacy document is in IRSG review.

Also discussed at the meeting was a framework for benchmarking mobility models. The purpose is to be able to evaluate models used in literature by using a common reference that outlines what to look for. In addition, there was lively debate over what approach to take for multicast mobility solutions. One camp argued that multicast needs to be extended, while the other contended that without extensions to mobility protocols, handover would not provide the performance necessary for multicast traffic.

Internet Congestion Control Research Group (iccrg)

The ICCRG meeting covered two main topics. The first was an exploration of alternative start-up mechanisms, and the second was the basis of Internet congestion control on TCP–friendly sending rates. Two sets of experimental data were presented on the evaluation of different start-up mechanisms that attempt to improve on standard TCP slow start. Matt Mathis then led a discussion on Rethinking TCP-Friendly, and the research group accepted that it should write a vision statement document to increase architectural discussion on whether TCP friendliness should still be used as a strict criterion for evaluation of new Internet protocols.

Routing Research Group (rrg)

The Routing Research Group met to continue to discuss next-generation routing architectures. The research group heard six separate presentations on different aspects of a new architecture and continued to discuss the many alternatives at hand. The group plans to draft a preliminary recommendation for an architecture in the coming months.

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

IETF 73 Acknowledgements and Interim Meetings

Many thanks to Google, Inc., and its staff, the VeriLAN Team, AMS, and the following volunteers for making IETF 73 a great success.

Volunteers: Chris Elliott, Cisco Systems Inc., David Farmer, University of Minnesota, Bill Fenner, Arastra, Inc., Joel Jaeggli, Nokia, Bill Jensen, University of Wisconsin Madison, Jim Martin, The Daedelus Group, Robert Nagy, DeepDive Networking, Karen O'Donoghue, U.S. Navy, Daniel Westacott, University of Minnesota

The Internet Society would like to express its deepest gratitude to its **Organization Members,** whose contributions directly support the work of the **Internet Engineering Task Force.** Organizations, businesses, and nonprofit organizations that are interested in offering financial support to the IETF, or in hosting an upcoming IETF meeting, are encouraged to contact **Drew Dvorshak** by phone at +1 703 439 2129 or by e-mail at dvorshak@isoc.org.

Interim Meetings

The following meetings were held between IETF 72 and IETF 73

- Behavior Engineering for Hindrance Avoidance, 1–2 October 2008
- Internet Area Open Meeting, 1–2 October 2008
- NETCONF Data Modeling Language, 8–10 October 2008
- Routing Over Low power and Lossy networks, 6 October 2008
- Softwires, 1–2 October 2008
- IPv6 Operations, 1–2 October 2008

Photos by Kevin Craemer, Peter Löthberg, and Wendy Rickard

IETF Meeting Calendar

IETF 74
22–27 March 2009
Host: Juniper Networks
Location: San Francisco, CA, USA

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

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

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

Register now for

IETF 74

22–27 March 2009

San Francisco, CA, USA
http://ietf.org/meetings/74/

Early bird registration: USD 635 (through Friday, 13 March 2009)
Regular registration: USD 785
Full-time students: USD 150 with on-site proof of ID

IETF 74 is being hosted by Juniper Networks

Special thanks to Google for hosting IETF 73

Special thanks to Juniper Networks for hosting IETF 74

The ISOC Fellowship to the IETF is sponsored by

Afilias Google intel Microsoft

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

Afilias Alcatel·Lucent ARIN Cisco Comcast RIPE NCC

IETF Journal
IETF 73
Volume 4, Issue 3
February 2009

Published three times a year by the Internet Society

4 rue des Falaises
CH–1205 Geneva
Switzerland

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

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

Internet Society