



Published by the Internet Society in cooperation with the Internet Engineering Task Force

Inside this issue

- Welcome 1
- Farewell from
the Retiring
IETF Chair 2
- IETF 68
Facts and Figures2
- Words from
the Incoming
IETF Chair.....3
- Farewell from
the Outgoing
IAB Chair 4
- ISOC Fellowship
Program4
- Words from the
Incoming IAB
Chair5
- New BoF
Meetings5
- Plenary Report6
- Update on DNS.....13
- Recent IESG
Document and
Protocol Actions 14
- More ROAP.....15
- IPv6 Update22
- DHCPv6
Bake Off22
- Not Being There.....24
- IRTF Report26
- Calendar28

Editor's Welcome

By *Mirjam Kühne*

This issue of the IETF Journal covers the meetings and discussions from IETF 68, which was held in March 2007 in Prague. The meeting was especially notable due to the many changes among the IETF leadership. Brian Carpenter, whose term as IETF chair ended in March, turned over the reins to veteran IETFer Russ Housley. Leslie Daigle, whose longtime position as chair of the Internet Architecture Board (IAB) also ended, handed the chairmanship over to Olaf Kolkman. As a result, in this issue you will hear from both the outgoing and incoming chairs, who reflect on their experiences and offer a few words about their view of the future. A number of other personnel changes were recognised at the meeting, including the naming of Kurtis Lindqvist as new chair of the IETF Administrative Oversight Committee (IAOC).

Perhaps the most significant topic of discussion at IETF 68 was the one that covered the problems associated with routing and addressing, known as ROAP, a subject that is taken very seriously by the IETF and the IAB. The Internet Research Task Force (IRTF) has stepped up its efforts with regard to ROAP and rechartered the Routing research group (rrg). A summary of all ROAP activities at IETF 68 can be found in the article "More ROAP" on page 15. A summary of the ROAP discussion at the plenary can be found on page 7. Another important topic of discussion during the technical part of the plenary session was internationalisation in the context of IETF work. A summary of that discussion appears on page 10 as part of the plenary report.



Atrium of the Hilton Hotel in Prague

Photo by Peter Lofberg

Here you will also find summaries of developments around the Domain Name System (DNS) (page 13) and IPv6 (page 22). It's interesting that most activities related to IPv6 are now incorporated in other working groups or have become operational issues. There are very few specific IPv6 working groups left.

Those who are not able to attend an IETF meeting but would still like to follow one remotely might be interested in reading Geoff Huston's "Not Being There" (page 24), an analysis and evaluation of the remote participation facilities.

As always, we wish you fun reading, and we welcome your comments as well as your contributions for future issues of this publication.



Words from the Retiring IETF Chair

By **Brian Carpenter**

IETF 68 was held in Prague, capital of the Czech Republic, in a very modern and spacious Hilton hotel. We were hosted by NeuStar, with additional support from CZNIC and CESNET. Local loop was provided by Dial Telecom, and for the first time, the site network was subcontracted to VeriLAN Networks. As always, the success of the event was due largely to an outstanding team of dedicated volunteers. We had excellent wireless networking throughout the week. Approximately 1,200 people from 45 countries attended. The week featured the usual mix of working group meetings, BoF (birds-of-a-feather) sessions, research groups, and formal and informal side meetings.

It was especially interesting to hear from Jon Lindberg, vice president of Secretariat Services at NeuStar, about his company’s motivation for supporting the IETF.

“As technology becomes more advanced, and protocols become more sophisticated, and service offerings continue to become more and more robust, and end-user expectations continue to increase, NeuStar’s reliance on IETF standards is absolutely essential in order to be able to continue to deliver successful and unified solutions,” Jon said.

Since IETF 67, three new WGs were chartered and six WGs were closed, leaving approximately 120 WGs currently chartered. Between the meetings, the WGs and their individual contributors produced 441 new drafts, not to mention 1,020 updates. The Internet Engineering Steering Group (IESG) approved 130 drafts for publication as RFCs, and the RFC publication queue was stable. The RFC Editor hit a new record by publishing 459 RFCs during 2006.

As I step down after two exciting and rewarding years as chair of the IETF, I’d like to thank all of the individuals who personally helped me do this job. I can easily identify well over 120 such people, and that’s without even counting the working group chairs and document authors who do so much to make the IETF productive. If we succeed, it’s in a spirit of open cooperation between hundreds of people. It remains only for me to wish every success to Russ Housley as he carries the work forward.

I look forward to seeing many of you in Chicago, 22–27 July 2007, and after that in Vancouver, Canada, 2–7 December 2007. 



Brian Carpenter, outgoing IETF Chair

IETF 68 Facts and Figures

Registered attendees	1129
Countries	45
New WGs	3
Closed WGs	6
New Internet-Drafts	441
Updated Internet-Drafts	1020
IETF Last Calls	119
Approvals	130

RFC Editor Actions (Nov. 2006–Feb. 2007)

- 95 RFC published of which
- 58 standards track or BCP
 - 27 Informational or Experimental
 - 10 from other sources

IANA Actions (Nov. 2006–Feb. 2007)

- Processed 1160 IETF-related requests of which:
- 796 Private Enterprise Number requests
 - 81 port requests
 - 16 MIME-type requests

Reviewed 300 I-Ds in Last Call or IESG Review



Russ Housley, IETF Chair

Words from the Incoming IETF Chair

By Russ Housley

At IETF 68, I was honoured to accept the position of IETF chair. I have the privilege of standing on the shoulders of the giants who came before me:

2005–2007	Brian Carpenter
2001–2005	Harald Alvestrand
1996–2001	Fred Baker
1994–1996	Paul Mockapetris
1986–1994	Phillip Gross
1986	Michael Corrigan

The mission of the IETF is to make the Internet work better. However, no one is “in charge” of the Internet. Instead, many people cooperate to make it work. Each person offers a unique perspective of the Internet, and such diversity of perspective sometimes makes it difficult to reach consensus. Yet once consensus has been achieved, the outcome is better, clearer, and more strongly supported than the initial position of any one participant.

As Security Area director, my focus was on the continuous incremental improvement of the security of the Internet. My focus as IETF chair must be broader. I will focus on continuous incremental improvement of all aspects of the Internet, as well as on continuous incremental improvement of the IETF standards development process.

I look forward to IETF 69 in Chicago on 22–27 July 2007 and to IETF 70 in Vancouver, Canada, on 2–7 December 2007. Scheduling information for the next IETF meetings may always be found via <http://www.ietf.org/meetings/meetings.html>. I hope to see you there. 

As a longtime researcher and as founder of Vigil Security, LLC, in Herndon, Virginia, Russ Housley is no stranger to Internet security or the standards-development process. As the new IETF chair, Russ brings 25 years' experience in security protocols, certificate management, cryptographic key distribution, and high-assurance design and development practices. Prior to accepting the IETF chair position, Russ served as Security Area director, and prior to that he chaired the Secure MIME (S/MIME) working group. His past work experiences include positions with the Air Force Data Services Center (AFDSC), Xerox Special Information Systems (XSYS), SPYRUS, and RSA Laboratories.

Russ served as editor for several cornerstone Internet public key infrastructure (PKI) standards, including RFC 3280. In November 2004, Russ was recognised by the IEEE 802.11 working group for his contributions to IEEE 802.11i-2004, which fixes the severe security shortcoming of the Wired Equivalent Privacy (WEP). He is co-author of Planning for PKI: Best Practices Guide for Deploying Public Key Infrastructure (John Wiley & Sons, 2001) and is listed as an author on 36 RFCs. Russ received a B.S. in computer science from Virginia Tech in 1982 and an M.S. in computer science from George Mason University in 1992.

Farewell from the Outgoing IAB Chair

By *Leslie Daigle*

As I've steadfastly refused to grow a grey beard, I will claim I'm too young to take full retirement! But, after five years of chairing the Internet Architecture Board (IAB), I am pleased to step down and pass the task into the capable hands of Olaf Kolkman.

Any number of jokes can be made about the A in IAB standing for administration or appeal—and it is a fact that the IAB has several roles to fulfil in those areas. However, the IAB has remained true to its mandate to provide oversight of the architecture for the protocols and procedures used by the Internet. As the scope of the Internet grows, so does the architectural oversight challenge. Over the past five years, we've worked to meet that growing challenge by leveraging the 13-member IAB to provide leadership and guidance of architectural discussions in the open consensus environment of the IETF.

The challenge of providing focused architectural leadership is not going to lessen in the coming years, particularly as the Internet hits some growing pains, such as the current discussions about routing and addressing issues and internationalisation. I am confident that Olaf will provide the thoughtful guidance needed for the IAB to take those issues on and provide oversight, even as the reality remains that the Internet is for, and built by, everyone: we can all be contributing to the discussions of sound Internet architecture. 



Leslie Daigle, outgoing IAB Chair

ISOC Fellowship Program

The five recipients of the Internet Society's Fellowship to the IETF completed their participation in March, following IETF 68, which was held in Prague.

The fellows selected to participate in IETF 68 were chosen from among more than 80 well-qualified applicants from the developing world. Each fellow received full funding to travel to Prague and participate in IETF 68. In addition, each fellow was paired with a mentor from the community of experienced IETF participants.

Fellow Jimmy M. Kimanzi is network administrator with Swift Global in Kenya. He was mentored by Dave Meyer of Cisco Systems. Fellow Hichem Maalaoui is head of technical innovation and new projects at the Tunisian Internet Agency in Tunisia. His mentor was Hesham Soliman of Elevate Technologies.

Fellow Alberto F. Martínez is research assistant at Tecnológico de Monterrey in Mexico. He was mentored by Shane Kerr of the Internet Systems Consortium. Fellow R. Lakmal Silva is from Sri Lanka and is a graduate student at Blekinge Tekniska Högskola in Sweden. He was mentored by Patrik Fältström of Cisco. Originally from Pakistan, fellow Tariq Rahim Soomoro is assistant professor at the Fujairah Colleges in the United Arab Emirates. He was mentored by Harald Alvestrand, an engineer at Google.

Five Fellows will participate in the program at IETF 69 in Chicago.

In addition to those who served as mentors to the Fellows at IETF 68, ISOC would like to thank James Galvin, Jaap Akkerhuis, Alain Patrick Aina, Sanjaya, and Frederico Neves for their participation as part of the Fellowship application review and selection committee.

The next call for applications will be announced shortly after IETF 69. ISOC is extending an opportunity for organisations to become sponsors of the program. Details can be found at <http://www.isoc.org/educpillar/fellowship/>.



IETF Fellows and mentors in Prague



Olaf Kolkman, IAB Chair

Words from the Incoming IAB Chair

By *Olaf Kolkman*

Writing this report, which is my first for the IETF Journal, is one of the many tasks I inherited from Leslie when I took over as Internet Architecture Board (IAB) chair during the IETF meeting in Prague. Only when I received the token did I begin to appreciate the extent of my predecessor's achievements. I will not make an attempt to enumerate all of Leslie's many achievements, but I think it is good to explicitly recognise her critical role in the reorganisation of the IETF over the past few years. In her role as IETF Administrative Oversight Committee (IAOC) member, she was a driving force in formalising the relations between the IETF and the RFC editor. In her role as IAB chair, she set the bar pretty high. It will be a challenge to match the quality and quantity of her efforts.

During this IETF there were two mutations in the IAB's membership. Both David Meyer and Bernard Aboba retired, while Danny McPherson and Barry Leiba joined the IAB.

Traditionally, shortly after the change of its members, the IAB organises a retreat to assess its previous year's work and to begin planning new activities. That retreat will take place 31 May–1 June in the Boston area. As of this writing, the agenda is not set, but, no doubt, two topics will get attention and both will be framed around the question of what the IAB can, will, and needs to do as a follow-up to the two IAB workshops held in 2006: the workshop on unwanted traffic and the routing and addressing workshop. (Both reports will be in the RFC editor queue when this publication goes to press.)

Based on activities during IETF 68, it is clear that the routing and addressing workshop unleashed a lot of energy. It is also clear that the community is still trying to get a grip on the exact problem and on the scope of the solution. The IAB and the Internet Engineering Steering Group (IESG) have established a routing and addressing directorate to assist in the coordination of efforts, as well as the establishment and maintenance of communications between the various stakeholders and the IETF leadership. (For a detailed charter, see <http://www.ietf.org/IESG/content/radir.html>.) While the IETF works on a better understanding of the problem and on scoping the solution space, there will no doubt emerge architectural issues that require better understanding and the determination of where the technical role for the IAB is. In the meantime, the IAB will also try to assess whether and how it can play a proactive role in this effort. 

Olaf Kolkman was born and raised in the Netherlands. He was trained as an astronomer but his interest in Internet technology took hold around 1996. He joined the RIPE NCC in 1997, where he became involved in the test-traffic project. That project brought him in contact with the IETF and he attended his first meeting in Munich. After acting as operations manager he became systems architect in 2000, responsible for DNSSEC deployment at the NCC. In 2005 he joined NLnet Labs, a R&D foundation, as chief executive. He is an IAB member since March 2006.

New BoF Meetings

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

Applications Area

fsm: Formal State Machines

General Area

sava: Source Address Validation Architecture

Internet Area

tictoc: Timing over IP Connection and Transfer of Clock

RAI Area

bliss: Basic Level of Interoperability for SIP Services
rtpssec: RTP Secure Keying

Security Area

ifare: IPsec FAllover and REdundancy

Plenary Report

By Mirjam Kühne

Note: This is not a complete report of the plenary sessions; rather, it is a summary of the highlights of the discussions. All IETF 68 presentations can be found at <http://www.ietf.org/meetings/past.meetings.html>

At IETF 68, the operations and administrative plenary session consisted of two parts: The first consisted of the usual updates on administrative and operational IETF issues. The second was dedicated to the status of the routing and addressing problem (ROAP).

IETF 68 attendees heard from Morgan Sackett—a representative of VeriLAN, the company providing network as well as network operations centre (NOC) services for the meeting—who gave a brief presentation that described the general network layout. Attendees also heard from Jon Lindberg, vice president of Secretariat Services, who gave an elevator speech about NeuStar's services. NeuStar Secretariat Services (NSS) organises the IETF meetings and is responsible for tools development and the overall IT infrastructure.

Brian Carpenter presented Jon with a plaque of contribution to NeuStar for hosting IETF 68, a gesture recently introduced at the meetings to recognise IETF meeting host organisations.

IETF Stats

Brian provided a summary report, including updates from the IETF Administrative Director (IAD), the Internet Assigned Numbers Authority (IANA), and the RFC Editor. The March 2007 meeting was attended by 1,200 people, which is roughly the same number of people who attended the meeting in Dallas one year ago. The number of countries increased, though, from 36 at the meeting in

Dallas to 45 at the meeting in Prague. Participants from the United States are still in the majority but constituted only about a third of the participants compared with about half last time.

Since IETF 67, three new working groups (WGs) have been created, and six WGs have been closed. This leaves approximately 120 WGs currently active. (See more statistics under Facts & Figures on page 2).

The RFC Editor published 95 RFCs since the last IETF meeting. Brian noted the improvements in the RFC editor queue and other improvements (see full report on the IAOC Web site) The new contract with the RFC Editor is imminent.

IANA processed 1,160 IETF-related requests since the last meeting, 81 of which were requests for port number. This is worrisome because the number of ports is finite. Since the service-level agreement (SLA) was signed between IANA and IAOC and the IAB, IANA is now ramping up its implementation. The IANA staff has worked especially hard to improve tools and metrics.

The IAD reported that the total meeting income for IETF 67 in San Diego was \$799,000. The total direct meeting expenses amounted to \$438,000. This leaves a surplus of



Mirjam Kühne

\$361,000 that was contributed to the IETF Administrative Support Activity (IASA). It is important to note that over the year, the IASA needs to build up about 50% surplus on meetings toward secretariat activities. The numbers for 2007 are on track so far.

The status of the IAOC action points from the last meeting is as follows:

- The new RFC Editor contract is close to be finalised.
- An SLA has been signed with IANA.
- An RFI on secretariat services has been issued.
- RFC copyright transfers from ISOC to the IETF Trust have been completed.
- An agreement has been reached to attach open-source license to Secretariat tools; source files have been received.
- IETF is a registered trademark.
- Document retention policy has been defined and implemented.

See <http://iaoc.ietf.org> for full reports.

News and Notes

Following Brian's reports, Andrew Lange, chair of the Nominations Committee (NomCom), gave an update on the results of the NomCom. All voting members and liaison members worked hard and kept the best interests of the IETF as a whole as their guiding principle.

As the NomCom removed Russ Housley for consideration as Security Area director by appointing him to the IETF chair position, it noted a dearth of candidates for the Security AD position. This resulted in an additional call for nominations for the Security Area position. From this call a large number of candidates was generated, with eight of them willing to serve.

Andrew discussed a few key issues being faced by the NomCom. There has been a significant decline in the number of volunteers for the NomCom in recent years, and the NomCom chair expressed concerns about this. Another problem is what Andrew called the Strong Incumbent Syndrome, wherein strong, well-respected incumbents appear to dominate and few other people are stepping forward. Finally, there seems to be a growing tension between openness and confidentiality. Andrew suggested that an amendment be made to the NomCom guidelines laid out in RFC 3777, which would make the list of willing candidates—or even all nominated candidates—public.

Daniel Karrenberg, chair of the ISOC Board of Trustees, reiterated a comment Andrew made: “You get the governance you deserve.” He recognised all of the outgoing IAB and Internet Engineering Steering Group (IESG) members. All of them will receive in the mail a plaque of recognition.

Brian and Russ celebrated the passing of the dots, a tradition in which Brian hands over his IETF chair dot—which appears on his badge—to the new IETF chair, Russ Housley.

In his presentation as incoming chair, Russ said he is on a steep learning curve and that he welcomes any input on how to do a good job as IETF chair. (See his welcome note on page 3.)

In closing, Russ said everyone has a slightly different view of the Internet. He welcomes suggestions on how to help figure out what his goals are and how to “achieve my goals in your part of the Internet.”

During the following IAOC open-mic session, Lucy Lynch, outgoing IAOC chair, recognised everyone who is leaving the IAOC and welcomed the new members. She further announced that Dave Crocker and Steve Crocker are the first two volunteers who signed their RFCs over to the IETF trust to ensure that the IETF continues to have as much scope as possible to extend and revise the protocols specified in the RFC series. Dave and Steve stepped up to the podium and each signed an open-ended license for the RFCs they authored. They encouraged others to do the same.

Daniel thanked the IAOC for an excellent and professional job over the past three years.

Leslie Daigle, former chair of the IAB and member of the IAOC, agreed and then reminded everyone that this conversation would have been inconceivable a few years ago. “Congratulations to us all!” she said.

Focus on ROAP

Part 2 of the plenary was dedicated entirely to the routing and addressing problem (ROAP), which was discussed during the last IETF plenary session and at many other meetings. This session was led by the Internet and Routing Area directors.

Brian Carpenter gave a presentation that illustrated where we are with this problem and what the IETF can and intends to do about it. To better understand the context for the recent activities, it is worth looking at the historical time line:

- 1962 Packet switching invented
- 1974 Internet concept invented
- ~1978 Internet Protocol designed
- ~1988 Border Gateway Protocol (BGP) designed
- 1992 Classless Inter-Domain Routing (CIDR) designed
- 1995 IPv6 designed

Since 1995 there has been growing concern about such issues as scaling, transparency, multihoming, renumbering, provider independence, traffic engineering, and the impact IPv6 has on the Internet and the routing system. In 2006, the IAB held a workshop to discuss routing and addressing issues (see the full report at <http://www.iab.org/about/workshops/routingandaddressing/index.html>).

An important architectural principle is that a network should be able to implement reasonable internetworking choices without unduly impacting



Photo by Peter Lotberg

Attendees interact at IETF 68

other networks' operations. The issue at hand on the architectural level is that today certain implementations need to be handled in ways that threaten that principle. This is the root cause of ISP problems and end-site dissatisfaction. The question is, What can be done to harmonise the network to that archi-

Continued on next page

Plenary Report, continued

A goal for 2050 could be to connect 10 billion end nodes with 10 million multihomed customers. Can we get there at a reasonable cost for vendors, Internet service providers, and end users? And what should the five-year goal be?

tectural principle? Brian mentions the “tragedy of the commons from provider-independent address usage” as an illustration of that problem.

There are a number of technical goals, both on the end-user level and on the ISP level. On one hand, end users want to connect to multiple ISPs while maintaining support for current traffic engineering capabilities. They also want to change ISPs without major switching costs. ISPs, on the other hand, want to keep the BGP table size and dynamics within their routers’ operational capabilities. ISPs also want the ability to match traffic engineering flows with their business needs. An overall technical goal is to support end-to-end session transparency.

Another issue is scaling: Today, 200,000 Internet BGP routes and several times more customers and virtual private network (VPN) routes are common. A goal for 2050 could be to connect 10 billion end nodes with 10 million multihomed customers. Can we get there at a reasonable cost for vendors, Internet service providers, and end users? And what should the five-year goal be?

A number of sessions at various recent meetings—such as at the regional operators meetings, the DARPA R&A workshop, NSF/OECD workshop, and the TER-ENA workshop—were held to raise awareness of the problem. Within the IETF community, activities are being organised to address the issue, including a Routing and Addressing Direc-

torate, which was formed recently. In addition, the Routing research group has been rechartered (see more details later in this article), and the Routing and Addressing mailing list ram@ietf.org has become quite active. During IETF 68, a number of meetings were organised to address the issue not only during the plenary session but also during the open Internet and the open Routing Area meetings.

With BGP4 in use with little change since 1994, there is growing concern about the growth of the BGP routing table. This is related primarily to the size and update rate as well as to the impact that multihoming has on the routing table. The problem exists in both IPv4 and IPv6. Attendees at the IAB workshop in October 2006 looked at hardware trends that raise economic and engineering concerns about the size of the Forwarding Information Base (FIB).

Another issue raised was a problem with transparency. Since 1981, the upper layers of the network stack have assumed that a thing that looks like an address is an address. For instance, application programmers often assume that an IP address is a valid end-system identifier that can also be passed on to third parties. Consequently, problems arise when addresses are viewed merely as locators. NAT, STUN and other applications attempt to deal with this problem in their own ways, which lead to newer problems. The historical reliance on address transparency creates specific difficulties for multihoming

and traffic engineering. There are a few areas for which solutions can be developed:

- Router and microelectronics designers can work on engineering to help solve the RIB/FIB scaling problem.
- BGP adjustments and better operational practices could help improve the update dynamics.
- Traffic engineering, multihoming, end-to-end transparency, and mobility would benefit from architectural changes. Identifier/locator split and/or multilevel locators form a hopeful approach.

All of those developments are orthogonal to both IPv6 deployment and application-level namespace issues.

What can the IETF do? It can provide a forum for open problem analysis and development of solutions by vendors, operators, and users working in concert. And the Internet Research Task Force (IRTF) can help with research. In the short term, routing-table growth is only an engineering issue. While routing dynamics need to be better understood, this is likely also an engineering issue and can possibly be addressed by stronger pushback in the ISP community as well as through implementation and protocol improvements. Thus, there is reason to believe that we do not have a short-term technology problem, even though it will continue to be hard work and will have an impact on business decisions. However, architectural problems remain. The IETF can help with short-term protocol work, such as better tuning BGP to meet today’s challenges. The IETF can also help by looking into architectural changes, such as the identifier/locator split and multilevel locators.

As an overall plan, it has been suggested that the IETF work on all of

the above levels (short-term incremental improvements as well as architectural changes) and continue the dialogue with the operators community.

The introductory presentation by Brian Carpenter was followed by a lively and animated discussion. Some people felt this problem is still not taken seriously enough and encouraged “the IETF to be bolder in tackling this problem,” as Ted Hardie put it.

Many people suggested that the problem be addressed on multiple levels. There are a number of things that can be done today to work on BGP dynamics, such as operational practices and certain routing policies or tweaks in implementations and protocols. And there are more-fundamental architectural changes—such as splitting



Photo by Peter Lothberg

Outgoing IAB Chair Leslie Daigle is honored at IETF 68

the locator from the identifier—that won’t come without costs. The trade-offs are not yet well understood. But “it is good to know that there are engineering solutions to keep the Internet running in the meantime while we start working on some fundamental changes,” said Ross Callon, one of the Routing Area directors.

Dave Ward, the other Routing AD recognised Dave Meyer, who has been working over the past few years

to bring this problem to light. Dave Meyer thanked the IETF for the applause but pointed out that we’ve already known about this development for 15 years and cannot wait another 15 years. He emphasised that this is not a discussion about incremental improvement to BGP, which needs to be done anyway. “This is a controversial topic, and it will be hard to get IETF consensus,” said Dave. “Do we have any way forward in our process to deal with this?”

Other people agreed, expressing concern that it is not clear what the right processes are to make architectural changes to the Internet. The general consensus was that this issue is too important to leave to the IESG or to the IAB; it should be worked on by the entire community. One participant mentioned that almost any architectural change will have costs, but this one might have significant impacts on security. Security implications need to be looked at early on in the process.

Summary of Progress Made

Much has been done in the past six months with regard to ROAP—a clear sign that the IETF now takes this problem seriously. Daniel Karrenberg suggested that in the discussion, the IETF should also look back and learn from past experiences. “There was a panic about routing-table growth that gave us CIDR,” he said during the plenary. “We needed a quick fix, and we did a quick fix. There was a panic about IP address space running out that gave us IPv6. And some of the ‘features’ IPv6 brings us is because we rushed it.”

However, it will take several years to develop, implement, and deploy any new architecture, and no one solution or tool will solve all problems; it might be necessary to apply different sets of tools in order to solve different parts

of the problem. Some have expressed concern that the IETF may have already constricted itself to a small problem space: routing (and small incremental changes) and addressing (big fundamental changes). In general, however, there is agreement that for short-term solutions, the focus should be on routing and that for the longer-term solutions, a lot more will be involved. Collaboration between all IETF areas will be necessary.

Additional concern about operators and users not being able to pay for continuing updates of hardware led to a discussion about the implication of architectural changes on business decisions, with an emphasis on the need for new architecture to take existing business models into account. “We have to make sure we deal with the architecture in some way that does not lead into high costs and short lifetime of equipment,” said Rüdiger Volk. On the other hand, it might be unrealistic to assume that one can make fundamental changes and keep all the existing business models.

“The Internet got started by breaking the then existing business models completely,” said Bob Hinden. “It could be that this will also happen today. We should not be constrained by the way people do business today.”

The IETF 68 Technical Plenary, held on the following evening, began with an update by Internet Research Task Force chair Aaron Falk on the activities of the IRTF (see page 26 for details).

Aaron’s presentation was followed by Aiko Pras of the University of Twente, The Netherlands. Aiko reported on a workshop held in October 2006 in Utrecht, which brought together researchers, operators, vendors,

Continued on next page

Plenary Report, continued

and technology developers to identify future directions of network and service-management research. The workshop was jointly organised by the Network Management research group (nmrg) of the IRTF and the European Network of Excellence for the Management of Internet Technologies and Complex Services (EMANICS). The work is expected to result in recommendations for research directions that are worth exploring over the next five years. The work is not intended to define the management standards that are needed now.

A number of challenges in this area were identified at the workshop:

1. Management models: The Manager-Agent approach (client-server) and hierarchical management (DisMan, TMN) are well understood. What is not yet understood are issues such as fully distributed management (P2P, ad-hoc) and self-technologies (auto-configuration, stability of control loops).
2. Distributed monitoring: A track number or quality of VoIP calls as well as a way to find the best proxies and peers (P2P) are needed. The goal is to develop a lightweight, distributed monitoring layer that offers aggregates of local information to applications.
3. Data analysis and visualisation: It is possible to create topology maps for small networks and statistic time series plots. There are difficulties with the creation of maps for large, multi-layer networks, the visualisation of anomalies and real-time, interactive visualisation techniques.
4. Economic aspects of network management: Most researchers tend to focus on technical solutions. There is limited research into the operational costs of such technologies. Network Management is risk management.
5. Uncertainty and probabilistic approaches: Many researchers focus on deterministic approaches and yet scalability problems might need more probabilistic approaches. How does one decide between probabilistic and deterministic approaches?
6. Ontologies: The data modelling is believed to be understood. There is still more research needed in areas such as how ontologies can be effectively used to automate the implementation of management interfaces as well as how ontologies can help check or enforce policies and behaviour.
7. Behaviour of managed systems: Management models usually represent a state (MIBs, CIM). More research is needed to model and manage behaviour, such as normal versus abnormal behaviour, detection of resource failure, and the design of self-stabilising systems.

As a follow-up to the nmrg workshop, an Internet-Draft will be written describing outcomes and an overview article will be submitted to the IEEE ComMag.

Internationalisation in an IETF Context

One of the more complex and challenging issues facing the IETF—and the Internet more generally—today served as the basis for discussion during the remainder of the technical plenary session: Internationalisation of the Internet in the context of the IETF. “There is an entire set of human Internet users who cannot use the Internet the way we do,” said Ted Hardie of the IETF’s interest in the issue. “We would like to change that.” As part of the technical plenary, a small team of experts—Leslie Daigle, Patrik Fältström, Ted Hardie,

John Klensin, Xiaodong Lee, and Pete Resnick—presented slides, explained the issues, and answered questions.

Tackling internationalisation is an ambitious undertaking—one that at times can be complicated by questions that are more philosophical or linguistic than they are technical. After all, it would be unrealistic to expect the IETF to attempt to develop a standard for helping one group understand the language of another if they don’t speak that language. Still, there is a lot that the IETF can do to improve international accessibility. As John has often said, “The IDN issues are tractable as long as we keep a clear focus on what problems we are trying to solve and what areas of the general topic actually need solving and can be solved.”

In a high-level introduction to the topic, Leslie shed light on the core set of issues associated with internationalisation and protocol design and offered a broader perspective on RFC 4690, Review and Recommendations for Internationalised Domain Names (IDN). She challenged a widely held assumption within the IETF that the topic does not concern those who are not involved in the applications area when, in fact, the technical issues associated with IDN touch many areas. She said that while there is no clear problem statement, there are, in fact, a variety of problems that promise to become bigger problems over time. “This portion of the technical plenary is meant as a heads-up on what may be coming into the IETF as new work,” she said.

The problems and concerns associated with IDN underscore the difficulties with which language is translated into code. In an attempt to put internationalisation into context for the IETF, Ted showed a clip from the movie *Charade*, in which Audrey Hepburn and Cary Grant are asked

to transfer an orange from under their chin to another person without using their hands. The clip demonstrates with humour how communication often relies on the ability to understand one or more languages as well as body language, context, and tone. In terms of IDN, the question is much larger and more complex than how characters are coded. The question is, How do we create character sequences and applications that depend on language use and context for accuracy and usability? As Ted pointed out to the audience, the clip's protocol description was in at least four different languages, but

it must be relevant on a local level. For many, that means a Domain Name System that is functional regardless of language and useful regardless of context. In RFC 4690, it is stated that "While IDNs have been advocated as the solutions to a wide range of problems. . . . They are no more and no less than DNS names, reflecting the same requirements for use, stability, and accuracy as traditional 'hostnames' but using a much larger collection of permitted characters. In particular, while IDNs represent a step toward an Internet that is equally accessible from all languages and scripts, they,

characters without any language identification or a requirement to conform to the syntax of any specific language. That would mean that in IDNAbis, character-to-character mappings become the responsibility of the user interface and not the protocol. "Most reasonable user interfaces won't need to be changed," he said. "The protocol should be able to work with different versions of Unicode and, hence, would not be restricted to specific Unicode versions." Some have suggested that more characters be added for functionality and that a larger number of characters be allowed to appear in registered strings. John expressed the hope that IDNAbis would allow for better treatment of bidirectional languages, such as Arabic languages that read from right to left. He also expressed confidence that the new work would lead to a protocol that would be easier to understand and explain.

"Accessibility from all languages is an important objective; hence it is important that our standards and definitions for IDNs be smoothly adaptable to additional scripts as they get added to the Unicode character set." — RFC 4690

only two were close to complete. To know that, though, one would have to recognise and understand all of them. And, in the end, you can't internationalise a chin or an orange.

As Ted said, there are specific instances when context is essential. Those include protocol descriptions, protocol elements, and human elements, the last of which is the most important and perhaps the most difficult. As spelled out in his slides, things that we may assume are protocol elements can sometimes be human elements and vice versa. That means that in some contexts, internationalisation is completely inappropriate; in others, it's necessary to understand how much context is available in order to do proper design. As the movie clip demonstrates, rapid language switching requires context switching. If you can't identify the context (or the language), you're going to have a problem.

If the Internet is to be truly global,

at best, address only a small part of that very broad objective. There has been controversy—since IDNs were first suggested—over how important they will actually turn out to be; that controversy will probably continue. Accessibility from all languages is an important objective; hence it is important that our standards and definitions for IDNs be smoothly adaptable to additional scripts as they get added to the Unicode character set."

In his presentation, John provided an overview of current IDN work within the IETF and emphasised the importance of advancing work on IDNA (Internationalising Domain Names in Applications). He said changes may need to be made in IDNAbis (the term for the successor to RFC3490 and related specifications), such as further clarification of terminology and separation of the user interface and language issues from the protocol, since the DNS deals only with strings of

The discussion that followed offered several examples of how confusion over internationalisation is evident even within the IETF. In one case, a participant asked a question in connection with the concept of Internationalised Resource Identifiers (IRI)s. Observing that the deployment community is not clear about whether it's possible to use non-ASCII characters in IRIs being passed in protocols, the participant asked if all the right things were done with regard to terminology. "You've just fallen into the same trap as the last questioner, which is, you used the word characters," said Pete. "IRIs are the things on the side of the bus with the funny characters. What goes in protocols are URIs [Uniform Resource Identifiers]. Generally speaking, those have octets in them; actually, they have 7bit filled octets in them. Well, sometimes they're filled, and sometimes they're not."

Continued on next page

Plenary Report, continued

“Maybe you now understand why we had headaches for the past three years,” said Patrik.

Unicode, as John pointed out, is concerned only with standardised code points; in other words, it does not specify how to draw characters or how to address fonts or context. For example, the criteria given in the Unicode Standard for assignment of code points excludes the assignment of codes to font variations on the same character. However, for the base Latin alphabetic characters, additional code points are assigned for mathematical uses; in other words, code points are assigned to characters that visually and within a specific mathematical context are the same as their base counterparts—except for such attributes as bold and italic, which meet any reasonable definition of font variations. As John said, one cannot distinguish those characters from others except that they are defined in the Unicode description as mathematical. “Given the context that Unicode essentially expects, these are not the base characters and a font difference at all, but mathematical symbols,” he said. “That’s a problem, because in our environment it causes complications.” For ordinary IETF text string purposes, one must either map these mathematical characters into the base ones or forbid them entirely. Either choice comes with problems, but in terms of IDNAbis, we concluded that it is better to forbid them in the protocol, leaving mappings, if any, to the user interfaces. John said in a separate discussion that there are other Unicode assignment rules that are not completely orthogonal to each other and that there are issues of some type with almost every script and language. “Difficulties arise when there is inconsistency in the ap-

plication of the Unicode Consortium’s own rules,” he said.

According to Harald Alvestrand, the IETF cannot claim expertise in character set coding and should not attempt to do so; that, he said, is the responsibility of the Unicode Consortium, even if its work doesn’t give us a complete solution to internationalisation issues.

Patrik responded that there is no need for the IETF to hold off on working on the IDN and IDNA issues if it is prepared to address or on internationalisation issues more generally. He believes there is a clear boundary between the Unicode Consortium and the IETF. “What we are doing is, relying on the data the Unicode Consortium is providing,” he said. “We are referencing their expertise for the classification of each of the Unicode characters.”

The conversation shifted to the potential for fragmentation of the Internet with IDN, particularly if different scripts are allowed. One participant raised the question of whether the worldwide operation of the Internet could be disrupted if e-mails are written in a script the recipient does not understand. In terms of content of messages, this situation has existed for more than a decade and has caused no more, and no fewer, problems than incompatibilities among languages traditionally do. The risks obviously increase if one is presented with domain names, IRIs, or e-mail addresses that one can’t understand or, worse, that one’s terminal device cannot render. Obviously, nothing the IETF can do about internationalisation will solve the problems of people communicating with each other by using languages that neither party speaks. However, the bigger and more relevant issue is the considerable number of people in the world who still cannot communi-

cate on the Internet, even with other people who share their language and writing system. By adding more functionality that would facilitate the use of other languages, many more people would be able to benefit from the Internet. “We are actually worried about that problem, and what we’re attempting to do in the internationalised e-mail address working group is to come up with a form of e-mail address that, if it gets to you and you’re not in on the game of fancy e-mail addresses with all sorts of interesting characters in them, you’ll at least be able to reply and get back to me,” said Pete.

Even so, other concerns persist, including the risk that different applications might deal with internationalisation in ways that will, once again, cause incompatibility and create problems for the users. “While those types of problems might exist today,” said Andrew Sullivan from the floor, “the scope of confusion will be significantly higher if people start trying to communicate across language and script borders.”

Another concern is the possibility that if the new IDNAbis becomes standard, names registered under the old IDNA may no longer be valid. “This might indeed happen,” said Patrik. “However, most of the names that will become invalid are corner cases, and only people who are interested in phishing use them.”

John agreed, adding that the vast number of labels that will not be valid in the new system were not valid under the old system—a system that included guidelines from the Internet Engineering Steering Group (IESG) and the Internet Corporation for Assigned Names and Numbers (ICANN)—either. “The old system was just more permissive,” he said. “The new system will standardise more and allow for fewer borderline cases.”

In reality, most of the issues are related to localisation and not internationalisation or multilingualism. “The purpose of our work is to include the people who cannot participate in the Internet today because we are failing to make it possible for them to localise properly,” said Ted.

IAB News and Notes

Following the long and lively discussion on internationalisation, Leslie

gave an update on recent IAB developments. A number of documents have been published or are in the publication queue, including reports on the IAB workshops on unwanted traffic and routing and addressing. There is also a document on the process for publication of IAB RFCs being published (draft-iab-publication-00.txt).

The IAB has appointed Bob Hinden to the IETF Administrative Oversight Committee (IAOC) and has renewed

Aaron Falk’s chairmanship of the IRTF. Aaron has been appointed to serve as chair for another two years.

At the conclusion of the technical plenary, Leslie handed the IAB chairmanship over to Olaf Kolkman, who thanked Leslie for her outstanding work and presented her with a gift. The session ended with a long round of applause and standing ovations for Leslie from the plenary attendees.



Update on DNS

By Jaap Akkerhuis and Peter Koch

Note: This article does not provide a complete summary of all IETF activities in this area. It reflects the author’s personal perspective on some current highlights.

This is an update on recent DNS activities based on current active working groups related to this area.

1. DNSEXT WG

Within the DNS Extensions working group, work is progressing steadily. A number of Internet-Drafts, including the ones on DNS Security Extensions (DNSSEC) Hashed Authenticated Denial of Existence (NSEC3) and DNS Name Server Identifier Option (NSID), have passed Last Call and are making their way to standard, experimental, or informational requests for comments (RFCs). For complete details, please consult the minutes and/or the standard tracker.

DNAME

Scott Rose gave a presentation on the DNAME clarification draft. This document is chartered by the WG to update RFC 2672, and it addresses issues that people have had with the original specification and implementation or operational experience. It also provides a clearer understanding of DNS and aliasing in general. The editors have started an issues tracker

and are looking for feedback on the issues.

DNS Hardening

An Internet-Draft was adopted as a work item that explains how resolvers can be made less vulnerable to spoofed DNS responses without adding protocol extensions such as DNSSEC. There were critical remarks about the lack of terminology and missing operational considerations. Several people stepped forward to improve the draft, and it is expected to be ready for Last Call this coming summer.

The 2929bis Template Argument

This is an experiment for allocation of new Resource Record (RR) types. The idea is that a lightweight process—for instance, a recommendation by a designated expert—might be sufficient to decide whether a new Resource Record should be allocated. One lesson learned is that 2929bis needs to be updated in order to establish what is expected of the expert and what the

boundaries of the expert are. In addition, it was decided that the process should be reviewed by the Internet Assigned Numbers Authority (IANA) and that the template may still need to be tweaked. Once the experiment has been completed, the area director will schedule it for an evaluation by the Internet Engineering Steering Group (IESG).

Is There a Future for DNSSEC?

Many drafts related to DNSSEC have been completed or are expected to be completed soon. The question is whether the DNSEXT WG should be closed or whether it should live on in some form. Typically, once the work is done, a WG is abandoned. One argument against abandoning the group is that the DNSEXT WG is often asked to comment on proposals done by other groups. Another is that it might be good to have the WG around to help advance the DNSSEC RFCs from proposed standard to the next level: draft standard. There have been cases where a WG is in a dormant state—for instance, the Point-to-Point Protocol (PPP) WG—or, in the case of the provreg WG, has maintained an active mail list for these purposes.

Continued on next page

Update on DNS, continued

2. DNSOP WG

The DNS Operations WG is still very active and in fact ran out of time during the meeting.

Some of the older Internet-Drafts, such as DNS Response Size Issues, are moving forward.

Reverse Mapping

Although lively discussions are still happening in this area, the discussion about the need for reversed mapping is expected to come to some close. The Internet-Draft is expected to be ready soon for WG Last Call.

AS112 Work in a Box Continues

AS112 is the popular term for how to deal with DNS queries that actually shouldn't happen—such as queries for the reverse mapping of private address space defined in RFC 1918, which managed to escape a local network (such as local top-level domains like .lan or .local) and make it into the Internet. The project, named after the origin Autonomous System 112, consists of a loosely coupled anycast cloud that responds to these queries to take load off the root name servers. (See <http://public.as112.net/> for a description.) To date, there is no detailed explanation available for end users or potential contributors. The WG is trying to fill that gap by creating a first document about it. The document “Help, I'm Attacked by prisoner.iana.org” is expected to go to the IESG soon. The need to explain how the current process surrounding the AS112 system actually works was identified.

New Work Items

Now that the WG has reached almost of all its milestones, there is still work to do in DNS operations. First, the management of large and distributed clusters of name servers is becoming more common but currently lacks automated, nonproprietary support for configuration and synchronization. A similar problem arises for the remote control of DNS secondary servers. The WG is now going to examine the need to address various DNS operational scenarios.

Two ICANN committees—the Root Server System Advisory Committee (RSSAC) and the Security and Stability Advisory Committee (SSAC)—jointly started an investigation earlier this year on how adding AAAA Resource Records for the root name servers would influence the DNS resolver priming process. While their results have been promising, it turned out that the priming process itself—although current practice—isn't fully specified and poses some questions related to DNSSEC.

Lixia Zhang presented some research on the effect of TTL (time-to-live) values for so-called infrastructure records—name server and address records—on the resolvability of a domain during longer periods of failure. While there is a general trend to treat DNS data more dynamically, there are side effects on both the infrastructure and the leaves in the DNS tree if the feature of caching is defeated by very low TTL values. This will be further investigated by the WG.

3. ENUM WG

The ENUM WG, dealing with the mapping of telephone numbers into the DNS, expected an interesting debate on the future of infrastructure ENUM, a supplement to the core ENUM protocol aimed at providers of telephony services. Just before the IETF meeting there was some confusion about the state of consensus as well as the political implications (for details, refer to Geoff Huston's excellent article on <http://ispcolumn.isoc.org/2007-03/infra-enum.html>), but this was resolved at least to the extent that the WG maintained its consensus, and other considerations will be taken into account as appropriate during the evaluation process.

Two other items are remaining for the ENUM WG: in response to various ENUM service registrations and in preparation of a closedown of the WG, guidelines are prepared on how to write and review ENUM service specifications.

The second major remaining task is an update of the base ENUM specification, for which a draft has already been published that tries to overcome some of the drawbacks of the Naming Authority Pointer (NAPTR)-based design in response to real-world deployment experiences. 

Note: For an in-depth look at on DNS infrastructure, see the March 2007 issue of the Internet Protocol Journal at http://www.cisco.com/web/about/ac123/ac147/about_cisco_the_internet_protocol_journal.html.

Recent IESG Document and Protocol Actions

A full list of recent IESG Document and Protocol Actions can be found at <http://ietfjournal.isoc.org/DocProtoActions0301.shtml>

More ROAP: Routing and Addressing at IETF68

By Geoff Huston

Note: This article does not provide a complete summary of all IETF activities in this area. It reflects the author's personal perspective on some current highlights.

Over the past year or so, we've seen a heightened level of interest in the topic of Internet routing and addressing. Continued intense examination of the IPv6 protocol together with associated speculation regarding the future role of the Internet raises the possibility of the Internet supporting a world of tens or hundreds of billions of chattering devices. What does such a future imply in terms of the core technologies of the Internet? Does what we use right now scale into such a possible tomorrow? Consideration of this topic has prompted critical examination of aspects of the architecture of the Internet, including the scaling properties of routing systems, the forms of interdependence between addressing plans and routing, and the roles of addresses within the architecture.

The Internet Architecture Board (IAB) has been active in facilitating discussion of this topic, both within the IETF and at various Internet operational gatherings around the world. This IAB effort culminated in a two-day workshop in October 2006 on routing and addressing to examine the characteristics of this space and to start to identify some of the interdependencies that appear to exist here. The workshop report (<http://tools.ietf.org/html/draft-iab-raws-report-01.html>) is close to completion, and there is also the author's informal report (<http://ispcolumn.isoc.org/2006-11/raw.html>) of impressions gained at the workshop.

IETF 68 saw some further steps in analysis of these issues, and during the week there was a plenary session on routing and addressing as well as meetings of the Internet and Routing Areas devoted to aspects of routing and addressing. This is a report of

these sessions plus some conjecture as to what lies ahead along this path.

Plenary ROAP: The Plenary Session on Routing and Addressing

The plenary session at IETF 68 presented an overview of the topic, looking at the previous initiatives in routing and addressing as well as providing some perspectives on the current status of work in this area. Routing and addressing, in the context of the Internet, has been visited on a number of occasions over the years, starting with the shift from the original 8/24 network and host part addressing to the Class A, B, and C addressing structures and the subsequent shift to the prefix-plus-length concepts of classless addressing. In the routing area, there was the adoption of a peer model of routing with the introduction of BGP (Border Gateway Protocol) and the shift in BGP to sup-

port classless addressing in the form of CIDR (Classless Interdomain Routing). And, of course, there has been the design of IPv6. However, there still remain the concerns that this is not completed work and that the technology is not in an ideal state to scale by further orders of magnitude without further refinement. There are concerns regarding the scalability of routing, the transparency of the network, renumbering issues, provider-based addressing and provider lock-in, service and traffic engineering, and routing capabilities, to name but a few issues that are relevant and challenging today and that appear to be even more so for the Internet of tomorrow.

Are there architectural principles that are relevant here? In the large, diverse but coupled set of networks that collectively define the Internet, it appears that each component network should operate within a general principle of containment or insulation of impact. The principle is that each network should be able to implement reasonable choices in its local configuration without undue impact on the operation or range of choices available to all other networks. In other words, each network should be able to make such local configuration choices relatively independently of the choices made by any other network. The relevant issue here lies in balancing this principle against the operation of the network as a whole, which can be seen as a binding of networks together as a coherent entity that supports consistent and robust communications paths through this collection of networks.

We do not use a routing technology that effectively isolates individual network elements from each other or even manages to localise the external impacts of local choices. On the contrary, far from being a protocol that

Continued on next page

More ROAP, continued

damps instability, BGP manages to be a highly effective amplifier of the noise components of routing events. So while it is a remarkably useful information dissemination protocol with considerable flexibility, the properties of BGP in an ever-more-connected world with ever-finer granularity of information raise some questions about BGP's scaling properties. Will the imposed noise of the protocol's behaviour completely swamp the underlying information content? Will we need to deploy significantly larger routers to support a much larger routing protocol load but need to route across a network of much the same size as today's network? The prospect here is that routing may become far less efficient because as we increase the degree of interconnection and the information load simultaneously, the inability to insulate network elements from each other and the inability to effectively localise information create a disproportionately higher load in network routing.

In addition to these observations about routing, there is the continuing suspicion that the semantic load of addresses in the Internet architecture, whereby an address conveys simultaneously the concepts of who, where, and how has some side effects that cause complexity in other aspects of the network—including routing, of course. To what extent the semantic intent of endpoint identity, (or "id") can be pulled apart from the semantic intent of network location and forwarding lookup token, (or "loc") is a question of considerable interest. While the current IP address semantics removes the need to support an explicit mapping operation between identity and location, the cost lies in the inability to support an address plan that is cleanly aligned to network

topology and in the inability to cleanly support functionality associated with device or network mobility. In the end, it's the routing system that carries the consequent load here. The issues in this area include evaluation of the extent to which identity can be separated from location, and the impact of such a measure on the operation of applications. How much of today's Internet architecture would be impacted

and remove that from the locator part of an address. In split id/loc terms, a mobile device is one that maintains a constant identity but changes locators. Multihoming can be expressed in id/loc terms as a single identity simultaneously associated with two or more locators; traffic engineering can be expressed in terms of locator attributes without reference to identifiers, and so on.

Behind this is the observation that the routing and addressing space is not infinitely flexible and that, on the contrary, it forms a highly constrained space. Part of the motivation behind the id/loc splits is to take some of the inflexibility of the id part of an address, in which persistence is a key attribute,

The prospect here is that routing may become far less efficient because as we increase the degree of interconnection and the information load simultaneously, the inability to insulate network elements from each other and the inability to effectively localise information create a disproportionately higher load in network routing.

by such a change, and what would be the resultant benefits if this were to be deployed? Would the benefits of such a deployment be realised directly by those actors who would be carrying the costs? Is deployment a complete and disruptive phase shift in the Internet, or are there mechanisms that support incremental deployment? Are we looking at one single model of such an id/loc split, or should we think about this in a more general manner with a number of potential id/loc splits?

Besides consideration of these general architectural principles and their application in routing and addressing, there are also more-specific sets of objectives that relate to Internet actors. For users, there are objectives here about maximising the user's service and provider choices without cost escalation; and for service providers, there are the objectives of using cost-effective technologies that can accommodate a broad diversity of both current and projected business needs, as well as the very real need to maxi-

mise the value of existing investments in network plant and operational capability.

Obviously, the study of the topic of routing and addressing and the related aspects of name space attributes and mapping and binding properties is one with a very broad scope. The larger question posed here is whether this is an issue whose resolution can be deferred to a comfortably distant future or whether we are seeing some of these issues affect the network of the here and now. Are we accelerating toward some form of near-term technical limit that will cause a significant disruptive event within the deployed Internet, and will volume-based network economics hold, or will bigger networks

start to experience disproportionate cost bloat or worse? Is it time to become alarmed? Well, there is the certainty of exhaustion of the unallocated IPv4 address pool in the coming years, but the sense of alarm in routing and addressing is more about whether there are real limits in the near future over the capability of continuing to route the Internet within the deployed platform by using the current technologies and by working within current cost performance relationships irrespective of whether the addresses in the packet headers are 32 bits or 128 bits in size. There was a strong sense of “Don’t panic!” in the plenary presentation, with the relatively confident expectation that BGP will be able to carry the Internet’s routing load over the next three to five years without the need for major protocol surgery and that Moore’s Law would continue to ensure that the capacity and speed of hardware would track the anticipated growth rates. There was the expectation that the current technologies and cost performance parameters would continue to prevail in this time frame. However, the subsequent plenary discussion exposed the viewpoint that such a prediction does not imply cause for complacency and that some sense of urgency is warranted given the criticality of this topic, the high level of uncertainty when looking at even near-term growth prospects, and the ease with which this industry adopts a comprehensive state of denial over pending events irrespective of their potential severity.

What we are up against as we consider these objectives as they relate to a future Internet is the relentless expansion of the network. Today the Internet sits in an order of size of dimension of around 10^9 . There are some 1.6×10^9 routed addresses in the Internet and an estimate of 10^8 to 10^9 attached devices. If we look out as far as four

decades to around 2050, we may be looking at 10^{11} to 10^{14} connected devices. (Yes, there’s a large uncertainty factor in such projections!) Can we take the Internet along such a trajectory from where we are today? And if that’s the objective, then how can we phrase our objectives over the next five years that are steps along this longer-term path?

The immediate steps at the IESG level have been to take the IAB’s initiative and work with a focus group—the Routing and Addressing Problem Directorate—to refine the broad space into a number of more-specific work areas, or “problem statements,” and undertake a role of coordination and communication across the related IETF activities. In addition, as there is a relatively significant research agenda posed by such long-term questions, the Routing Research Group of the Internet Research Task Force has been rechartered and, judging by participation at its most recent meeting just prior to IETF 68, effectively reinvigorated to investigate various approaches to routing that take us well beyond tweaking the existing routing tool set.

Internet ROAP: The Internet Area Meeting

The Internet Area meeting concentrated on aspects of this approach of supporting an identifier/locator split within the architecture of the Internet, and, specifically, on the internet-working layer of the protocol stack, and on gathering some understanding as to whether this approach would assist with routing scaling. One of the key considerations in this area involves working through what could be called boundary conditions of the study. For example, is this purely a matter for protocol stacks within an endpoint, or are distributed approaches that have active elements within the network

also part of the consideration? To what extent should a study consider mobility, traffic engineering, network address translations, and minimum-transmission-unit (MTU) behaviour? What appears to be clear at the outset is that this is not a clean-slate network, and any approach should be deployable on the existing infrastructure; should use capability negotiation to trigger behaviours so that deployment can be incremental and piecemeal and allow existing applications and their identity referential models to operate with no changes; and, hopefully, should have a direct benefit to those parties who decide to deploy the technology.

From the routing perspective, the overall desire is to reduce the growth rates of the interdomain routing space. The desired intent is to reduce the amount of information associated with locators so that locators reflect primarily network topology in such a way that the locators can be efficiently aggregated within the routing system that attempts to maintain a highly stable view of the network’s topology.

The resultant system must be able to express, in routing terms, most of the flexibility we see in today’s system, perhaps on a more ubiquitous scale. This includes site multihoming across multiple providers, ease of provider switching and locator renumbering (assuming that locators may include some provider-based hierarchy), support for mobility, roaming and traffic engineering, and allowing for session resilience across various locator switch events. In and of themselves, these objectives form a challenging set—but not the complete set—of objectives. In addition, it is necessary that these outcomes be achieved within tight cost constraints and volume economics that allow for scaling without disproportionate cost escalation. Plus, of

Continued on next page

More ROAP, continued

course, such systems should be resilient to various currently known—and currently unknown—forms of hostile attack.

Today's system uses two critical mapping databases to support the discovery of the binding between identifiers and addresses. The Domain Name System (DNS) is used for mapping between a human-oriented name space used at the application level (domain names) and IP addresses, and the routing database in each router is used for mapping from addresses to particular local forwarding decisions (the forwarding mapping from the RIB to the FIB¹ data structures). The current mapping system assumes stable endpoints with simple resource requirements and rudimentary security.

When we consider in further detail the implications of disambiguating aspects of identity from those of network location, we must recognise that there are a number of dimensions to such a study, including the structure of the spaces, the mapping functions, and the practicalities of any form of deployment of such a technology.

The first of these topics consists of the desired properties and structure of these distinct identification and locator spaces. Should the identity space be a flat space of token values, or should it use some internal structure within the token that matches some distribution hierarchy? Is identity something that is embedded into a device at the point

of manufacture (such as IEEE-48 MAC addresses)—or at the point of deployment (such as domain names)? Is uniqueness a statistically likely outcome—or one that is ensured through the structure of the token space? Are there properties of the identity space that aid or hinder the security properties of the use functions in terms of mapping and referral operations? Is there necessarily one identifier space or are there potentially many such spaces? There are similar questions regarding a dedicated locator space, particularly related to the time and space properties of locator tokens.

The next critical topic appears to be how an identity-mapping function relates to the forwarding-mapping function. Assuming that the existing name spaces remain unaltered, the resultant framework appears to require distinct name-to-identifier mappings, identifier-to-locator mappings, and locator-to-forwarding mappings. Where these mapping functions should be performed, who should perform these functions, when they should be performed, what should be the duration of the validity of the outcomes, whether the mapping-function outcomes are relative or universal, the scope and level of granularity in time and space of the map elements, the security of these mapping functions, and whether there is a simple operation or multiple operations in each mapping function all remain undefined at this point. There are also the issues of whether the mapping is explicit or im-

PLICIT, of what evidence of a previous mapping operation is held in a packet in a visible manner, and of what is occluded from further inspection once the mapping operation has been performed. What level of state is required in each host, and is there true end-to-end transparency and at what level? To illustrate some of the dimensions here, a particular approach to an identifier/locator split could see identifiers in the role of the end-to-end-tokens that are used by upper levels of the protocol stack, where identifiers are preserved in such a manner that both parties to a packet exchange use the same identifier pair for each transmitted packet, while locators would have to be more elastic in intent and various identifier-to-locator and even locator-to-locator mappings could be performed while the packet is in transit. Another approach would take a more constrained view of locators and attempt to protect the initial locator value in such a way that any attempts to alter that value during transit would be detected and discarded by the receiver.

The other aspect to consider here is what one presentation termed the incentive structure, where it was advocated that the most-effective incentives are those in which local change is performed as a means of alleviating local pain. This would indicate that routing scalability is predominantly a concern of service providers, whereas host mobility and service multihoming and session resilience are matters of concern to the host and service pro-

Footnote

1. The RIB, or routing information base, is a router's internal data structure that stores the current state of reachability information (or "routes," where a route is defined as a unit of information that pairs a destination address prefix with the attributes of a path to that destination) as provided by the operation of the routing protocol and local policies. In BGP there are three notional RIB sets: an Adj-Rib-In, used for storing routes received from BGP; a peer, the Loc-RIB, representing the routes selected for use by the local BGP instance; and an Adj-RIB-Out, which contains routes that are announced to a BGP peer. RIBs contain bindings of next-hop addresses to routes. The FIB, or forwarding information base, is the data structure for making local forwarding decisions; it contains a set of address prefixes and the associated local forwarding action, conventionally denoted as an interface identifier. A FIB contains bindings of addresses (or address prefixes) to interfaces.

vider and consumer. It's also useful in an incentive structure that benefit be realised unilaterally, in that one party's efforts at deployment provide local benefit for that party without regard to the actions of others, so that the problems of initial deployer penalties and lockstep are avoided.

It is likely, at least at this stage of the study, that there is a diversity of approaches to such a split both in the intended roles of identifier and location tokens and in their means of binding. Already in the HIP (host-identity-protocol) and SHIM6 approaches we've seen a difference of approach, wherein the SHIM6 approach coops locators as identifiers on a per-host-pair basis, while the HIP approach uses a persistent identity value that cannot assume the role of a locator. The expectations at this stage of the study are that further ideas will surface here and that such ideas are helpful rather than distracting. It is unclear whether a single solution can emerge from this activity or whether different actors have sufficiently different sets of relative priorities so that multiple approaches, each of which expresses different prioritisation of functionality, are viable longer-term outcomes.

The critical consideration here is that it is unlikely that scaling routing over the longer term to a very much larger network is simply a matter of just changing the operation of the routing system itself. Real leverage in this area appears to also require an understanding of the meaning of the objects, or addresses, that are being passed within the routing system. The motivation for opening up the identifier/locator space within the Internet Area appear to be strongly tied to the notion that if you can unburden some of the roles of the addresses used in routing and can treat these routed tokens as unadorned network locality tokens, then

you gain some additional capability in routing. The intended outcomes include being able to group 'equivalent' locators together and thereby reduce the number of elements being passed within the routing system, ensure that the locator set readily maps into local forwarding actions, and hopefully, reduce the amount of dynamic change that is propagated in routing. It would also be useful if such an approach facilitates traffic engineering, site multihoming, various forms of mobility, and roaming. It might also be possible to remove from the application's end-to-end model the consideration of not just endpoint locality but also the tokens used in the transport protocol, thereby proving a different approach to IPv4 and IPv6 interoperability.

overall trend appears to be a system that is increasingly densely interconnected and carrying more information elements, each of which expresses finer levels of granularity in reachability. As an example of some of the relativities here, it was reported that the amount of address space advertised in 2006 increased by 12% from January 2006 to December 2006, while the number of advertised Autonomous Systems increased by 13% and the number of advertised prefixes increased by 17% over the same period. The report also considered the dynamic behaviour of the routing space, looking at various distributions of the 90 million prefix updates that had been recorded for the year. One of the major aspects of BGP updates in both 2005 and 2006

As an example of some of the relativities here, it was reported that the amount of address space advertised in 2006 increased by 12% from January 2006 to December 2006, while the number of advertised Autonomous Systems increased by 13% and the number of advertised prefixes increased by 17% over the same period.

At this juncture there is no unity or even clarity of the exact requirements of system design, let alone solutions for this work. Exploration of the interdependencies of mapping functions, the properties of identity and locator spaces, and the ways in which mapping functions can be supported in this environment is still at an early stage.

Routing ROAP: The Routing Area Meeting

The last of these ROAP sessions at IETF 68 was that of the Routing Area. The first part of the Routing ROAP session looked at trends in the routing system during 2005 and 2006. The

is the skewed distribution of updates, whereby, in 2006, 10% of the announced prefixes are the subject of 60% of the BGP updates, and 60% of the announced prefixes generate just 10% of all updates. By using known control prefixes, it appears that BGP appears to be an effective noise amplifier, whereby a single origin event can generate a considerably larger set of updates at the measurement point.

There appear to be two forms of dynamic BGP load: the BGP "supernova", which burst with an intense BGP update load over some weeks and then disappeared, and "background ra-

Continued on next page

More ROAP, continued

“mediation” generators that appear to be unstable at a steady update rate for months or even the entire year.

With respect to scaling of the BGP routing environment, it appears that one form of approach is to look in further detail at this subset of prefixes and ASs that are associated with the overall majority of BGP updates. One approach is to investigate whether damping of unstable prefixes in some fashion, or detecting routing instability that is an artefact of origination withdrawal, or deployment of propagation controls on advertisements would be effective in reducing the overall dynamic load of BGP updates. This approach represents a behavioural change in local instances of BGP that reduce the potential for unnecessary updates to be propagated beyond a need-to-know-now radius. Another approach is to consider changes to BGP in terms of additional attributes to BGP updates, such as a withdrawal-at-origin flag, or selective advertisement of next-best path, both of which are intended to limit the span of advertised intermediate transitions while the BGP distance vector algorithm converges to a stable state.

Again, the considerations of deployment were noted, where the Internet’s routing system is now a large system with considerable inertia. The implication is that any change to the routing system needs to use mechanisms that allow for piecemeal incremental deployment and whose incremental benefit is realised by those who deploy. One potential case study of such a change is the 4-Byte AS Number deployment.

It appears that we could improve our understanding of the operational profile of the routing space—particularly by looking at the various forms

of pathological routing behaviours and comparing these against the observations of known control points. Such a study may also lead to more-effective models of projections of the size of the routing space in the near-term and medium-term future and allow some level of quantification as to what the concept “scaling of the routing space” actually implies.

The second part of the Routing ROAP session took a look at the current status of the routing world, updating some of the observations made at the IAB Routing Workshop and outlining some further perspectives on this space.

One critical perspective on BGP is the behaviour of BGP under load. BGP uses TCP (transmission control protocol) as its transport protocol. This is a flow-controlled protocol, whereby the sender must await an advertisement of reception capability from the receiver (an advertised “window”) before being able to send data. When this session is uncongested, a BGP speaker sends updates as fast as they are locally generated (depending on the Minimum Route Advertisement Interval (MRAI) timer). When the transmission is congested, a local send buffer of queued updates forms. Unlike conventional applications that treat TCP as a simple black box, most BGP implementations use state compression on these update queues. As a simple example, the queuing of a prefix withdrawal should remove any already queued but as yet unsent prefix attribute updates for this prefix. This state compression of the advertisement queue should be on a peer-by-peer basis, so that a congested BGP peer does not slow down an uncongested peer. The implication is that the load characteristics of BGP alter as the load level increases, and BGP attempts to ensure that its peer receives the latest

state information only when the peer signals (via TCP flow control) that it is not keeping pace with the update rate.

Another critical factor is the nature of “convergence” in BGP. Convergence is at least an $O(n)$ -sized issue, where n is the number of discrete routing entries. This may appear daunting, but the real question is: How important is convergence? The presentation included the claim that this was BGP’s biggest, yet least important, problem. Convergence delays can be mitigated by graceful restart, nonstop routing, and fast reroute. One of the measures that exacerbates convergence is the use of route reflectors, whose model of information hiding is intended to reduce the number of BGP peer sessions and the total BGP update load, but what benefits they achieve come at the cost of slower convergence, with a higher message rate during the intermediate-state transitions. Perhaps it is appropriate to consider small-scale changes to BGP behaviour so as to mitigate the transient BGP update bursts caused by path hunting, including the already mentioned withdrawal-at-origin notification and propagation of backup paths.

One approach is to take the current set of potential tools that are proposed to address or that mitigate various BGP pathologies and prune this set by looking at those that align cost and benefit in deployment, allow piecemeal incremental deployment, and have beneficial changes to the load properties of BGP.

The approach advocated here is based on the perspectives that BGP is not in danger of imminent collapse and that there is still considerable headroom for BGP operation in today’s Internet. This allows the IDR Working Group of the IETF to focus on measures that include tools and

behaviours that tweak the current behaviour of BGP in ways that could mitigate some of the more excessive behaviours of BGP. And it gives the Routing Research Group the latitude to study the broader topics of fundamental changes that may be associated with novel routing and addressing architectures.

More ROAP?

So, is there some urgency here in looking at this problem? It's not clear that the problem is pressing, in that it is likely that the Internet will still be around tomorrow and probably the day after tomorrow as well. However, like many other issues in which there are complex feedback loops with internal amplification factors, it may not be apparent that there is a near-term problem with the health of the routing system until such time as the problems have already surfaced—and by then, dire warnings of impending trouble are just too late! Also, by the time that that stage is arrived at, there is no time to think about the various approaches to the space and the relative drawbacks and merits of each, because the pressure to simply deploy any measure to mitigate the issue becomes overwhelming.

The routing space is a classic example of the commons, where each party is free to generate as many or as few routing entries as it sees fit and is also free to adjust these entries as often as it sees fit. This allows each party to use routing to solve a multitude of business issues, including, for example, using routing to perform load balancing of traffic over a set of transit providers, or using a spot market in Internet transit services, or creating differentiated transit offerings by using more-specific routes and selective advertisements. The ultimate cost of these local efforts in optimising business

outcomes through the loading of the routing system is not necessarily a cost that is imposed back on the originating party. The ultimate cost lies in the increasing bloat in the routing system and the consequent escalation in costs across the entire network in supporting the routing system. There are no routing police, nor is there a routing market. There is no way to impose administrative controls on the global routing system, nor have we been able to devise an economic model of routing wherein the incremental costs of local routing decisions are visible to the originator as true economic costs for the business and wherein the benefit of conservative and prudent use of the routing system reaps economic dividends in terms of relatively lower costs for the business. Like the commons, there are no effective feedback mechanisms to impose constraint on actors in the routing space, and also like the commons, there is the distinct risk that the cumulative effect of local actions in routing creates a situation that pushes the routing system—either as a whole or in various locales—into a nonfunctioning state.

It appears that there are a number of avenues of approach here for efforts to place some constraints on the potential expansion of the routing system. What is less than clear is the ultimate value of such approaches in the context of the future Internet. Is making a functionally richer endpoint protocol stack a course of action that sits comfortably within a world of communicating RFID (radio-frequency-identification) labels? Is the lack of a routing market and an associated routing economy such a fundamental weakness that no technical efforts to alleviate the situation can gain traction in a world dominated by the desire to perform local optimisations in the cheapest possible manner? Have we already constructed a massive multi-trillion-dollar indus-

try that now uses business models that assume particular routing behaviours, and would efforts to alter those behaviours simply founder because of trenchant resistance to change in the business models within the communications industry?

Whether it needs a sense of urgency to motivate the work or a sense that there can and should be a better way of planning a future than via crude crisis management, the underlying observation is that the routing and address world is fundamental to tomorrow's Internet. Unless we make a concerted effort to understand the various interdependencies and feedback systems that exist in the current environment and understand the interdependences that exist between network behaviours and routing and addressing models, then I'm afraid the true potential of the Internet will always lie within our vision but frustratingly just beyond our grasp.

Yes, more ROAP, please! 

Further Reading

This is the set of references to further material on this topic, as presented in the plenary session.

<http://www.ietf.org/internet-drafts/draft-iab-raws-report-01.txt>

http://submission.apricot.net/chat07/slides/future_of_routing/apia-future-routing-john-scudder.pdf

http://submission.apricot.net/chat07/slides/future_of_routing/apia-future-routing-jari-arkko.pdf

<http://www3.ietf.org/proceedings/07mar/agenda/intarea.txt>

<http://www3.ietf.org/proceedings/07mar/agenda/rtgarea.txt>

<http://www1.tools.ietf.org/group/irtf/trac/wiki/RRG>

<http://www.ietf.org/IESG/content/radir.html>

IPv6 Update

By Mikael Lind

Note: This article does not provide a complete summary of all IETF activities in this area. It reflects the author's personal perspective on some current highlights.

IPv6 is quietly fading away as a topic. Instead, it is becoming an integrated part of all of the work within the IETF. There are still many remaining IPv6-specific problems to solve, but instead of being treated on the side by the IPv6 community, the majority of the problems are now part of different working groups covering both IPv4 and IPv6.

This is a good sign for IPv6, and it shows that the closing of the IPv6 working group was a sensible move and that IPv6 is becoming mature. IPv6 operations—one of the few remaining IPv6-specific groups—discussed the issue of address selection in a multi-prefix environment as the only IPv6 technical issue. The rest of the topics covered operational guidelines and experiences. One example is experience from network scanning that has been tracked in a university network that is running IPv6. A conclusion was that general network scanning to find active IP addresses doesn't exist but that port scanning of known hosts occurs in much the same way as with IPv4.

with IPv6-related problems is DHC (the Dynamic Host Control Working Group). Both operational issues—like prefix delegation—and new additions for IPv6 are on the agenda. The same can be seen in a large variety of working groups from mobility to routing. Trying to address the problem with the future requirements on routing is perhaps the best example of IPv6 being integrated. Even if IPv6 initially was intended to solve the problem of growing routing tables by aggregating prefixes, such is no longer the case. The introduction of provider-independent IPv6 addresses—due to the wish to maintain the operational models—has moved routing table growth issues

One working group that is busy

higher up on the agenda because aggregation no longer is a viable option. A lot of attention was placed on trying to initiate work on that problem during IETF 68 and the topic was brought up in several working groups and even at the plenary. The routing and addressing problem is not a new issue, but with latest developments it has come into a new light, and the work at this meeting is the first step into creating a ROAP working group. There are several views about the issue and if it is even an issue at all. Some think this won't create any problems because the development of hardware will keep up with growing demands, but even if that's the case, there is major concern regarding the cost involved for the operators who eventually will have to pay for this new equipment. Even if there is disagreement about the topic, it seems clear that it has to be studied in more detail. The fact that the focus for IPv6 is on operational issues and integrated IPv4 and IPv6 issues shows that IPv6 has moved into a new phase and is now starting to become a natural part of the daily life of the IETF and the network community as a whole. 

DHCPv6 Bake Off

By Alain Durand

Note: This article does not provide a complete summary of all IETF activities in this area. It reflects the author's personal perspective on some current highlights.

The DHCPv6 protocol has been in existence for several years, but only recently have a number of independent, production-ready implementations been available on the market. Initial lab tests pointed to some interoperability issues between codes of different origins. With the help of the Internet Systems Consortium (ISC), the office of the chief technology officer of Comcast decided to organise a DHCPv6 bake off (see sidebar, page 23) for the purpose

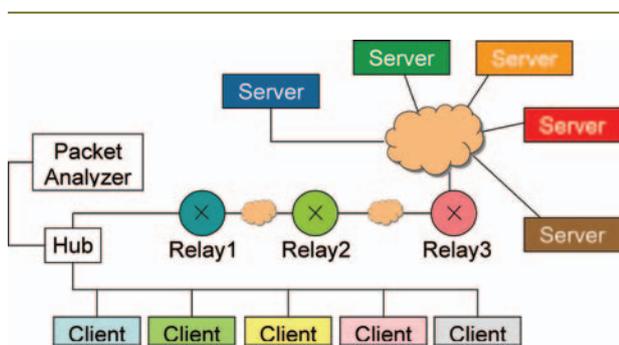
of testing as many implementations as possible for interoperability, operational impact, and usability.

The DHCPv6 bake off was held 14–16 March, just prior to IETF 68, at the RIPE NCC in Amsterdam. Fourteen engineers, representing seven vendors and/or open-source providers developing DHCPv6 software, gathered for two and a half days of intensive testing. One participant was connected remotely from California and was in constant contact via e-mail and instant messaging. The tests involved 13 independent DHCPv6 implemen-

tations: five clients, five servers, and three relays. The workshop enabled the participants to meet and discuss their software, assess whether it works with other implementations, find bugs, and exchange experiences and ideas.

The test plans were articulated around three axes:

- The first axis covered X-versus-Y tests, whereby each client was put in front of each server to verify basic interoperability.
- The second axis covered client tests in which the host interactions with IPv6 stateless autoconfiguration and router advertisements were explored.
- The third axis targeted relays/servers tests building increasingly complex network topologies, including a series of up to three relays, multicast relaying and high availability, and redundancy achieved through a set of anycast DHCPv6 servers.



The Tests

The key take away from the bake off is that the basic DHCPv6 technology works very well. One implementation, designed by someone who'd never attended an IETF meeting before, interoperated well with implementations performed by more-seasoned IETF participants, experiencing only minor bugs, which were quickly corrected on-site. This is certainly a trib-

ute to the quality of RFC 3315, the DHCPv6 specification.

However, problems with RFC 3315 arose, including 16 issues that were discovered and discussed in the course of the bake off. While the issues were varied and diverse, the most significant ones are described here.

- Interaction with router advertisement. An option describing the length of the prefix associated with the assigned address would be helpful.
- Interaction with the Domain Name System (DNS). The DHCPv6 client would benefit from feedback from the DHCPv6 server performing dynamic DNS update on its behalf.
- Client/server interactions and the semantics of the negotiation of certain parameters, which raises the question: What should happen when a client requests a particular address and the server does

not agree to the request?

- Relay/server interaction. What is the best way to keep track of the different levels of relaying, and how and when—if at all—should multicast be used?

The issues were again discussed with the DHC

working group during the IETF meeting in Prague, and the discussions will be documented through an Internet-Draft that could become the basis for an update to RFC3315.

Many thanks go to the RIPE NCC for organisational support of the workshop (networking hardware and staff to help run tests) as well as for the meeting venue; to ISC for the organi-

sation of the test plan; to the Comcast crew that made the bake off possible; and, of course, to the participants, who came from three continents, and to their employers.

Given the success of the bake off and the unanimous feedback from the participants, we are now considering organising a second DHCPv6 bake off to be held this coming fall, tentatively scheduled for the week before IETF 70 in Vancouver, Canada. The exact location is still to be determined.



Bake Off

The term bake off is jargon and stands for interoperability testing. See RFC1025 by Jon Postel for a more detailed explanation of the term:

In the early days of the development of TCP and IP, when there were very few implementations and specifications were still evolving, the only way to determine if an implementation was correct was to test it against other implementations and then argue that the results demonstrated that your implementation worked. These tests and discussions could, in those early days, as likely change the specification as change the implementation. There were a few times when this testing was focused, bringing together all known implementations and running through a set of tests in hopes of demonstrating the N squared connectivity and correct implementation of the various tricky cases. These events were called bake offs.

For more information, see <http://www.rfc-archive.org/getrfc.php?rfc=1025>.

Not Being There

By Geoff Huston

Note: This article does not provide a complete summary of all IETF activities in this area. It reflects the author's personal perspective on some current highlights.

Within the IETF, as much work as possible happens on its various mailing lists, so that attendance at IETF meetings is not an essential part of being an effective contributor to the work of the IETF.

Well, that's the objective, and, in general, that's been the case: critical decisions involving the milestones of the progress of IETF documents and the procedures we use are proposed, debated, and, as much as possible, concluded through the use of these mailing lists. But face-to-face meetings are still important. As one IETF attendee pointed out (on an IETF mailing list, by the way!) after IETF 68: "It's hard to be consistently effective in the IETF without attending the face-to-face meetings. A lot happens during IETF week."

But what if you just can't be there?

While I cannot claim to have attended all 68 IETF meetings, I have been able to attend most of the recent ones since the mid '90s, and I've found them to be very useful not only as a hard deadline for document editing but also as an opportunity to exchange ideas with others who share a common interest. However, I was unable to make it to IETF 68 in Prague, so I was using the IETF meeting's facilities for remote participation for the first time in many years.

The IETF has supported remote participation in its meetings for many years. At IETF 23, in March 1992, the IETF was experimenting in multicasting the meetings over the Internet. From the proceedings of this meeting (<http://www3.ietf.org/proceedings/prior29/IETF23.pdf>) I found this:

Packet Audio Experiment

Thanks to the organizing efforts of Steve Casner (ISI) and Steve Deering (Xerox), and the behind-the-scenes efforts of Van Jacobson and others, we had a very exciting demonstration of the DARPA packet audio experiment in San Diego. These stalwart experimenters set up IP multicast tunnels through the NSFnet backbone, and broadcast the Plenary proceedings of the IETF to multiple sites across the Internet. Sweden, UK, and Australia all took part in this exercise. We even had a brief 2-way communication, in which several remote listeners spoke to the assembled Plenary. The quality was not perfect. Some sites had much better reception than others. For some sites, the broadcast was apparently unintelligible at times. Still, for all its imperfections, this demonstration was an impressive promise of services to come. Some of us speculated that this new technology might play an important role in helping to deal with our future growth. For example, if the Proceedings of the Plenary (and perhaps even certain Working Groups) could be made available as a reliable Internet service, especially if it provided a robust 2-way interaction, it might give prospective attendees an alternative way to participate, rather than flying to attend the meeting in person. The "Information Age" could truly

be at hand! One of the traditional strengths of the Internet community is that we use the technology we are developing to assist that very development. This exciting packet audio demonstration offers the promise of adding to that tradition. We hope to see more packet audio, and perhaps even packet video, experiments at the IETF in the future.

So far, we haven't reached the level of support that creates an immersive telepresence for remote participants, but we have managed to make some progress with practical measures to support remote participation.

Multicast audio feeds, which started in 1992, continued for some years, to be replaced by a two-channel video and audio multicast feed by the late '90s. Beginning with IETF 63 in 2005, the streaming media moved to an eight-channel audio feed, also moving away from multicast to a unicast audio-streaming service. This means that all IETF working group sessions, BoF sessions, and plenary sessions are now part of the streaming audio service.

In recent times, Jabber has started to extend its pervasive reach into IETF meeting rooms, and these days—in addition to using a scribe to take the minutes of each working group session—the working group chairs are asked to find a willing volunteer to act as a jabber scribe, noting the meeting in real time. This was taken one step further in the SIPPING working group at IETF 68, with a stenographer providing a full transcript of the session in the jabber room in real time. Not only does this make the jabber service really useful as both a real-time feed and a meeting record, but also it is invaluable for non-native English speakers, because often, the written transcript provides a means to readily resolve some of the uncertainties that are caused by the variety of

accents, acronyms, unfamiliar use of terms, local dialects, and the speed of verbal delivery.

And working group chairs have become more aware of the need to publish both session agendas and the presentation packs to be presented in the session well before the scheduled time of the session, thereby giving the remote audience the opportunity to match the audio feed and the jabber log to the material being presented at the face-to-face meeting.

If you are familiar with the IETF format and familiar with the accents and colloquialisms of the usual contributors in your favourite working groups, then in terms of your being part of the audience, these tools are good enough to give you an excellent feel for what is happening at the

rather than a participant on the whole. Yes, you can attempt to ask questions into the jabber room, and—depending on the session and the state of the microphone queue in the room—you may be fortunate enough to have your question read out to the microphone and to have an answer provided. Simple questions can be posed—and answered—in this manner, but as an effective form of interaction with a dialogue between the remote participants and those in the room, this is not an ideal solution.

And then there's the other important part of the IETF face-to-face meeting: face-to-face interactions throughout the week. No, it's just not possible to participate remotely at the Social Event, which, I understand, was an especially notable event in Prague! And corridor conversations and of

actually the rich social fabric and the strong, even tribal, sense of "us" as a community of people with common interests that keep many of us contributing to the IETF for far longer than we may have originally envisaged. And for that, while not being there from time to time is probably unavoidable, it is the being there that makes working in the IETF all the more valuable and enjoyable. 

Remote Resources for IETF 68

- The Working Group and Plenary session agendas and presentations for IETF 68 can be found at https://datatracker.ietf.org/public/meeting_materials.cgi?meeting_num=68 while the material is being assembled during April 2007, and then the material will be published as proceedings at: <http://www3.ietf.org/proceedings/07mar/index.html>.
- The audio streaming for IETF 68 is described at <http://videolab.uoregon.edu/events/ietf/ietf68.html>, and the recordings for each of the eight channels are archived at <http://limestone.uoregon.edu/ftp/pub/videolab/media/ietf68>.
- The logs of the jabber chat rooms are referenced at the IETF text conferencing home page, at http://www.ietf.org/meetings/text_conf.html.
- And of course a large set of really useful tools can be found at <http://tools.ietf.org>.

If you are familiar with the IETF format and familiar with the accents and colloquialisms. . . . then in terms of your being part of the audience, these tools are good enough to give you an excellent feel for what is happening at the face-to-face meeting.

face-to-face meeting. Indeed it probably matches the level of sensory input you have while being in the room with your head down making notes on your laptop. These days the Internet is clearly robust enough to support an excellent quality of audio streaming, and the diligence of working group participants to use the microphones in the room and announce their names before speaking is really helpful for remote listeners. The online presentations allow you to keep track of the progress of the presentation, and the experience is pretty much the equivalent of being there.

What is still somewhat frustrating is that the sense of being there is limited to being a member of the audience

course the many fine lunches and dinners are not directly accessible either, when you're sitting over a laptop in the middle of the night.

Yes, the work of the IETF is meant to happen on the mailing lists, and physically being there at IETF meetings is not intended to be an essential precondition for effective participation in the IETF. And we do try hard to actually make that the case. But the IETF is as much about "us" as a community of people with a common interest in Internet technology as it is about mailing lists and processes, and while it's the unique and fascinating technical agenda that draws many of us into the IETF in the first place, it's

IRTF Report

By Aaron Falk

Below are summaries of several updates on the Internet Research Groups (RGs), as reported during the Technical Plenary at IETF 68.

Anti-Spam RG (asrg)

The asrg published a couple of drafts and had some energetic discussions about them on the mailing list:

- BCP on DNS-Based Blacklists
- Criteria for Proposed Techniques for the Management of Spam

Delay-Tolerant Networking Research Group (dtnrg)

There is a meeting planned in Dublin on 21–22 May 2007 for the purpose of making progress in the areas of routing, key management, and applications. The document titled DTN Architecture is in the RFC Editor queue. Another document, titled DTN Bundle Protocol Spec, has finished Internet Research Steering Group (IRSG) review. In addition, the RG review on LTP, a transport protocol for interplanetary-scale delays, has been finished.

End-Middle-End Research Group (eme)

The eme RG has circulated the first draft of a requirements document and is planning to meet at IETF 69 in Chicago.

End-to-End Research Group (end2end)

The End-to-End Research Group met at University College London, on 26–27 February 2007. The agenda included a half day revisiting the topic of congestion control. Again, some of the leadership of the Congestion Control Research Group and other experts on the subject joined and led a discussion on the most challenging problems in this space. The group spent another half day in discussion of future network architecture, including reports on the NSF FIND program and the Eiffel effort in Europe, and held additional discussion on the nature of architecture, the scope of the term network, the question of what it means (or whether one ought) to be doing so-called architecture research, and other related questions. The meeting concluded with several talks on miscellaneous topics, including revisiting an extension to the end-to-end questions.

Host Identity Protocol Research Group (hip)

The hip RG met in Prague during IETF 68 to discuss such issues as Lightweight HIP and HIP & Mobility. A number of topics—including Application Programming Interface (API), Network Address Translation (NAT), and other application drafts—have been moved to the IETF HIP WG. The stability of HIP implementations is improving. The RG is now concentrating more on experiments.

Internet Congestion Control Research Group (icrg)

The icrg met in Los Angeles on 12–13 February 2007. The RG is working with the IETF Transport Area directors on evaluating and publishing



Aaron Falk, IRTF Chair

the congestion control proposals. It is also developing a draft providing an overview of existing congestion-control-related RFCs. Many other issues are currently being discussed on the mailing list, such as fairness concepts and corruption loss. A meeting on these topics is planned for IETF 69 in Chicago in June.

Internet Measurement Research Group (imrg)

The imrg is planning to hold a workshop on 4 October 2007 to discuss classifying traffic by application or application type regardless of whether standard ports, tunnelling, or encryption is employed in an attempt to evade classification. Watch the IMRG and IPPM mailing lists for an announcement.

IP Mobility Optimisation Research Group (mobopt)

The IRTF RFC 4651—A Taxonomy and Analysis of Enhancements to Mobile IPv6 Route Optimisation—was published recently. The document, titled Unified L2 Abstractions for Fast Handovers, has completed RG review. The group is now working on location privacy, multicast, and mobility as well as policy implications on mobility.

Network Management Research Group (nmrg)

The RG organised a workshop in October 2006 to identify network management research challenges for the next five years. See a workshop summary in the plenary report on page 6. The group is also finalising a draft on Simple Network Management Protocol (SNMP) trace exchange formats and considerations. During its meeting at IETF 68, the group discussed adaptive monitoring work that permits trade-offs between overhead and accuracy.

Peer-to-Peer Research Group (p2prg)

A group member has authored a Survey of Research toward Robust Peer-to-Peer Networks. Also, the RG is involved in a very active discussion on a P2P session initiation protocol (SIP).

Routing Research Group (rrg)

The Routing RG has been rechartered and has two new co-chairs: Lixia Zhang and Tony Li. The RG met during IETF 68 and is now discussing goals and proposals for new routing and addressing architectures. (See also More ROAP on page 6.)

Scalable, Adaptive Multicast Research Group (samrg)

The samrg held an interim meeting in January 2007 at the workshop on Peer-to-Peer Multicasting (P2PM 07). There are currently three active drafts: problem statement, requirements, and hybrid framework. The group is also doing some preliminary experimental work on PlanetLab.

Transport Modelling Research Group (tmrg)

The tmrg submitted a Metrics document to the IRSG to be considered as an Informational RFC.

The IRTF as a whole is currently revising the IRTF RFC review and publication process. 

For more information about the IRTF, please see <http://www.irtf.org/>.

NomCom Results

Incoming IESG Members

Russ Housley**	IETF Chair/ General Area
Mark Townsley*	Internet Area
Ron Bonica	Ops & Mgmt Area
Jon Peterson*	RAI Area
Dave Ward	Routing Area
Tim Polk	Security Area
Lars Eggert*	Transport Area

Incoming IAB Members

Barry Leiba
Loa Andersson*
Kurtis Lindqvist*
Danny McPherson
Dave Thaler*
Lixia Zhang*

Incoming IAOC Members

Jonne Soininen*

* returning incumbent

** moved from Security Area

IETF Meeting Calendar

Summer 2007—IETF 69

22–27 July 2007

Host: Motorola

Location: Chicago, IL, USA

Spring 2008—IETF 71

9–14 March 2008

Host: Comcast

Location: Philadelphia, PA, USA

Fall 2007—IETF 70

2–7 December 2007

Host: TBD

Location: Vancouver, BC, Canada

IETF 72

27 July–1 August, 2008

Host: TBD

Location: Asia (Provisional)

Register now for

IETF 69

22–27 July 2007

Chicago, IL, USA

<http://www3.ietf.org/meetings/69-IETF.html>

Early bird registration: 600 USD (through 13 July 2007)

Regular registration: 750 USD

IETF 68 has been sponsored by:



The ISOC Fellowship to the IETF is sponsored by



IETF Journal

IETF 68
March 2007
Volume 3, Issue 1

Published three times
a year by the
Internet Society

4 rue des Falaises
CH-1205 Geneva
Switzerland

Managing Editor
Mirjam Kühne

Associate Editor
Wendy Rickard

Editorial and Design
The Rickard Group, Inc.

Editorial Board
Brian Carpenter
Leslie Daigle
Peter Godwin
Russ Housley
Olaf Kolkman

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

