IETF

# IETF Journal

IETF 70 • Vancouver
Volume 3, Issue 3
December 2007

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

***Inside this issue***



Internet Society

## Security and Protocols

***From the Editor's Desk, by Mirjam Kühne***

Even though IPv6 and related topics were, once again, being discussed at the most recent meeting of the IETF, another, higher-level discussion took place at IETF 70 in Vancouver, Canada, last December: How can one measure the success of a standard and protocol developed within the IETF? Dave Thaler and Bernard Aboba have studied that issue, and they offered a number of answers, some of which are surprising. A summary of their presentation appears on page 20.

The subject stimulated good discussions at IETF 70 and raises some interesting issues, particularly as it relates to Internet security and security protocols. While opinions may vary about whether security protocols developed by the IETF are successful, security remains a topic close to IETF's heart. For more than 10 years, every document has been required to include a section on security considerations. Still, the enormous amounts of unwanted traffic on the Internet cause concern. A few years ago, the Internet Architecture Board held a workshop on the subject. In this issue of the *IETF Journal,* we feature an updated summary of the workshop, including a number of important facts and notable observations. Also in this issue you'll read about João Damas's and Frederico Neves's solution to a long-standing security hole in the Domain Name System, which is described in their article "The Perfect Attack," on page 27.

Photo by Mirjam Kühne

Vancouver, site of IETF 70

> As with most success stories, to be wildly successful can be both good and bad. A wildly successful protocol is one that solves more problems or that addresses more scenarios or devices than originally intended or envisioned. However, if a protocol is used for purposes other than the one for which it was designed, there can be undesirable side effects—such as performance problems.
>
> — Dave Thaler and Bernard Aboba, page 20

Typically, the *IETF Journal* features short updates of the ongoing activities of Internet Research Task Force research groups. In this issue, we are pleased to offer more-detailed reports of those activities, including current work, achievements, and future plans.

We would also like to call attention to a number of newcomers who have contributed to this issue of the *IETF Journal.* One is Tomas Carlsson, who, in addition to an in-depth report on the IETF 70 fellows, offers an analysis of IETF culture (see page 18). Another is Bryan Ford, an MIT student who reports on new directions in the Transport Area (see page 29).

We thank all of our contributors to this issue, and we wish you fun reading. And, as always, we welcome both your comments and your contributions for future issues.

# Message from the IETF Chair

***By Russ Housley***

Russ Housley, IETF Chair

The IETF returned to Vancouver, Canada, in December 2007 for IETF 70. The Westin Bayshore, site of our previous visit to Vancouver, has excellent facilities for the IETF. With 1,114 people attending from 37 countries, the meeting was, by all accounts, successful, with progress made in many working groups (WGs). Cisco Research and Microsoft served as hosts for the event, and the site network was subcontracted to VeriLAN Networks. Sponsors included BC.NET, Eyeball, Huawei, and Telus. On behalf of the IETF, I'd like to express my gratitude and appreciation to our hosts and sponsors for their outstanding contributions. As usual, the IETF depends on a team of dedicated volunteers, which this time included a group of programmers who helped with the development of software tools that are used by the IETF on the Saturday before the meeting. The week was filled with the usual mixture of working group meetings, BoF (birds-of-a-feather) sessions, research group meetings, and, as always, many side meetings.

It was interesting to hear from Stephen Wolff of Cisco Research Center, who talked about the early days of the IETF. He was one of the 21 people who attended IETF 2! He recalled a time when 160 million packets per week was considered significant and when the first gigabit research networks were set up. He expressed hope that today's network research initiatives, such as GENI and FIND, will lead to similar advancements. He also mentioned that unsolicited proposals for network research are welcomed by Cisco.

Since IETF 69, 3 new WGs were chartered and 15 WGs were closed. Approximately 115 WGs are currently chartered. Since July 2007, the WGs and their individual contributors produced 421 new Internet-Drafts and generated 967 updated Internet-Drafts. The Internet Engineering Steering Group approved 106 Internet-Drafts for publication as RFCs. The RFC Editor published 103 new RFCs.

I'm happy to announce the winner of the IETF Secretariat services RFP. At the Wednesday evening plenary session, staff members from NeuStar Secretariat Services (NSS) received a standing ovation for their years of dedicated service. The winner was named too. Association Management Solutions (AMS) will begin providing Secretariat services in early 2008.

On a sad note, we recently lost a longtime IETF participant. Jun-ichiro "Itojun" Hagino passed away on 29 October 2007 at the age of 37. The "IPv6 Samurai" will be missed. He will be remembered for many things, including his contributions to the KAME project, which developed the IPv6 implementation that is now in FreeBSD, NetBSD, OpenBSD, and MacOS X.

On a happier note, at the Wednesday evening plenary session the IETF community offered thanks to Mark Foster, chief technology officer of NeuStar, for the pivotal role he played in the administrative restructuring of the IETF. Without Mark's assistance, the restructuring would have taken much longer and would have been much more painful. I personally appreciate his dedicated support to the IETF community. Again, thanks, Mark.

During IETF 69, one of the hot topics in the several sessions and many hallway discussions was IPv6 adoption, which remained a hot topic at IETF 70. The hope was to identify tasks the IETF can do to facilitate a smooth adoption.

*Continued on page 4*

## New BoF Meetings

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

*Internet Area*
csi: Cga & Send Extensions
savi: Source Address Validation Improvements
tictoc: Timing over IP Connection and Transfer of Clock

*Real-Time Applications and Internet Area*
peppermint: Provisioning Extensions in Peering Registries for Multimedia Interconnection

*Routing Area*
rl2n: Routing for Low Power and Lossy Networks

*Transport Area*
safe: Self-Address Fixing Evolution


Olaf Kolkman, IAB Chair

# Words from the IAB Chair

***By Olaf Kolkman***

The technical plenary during IETF 69 in Chicago did not include a technical presentation. During the open-microphone session at that IETF, a number of people in the audience expressed dissatisfaction with the lack of a substantive technical presentation. In contrast, IETF 70 in Vancouver featured two technical topics that were sufficiently thought provoking to stimulate lively debate.

The technical plenary serves as a forum in which the Internet Architecture Board (IAB) reports and receives feedback from the community. The IAB chair's report, the IRTF (Internet Research Task Force) chair's report, and the open-microphone session are fixed agenda items. These agenda items serve a purpose similar to those of the IAOC (Internet Administrative Oversight Committee) and the IESG (Internet Engineering Steering Group) administrative plenary sessions: they serve as a meaningful and effective way for the IAB to receive feedback from and undergo scrutiny by the community.

In addition to those agenda items, the IAB looks for presentations that inform the IETF of technical topics or developments the community should be aware of or that require further discussion.

Often, these sorts of presentations relate directly to work that is ongoing within the IAB. One example is the session on internationalisation that was featured at IETF 68 in Prague. (See http://www3.ietf.org/proceedings/07mar/index.html.) It was inspired by the earlier publication of RFC 4690.[1] Dave Thaler's presentation on protocol successes in Vancouver is another example; clearly, some of the protocols that we design have been much more successful than others. The IAB has been working to try to understand which factors lead to success. Its goal is to help make current and future protocol work more successful. The paper titled "What Makes For a Successful Protocol?"[2] represents the current state of our thinking (see page 20). We welcome yours.

Sometimes the topics presented at the technical plenary are of wide interest but not directly related to the IETF's or the IAB's agenda. An example of that was the presentation in Vancouver on power consumption of network elements, a topic that is starting to attract a lot of interest but that is not something that we in the IETF have heretofore considered when designing protocols or systems. We hope it was useful to hear someone who is actively researching power consumption reflect on the power consumption issues in IETF protocols and networks in general.

When selecting plenary topics and speakers, we always aim for presentations that are entertaining, informative, and thought provoking and that will lead to healthy group discussion. Of course, defining suitable topics and finding good speakers are continuing challenges for which the IAB welcomes suggestions.

## IETF 70 Facts and Figures

| | |
|---|---|
| Registered attendees | 1,114 |
| Countries | 37 |
| New WGs | 3 |
| Closed WGs | 15 |
| New Internet-Drafts | 421 |
| Updated Internet-Drafts | 967 |
| IETF Last Calls | 116 |
| Approvals | 106 |

*RFC Editor Actions (March–June 2007)*

138 RFCs published of which

- 76 standards tracks
- 4 BCP

*IANA Actions (July–October 2007)*

Processed ~1,600 IETF-related requests of which:

- 844 Private Enterprise Numbers
- 93 Port Numbers
- 126 TRIP ITAD Numbers
- 24 media-type requests

1. RFC 4690: "Review and Recommendations for Internationalized Domain Names (IDNs)," Klensin, Fältström, Karp, and the IAB. See http://www.ietf.org/rfc/rfc4690.txt.
2. "What Makes For a Successful Protocol?" D. Thaler, B. Aboba, and the IAB. See http://tools.ietf.org/html/draft-iab-protocol-success for the latest version.

# Plenary Report

***By Mirjam Kühne***

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 70 presentations can be found at http://www.ietf.org/meetings/past.meetings.html.*

Following a warm welcome by IETF chair Russ Housley, Stephen Wolff of the Cisco Research Center, one of the hosts of IETF 70 together with Microsoft, gave a presentation in which he reflected on Internet research.

Stephen's participation in the IETF goes back to its beginnings. Stephen recalled that at the second meeting of the IETF in April 1986—which was considerably smaller than today's meetings and which had a much smaller network—a presentation by Bob Hinden showed the actual size of the Internet: "131 Networks, 85+ Gateways, 160,000,000 packets/week."

At the time, there was not a lot of Internet-related research and there weren't many textbooks on networking. In fact, the entire library of Internet-related books would most likely have fit on one shelf.

However, even at that time, Stephen said, the Internet was a rich source of problems. On the technical side there were routing failures, collapses through congestion, fast long-distance networks, and lack of security. But there were non-technical problems as well.

In his presentation Stephen described one particular research project—the Gigabit Testbeds (1990–1995)—that was of particular interest. Done in cooperation with CNRI (the Corporation for National Research Initiatives) and funded by 20 million USD over five years, Craig Partridge called it "How Slow Is One Gigabit per Second?" With a total of five testbeds, the project had mixed success, but according to Stephen, the community gained a much better understanding of the challenges of speed over long distances.

Much has happened since. Today there are two big research projects: (1) GENI (Global Environment for Network Innovation), which, with a budget of approximately 367 million USD, is much larger in scope than the Gigabit Testbeds project but similarly organised, and (2) FIND (Future Internet Network Design), a major, new, long-term initiative of the NSF NETS research programme and that has a budget of 30 million–40 million USD per year. Together they add up to more than 40 different projects. "Many of them are good," said Stephen, who encouraged everyone who is interested in research to look over the programmes at http://find.isi.edu/. There also is the work of the Internet Research Task Force (IRTF). A number of IRTF research groups (RGs) have funded FIND proposals, such as dtnrg, eme, end2end, mrg, p2prg, and rrg. The Crypto Forum RG (cfrg) seeded the GCM mode for IPSec (RFC 4106) and UMAC message authentication (RFC 4418).

Stephen closed his presentation by welcoming unsolicited proposals to the Cisco Research Center. For more information, see http://www.cisco.com/research.

Following Stephen's presentation was the Network Operations Center (NOC) report, which was presented by Morgan Sackett of VeriLAN. The NOC was again run by VeriLAN staff and volunteers. Upstream connectivity was provided by Telus and BCNET. An IPv6 tunnel was set up to ISC. Once again, the NOC did a fantastic job. The network operated without disruptions during the entire meeting.

Lakshminath Dondeti, chair of the Nominations Committee (NomCom), introduced the new NomCom members and gave a status report. The NomCom regularly sends requests for feedback about the various candidates. The return rate is, at best, 12 percent and, at worst, 4 percent. This is not good enough. More

---

*Message from the IETF Chair, continued from page 2*

In the past, the IETF has taken the approach that IPv6 adoption would happen naturally, before the IPv4 address space was exhausted. However, there is an increasing realisation that this is not the case. The last IPv4 address block will most likely be allocated by the Internet Assigned Numbers Authority before widespread IPv6 deployment occurs. There were a number of varying opinions and lively discussions on the topic. Ultimately, no consensus was reached on what the IETF can do right now to expedite IPv6 deployment. It is clear to me that there will continue to be much speculation and energetic debate, but I continue to believe the IETF has a valuable contribution to make in this area.

I look forward to seeing all of you at IETF 71 in Philadelphia on 9–14 March 2008 and at IETF 72 in Dublin, which is scheduled for 27 July–1 August 2008. Scheduling information for upcoming IETF meetings may always be found at http://www.ietf.org/meetings/meetings.html.

feedback is needed in the future. Community input is crucial for this process to work effectively.

Henrik Lewkovetz put together some extremely useful tools. Lakshminath thanked Henrik and the NomCom members for all the work they have put into this process.

## Recognitions

A number of people were recognised for their contributions to the IETF and the Internet. Mark Foster, chief technology officer of NeuStar, was recognised for his pivotal role in the administrative restructuring of the IETF. Without his assistance, the restructuring would have taken much longer and would have been much more difficult.

## Jon Postel Award

The Jon Postel Award Committee announced that the 2007 Jon Postel Award was given to Nii Quaynor for "his vision and pioneering work that helped countless others to spread the Internet across Africa."

The award is traditionally presented to an individual who has made outstanding contributions in service to the data communications community and to honor an individual who, like Jon Postel, has provided sustained and substantial technical contributions, service to the community, and leadership. With respect to leadership, the committee places particular emphasis on candidates who, in addition to their own individual accomplishments, have supported and enabled others to achieve success.

Nii established the first Internet services in West Africa in 1993. He is the founding chair of AfriNIC, the African numbers registry, and has been convener of the African Network Operators Group (AfNOG) since 2000. Earlier in his career, Nii established the computer science department at the University of Cape Coast in Ghana in 1979. He was awarded a Ph.D. in computer science from the State University of New York at Stony Brook in 1977 and worked at Digital Equipment Corporation from 1977 to 1992. Currently, Nii is chair of Network Computer Systems in Ghana and professor of computer science at the University of Cape Coast.

Nii thanked ISOC and the IETF, saying he felt humbled by the award and by what it represents. "Africa thanks ISOC and the IETF for this recognition. Africa will be very pleased with this contribution." He thanked his colleagues in Africa who supported his efforts and pushed him along. He also recognised the IETF community for contributing in such areas as how the number and name resources have been defined, which has helped Africa's underlying understanding and its self-organisation.

Nii plans to use the award of 20,000 USD to establish a new fund for technical engineers in Africa. The fund will be managed by AfriNIC and AfNOG.

## Stats and Updates

In his IETF chair report, Russ Housley provided some meeting statistics as well as an update on IANA and RFC activities. He also thanked the team of volunteers responsible for the audio streaming. The IETF received outstanding support from the Network Resource Startup Center and the University of Oregon.

Kurtis Lindqvist, chair of the IETF Administrative Oversight Committee (IAOC), reported that the Association Management Solutions (AMS) secretariat has won the RFP for the IETF Secretariat services. A transition plan is being worked out. The IAOC thanked the staff of NeuStar Secretariat Services. The 2008 IETF budget was submitted to the ISOC Board of Trustees and was approved shortly after IETF 70.

Ray Pelletier, IETF administrative director, announced that contracts with venues are already in place for meetings in 2008. In fact, the entire meeting-planning process is now happening with much longer lead times. After a survey of the IETF community, it was decided to follow a 3-2-1 model with respect to meeting locations: Within two years there will be three meetings in North America, two meetings in Europe, and one meeting in Asia. This seems to be appropriate for meeting the needs of attendees.

At the end of his presentation, Ray acknowledged the ISOC Fellowship Programme, which brings engineers from developing countries to IETF meetings. All costs are covered by ISOC. Each fellow is paired with an experienced IETF participant, who acts as a mentor for that fellow. Many thanks to the mentors and sponsors of the programme. (See page 9 for more information about the ISOC 70 fellows and mentors.)

Photo by Peter Lothberg

Jon Postel Award Winner Nii Quaynor

## Open Mic

A short discussion regarding tools development took place at the beginning of the open-mic session, which was directed mainly toward the IAOC. While tools development and maintenance falls within the responsibility of the IETF Secretariat, the volunteer effort to develop tools is still seen as critically important both to save money and because it is a hands-on community effort. Some

*Continued on next page*

*Plenary Report, continued from page 5*

people would like to see a plan for moving forward with tools development and maintenance and learn more about how the plan will support the community. The IAOC is working on both the plan and a license agreement.

With regard to the IETF Trust, Ray said that while the IAOC has not kept an inventory of the nearly 2,000 RFCs that have by now been signed to the Trust, all names are listed on the IAOC Web site. In addition, businesses have signed their RFCs over to the Trust, which means that all documents that have been published by employees of a company are automatically signed over to the Trust. It is estimated that approximately half of all RFCs are by now signed over to the IETF Trust, which is a positive development.

## Discussion on NAT and IPv6 Continues

Network address translation (NAT) was a topic again raised during the plenary session. The behave working group (WG) was chartered to define how NATs can behave more reasonably and according to specification. One of the properties would be incremental deployability. One speaker was concerned that incremental changes to NATs would be the wrong approach. On the other hand, there are currently ongoing discussions within the STUN (Simple Traversal of UDP through NATs) and ICE (Interactive Connectivity Establishment) communities that describe why a general solution will not work. Those discussions will have to be continued by the appropriate working groups.

IPv6 also remains a big topic for the IETF. Even though at this IETF meeting, IPv6 deployment was not officially on the IESG or IAB plenary agenda, there were several discussions both during various working group meetings and in the hallways. IESG member Ross Callon said he hopes that at some point "the pain will be high enough to deploy IPv6."

Sam Hartman, one of the two security area directors, disagreed, saying, "IPv6 is being incrementally deployed and is catching on where it has value." He wondered whether it really makes that much sense for everyone, at some point, to switch over to IPv6. Another attendee added that deployment of IPv6 is not always straightforward, leading to agreement that better documentation is needed and that the IETF community could help with that.

Jari Arkko, an IESG member who is active in the area of IPv6, reiterated that the IETF can also help by making sure all the necessary pieces are in place so that users can deploy IPv6. The v6ops WG is looking into whether all transition mechanisms are in place. Outside the IETF, education is needed, and the IETF is working with ISOC to address that issue. Overall, it was expressed that it is necessary to understand that no true transition is possible. IPv4 and IPv6 are disjoint address spaces. Proper mechanisms for moving between the two versions are essential.

What followed was a discussion about what would motivate people to use IPv6 in their networks. Some people believe only a killer application or more features will help. Others disagree, saying nobody is going to develop an application that runs only on an infrastructure that hardly anyone uses yet. The biggest benefit of IPv6 is a much bigger address space. Alain Durand summarised it as follows: "The motto for many years has been 'Bandwidth, bandwidth, bandwidth.' Now the motto has changed to 'Address space, address space, address space.' It's as simple as that." However, the immediate address shortage of IPv4 has been fixed by introducing NAT. Large corporations that need a lot of address space are starting to use IPv6, now that there aren't enough IPv4 addresses.

Adiel Aklogan, CEO of AfriNIC, the African numbers registry, agrees that IPv6 is indeed happening and that the IETF has been doing a lot to encourage people to use it. "All RIRs are working with their communities on that," he said. "Maybe the IETF can help by sending the right message to the operators community that the protocol is ready."

Echoing much of the discussion at IETF 69 in Chicago, it was concluded in Vancouver that most of the IETF's work on IPv6 has been completed. What is left to do is education. And vendors need to be encouraged to implement IPv6. "There are some bugs and issues with IPv6 equipment," said Jari. "With more users, those will be fixed faster. It's a matter of time."

## Technical Plenary

The technical part of the plenary session at IETF 70 was devoted primarily to two technical presentations. The first one, called What Makes For a Successful Protocol, was presented by Dave Thaler (see page 20). The authors were applauded by attendees for their excellent work in this area, and the presentation was followed by a constructive and lively discussion. It was suggested that not only the IESG but also working groups need to be paying attention to this work so that newer protocols have a better chance at becoming successful. Economical alignment and deployability are now already being used as criteria for successful protocol design in some WGs. This is a positive development, as Olaf Kolkman noted.

Leslie Daigle warned that technical superiority is not necessarily a factor for becoming successful. As she explained, older protocols that were brought into the IETF could also be more successful because at that time new work was more elastic and more experimental in nature. It was easier to bring in new ideas. "Nowadays we have to check if things are successful before we start, looking

"The motto for many years has been 'Bandwidth, bandwidth, bandwidth.' Now the motto has changed to 'Address space, address space, address space.' It's as simple as that."

— Alain Durand

at them in order to not bring the whole industry down," she said. "This is a challenge."

Dave Thaler emphasised that success often doesn't become clear until sufficient time has passed. "In hindsight it's easy to tell what is successful," he said.

"It's much more difficult when you're designing it. Sometimes you just can't tell. It may be successful later for other reasons," Bob Hinden added.

Some interesting suggestions were made during the discussion: For example, one could identify those protocols that were developed outside the IETF and what factors influenced the decision to not take them on as IETF work items. Some of them became successes. One could look at those cases and see why they became successful and why they were developed elsewhere.

Another speaker mentioned that one reason protocols are successful outside the IETF could be that protocols have a high turnover rate and suggested that perhaps one should look at how to adopt protocols into the IETF.

There are also cases where efforts were made to kill a protocol but it survived nonetheless. What were the reasons for that? Is it possible that working groups sometimes try too hard to predict what is going to be successful?

Olaf gave DNSSEC some thought in that context. "This has been in the IETF for a long time," he said. "One success criterion is whether there is a perceived benefit. This is very hard to sell for a security mechanism. There is very little the IETF can do except make the case that some things are important. Then the marketplace decides."

"There are always various demands on a protocol," said Mark Crispin in closing. "The processes of the past cannot be applied today. Are there other organisations that have a faster turnaround and the same diversity as we have? Our diversity is our strength."

The second technical presentation covered a topic that is a bit unusual for the IETF but was received positively. It was called Energy Engineering for Protocols and Networks, and it was presented by Bruce Nordman, a researcher at the Environmental Energy Technologies Division of Lawrence Berkeley National Laboratory.

"Why might we care about energy engineering?" asked Elwyn Davies as he introduced Bruce. "Is there a way—when we design protocols—to keep the amount of total energy use down?"

Bruce presented a number of statistics demonstrating that most energy is used at the edges and that energy use is affected by applications and protocols—and not just hardware. A number of research projects are exploring that topic, including a project called Energy-Efficient Ethernet and one called Network Presence Proxying, which focuses on the significant amount of energy that is used when devices are idle. Bruce is currently working with the industry to draft the content of a proxying standard. For more information, see http://www.ethernetalliance.org/technology/white_papers/. A related initiative is an NSF/FIND project called Energy-Aware Network Architecture (http://www.icir.org/mallman/research/proj-energy-arch.html).

What can the IETF do to help reduce energy usage? Bruce made a number of suggestions that IETF engineers and network operators could act on to save energy:

- Facilitate multiple forms of reduced presence (instead of always assuming full presence of all edge devices)
- Enable optional reduced speeds (instead of assuming network links always run at maximum speed)
- Expose knowledge of acceptable latencies
- Determine when and how to facilitate slower acceleration
- Facilitate powering-down links when capacity is not needed (instead of maximising interconnections)

The IETF can also ask itself which existing and developing protocols have features that inadvertently work against energy savings and, conversely, which protocols facilitate energy savings. "Could there be some guiding principles that might ensure that protocols maximise energy?" he asked. "And could existing protocols be modified to follow those principles in future revisions?" In closing, Bruce suggested the IETF community make energy savings an integral part of protocol design, just like security.

The IETF has two WGs dealing with related issues: 6lowpan (IPv6 over Low-Power WPAN) and, possibly, rl2n (Routing for Low Power and Lossy Networks). The latter was a BoF at IETF 70 and might soon become a WG.


Photo by Mirjam Kühne

Sidewalk signage in Vancouver

# An Interview with ISOC Fellow Subramanian Moonesamy

***By Tomas Carlsson***

Tomas Carlsson

"Now I got to put a face to all strong people from the mailing lists," said Subramanian Moonesamy, the man from exotic Mauritius. While the rest of the world dreams of visiting this mythic island country, SM, as he is called, had one of his dreams fulfilled when he attended IETF 70 in Vancouver as part of the ISOC Fellowship to the IETF programme.

Think of Earth as a ball of wool. Then think about pushing a knitting needle in where Vancouver would be. At the other side of the ball is where Mauritius is located, a 2,000-square-kilometre island east of Madagascar. The knitting needle is a good metaphor, representing the ideal connection to the U.S. IP backbone. For more than five years, SM has been active in several IETF mailing lists. And while he has made comments on and suggestions for drafts, he never thought he would ever attend a meeting. "It seemed very far away," SM said.

Photo by Tomas Carlsson

ISOC Fellow Subramanian Moonesamy

"I have to travel for 24 hours to get to the United States. And such a trip is also very expensive."

The ISOC fellowship programme made the trip possible, and the *IETF Journal* took the opportunity to interview a person from a country generally unknown within the IP world.

SM is aware of his uniqueness. One of the first questions he would ask was, "Do you know something about Mauritius?" Most of us would have to admit that all we knew was gleaned from vacation advertisements, where the country is described as an island paradise.

For SM, the visit to Canada was quite a contrast. He had encounters with snow, with cold days, and with ice on the pavements. Inevitably, SM caught a cold after a few days.

We met at an Indian restaurant on a main street of the city's center where we talked a bit about Mauritian culture, including the country's food traditions. The mix of African and Asian food sounded promising. SM served as a good ambassador for Mauritius.

Unavoidably, the talk turned to the Internet and its current state in Mauritius. Even if Internet penetration among the 1.3 million people living in Mauritius seems substantial, less than 23 percent of the population has broadband connections.

According to SM, there are three Internet service providers and one national gateway. The administrator of the top-level domain is a private company called Internet Direct. For each name registered, it charges 65 USD, a price that is not likely to speed up Internet usage. Actually, when we checked the registration fee of other registrars, it was even 267 USD. SM was reluctant to talk about his opinion on this and said he better not comment on it.

SM is a freelance consultant for Eland Systems. He has chosen not to become an employee, because he's in favor of the freedom that being independent brings.

"It means that I can be 'nice' to my customers," he said. "I do not have to sell expensive solutions, but, rather, the best one. Then customers come back, and in the long run it is a better business."

SM designs appliances for virtual private networks, firewalls, antispam, and e-mail used by bigger companies. His mission, he said, is to assemble software and hardware, do some code authoring, and then adjust it to the requirements of the customer. He finds it very important to be up-to-date with the latest standards, and he follows with great interest certain working groups, including SMTP, DKIM, EAI, DNSOP, IPR, SIEVE, and SASL.

"I found that a lot of work gets done when people meet in the corridors. Sometimes I found people to be younger than they seemed to be when I was only reading their mail," SM observed.

He talked about the lack of body language in e-mail and how a physical meeting makes it possible to see that people are not angry even when they disagree.

"I found the people in the IETF working group meetings open, sharing, and cooperative," he said, "just like Internet people used to be. They were also surprisingly open to different cultures."

In the past, SM travelled mostly to African countries for business. For him, the trip to Vancouver opened up a wider working field. If the world of the Internet doesn't currently run to Mauritius, at least Mauritians can reach out to the world.

*Tomas Carlsson is freelance journalist who specialises in writing about Internet technologies in Sweden.*

**Report from ISOC Fellowship Programme**

# Fellows Motivated to Become More Involved

***By Tomas Carlsson***

Becoming more active in IETF working groups is the goal of all five participants in the ISOC Fellowship to the IETF programme after their visit to Vancouver for IETF 70. "I will spread more information about the IETF to colleagues at home," said Fellow Pedro Torres.

Photo by Tomas Carlsson

IETF fellows and mentors in Vancouver

The *IETF Journal* chatted with each of the fellows, focusing on their perceptions of the meeting and on the status of the Internet in their countries.

Pedro manages an Internet-based academic backbone as well as an Internet exchange point and a metropolitan area network. His home city of Curitiba, Brazil, is in the Parana region, which is smaller in poulation than all of São Paulo. Pedro is concerned about the relatively few IETF-meeting participants from Africa and South America. "We use the standards but do not participate in creating the solutions," he said.

He hopes to increase his knowledge of the work of the IETF among his colleagues as well as in his geographic region. He also hopes to develop better relations with other countries, both through person-to-person contacts within the industry and through the creation of more and better access points. Those ambitions explain his interest in LAPLA/LACNIC—the NIC of Latin America and the Caribbean—and its associated mailing list.

Eduardo Ascenço Reis of São Paulo is a network analyst at CBTC, a telecommunications company in Brazil. He faces considerable challenges at a higher level of networking, which explains his interest in routing, IPv6, and TCP management. "I have been fairly passive in the working groups until now," he said, "but with the help of my mentor, Scott Brim, I feel more integrated in the groups."

Frederico Faria works with customers throughout South America and the Caribbean. "Not many people in Brazil are aware of how the IETF working groups gather folks from different parts of the world to discuss Internet development," he said. "The discussions are open and mature, and the people are welcoming and encouraging. However, I have noticed that the IETF lacks strong links between working groups. Some are trying to solve the same problems, which could be avoided with more cooperation among the groups."

Veaceslav Sidorenco's home country of Moldova has only 22,000 broadband connections. Internet penetration is near 20 percent, an exceptionally high number considering that only 30 percent of the population has wired telephones. Veaceslav has been committed to realising the potential of the Internet since the 1980s. While he is now a UNDP-expert in the government and has given classes in RFC theory, he has never before participated in an IETF working group. He is now helping with the management of RENAM, the Research and Educational Networking Association for 8 universities and 20 research institutions. He is also involved in building a network operation centre and a computer emergency incident response centre, and he is actively participating in an effort to start an ISOC chapter in Moldova.

In addition to the challenges of understanding IETF processes and culture as a newcomer, some of the fellows face challenges even in just getting themselves to the meeting. Processing times and requirements for obtaining travel visas can be long and onerous, especially for individuals from developing countries. Since the inception of the programme, a few fellows have needed to reschedule their meeting attendance because of such delays.

"ISOC assists the fellows before and during the visa application process, but sometimes the consulates are very slow in responding," said Karen Rose, ISOC's director of education and programmes. "We have added more time into our programme schedule to minimise this kind of disappointment, but if a problem arises, we always offer the fellow an opportunity to attend the next IETF meeting."

### ISOC 70 Fellows and Mentors

Frederico Faria, Brazil (mentor: Frederico A. C. Neves, Chief Technical Officer at Nic.br, Brazil)

Subramanian Moonesamy, Mauritius (mentor: John Klensin, independent consultant, U.S.)

Eduardo Ascenço Reis—Brazil (mentor: Scott Brim, Senior Consulting Engineer, Cisco, U.S.)

Veaceslav Sidorenco, Moldova (mentor: Jaap Akkerhuis, Network Research Engineer, NLnetLabs, the Netherlands)

Pedro Rodrigues Torres-Júnior, Brazil (mentor: Henk Uijterwaal, Senior Project Manager, RIPE NCC, the Netherlands)

# Unwanted Traffic

*By Elwyn Davies*

The Internet carries a lot of unwanted traffic today. At its most fundamental, unwanted traffic is made up of packets that consume network and computing resources in ways that do not benefit the owners of the resources. To gain a better understanding of the driving forces behind such unwanted traffic and to assess existing countermeasures, the Internet Architecture Board (IAB) organised a workshop in March 2006 called Unwanted Internet Traffic. At the workshop, a number of experts—including operators, vendors, and researchers—exchanged experiences, views, and ideas on this important topic (the full report of the workshop was published in RFC 4948). This article presents the findings of the workshop and looks at some developments that have occurred since the workshop.

## The Underground Network Economy

The most important message from the Unwanted Internet Traffic workshop was that the enormous volume of unwanted traffic is a symptom of a vast criminal underground economy that is a parasite on both open technology and the innovative culture of the Internet as it has developed over the past 20 years.

### *From Anarchy to Criminality*

Early in the life of the Internet, unwanted traffic was largely an expensive nuisance. Much of it was generated by so-called script kiddies, who had no clear motive beyond demonstrating to their equally mindless peers their ability to cause mayhem. While the consequences for the networks and hosts that were targeted were generally immediate and catastrophic, often resulting in significant economic loss for the victims, the attackers profited little or, in most cases, not at all.

Over the past few years, the situation has altered dramatically. The anarchic hackers of the past have been harnessed or have been displaced by criminals who seek to use the Internet for illicit gain.

The underground network economy that has developed within the Internet mirrors the underground economy in the physical world: tools of the [criminal] trade are created and sold to other criminals; stolen information is fenced for use in further criminal activity; and routes are created through which the illicit proceeds can be laundered to enable the criminals to benefit from their activities.

The underground network economy has evolved quickly, changing from an initial barter system into a gigantic shopping mall for tools and information. This has led to a rapid shift in the nature of unwanted traffic and the ways in which the traffic affects the network. It is now a fully integrated and persistent subculture that sucks many billions of dollars out of the legitimate network economy by exploiting the commercial growth of e-business. It is no longer in the interests of these types of criminals to destroy or significantly damage the network; as with any parasite, the parties responsible are absolutely dependent on the continued existence and availability of the network to supply their income.

### *Subverting the Network*

The marketplace for the underground network economy is typically hosted on IRC (Internet Relay Chat) servers that provide access to "stores" that sell the tools that are needed to operate in the underground economy. Easily available is strong encryption software for e-mail and other communications tools, both of which allow deals to be closed with little risk of detection. Consequently, it is no longer necessary to be a skilled programmer to be a successful miscreant in the underground economy. The malware, bot code, and access to compromised hosts or Web servers can be bought off the shelf, and some of the profits can be used to finance new tools and to set up "dirty" Internet service providers to host IRC servers and fraudulent Web sites.

The network itself provides the means to turn the available tools and stolen information into real assets. In the simplest case, electronic funds transfer can be used to drain money from online bank accounts directly into short-lived accounts—often in another country, which makes it difficult to trace or recover the money. More-sophisticated schemes use stolen credit cards to purchase goods that are redirected or resold through money-laundering services that obscure the trail that leads to the beneficiary. The international nature of the Internet, the absence of audit trails, and the ease with which anonymity can be achieved are important features of the network, but they also facilitate misuse.

One of the key weapons used by criminals consists of compromised hosts, also known as *bots* or *zombies*. Networks of bots (*botnets,* for short) are created by exploiting security flaws in networked machines or by inducing naive users to install in their machines certain backdoor remote control capabilities of which they are unaware. Remotely controlled bots can then be used either as means of capturing valuable personal or financial information from the users of the machine or as ways of generating further unwanted traffic, such as e-mail spam or distributed denial-of-service (DDoS) attacks that cannot easily be traced to their true origins. In most cases, bots do not cause major disruption to the hosting machine by either obviously disrupting operations or clogging the machine's network connection with large amounts of unwanted traffic. The objective in most cases is to provide a resource that can be used by the miscreants for as long as possible. To make a medical analogy, unwanted traffic no longer creates an acute disease in the compromised host; rather, it creates chronic carriers that

may go undiagnosed for a long time and that act as sources of infection that can perpetuate the problem.

A major reason that the underground economy is so successful is the ease with which botnets can be created. Miscreants view them as expendable resources, and they are rarely bothered by operators who may see what they're doing. As long as their cash flow is not significantly impacted, miscreants simply move on to new venues when ISPs take action to clean up bots and protect their customers. However, taking out one of the IRC servers might provoke a severe and ruthless attack on the ISP, typically through the use of botnets to launch a DDoS attack targeting the ISP's network. In this way, the attackers create an example that might intimidate other ISPs into leaving them alone.

### *Simplicity and Power versus Vulnerability and Ignorance*

The end-to-end architecture of the Internet emphasises the flexibility of implementing new applications in the end system while keeping the network itself as simple as possible. The network neither enhances nor interferes with end system data flows. The success and adaptability of the Internet demonstrate the power of this model but can also make life easy for those who operate in the underground economy.

The concentration of capabilities in a large number of end hosts means there is an enormous field of complex systems available for launching an attack. Inevitably, complex systems are difficult to analyse and protect. Consequently, it is not surprising that the majority of hosts are to a greater or lesser extent vulnerable to compromise. Miscreants maximise the return on their investment by exploting vulnerabilities in the most common platforms, such as Microsoft Windows; the volume of exploits reported is a measure of the system's market penetration rather than its lack of security.

Many of these complex systems are owned and controlled by ordinary people who come from all walks of life and who eagerly jump into the exciting online world but are rarely given the training to fully understand the implications of the systems they own. The operating systems and applications they are using are generally designed to hide the complexities of the system so that the users are not deterred from making use of the system. As a result, a large proportion of users fail to anticipate how such a great invention as the Internet can be readily abused, and they do not understand that their system can be compromised without their being aware.

It is therefore not surprising that the Internet now has a considerable number of compromised hosts where the owners are not aware that a compromise has happened. Although a large percentage of those machines are home PCs, evidence shows that corporate servers or backbone routers—even government firewalls—have also fallen victim to compromise.

### *Running under the Radar*

Although some of the consequences of the flood of unwanted traffic—such as spam e-mails and DDoS attacks—are all too visible, many other types of unwanted traffic are hard to detect and counter.

Hosts are now quietly subverted and linked to botnets while leaving their normal functionality and connectivity essentially unimpaired. Bots and the functions they perform are often hard to detect—especially since owners and operators are oblivious to their presence. And detection may well come too late, because the bot may have already carried out the intended (mal)function.

The presence of large numbers of quiet bots in compromised hosts is a particularly challenging problem for the security of the Internet. Not only does the resulting stolen (financial) information lead to enormous economic losses, but also there does not appear to be a quick fix for the problem. The fix needs to be applied at places that see little or no local benefit from the solution. For example, the owner of a machine infected with a bot may not care about fixing the problem if the bot has negligible impact on the way the machine performs for the owner. As long as the owner can keep playing online games, the owner may not be interested in applying a time-consuming and potentially technically complex fix, even though the public interest is endangered.

#### Editor's Note

This article is based on discussions at an IAB workshop held in March 2006. The full report of that workshop has been published as RFC 4948.

Work to address the issues here is active and on-going.

The IETF has the following working groups addressing some of the problems identified here:

- Operational Security (OPSEC) WG
- Routing Protocol Security (RPSEC) WG
- Secure Inter-domain Routing (SIDR) WG

Much of it extends well beyond the technical sphere of IETF specifications. See, for example, some efforts by the U.S. Federal Bureau of Investigation (http://www.fbi.gov/pressrel/pressrel07/botnet061307.htm). A commentary on that effort is available from Arbor Networks, which is engaged in measuring, researching, and proposing paths forward (http://asert.arbornetworks.com/2007/06/who-ya-gonna-call/).

Further information and resources are available from Team Cymru Resources (http://www.cymru.com/).

*Continued on next page*

*Unwanted Traffic, continued from page 11*

Simplicity at the core of the network and the nature of the routing system can also make life easier for attackers. IP is specifically designed to minimise the amount of state information needed in the data plane to forward traffic from one end to the other. The network core does not record audit trails for individual traffic streams unless special measures have been planned in advance, such as when the police request lawful interception of some particular traffic.

A major strength of the Internet is its ability to provide seamless interconnection among an effectively unlimited number of parties and with no constraints on where the parties are located geographically. The simplicity of the core combined with worldwide access means not only that there is essentially no limit on what a host can use the network to do, but also that there is no trace—after the event—of what a host may have done. Currently, there is virtually no effective tool available to provide either problem diagnosis or packet traceback. This makes tracking DDoS attacks and other generators of unwanted traffic launched from multiple compromised hosts labor-intensive, requiring sophisticated skills. Even if the compromised hosts and the controller of the botnet can be located, it is likely that more than one organisation has responsibility for the machines and networks involved, which makes investigation difficult. Compounding the problems associated with the high cost and the lack of incentive to report security attacks (see below) is the fact that attacks are rarely traced to their real roots.

## The On-Ramp

The Internet is designed to be both friendly and flexible so that it does not constrain new applications that could be developed for and deployed in end systems. Such a design is, of course, a double-edged sword: capabilities that make it easy to develop useful new applications can be just as easily misused to create unwanted traffic. The aspects of Internet architecture that can be exploited to insinuate unwanted traffic onto the Internet are quite complex. Trying to ensure that the Internet remains open to innovation while denying access to unwanted traffic requires a deep understanding of the ways the Internet is intended to work and of the complex value judgments that need to be applied in order to balance the ease of use with the danger of misuse.

### Known Vulnerabilities

According to a survey conducted by Arbor Networks, the first two vulnerabilities discussed here are currently believed to be the most critical for the Internet. Other possibilities certainly exist, and the ones that are most commonly exploited shift in the continuing tussle between miscreants and security experts.

***Lying about Traffic Source Addresses.*** In the past, many attacks on networks using unwanted traffic relied on injecting packets with a forged IP source address. Receivers might then be deceived about the source of questionable packets and might therefore accept packets they would not have accepted if the packets' true source were known, or they may direct return traffic to the forged source address, making them part of a DDoS attack (reflection attack). This process is called *address spoofing*. The prevalence of botnets that can launch various attacks using the real address of the bot means that address spoofing is no longer as important a technique as it used to be, but many attacks—especially reflection attacks—still use spoofed addresses.

***Hijacking Inter-Domain Routing.*** Attacks can be launched on the Border Gateway Protocol (BGP), which routes Internet traffic between administrative domains. Various attacks can lead to traffic that gets misrouted, but a particularly insidious attack injects routes for IP addresses that are not in genuine use. Because the existence of these routes provides a measure of acceptability for packets sourced from the bogus IP addresses, attackers can use these addresses to source spam messages. Since the additional routes do not affect normal packet delivery and since careful selection of the address prefix used can hide the bogus route among genuine ones, the bogus routes often have little chance of being noticed.

***Misusing Web Protocols.*** The HTTP (Hypertext Transfer Protocol) used for accessing Web servers is now frequently used as a general-purpose transport protocol for applications that have little or nothing to do with the World Wide Web. The reason is that one of the ways attackers identify vulnerable systems is to perform a port scan. The standard transport protocols—UDP and TCP—used in the Internet identify communication end points on a host with a 16-bit port number. Targeted systems are challenged by trying to start a communication using every possible UDP and TCP port number in turn. If the communication can be started, it may give the attacker a wedge with which to pry open the security on the system. The system management reacts by closing down all unused ports to incoming communications, especially at firewalls. This has, in turn, led to difficulties for new applications that use previously unused ports and that need to have packets traverse firewalls. Applications designers have responded by reusing the HTTP communication channel, which can be pretty much relied on to be open in any firewall. However, transporting everything over HTTP does not block attacks; it simply moves the vulnerability from one place to another, and the miscreants are following.

***Everyone Comes from Everywhere.*** On the Internet it used to be possible to get some indication of the authenticity of traffic coming from a specific sender based, for example, on the number of

hops between routers that had been traversed. Each arriving packet contains a Time to Live (TTL) count, and packets that have followed the same route from a static source would have the same original TTL value decremented by the same amount, resulting in an almost constant value of TTL on arrival. A change in the TTL value for a source without a corresponding change in routing could be interpreted as meaning that the traffic with a different TTL was potentially bogus. More recently, hosts have become mobile, and a change in TTL value may simply indicate that the host has moved, with the roaming putting more or fewer hops between the source and the destination. Similarly, multihoming of a network can mean that two or more different values for the TTL are equally valid. Thus, changes in TTL value can no longer be seen as indications that traffic has been subverted, even if the underlying routing is unchanged.

***Difficulties Authenticating Identities.*** Authentication of users and machines attaching to networks as it is used today is far too complex to be feasible for users to use effectively. Consider a scenario in which a customer's handset is initially on a corporate wireless network. If that customer steps out of the corporate building, the handset may get connected to the corporate network through a GPRS cellular telephone network. The handset may then roam to a wireless LAN when the user enters a public area with a hotspot. The authentication mechanisms are usually tied to the type of data link layer used; the mechanisms use different credentials for each type, with different semantics; and there is little or no linkage between the authentication databases used with the different technologies or with policy databases that control what a user may do when attached to a network. Consequently, we need authentication tools for unifying and simplifying this authentication infrastructure and that can cope with cases when the underlying data link layer technology changes quickly—possibly during a single application session—to ensure that users and applications will not be surprised when operations that are allowed at one moment fail a little later, once the attachment point has changed.

Attackers are interested in finding targets that offer maximal returns with minimal efforts. Regions with lots of high-speed, high-bandwidth user connections but poorly managed end hosts are ideal targets for originating DDoS traffic.

## The Scale of the Problem

Unwanted traffic is a major problem for network owners and operators today both because of the volume and because of the ubiquitous adverse impact of the traffic on normal operations. The workshop did not look in any detail at the actual volumes of traffic: a look at almost any e-mail in-box is evidence enough that the volumes of spam alone are very large. This section looks briefly at how specific types of network are affected.

### *Everywhere Is Affected*

There are a variety of types of unwanted traffic on the Internet today. The IAB workshop concentrated on DDoS and spam. The impact of unwanted traffic depends on the nature of the network domain through which it is flowing, but it affects almost every part of the network adversely.

The global nature of the Internet and the ease of ubiquitous connectivity allow miscreants to originate unwanted traffic from almost anywhere in the network and to target victims who are equally widely distributed. Attackers are interested in finding targets that offer maximal returns with minimal efforts. Regions with lots of high-speed, high-bandwidth user connections but poorly managed end hosts are ideal targets for originating DDoS traffic.

### *Effects on Specific Domains*

***Backbone Providers.*** Backbone providers are primarily in the commodity business of packet forwarding. Since they do not support end users directly, spam and malware are not major concerns. Sometimes backbone routers become compromised, but this is not currently a major problem. Thus the impact of unwanted traffic is measured chiefly through the effect of DDoS traffic on network availability.

Backbone networks are generally well provisioned with high-capacity links and are therefore not normally affected by DDoS attacks. A 5 Gbps attack that would challenge most access networks can usually be absorbed without noticeable impact. On the other hand, the fact that the backbone can handle this traffic amplifies the effect on the backbone's access customers. A multihomed customer is highly likely to suffer from aggregated DDoS traffic arriving from all directions through its multiplicity of connections.

***Access Providers.*** From the access providers' viewpoint, the most severe impact of unwanted traffic is on their customer support load. Access providers have to deal directly with end users. Residential customers in particular see the access provider as their IT help desk, and the competitive nature of the business means that a single call can possibly wipe out any profits the provider might have made from the customer.

***Enterprise Networks.*** Enterprises perceive many different categories of unwanted traffic. Apart from accidentally

Continued on next page

*Unwanted Traffic, continued from page 13*

created traffic resulting from misconfiguration, a large part of the deliberately created unwanted traffic is usually just a background nuisance for enterprises because such traffic absorbs bandwidth, computing, and storage resources. Spam and peer-to-peer traffic that is not related to company business are good examples. Some of the remaining unwanted traffic may have an unknown purpose, but the big problems are caused by what is often a small volume of malicious traffic, such as traffic that spreads malware. The damage that results from undetected malicious traffic can be very costly and can take a lot of highly skilled effort to remedy.

Today, malicious traffic is often stealthy and can be obscured by encryption or can masquerade as legitimate traffic. Existing detection tools may be ineffective against this kind of traffic, and as with bots, stealth worms may open backdoors on hosts but remain dormant for long periods without causing any noticeable detrimental effects. This kind of traffic has the potential to be the gravest threat to an enterprise.

On the other hand, an enterprise may become the target of a DDoS attack, often focusing on its customer-facing Web servers. Such an attack can transform unwanted traffic from a background nuisance to a critical constraint on the enterprise's ability to do business for the duration of the attack. For civilian businesses, this risks loss of customer confidence and in addition, has longer-term implications for the business, but for infrastructure and government services there can be political or terrorist motivations with the intention to affect the stability of the state. Detecting such an attack and dealing with it as soon as possible can be vital to the survival of the enterprise: advance planning is key to managing a DDoS attack because there is little time to react once an attack starts, and the traffic has to be suppressed before it concentrates on the target resources, which may mean having tools installed by the service providers feeding the enterprise.

Network reputation is key to gaining new customers, and so, minimising the amount of publicity given to security incidents is important to service providers' survival.

### *Unwanted Traffic and Internet Infrastructure Services*

The Internet needs certain infrastructure services—such as provision of the Domain Name System (DNS)—that are potentially vulnerable to DDoS attacks. Participants at the workshop heard reports of increasingly significant DDoS attacks on the servers that handle the root of the domain hierarchy as well as the .com and .net top-level domains.

Those attacks lead to disruption of critical services, and the situation is likely to get worse because the daily peaks of DNS usage have been growing at a much faster rate than the number of Internet users. This trend is expected to continue. The increasing load on the DNS infrastructure has led to an increase in complexity that potentially makes greater targets for attacks.

## Defenses: Available but Relatively Ineffective

The Internet is not totally defenseless against the attacks from the underground economy. It is unfortunate that for a variety of reasons, many of the defenses are not as effective as they might be. Many of the reasons are economic and political rather than technical, including lack of resources, a perception that the benefits of deployment are felt by organisations other than those that have to bear the costs, and the need for coordination between competing organisations to achieve best results.

Analysis of the reasons for the ineffectiveness of the Internet's defenses is critical to the design of future effective approaches to the unwanted-traffic problem.

### *Problems for Today's Defenses*

Although there are some techniques available to protect against the known vulnerabilities, a number of inadequacies exist in the tools themselves; more critically, a number of the tools that vendors and standards organisations have produced do not get used, and the scale of deployment of the tools of the remainder is inadequate, as is education of users and operators in the secure usage and operation of the Internet.

Generally, operators do not have adequate tools for diagnosing network problems. Current approaches rely primarily on the skills and experience of operators that use time-consuming manual operations. Better and automated tools would help; the same is true for tools that help by mitigating attacks.

### *Lack of Incentives for Countering Unwanted Traffic*

A common theme that runs through the analysis of how unwanted traffic affects networks outside the enterprise is the lack of incentives for network operators to deploy security measures. That lack is due mainly to the low return on investment from what are essentially preventive measures.

Expressed in the workshop discussion of the underground economy was an unwillingness to report fraud due to commercial sensitivity. That sensitivity also applies to the reporting of security incidents by network operators who fear that their reputations—or the reputations of their customers—would

be damaged. Network reputation is key to gaining new customers, and so, minimising the amount of publicity given to security incidents is important to service providers' survival. As a result, investment in prevention is minimal, and mitigation work tends to be local so as to avoid releasing commercially sensitive information, thereby hamstringing efforts to coordinate responses to attacks or to track malicious activity.

Notwithstanding the inadequacies of the available techniques, the view of the IAB workshop was that a significant reduction of unwanted traffic could be achieved with the limited tools available if those tools were deployed extensively and were operated correctly. Educating users to be more demanding and to lobby for judicious application of government regulation may assist in the incentivisation of providers to deploy the tools.

### Available Defensive Techniques

***Countering DDoS in the Backbone.*** At the time of the workshop there was no effective diagnosis and there was only a limited supply of mitigation tools that could help backbone providers fight DDoS attacks. That situation has changed over the past two years, and many providers are offering managed DDoS security services that deliver cleaned traffic to attached customer or lower-level provider sites based on traffic pattern learning, which allows recognition and filtering of abnormal patterns that signal a DDoS attack before they concentrate on the target. On the other hand, these solutions are designed to aid particular customers who are willing to pay for the extra service, and because of the perceived low return on investment, there is still little incentive for the backbone provider to deploy these solutions for every connection.

***Know Your Sources.*** The IETF documented current best practices for filtering out incoming traffic with spoofed-source addresses in BCP 38 (RFC 2827), "Network Ingress Filtering: Defeating Denial of Service Attacks Which Employ IP Source Address Spoofing." Many routers support this type of filtering as well as the updated version for multihomed networks in BCP 84 (RFC 3704).

Network operators have not deployed these techniques universally—at least partially because of the lack of incentive resulting from the heavy management costs of maintaining the filtering and because of the need to ensure that legitimate traffic is not accidentally filtered out. Although source spoofing is no longer the indispensable tool of the underground economy that it once was, more widespread use of BCP 38 and 84 filtering can still make attacks using spoofed addresses unprofitable and facilitate traceback of attacks.

***Managing Access: Customer Behaviour.*** Access providers routinely offer free security software to customers in the hope of avoiding future help calls after a security break-in. Unfortunately, customers are often not educated about the need to install security software, and even when they are, they may lack the skills to correctly configure a complex system.

Customer behaviour in the face of security breaches is depressing:

- All customers behave in essentially the same way.
- Notifying customers that they have a problem has little effect on whether they take action to repair the breach.
- Patching of breaches works in the same way as radioactive decay. A fixed proportion (about 40 percent) of remaining vulnerable systems are patched every month after the patch becomes available. In the large population of Internet hosts, this leaves a significant number that will be vulnerable for the rest of their working lives.
- Lack of understanding often leads to compromised systems' being replaced rather than being repaired, but this ignorance often leads to the occurrence of infections during installation of the replacement.

***Maintaining Profitability in Enterprises.*** Enterprises, particularly large ones, are more willing to investigate security breaches than backbone or access providers are, because they can directly impact the enterprise's operations and profitability. However, enterprise network operators are very wary of security solutions that generate false-positive alerts, because such alerts can be very costly to the enterprise if parts of the network have to be shut down unnecessarily. Most prefer prevention solutions to detection solutions because of this and are often willing to accept some missed alerts rather than significant false positives.

Enterprises are motivated by potential losses to spend money on security tools. Consequently, a thriving market has emerged to meet the demand. Unfortunately, the tools offered provide mostly reactive solutions, such as regularly updated virus scanner databases for countering newly emerging vulnerability exploits, which leads to an ongoing arms race between security exploits and patching solutions. Workshop participants expressed concerns that this was not a sustainable situation because it does not enable us to get ahead of the attackers.

***Over-engineering the Infrastructure.*** At present, the only effective mitigation strategy for DDoS attacks on critical infrastructure services is over-engineering. There is some concern that the runaway growth of demand especially for DNS services is eroding the safety margins. The expected widespread deployment of IPv6 and deployment of the new DNS security extensions (DNSSEC) in the near future will bring new and potentially flawed software into widespread use that could be abused to generate new DDoS attacks.

Continued on next page

*Unwanted Traffic, continued from page 15*

### *Law and Regulation Playing Catch-up*

In human society, legal systems provide protection from and deterrence for criminals. Laws and regulations aim to penalise criminal conduct after the fact, but if the likelihood of detection is low, the deterrent effect is also minimal. At present, the development of legal systems aimed at cyberspace crime is lagging behind the development of the crime that the legal systems are intended to deter, and the likelihood of detection of the real criminals is low.

Some of the reasons for the ineffectiveness and slow development of the law of cyberspace include:

- **The international scope of the problem.** The Internet spans the globe, and crimes masterminded in one national jurisdiction may be executed by machines in one or more other countries, with victims in yet other jurisdictions. While some countries, particularly in the developed world, criminalise computer fraud and abuse, regulate unauthorised use of government and other critical infrastructure, and prohibit access to confidential information on protected computers, the laws are not uniform, which makes it difficult to prosecute criminals for offences carried out from other jurisdictions. There is also little political incentive to pursue criminals when the victims are not in the same national jurisdiction. Although there is a coalition between countries on collecting evidence of cybercrime worldwide, there is no rigorous way to trace unwanted traffic or to measure the consequences of cybercrime across national borders.
- **Pinning down the responsible organisation.** A single episode of unwanted traffic and the botnets that are responsible for much of the traffic can involve many different organisations, such as owners of hosts, enterprise networks, and service providers of various kinds. Many of these organisations would see themselves as innocent parties, and others, such as the owners of compromised hosts, see no incentive to take action. This makes it extremely difficult to either regulate effectively in advance to make life difficult for the criminals or to make any organisation responsible for cleaning up after an attack has been detected.
- **Getting the legal definitions right.** Lawmakers are generally unfamiliar with the new world of cyberspace, and therefore they often lack the technical understanding necessary to specify laws precisely and in such a way that they will actually target undesirable acts without limiting legitimate use of the network. As in many areas where there are active innovation and financial incentive, the underground economy will always be seeking to push the limits by using techniques that are borderline legal and conceal evidence through complexity. The lawmakers are inevitably playing catch-up in cyberspace.
- **Quantifying the damage.** Investigative authorities are already stretched, and so, active legal action tends to be restricted to cases where the harm caused exceeds a fairly high threshold. In the case of unwanted traffic, this generally means either significant damage to national infrastructure or a large, quantifiable monetary loss. Unfortunately, either (1) it is often difficult to quantify the loss, or, when financial institutions are involved, (2) there is a reluctance to admit the scale of the losses for fear of ongoing commercial damage. Consequently, much cybercrime is either not reported to the authorities or not investigated.
- **Defining unwanted traffic.** Creating capabilities to limit unwanted traffic can have unwanted side effects. It needs only a shift in the definition of *unwanted* to move from constraining the underground economy to facilitating censorship and limiting open access. Countries already differ over what is defined as unwanted traffic; and traffic that would be seen as wholly legitimate in many countries may result in criminal prosecutions elsewhere. There is a trade-off between having audit trails to facilitate forensic analysis and providing the means to enforce censorship. Building monitoring capabilities into the network will surely result in stronger pressure from legislators, requiring that operators actually carry out monitoring.

The workshop also emphasised that, while an effective legal system is necessary to create effective deterrence for and sanctions against the parasites, it is by no means sufficient on its own. It can work only in conjunction with effective user education as well as technical solutions to unwanted traffic prevention and detection. Only a well-informed and motivated user community can collectively establish a defense against unwanted traffic in cyberspace.

## Consequences

The consequences of the large volumes of unwanted traffic on the Internet to-

**Recent IESG Document and Protocol Actions**

A full list of recent IESG Document and Protocol Actions can be found at http://ietfjournal.isoc.org/DocProtoActions0303.shtml.

day are highly detrimental. The health of the network presents a picture that is far from rosy.

- There are big economic incentives and a rich environment to exploit.
- There is no specific party to carry responsibility.
- There are problems of underdeployment of the limited defensive tools that are available.
- There are no auditing systems to trace back to the sources of attacks.
- There are no well-established legal regulations to punish offenders.

The combination of these factors inevitably leads to ever-increasing types and volumes of unwanted traffic. However, the real threats are not the bots or DDoS attacks but the parasitic criminals behind them. Unwanted traffic is no longer aiming only for maximal disruption; in many cases, it is now a means to illicit ends, and its specific purpose is to generate financial gains for the miscreants. Their crimes cause huge economic losses, counted in multiple billions of dollars and growing.

The Internet community needs to increase its awareness of the problem of unwanted traffic and take action to make the network less friendly to this type of traffic. And it needs to do so without significantly reducing the flexibility of the network that has been the key factor in the economic success of the Internet.

All Internet stakeholders can potentially contribute to the reduction of unwanted traffic. At a high level, actions should include the following.

- Research into specific problems resulting from unwanted traffic, involving:
  - Sponsoring and funding agencies that prioritise this kind of research
  - Network operators, equipment vendors, and users who can identify the most important problems that require research effort and who can make sure that researchers are aware of them
  - Standards organisations, which should help coordinate communication between researchers and the rest of the community to identify the fundamental problems and standardise any solutions that may be found.
  - Development of a uniform global legal framework that will facilitate successful legal pursuit of the miscreants in the underground network economy across national borders. This work needs to be informed by the best possible technical expertise to ensure that it leaves Internet flexibility intact so far as is possible.
- Appropriate regulation to require that network operators take action to minimise the effects of unwanted traffic and that they share information that will lead to mitigation of attacks and will drive miscreants out of business
- Increased deployment of available tools, possibly aided by incentivisation through regulation or customer demand
- Vendors that provide more-appropriate default security settings in equipment so that end hosts are less vulnerable to subversion from the moment they are deployed and without the need for sophisticated configuration by users
- Vitally, improved education of users to make them more aware of the risks to their systems, to make them aware of the ways those risks can be mitigated, and to mobilise them so they'll demand action from network operators when action is needed to support network security in both enterprises and homes

> At present, the development of legal systems aimed at cyberspace crime is lagging behind the development of the crime that the legal systems are intended to deter, and the likelihood of detection of the real criminals is low.

Above all, the Internet community needs to get ahead of the miscreants. At present, almost all activity for countering unwanted traffic is reactive, by ex post facto identification of malware and retroactive patching of security holes. Recently, there have been improvements in the use of traffic pattern analysis to identify attacks as they happen, but future work needs to be intelligence led, and it must concentrate on eliminating opportunities for miscreants before such opportunities are deployed.

*Many thanks to Lixia Zhang, Loa Andersson, and Danny McPherson for their feedback and review.*

Photo by Peter Lothberg

*IETF Journal* editor Mirjam Kühne and friends attend the IETF 70 plenary session.

# The Unique Political Soul of the IETF

***By Tomas Carlsson***

Known to the world are two different political systems. Known to the IETF community is a third system. Whether we call it *IETF democracy* or *Majhum* (majority by humming), IETF meeting attendees will know what I mean. Everyone else will have to fight through several levels of abstractions to get a sense of it.

Perhaps the IETF community hasn't thought of its procedures and processes in terms of a political system. Regardless, I will explain why it is the third form that is relevant. Welcome to the world of mash-up politics.

First, we have democracy. Whether you have it or you don't have it, you usually want it. In a democracy, decisions are made by representatives who are elected by the voting population. Sometimes the system is referred to as *parliamentarism*. Whatever you call it, in a democracy, different opinions on a subject are allowed to be expressed prior to a decision's being made.

Then we have dictatorship. Within this category we also place Muammar al-Gadhafi of Libya and his so-called Third International Theory. In a dictatorship, one person or one group makes decisions without requiring a registered mandate from those affected by the decisions. In most cases, that person or group maintains power with the help of armed forces. Some monarchies operate as dictatorships. In those cases, power is inherited.

In the past, I have studied and written about the procedures of the IETF, but I couldn't—even in my imagination—have believed it to be so close to direct democracy until I experienced it myself during the IETF meeting in Vancouver.

Politically speaking, the IETF has no equivalent. Its power is intricately tied to the expectation that participants understand the IETF's rather complex and abstract culture—a culture that, above all, demands that interaction among its participants be handled properly. When accustomed to working among directors, executives, and board members—a culture that embraces a clear organisational hierarchy—one can at first feel annoyed at having to listen to the diverse opinions that are allowed to float up as part of the IETF's democratic system. I found this to be especially true when the opinions were expressed in the later stages of a process to settle a developed proposal for an Internet standard. It is said the IETF has no members and no voting, but in my opinion there are both members and voting—in the same way that an ant is a member of a colony and voting is a means for determining which way the crowd will head next. The IETF may not be a legal entity, but it offers power to the masses and confidence that the rules of interaction will result in the right decisions.

In both ant colonies and the IETF, decisions are made every moment. Small decisions become bigger decisions. They say that practicing democracy is time-consuming. I'd say with regard to the IETF that that is an understatement. If I can identify one single factor that affects the time it takes for a group to move forward in the process, it would be the rigor with which working group chairs demand that participants follow the correct, stated procedures when giving their say.

The word *correct* is key within the IETF. This is not to say that the goal is for everything that is said to be correct; the goal is to achieve the best technical standards, and within the IETF, doing so means adhering to the rules of engagement and accepted procedures that working group participants have long followed. "This is the correct way of doing it" is an oft-repeated mantra. Correct standards wouldn't be engineered without the help of an overarching authority. Such authorities

Photo by Peter Lothberg

IETF 70 attendees take a break to chat and catch up on e-mail.

used to be Internet veterans with beards that have growth rates that are proportional to the cumulative list of assigned IP numbers. But I have noticed that the fashion has changed lately to one of more bald chins and cheeks. In the same way, over time, the focus changed in the different organisational elements. This is a healthy sign and one that demonstrates a changing reality.

One could ask why it is necessary for the IETF to meet in person three times a year. Isn't all the hard work that is done in each working group's mailing list enough? The answer is no. I have seen research on the effectiveness of distance education. The result is that face-to-face interaction is necessary

to keep the motivation, the passion, and the understanding among people strong. Face-to-face meetings are also the places people discover the extent to which chemistry translates into well-functioning groups. As ad hoc as they may seem, the personal connection that is found at IETF meetings makes it more likely that these groups will survive.

At first, I found the IETF's insistence on consensus and the humming as a method to determine rough consensus a bit silly. Eventually, though, the psychological effect grew on me. One can feel the strong hum of a majority in the chest, and no matter how logical your objections, that feeling cannot be erased. It will hold back every not-very-well-grounded opinion. It may not prevent situations where participants are objecting for the sake of objecting, but a good working group chair will in that case make sure the meeting proceeds.

Within the IETF's system, if I crave the cult status of having initiated, written, and published an IETF standard in the form of a finished RFC, I first have to convince an area director that we need to have a meeting—known as a birds-of-a-feather meeting—to discuss it. Even if I think it is a splendid idea, there will be no working group, no draft, and no nothing if I can't come up with enough support to keep it going. The best way to get support for your ideas is to first gain respect for your knowledge. You will probably not get that respect in the short time you have at the microphone at the meetings. You earn it in the corridors, or at the late-night get-togethers in the lobby, or in the bar, or on the mailing lists. On the mailing lists in particular, concrete and clever comments and contributions will result in people fighting to hear your opinion.

> The IETF's power is intricately tied to the expectation that participants understand its rather complex and abstract culture—a culture that, above all, demands that interaction among its participants be handled properly.

This is the essence of the third political system: Anyone—no matter their social or cultural background—can take a leadership position within and make a contribution to the IETF system. If you earn respect, if you demonstrate that you are knowledgable, then you will be heard. But it takes time, commitment, and a willingness to participate in a direct democratic system. The entire IETF standards-building process is based on individual contributions that ultimately lead to teamwork. In other words, if you demonstrate wisdom, others will team up around your idea.

## IETF Loses Participant and Former IAB Member

The Internet Society and the IETF community mourn the loss of Jun-ichiro "Itojun" Hagino, who died on 29 October 2007. He was 37 years old.

Photo by Diane Bruce

Jun-ichiro "Itojun" Hagino, pictured at BSDCan in May 2004

Itojun was an active participant in the IETF and a member of the Internet Architecture Board (IAB) from 2003 to 2005. He was a senior researcher at the Internet Initiative Japan (IIJ) and a member of the board of the Widely Integrated Distributed Environment (WIDE) project. Itojun was a strong supporter of open standards development and open software, working as a core researcher from 1998 to 2006 at the KAME project, a joint effort of six companies in Japan to provide a free stack of IPv6, IPsec, and Mobile IPv6 for BSD variants.

Itojun will also be remembered as a kind and mild-mannered friend to many, a very helpful cross-cultural bridge, and a knowledgeable international foodie. "He seemed to be at his happiest when programming and when sharing a good meal," said Randy Bush, a friend and colleague.

In a brief, joint statement, IETF chair Russ Housley and IAB chair Olaf Kolkman recognised the valuable contributions Itojun had made to the IETF—particularly through his work in IPv6-related working groups. "He inspired many and will be missed," the statement read.

# What Makes For a Successful Protocol?

*By Dave Thaler and Bernard Aboba*

HTTP/HTML versus Gopher. IPv4 versus IPX. Interdomain IP Multicast versus application-layer overlays. As we learned from the more mainstream VHS-versus-Betamax-format war, the reasons that one technology or protocol takes off while another one crashes and burns are obvious only in retrospect. Success may not be easy to predict, but it's rarely if ever an accident or simply a matter of luck or timing (though timing can be a critical ingredient in achieving success). More often than not, success happens when a problem gets solved or a need gets addressed in a manner that is cost-efficient, easy to deploy, and useful for more than a minute and a half.

Sound simple? It is, as long as it's understood that simple and easy are not the same things. Even if a formula existed for designing the perfect protocol, the Internet—together with all that is layered on top of it—is too vast, too changeable, and too complex to make any proposed solution or fix a sure thing. Fortunately, though, the sheer number of Internet protocols developed, published, and deployed in the past few decades offers valuable opportunities for determining the factors that could stack the deck in favor of success.

## Defining success

What does it mean for a protocol to be successful? Is a protocol successful if it has met its original goals but is not widely deployed? Perhaps, but for purposes of this article, we define a successful protocol as one that both meets its original goals and is widely deployed. Perhaps the best examples of successful protocols are IPv4 (RFC 791), TCP (RFC 793), HTTP (RFC 2616), and DNS (RFC 1035).

Success, however, is multidimensional. When designed, a protocol is not intended only for some range of purposes; it is also designed to be used on a particular scale. Therefore, the two most important measurements by which a protocol can be evaluated, as shown in Figure 1, are purpose and scale.

According to those metrics, a *successful* protocol is one that is used for its original purpose and at its originally intended scale. A *wildly successful* protocol is one that exceeds its original goals either in terms of purpose (it is used in scenarios that extend beyond the initial design) or in terms of scale (it is deployed on a scale much greater than originally envisaged) or in terms of both; that is, the protocol has overgrown its bounds and has ventured out into the wild.

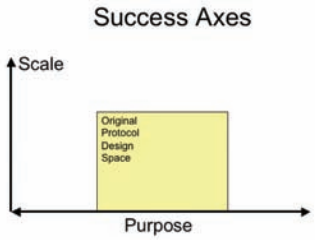


Figure 1

If we apply those definitions, then a protocol such as HTTP is defined as wildly successful because it exceeded its design in both purpose and scale. Another example of a wildly successful protocol is IPv4. Although it was designed for all purposes ("Everything over IP and IP over Everything"), it has been deployed on a far greater scale than it was originally designed to meet. Still another example is ARP (Address Resolution Protocol). Originally designed for a more general purpose (namely, resolving network-layer addresses to link layer addresses regardless of media type or network-layer protocol), ARP was widely deployed for a narrower scope of uses (resolution of IPv4 addresses to Ethernet MAC addresses). More recently, it has been adopted for other uses, such as detecting network attachment (DNAv4 [RFC 4436]). Like IPv4, ARP is deployed on a much greater scale (in terms of number of machines but not in terms of numbers on the same subnet) than originally expected.

As with most success stories, to be wildly successful can be both good and bad. A wildly successful protocol is one that solves more problems or that addresses more scenarios or devices than originally intended or envisioned. When this happens, it may mean it's time to revise the protocol to better accommodate the new space. However, if a protocol is used for purposes other than the one for which it was designed, there can be undesirable side effects—such as performance problems. The design decisions that are appropriate for the intended purpose may be inappropriate for another purpose. Worse, wildly successful protocols tend to become popular, which means they can be attractive targets for attackers.

## When failure becomes an option

Unlike a major motion picture, which can be dubbed a failure at the box office within a week or two of theatrical release, the failure of a protocol can be determined only after a sufficient amount of time has passed—generally 5 to 10 years for an average protocol. To be considered a failure, a protocol must be lacking in three key areas: (1) mainstream implementation (little or no support in hosts, routers, or other classes of relevant devices), (2) deployment (devices that support the protocol are not deployed, or, if they are, the protocol is not enabled), and (3) use (the protocol may be deployed but there are no applications or scenarios that actually use

the protocol). It's important to note that at the time a protocol is first designed, there is of course no implementation, deployment, or use, which is why it's important to allow sufficient time to pass before evaluating the success or failure of a protocol.

## Identifying success factors

A series of case studies examined by the authors laid the groundwork for determining the key factors that contribute to a successful or a wildly successful protocol as well as the relative importance of those factors. Note that just as a successful protocol may not necessarily include all of the factors, a failed protocol could very well include some of the factors that determine success.

*Positive net value (meeting a real need).* The success of a protocol depends largely on the notion that the benefits of deploying the protocol (monetary or otherwise) outweigh the costs—such as the costs of hardware, operations, configuration, and management—as well as costs associated with any changes to the business model that might be required. A few key benefits might include pain relief (lower cost than before), opportunities to enable new scenarios (though this type has a higher risk of failure than the other types), and incremental improvements (for example, better video quality).

Success seems more likely when the natural incentive structure is aligned with the deployment requirements—that is, when those who are required to deploy, manage, or configure a protocol are the same as those who gain the most benefit. In other words, it's best if there is significant positive net value at each organisation where a change is required.

*Incremental deployability.* A protocol is incrementally deployable if early adopters gain some benefit even if the rest of the Internet does not support the protocol. It also appears that protocols that can be deployed by a single group or team have a greater chance of success than do those that require cooperation across organisations (or, worse, those that require a flag day where everyone has to change simultaneously).

*Open code availability.* Perhaps the next most important technical consideration is that a protocol have freely available implementation code. This may have been the case when deciding between IPv4 and IPX, the latter of which at the time was, in many ways, the technically superior of the two.

*Freedom from usage restrictions.* A protocol is far more likely to succeed if anyone who wishes to implement or deploy it can do so without legal or financial hindrance. Within the IETF, this point often comes up when choosing among competing technologies; for example, the one with no known intellectual property restrictions is the one most likely to be chosen even if it's technically inferior.

*Open specification availability.* What remains true for all RFCs—and has contributed to the success of protocol specifications both within and outside the IETF—are protocol specifications that are made available to anyone who wishes to use them. This might include worldwide distribution (accessible from anywhere in the world), unrestricted distribution (no legal restrictions on getting the specification), permanence (remains even after the creator is gone), and stability (doesn't change).

*Open maintenance processes.* The protocol is maintained by open processes; mechanisms exist for public comment; and participation by all constituencies affected by the protocol is possible.

*Good technical design.* The protocol follows good design principles that lead to ease of implementation and interoperability.

## What makes a protocol wildly successful?

The following factors do not seem to significantly affect initial success, but they can affect whether a protocol is wildly successful.

***Extensible.*** An extensible protocol is one that carries general-purpose payloads and options or easily accommodates the addition of new payload and option types. Such extensibility is desirable for protocols that are intended for application to all purposes, such as IP. However, for protocols designed for a specialised purpose, extensibility should be considered carefully.

***No hard scalability bound.*** Protocols that have no inherent limit near the edge of the originally envisioned scale are more likely to be wildly successful in terms of scale.

***Threats sufficiently mitigated.*** Protocols with security flaws may still become wildly successful provided they are extensible enough to allow the flaws to be addressed in subsequent revisions. However, the combination of security flaws and limited extensibility tends to be deadly.

## Conclusion

It can't be emphasised enough that the most important factor that contributes to the success of a protocol is that the protocol fill a real need. It also helps if the protocol can be deployed incrementally. When there are competing proposals of comparable benefit and deployability, open specifications and code become increasingly significant success factors. Open source availability is initially more important to success than is open specification maintenance.

In most cases, technical quality was not a primary factor with regard to initial success. The initial design of many protocols that have become successful would not pass IESG review today. Technically inferior proposals can win if they are openly available. Factors that do not seem to be significant in determin-

Continued on next page

*What Makes For a Successful Protocol? continued from page 21*

ing initial success (but that may affect wild success) include good design, security, and having an open-specification-maintenance process.

Many of the case studies we evaluated concern protocols originally developed outside the IETF but that the IETF played a role in in improving after initial success was certain. While the IETF focuses on design quality, which is not a significant factor in determining initial protocol success, once a protocol succeeds, a good technical design may be key to its continuing success. Allowing extensibility in an initial design enables initial shortcomings to be addressed.

Security vulnerabilities do not seem to limit initial success, most likely because vulnerabilities often attract attackers only after the protocol becomes deployed widely enough to become a useful target.

Finally, open specification maintenance is not very important to initial success, because many successful protocols were initially developed outside the IETF or other standards bodies; they were, in fact, standardised later.

In light of our conclusions, we recommend that the following questions be asked during the evaluation of a protocol design:

- Does the protocol exhibit the critical initial success factors?
- Are there customers (especially high-profile customers) that are ready to deploy the technology?
- Are there potential niches where the technology is compelling? If so, can complexity be removed to reduce cost?
- Is there a potential killer application? Or can the technology work underneath existing, unmodified applications?
- Is the protocol sufficiently extensible to allow potential deficiencies to be addressed in the future?

If it is not known whether the protocol will be successful, should the market decide first? Or should the IETF work on multiple alternatives and let the market decide among them?

Are there success factors that may predict which among multiple alternatives is most likely to succeed?

In the early stages of protocol design, evaluating the factors that may influence initial success is important in facilitating success. Similarly, efforts to revise or revive unsuccessful protocols should include an evaluation of whether the initial success factors (or enough of them) were present rather than focusing on wild success, which is not yet a problem. For a revision of a successful protocol, on the other hand, focusing on the wild-success factors is more appropriate.

# Security Protocol Failures

***By Phillip Hallam-Baker***

*This article is a condensed version of the argument made in* The dotCrime Manifesto: How to Stop Internet Crime, *in which the question of how to fix these problems is considered.*

The Internet is insecure, so what went wrong? Contrary to widely held belief, the reasons for Internet security protocol failure are not primarily technical. Failure to understand the risk model and to meet the actual user requirements are much more significant causes of security failure. The economics of security protocol deployment and security usability engineering are also key: a protocol might as well not exist if it is not used.

## Is It Safe?

Is the Internet safe? To paraphrase Douglas Adams, yes, the Internet is perfectly safe: it's the rest of us who have to worry.

The Internet was built to meet a specific set of needs and be adaptable beyond those needs. Contrary to common assertion, security was a consideration early in the design of Internet architecture and protocols. Saltzer et al. addressed security at some length in their seminal end-to-end-arguments paper of 1981.[1]

There are many reasons why cryptographic security was not embedded into the Internet from the first, not the least of them the limited computing power available. But even if more powerful machines had existed, the risks did not. There were no shops or banks on the primordial Internet. Military secrets were isolated on an essentially separate network—albeit not isolated enough, as subsequent events would prove.[2]

Although the primordial Internet lacked cryptographic security, it did have a strong and effective accountability mechanism. Access to the Internet required access to one of the tiny numbers of computers connected to it. Miscreants faced a real risk of consequences; a visit to the dean's office, loss of computing privileges, and in extreme cases, expulsion.

The Internet protocols were capable of scaling to support a billion users; the accountability mechanisms were not. At the same time, the Internet became, in Willie Sutton's infamous phrase "where the money is." Consequently, the need to urgently retrofit security to the Internet

became sharply apparent—in particular, with the rise of the Web beginning in 1993.

Yet here we are, 15 years later. Internet crime is a multibillion-dollar nuisance, and cybersecurity is a campaign issue in the U.S. presidential campaign.[3]

What went wrong?

## Systems Failure

According to the traditional view, the first concern in security protocol design is to get the job done right. "Bad security is worse than no security." But while this may have been true for Mary Queen of Scots and the Rosenbergs—executed as a result of misplaced faith in a faulty cipher—it is certainly not the major cause of Internet security failures.

Mistakes matter rather less than is often supposed. The most elementary of errors—complete lack of any authentication capability—was discovered in SSL (Secure Sockets Layer) 1.0 just 10 minutes into the first public presentation on the design. That error was fixed in SSL 2.0, but this time the designers made no effort to obtain public review prior to release, and further design errors were identified. It wasn't until the design of SSL 3.0 that an experienced designer of cryptographic protocols was engaged to evaluate the design—but for only 10 days.[4]

Rather more significant than the making of the mistake itself is an architecture that allows the mistake to matter. Lampson's security reference monitor[5] does not make it less likely that the programmer will make a mistake but does reduce the number of places where a mistake is likely to matter.

## Failure Commitment

Fear of making a mistake has frequently led to security protocol design that takes far longer than it should.

Despite breaking every accepted rule of open standards design, SSL and its IETF successor TLS (Transport Layer Security) are the only Internet security protocols to have achieved ubiquitous use. Getting the protocol specification as correct as possible should certainly be the first concern of the protocol designer who wants to find future employment, but nobody is served by the designer who is perpetually unable to commit to a design that can be tried in the real world.

## Requirements Failure

The Internet is a work in progress, not an absolute truth. It was not the original purpose of the Internet to provide a communication network; it was to provide an environment for the research and development of computer networks. The World Wide Web was not imagined in 1980, nor were the security requirements for employing the Web as the ubiquitous engine of electronic commerce understood in 1995. It is only with experience of use that these requirements have become better understood.

The IETF has produced four specifications for an end-to-end e-mail security protocol: PEM, MOSS, Open-PGP, and S/MIME. None is widely used. For many years it has been asserted that the lack of use of S/MIME was due to the inadequate deployment base of capable clients—despite the fact that Outlook, Thunderbird, Notes, AOL, and express variants thereof have all supported S/MIME for almost a decade.

It is time to admit that one of the many reasons for this failure is that none of the end-to-end mail security protocols actually met users' real requirements. Ease of deployment and use were far higher in most users' list of real requirements than was the theoretical possibility of an attack by the mail server administration. And today, users who are concerned about the need for end-to-end security consider it in terms of the end-to-end life cycle of their confidentiality-sensitive documents.

## Political Failure

Another reason for the failure of end-to-end e-mail security is political failure. S/MIME has widespread deployment, but Open-PGP is still the leader in mindshare.

## Infrastructure Failure

Another commonly cited reason for the failure of end-to-end e-mail security protocols is the lack of a deployed, public key infrastructure (PKI), but this explanation may confuse cause for effect. There is a large and robust market providing PKI infrastructure for SSL—albeit a commercial, for-profit infrastructure rather than a free one.

A more convincing explanation of the failure to establish an end-user PKI infrastructure is that both Open-PGP and S/MIME resort to what can only be described as hand-waving arguments wherein the question of public key discovery is concerned. If the Open-PGP web of trust is to be taken seriously, we should expect to see a rich maintenance protocol offering features similar to those being discussed in the areas of social networking and Identity 2.0. S/MIME lays the responsibility off onto PKIX, which in turn lays it off onto an entirely underspecified Lightweight Directory Access Protocol (LDAP) instantiation.

The problem, then, is not merely the failure of the necessary PKI infrastructure to deploy, as is often claimed; we lack the necessary S/MIME infrastructure to make use of it even if it did.

## Context Failure

Many security failures result from security by analogy. If a security control is adequate in one context, then it should be adequate in another context. If a four-digit PIN is good enough for securing an automatic teller machine (ATM) transaction, then it's good enough for online banking. If sending passwords in the clear is good enough for FTP, then

*Continued on next page*

*Security Protocol Failures, continued from page 23*

it's good enough for HTTP.

The problem with security by analogy is that while it can certainly be effective in identifying possible risks (i.e., if protocol A fails due to X, then look for the potential for X in protocol B), it is rather too easy to overlook significant differences in the context in which the protocols are applied. The PIN is only one factor in a two-factor authentication scheme in ATMs. Moreover, there is a maximum daily limit on withdrawal. In an online brokerage application, the PIN is the only authentication factor, and there is no transaction limit.

The name of the Wired Equivalent Privacy (WEP) protocol used in securing IEEE 802.11b wireless Ethernet demonstrates another form of context failure. The designers of WEP assumed that the principal change in the security context of moving from a wired to a wireless LAN was the risk of disclosure. Consequently, they designed a protocol intended to provide a strong confidentiality protection wherein the authentication component consisted of a single secret key shared by every machine in the network. The practical security implications of this model included terminated employees' surfing the corporation from the parking lot, among others.

## Experience Failure

It is an old but true saying that familiarity breeds contempt. While almost everyone has an Internet security story to tell, the telephone network raises rather fewer concerns than it should. The security posture of the telephone system in virtually every industrialised country is predicated on the assumption that there is a single, monopoly operator whose employees are absolutely trustworthy.

I have a fax machine in my office because some people insist that the Internet is not secure enough. The fax is served by a VoIP connection and forwards the messages to me by e-mail.

## Usability Failure

Designing security protocols is not enough. If we wish to secure the Internet, we need people to use those protocols. Until recently, the field of security engineering usability was virtually ignored. Today it is much more widely appreciated as a security protocol that people do not use because people do not use what they cannot either use or understand.

Much has been written about the need for end-to-end security. On the Internet the ends of the communication are the user's brain and the person or organisation the user is interacting with. To provide end-to-end security, we must secure the last two feet between the user's eyeballs and the screen. Secure Internet Letterhead[6] was proposed with a view to that end: the customer recognises the bank by the bank's brand on the ATM, the bank card, the branch, and so on. We should adopt the same cues on the Internet (e.g., via RFC 3709).

Recognising the need for usability is much easier than achieving usability. The entire computing field is facing a usability crisis, and such techniques as exist tend to be designed by and for usability experts. Much work remains to be done before the techniques are part of every security engineer's tool kit.

## Accountability Failure

When the Internet crime wave first hit, a great deal of effort was put into consumer education. Such efforts frequently missed the point that Internet crime is neither the consumers' fault nor the consumers' responsibility. The design flaws are in the financial infrastructures and the Internet. Consumers were not responsible for the security design of either.

Responsibility for security must lie with the party best able to provide it. Customers put their money in a bank because they believe that the bank is better able to keep the money safe than they are. If the bank starts telling customers that safety is customers' responsibility, the bank undermines its own purpose.

We cannot hope to hold a billion Internet users accountable for their actions, but we can hold ISPs accountable for allowing SYN floods and spoofed source address packets onto their networks, just as we now hold them accountable for spam. We cannot hold application providers accountable for every last bug in their systems, but we can hold them accountable for allowing the bugs to matter, and we can hold them accountable for systems whose default behaviour is to automatically run unknown programmes from unknown sources with full user privileges.

## Deployment Failure

Perhaps the most common reason for Internet protocol failure is that the protocol is never used. Security specialists have been considering the economics of profit-driven Internet crime for some time. Recently, attention has focused on the economics of protocol deployment.[7] A study of deployment of the SSH protocol by Rosasco and Larochelle[8] concluded that the reason for the protocol's success lay not in the specific security features supported in the SSH protocol itself but in the additional, nonsecurity functionality that the SSH application made possible—in particular, the ability to tunnel X-Windows sessions using SSH.

Many opportunities for applying this codeployment strategy remain untapped. Establishing an Internet Webcam session in the presence of firewalls and NAT remains a largely futile effort. In *The dotCrime Manifesto*[9] I make the case that simplified network administration could be the killer application for Default Deny Infrastructure.

## Conclusions

Despite our best efforts in applying our core skill sets, the Internet remains

unacceptably insecure. Internet crime is a large and growing problem. Neither Internet crime nor the failure to deploy the necessary effective security protocols is an exclusively technical problem. We must therefore look beyond the narrow focus of our own expertise to other communities of experts that can help us.

The list of problems to be addressed is reassuringly large: if we had no idea what the cause of the problem might be, we would have no way to fix it. While each cause of failure is significant, all are readily fixed once the problem is identified. All we need is the will to do so.

Some have objected that these concerns are not engineering and thus lie outside the scope of the IETF. This is not my view. In Europe a person with mere domain expertise is known as a technician. Only once a candidate has demonstrated the ability to combine personal expertise with whatever other expertise is necessary (managerial, legal, commercial) to solve problems does the candidate qualify for the title *engineer.*

## References

1. Jerome H. Saltzer, David P. Reed, and David D. Clark, s.l. *End-to-End Arguments in System Design*. IEEE Computer Society, 1981, Proceedings of the 2nd International Conference on Distributed Computing Systems, Paris, 1981, pp. 509-512.

2. Clifford Stoll. *The Cuckoo's Egg: Tracking a Spy through the Maze of Computer Espionage.* s.l.: Pocket, 2000. 0743411463.

3. Rudolph Giuliani. "The Resilient Society: A Blueprint for Homeland Security." *Wall Street Journal* [online]. January 9, 2008. http://www.opinionjournal.com/federation/feature/?id=110011099.

4. Paul C. Kocher. Personal communication.

5. Butler Lampson. *Hints for Computer System Design*. 5, 1983, *ACM Operating Systems Review,* SIGOPS, vol. 17, pp. 33-48.

6. Phillip Hallam-Baker, Secure Internet Letterhead. Presented at W3C Workshop on Transparency and Usability of Web Authentication. http://www.w3.org/2005/Security/usability-ws/papers/27-phbaker-letterhead.

7. Ross Anderson and Tyler Moore. Information Security Economics—and Beyond [online]. August 21, 2007. http://www.cl.cam.ac.uk/~rja14/Papers/econ_crypto.pdf.

8. Nicholas Rosasco and David Larochelle. How and Why More Secure Technologies Succeed in Legacy Markets: Lessons from the Success of SSH [online]. 2004. http://www.cs.virginia.edu/~drl7x/sshVsTelnetWeb3.pdf.

9. Phillip M. Hallam-Baker. *The dotCrime Manifesto: How to Stop Internet Crime.* s.l. Addison-Wesley Professional, 2008.

# Doing Business Securely in an Insecure World

***By Randy Bush***

What real improvements in Internet security have we achieved? The Net certainly is not a safe place as long as there are phishing, DDoS attacks, and cross-site script HTML attacks. However, although those are serious problems, we should not ignore the ways in which we have been successful in making the Net a safer place, and we should try to take some lessons from those successes.

## The Net Is a Dangerous Place, but We Can Do Things Safely

Twenty years ago, it was considered rude to connect a UNIX machine to the Net without offering a password-free account named *guest* so that any passerby could use it. When I tell this to people who have less than 10 years of experience working with the Internet, they don't believe me.

Times have certainly changed. Today the Net is much less secure. In an environment where attacks are recurring events, where operating system and application vulnerabilities are discovered daily, and where bot-nets of 100,000 zombies attack, we still feel comfortable conducting financial and other private transactions worth billions of dollars on the Net.

How is this possible? It's possible because we have deployed tools and protocols that enable secure transactions in an insecure world. This philosophy is similar to that which explains how we were able to build a reliable network from a set of unreliable components: circuits may fail and equipment may have errors, but the packets route around these problems.

So, what are the successful protocols and tools that have made the Net a safer place? Here are some of what I consider the most significant.

## Secure Sockets Layer/Transport Layer Security (SSL/TLS)

One would not consider sending credit card or other personal information over an unencrypted link. Encrypted and authenticated browsing, using https as opposed to http, are the bases on which almost all Internet commerce is founded. Even when a transaction is not facilitated by a browser, TLS (Trans-

*Continued on next page*

*Doing Business Securely in an Insecure World, continued from page 25*

port Layer Security)—the new name for SSL (so that the IETF could make a "contribution")—is used underneath most client-server exchanges.

### SSH (Secure Shell)

Can you imagine having to telnet to a remote system today? The exposure to attack, password interception, and other potential dangers have made telnet a thing of the distant past, along with rsh (remote shell) and rpc (remote procedure call), among others.

The SSH protocol and tool set—and SSHv2 in particular—now dominate this niche.

### IPsec

VPN technology allows safe business transactions among branch offices, trusted vendors, road warriors, and telecommuters. IP Security (IPsec) in particular offers more than just seemingly private channels; it encrypts the data flowing over those channels, which multiprotocol label switching (MPLS), asynchronous transfer mode (ATM), and others do not. Aside from being complex and fragile, circuit emulators such as MPLS and ATM are unencrypted and therefore vulnerable to tapping because they are virtual networks, not virtual private networks.

Unfortunately, IPsec is only half a win. It is widely deployed in prepackaged and configured VPN devices, and it is usually managed by an ISP, security company, or local guru. It is the last that is the sore spot in IPsec. It requires a guru to configure. This is inexcusable and unnecessary, and until setup and maintenance become a lot easier, IPsec will remain a specialised corner and its promise will be only partially realised.

Photo by Peter Lothberg

Russ Housley and members of the IESG at the IETF 70 plenary

### Pretty Good Privacy (PGP)

PGP allows us to exchange signed and strongly encrypted e-mail whose content cannot be repudiated; in other words, the sender cannot deny sending it. PGP can also be used to encrypt files on one's hard drive. This free tool is so powerful that the U.S. government tried to suppress its export even harder than the efforts it usually makes.

It's also worth noting that trust in PGP is nonhierarchic (meaning, there is no central authority). PGP entities attest to each other's identities in a web of trust, as opposed to adhering to a particular hierarchy. Therefore, it is quite decentralised and immune to compromise of root trust anchors.

### X.509

X.509 certificates and the public key infrastructure to support them are used in browser authentication. The problem here is that because they are totally hierarchic, they are only as reliable as the Certifying Authority (CA) that issues them; and the commercial CAs that issue certificates have little financial incentive to make the effort to truly validate identity. There have been notable compromises of the X.509 certificate hierarchy.

Use of X.509 certificates for attesting to IP address space ownership will be coming into use at the ARIN, APNIC, etc., and ISP levels in 2008.

### S/MIME

There is a second, far less used, e-mail signing method that is used to some extent in the corporate world. Its function is similar to that of PGP, but it relies on an X.509 certificate hierarchy.

In an environment where attacks are recurring events, where operating system and application vulnerabilities are discovered daily, and where bot-nets of 100,000 zombies attack, we still feel comfortable conducting financial and other private transactions worth billions of dollars on the Net.

### Summary

There are free, open-source tool kits for all of the tools and protocols described here, all of which are incorporated in browsers and e-mail packages.

This is not to say there are no longer security threats on the Internet. Certainly, spam, DDoS, the Estonian DDoS, and the root DNS attack of 7 February 2007 are all very real and serious problems, but thanks to the good folks who gave us the protocols and tools described here and the applications that use them, we can walk safely through a dangerous city.

Again, in the same way that we can build a reliable Internet out of unreliable components, we can build secure applications and services that work well in today's highly insecure environment. This is a big win.

# The Perfect Attack

***By João Damas and Frederico Neves***

Attacks of all types have existed on the Internet for a long time. They have targeted individual users, servers, client machines or applications, and the infrastructure itself. They have had different degrees of success in achieving their goals, many of which have not always been clear for everyone to see. While attacks initially may have been motivated by clever coders seeking attention, over time the reasons behind attacks have become more varied. The most popular of them seems to be economic gain but not necessarily legal in any jurisdiction.

In this article we visit one type of attack: a type that may not be the most directly profitable for the attackers but one that has proved especially threatening to potential victims. What is especially insidious about this type of attack is that it is often launched for the purpose of attracting the attention of potential customers for the attackers, who initiate these attacks to demonstrate their skill. In fact, there is general consensus within the Internet security community that attracting business is the key motivator behind attacks on root servers, since such attacks always make headlines even if the damage is minimal.

It is extremely difficult, if not impossible, to defend oneself from this type of attack because the traffic it generates appears to be nearly identical to legitimate traffic; therefore, a direct defense is in itself a component of the attacker's success.

## The Perfect Attack

The attack on Internet infrastructure that we are describing here makes use of the Domain Name System (DNS). The attack uses some of DNS's necessary features for the profit of the attacker or its customer.

The DNS is designed to use the UDP part of TCP/IP as its main transport. UDP is a perfect match for DNS due to the short question/answer exchange that is involved in most DNS queries and usually completed with the involvement of only one packet from client to server and one packet from server to client. This feature is one of the characteristics of DNS that makes it especially scalable, as there is usually none of the overhead of a session establishment during a DNS query/answer interaction. Because there is no state referring to the client in a DNS authoritative server, it's possible for an authoritative DNS server to answer a high rate of incoming queries. Recursive servers—those that perform queries on behalf of their clients to walk the DNS tree and find the requested information—do have to maintain state while they issue the various queries that may be required to get the answer the client initially asked. Even with this added burden, a small number of recursive servers can handle large populations of clients because only a small portion will be performing queries at any given time. Recursive servers implement the caching mechanism (as described in DNS RFCs) and therefore can reduce the amount of external traffic required, as long as results from previous queries are readily at hand. This caching mechanism has worked very well over the years.

On the other hand, the connectionless/stateless characteristics of DNS and its preferred underlying UDP transport are, as it happens, Achilles' heels when it comes to differentiation of traffic arriving at a server. With DNS, a client is required to send a single packet to a server only in order to trigger a response by the server, with the corresponding work on the server side being performed. The effects on the server and network are therefore much more severe than the one-packet interaction of SYN attacks against TCP stacks popular some years ago.

In addition, even though IP packets carry both a destination and a source address as part of their header information, the source address is generally looked at only once the packet arrives at its final destination and is used only as a means to address the reply sent by the destination server, if any is to be sent.

The combination of these characteristics opens up the possibility for what are now called *reflector attacks*. In reflector attacks, a series of compromised hosts sends correctly formed DNS queries to recursive resolvers to which they have access over the network, but queries are crafted in such a way that the source IP address of the UDP packet is not that of the actual sending host but that of a victim. When the packet arrives at a recursive resolver, the answer is eventually sent out not to the host that sent the original packet but to the host with the faked source (now destination) address. By using a few of these recursive resolvers spread around the network, one can concentrate the streams of network traffic from the individual recursive resolver onto the victim in a way that can generate traffic levels that are unbearable for

*Continued on next page*

## Comic BoF

*The Perfect Attack, continued from page 27*

the target host or the networks it lives in, thereby causing the service to effectively collapse.

The number of hosts used in the attack and their bandwidth can vary depending on what the attackers have at their disposal—sometimes preferring a few well-connected machines located on campus LANs with good Internet access and other times using botnets of hosts behind domestic broadband links. In all cases, careful distribution of traffic among originating hosts and recursive resolvers can make the traffic levels go undetected until the traffic gets close to the victim's network and becomes focused on a single point.

The victim always has a hard time because it is usually the network itself that gets saturated, and even if the network administrators were to block incoming traffic at their border routers, it is probably too late; by this point, their incoming links are probably saturated.

The attack becomes a perfect attack when the victim is itself an authoritative name server. The server administrator is then faced with the problem of not being able to distinguish between real queries and attack traffic. Each of the recursive resolvers being used as reflectors is likely to also provide service for a community, and simply blocking traffic from one of the reflectors will render the service unavailable for the entire community served by it.

Mitigating these attacks usually requires collaboration among the organisations responsible for the victim servers and the ISPs that carry the traffic, which attempt to trace traffic back to its origin and to find the command and control centre that coordinated the attacking hosts. So far, the usual reaction has been to increase capacity in order to increase the chances of surviving such an attack.

While there are a number of forms these attacks may generate, only a few measures will be mentioned here that can be taken to at least help with closing the door to some of the possibilities. The first and easiest measure is to have the administrators of recursive resolvers configure their servers so that they provide service only for their intended audience and not the entire Internet. It is still quite common today to see a recursive resolver that will answer to any machine on the Internet—a holdover from older, gentler times. The IETF is trying to make a recommendation for these administrators in hopes of seeing some improvements. The recommendation is also geared toward DNS software vendors to alter their default parameters, so that service is by default provided only for a relevant default population—for example, those machines using the same IP prefixes as the server interfaces. It is good to see vendors taking action in this area already.

> It is extremely difficult to defend oneself from this type of attack because the traffic it generates appears to be nearly identical to legitimate traffic.

The second and more difficult option is for ISPs to check the source IP address of packets in their networks and weed out the ones that shouldn't be there. This can itself be tricky.

While ISPs use knobs in their routing configurations, referred to as *policy routing*, which may inspect the source address in the packets for a variety of reasons, these source addresses are used mostly for traffic categorisation—for instance, for sending some source addresses or UDP/TCP ports via connections with more-controlled timing characteristics than other source addresses or ports have.

A less frequent event is the utilisation of router capabilities that look at the source addresses of the packets to verify that they are within the set of addresses that should be seen coming in via a given router interface. In principle, only networks behind that interface should be sending packets to that interface, with their addresses as sources. This is generally the case for customer networks behind their ISP's access routers or enterprises sending traffic through their office or campus routers to their upstream providers. The picture gets harder when ISPs with multiple peers are involved. When multiple paths are made possible for traffic exchange, asymmetric traffic is a definite possibility. Any such network is likely to see traffic intended for one destination exit via one path and the response enter the network via an entirely different path. This can be due to engineering considerations at the ISP in question or as a result of similar decisions made elsewhere in the network. Whatever the case, this is a feature that provides one of the pillars of resilience to the Internet and therefore should not been seen as a problem.

Encouraging this sort of check and balance at the internal edges of the networks, where the core of the ISP network faces its customers, has the best chance of success. Most current router software for ISPs already includes features that allow this check to be performed without complex configuration and based on dynamic—rather than static and therefore harder-to-maintain—data.

As can be seen, the involvement and cooperation of all parties are required for a complete solution to this sort of problem. While many will think this kind of coordination is a utopia, it is also one of the basic features of the Internet—the interconnected mesh of disparate networks out in the world.

# Directions in Internet Transport Evolution

*By Bryan Ford*

At the Transport Area Open Meeting at IETF 70, area directors Magnus Westerlund and Lars Eggert noted that their queues are empty and that several Transport Area working groups are nearing completion. It may therefore be an opportune time to step back and consider possible directions for new projects in the Transport Area. This article summarises some of the ideas presented and discussed at that meeting, and it attempts to synthesise those ideas into an outline that describes a number of potentially promising directions for future Transport Area work. Identifying a complete shopping list of issues or work areas was not the goal of the meeting, nor is it the goal of this article. Instead, what follows can be best described as a cross section of possibilities.

The topics discussed at the meeting and summarised here fall into three main categories: *transport semantics*, *traffic management*, and *end/middle interaction*.

- *Transport semantics* is concerned with what transport abstraction the application writer sees or would like to build on. Described here is a new experimental transport protocol I presented at the meeting and how it compares with several existing transports.
- *Traffic management* is concerned with mechanisms that allow transport endpoints to take advantage of available network bandwidth while being fair to other applications and users competing for network resources. This section summarises some of the key issues in defining fairness, as identified by Bob Briscoe in his presentation.
- *End/middle interaction* is concerned with the interaction between the transport protocols at the endpoints, which traditionally assume that they have a clear end-to-end path provided by the IP layer, with middleboxes (such as NATs) and firewalls that intentionally obstruct or otherwise complicate this end-to-end path in various ways.

## Transport Semantics

TCP introduced the ordered byte stream abstraction on which most Internet applications have been built. The conceptual simplicity, elegance, and flexibility of this minimalistic stream abstraction continue to be among TCP's greatest strengths. Unfortunately, one of the basic assumptions embodied in this abstraction—that all bytes communicated in one direction as part of a given stream must always be delivered in order—creates practical performance problems for modern Internet applications that did not exist when TCP was designed. Because one lost packet in a TCP stream holds up all data queued behind it until the lost packet has been successfully retransmitted, TCP is almost unusable for real-time audio or video, in which it is much better just to drop and interpolate over isolated lost frames than to delay delivery of a whole series of frames. Modern Web browsers and other transaction-oriented applications similarly challenge TCP's simple, totally ordered stream abstraction, with their need to submit many logically independent or parallel requests to one or more servers efficiently (for instance, to load all of the images and other embedded objects on a complex Web page). TCP forces applications to choose between (1) using one stream per transaction, as in HTTP 1.0, which can be inefficient due to the costs of creating and destroying short-lived streams, and (2) multiplexing many logical transactions onto a smaller pool of streams, as in HTTP 1.1, which creates the same head-of-line blocking problem, as in real-time media applications, where one lost packet delays delivery of all of the (potentially unrelated) transactions queued behind it on the same stream.

TCP's limitations have been recognised for years, and they constitute the motivating force behind the development of several alternative transport protocols, such as RDP, SCTP, and DCCP. While each one of the alternative transports tends to address heavily overlapping needs and problem areas, each one also takes a different approach to solving them. A common feature of all existing alternative transports is that they move away from TCP's conceptually simple *byte stream* semantics and toward transport abstractions. DCCP provides unreliable, unordered delivery equivalent to UDP but with congestion control. RDP provides reliable, optionally sequenced message delivery with congestion control. SCTP provides reliable, optionally sequenced message delivery similar to RDP, but it also allows the application to associate each message with one of several *logical streams* within the application's transport connection, thereby permitting messages on different streams to be delivered out of order. A common issue with all of these message-oriented transports is that their message abstraction does not scale arbitrarily the way that TCP's byte stream abstraction does; instead, the application must break up large transactions into reasonably sized messages to avoid subtle performance problems or outright transport failures. This reasonable-size threshold is not generally well-defined and often varies with network conditions.

*Continued on next page*

*Directions in Internet Transport Evolution, continued from page 29*

At the Transport Area Open Meeting, I presented a new experimental transport called *Structured Stream Transport* (SST). Instead of moving away from TCP's familiar, conceptually simple, and scalable byte-oriented stream abstraction like other alternative transports do, SST enhances TCP's byte stream abstraction to permit applications to use streams in larger numbers easily and efficiently. With SST, for example, transaction-oriented applications like Web browsers need not either multiplex many transactions onto one TCP stream, as in HTTP 1.1, or retool to run on a message-oriented transport; instead, the application simply opens one new stream per transaction and relies on SST to implement those streams efficiently enough, whether it needs a few large streams, or a large number of short-lived streams, or some of each. An audio-streaming or video-streaming application on SST can preserve the transmission independence of separate frames simply by (1) opening a new (rather short-lived) stream for each frame it wishes to transmit and (2) using SST to take on the challenge of transmitting those ephemeral streams efficiently. Thus, SST's philosophy is not to discard TCP's serial byte stream abstraction but to adapt it to the demands of modern applications that demand nonserialised communication.

SST's main application-visible enhancement to TCP's stream abstraction is what amounts to a fork operation, meaning that given any existing SST stream, either endpoint can initiate a new stream as a child of that existing stream. The other endpoint accepts this child stream by performing a listen and accept on its corresponding end of the parent stream rather than on a traditional listen socket. For example, a Web browser using SST might open a top-level stream to communicate with a particular Web server, and then open a child of that stream to fetch the HTML for a given page on that server, and then further fork a Web page's HTML stream to load each of the embedded objects on that page. SST thus organises streams into a heredity structure—hence the term *structured streams*. Because SST preserves and communicates this heredity structure between the participating hosts, applications do not have to bind port numbers or authenticate each new stream: a child stream always starts with a clear communication context defined by the parent stream it is derived from.

Once created, each SST stream is independent and provides semantics essentially identical to TCP streams, including reliable delivery, ordering, and flow control independent of all other streams. The SST protocol contains optimisations that allow the application to create and start sending data on new streams, with no three-way handshake delay as in TCP, and SST can destroy streams without maintaining their state for a four-minute TIME-WAIT period as in TCP. All SST streams between given pairs of endpoints automatically share congestion control state, thereby avoiding the performance costs of a separate slow start for each new stream. The application can limit the length of time a stream's data is buffered for retransmission, which permits SST to be used for unreliable delivery when needed, such as when streaming media. For further details about the protocol, see my SIGCOMM paper.[1]

A natural question, then, is, How does SST differ in practice from SCTP, which also supports the sharing of congestion control and some other transport state among multiple logical streams? In addition to the basic semantic difference between SST's TCP-like *streams of bytes* and SCTP's RDP-like *streams of messages*, there are two, key, pragmatic differences. First, an SCTP application cannot dynamically open or close individual streams; instead, it opens a *connection* representing all the streams it will need, and SCTP negotiates the number of streams to be multiplexed onto that connection only once during connection setup. Second, SCTP provides receiver-directed flow control only for the entire connection and not independently for individual streams. The application can receive only the next message available on *any* stream and cannot pick a particular stream on which to receive. This means that the receiver cannot hold off the sender's transmission on one stream—such as in the case of a video file being downloaded for playing at a constant frame rate—while continuing to accept data on another stream, such as in the case of a file being downloaded to disk as quickly as the disk will accept it. SST streams, in contrast, work like fully independent TCP streams, only implemented more efficiently: the application can open and close them at any time and can read from some streams while holding off the sender on others.

> One of the basic assumptions embodied in the ordered byte stream abstraction—that all bytes communicated in one direction as part of a given stream must always be delivered in order—creates practical performance problems for modern Internet applications that did not exist when TCP was designed.

Could SST be implemented as a layer on top of SCTP? Yes and no. It may be relatively straightforward to imple-

1. "Structured Streams: A New Transport Abstraction," Bryan Ford. ACM SIGCOMM, August 2007. http://www.bford.info/pub/net/sst-abs.html.

ment SST's hierarchical stream abstraction with dynamic open/close on top of SCTP, making use of SCTP's fixed set of streams negotiated at connection time as a pool of low-level message streams on which to multiplex SST's TCP-like byte streams. Making up for SCTP's lack of independent per-stream flow control may be more difficult to do this way, however, because it would require the adaptation layer to maintain an additional set of send and receive buffers between SCTP's and the application's, thereby subjecting data to additional copying on the critical path.

SST is still in an early experimental stage, and all transports other than TCP and UDP face serious deployment challenges due in part to the end/middle interaction issues discussed later. Nevertheless, the amount of effort expended over the past decade on alternative transports with relaxed ordering and delivery semantics suggests that there is clearly a widely perceived need for such alternatives to TCP, even if the best approach is not yet clear.

### Traffic Management

A large portion of ongoing transport-related work both in the IETF and elsewhere is concerned with mechanisms for controlling the flow rate of network traffic. The goal is to permit applications to take full advantage of whatever bandwidth is available over a given path while ensuring that different applications and users share available bandwidth fairly and avoiding congestive collapse. The Internet's traditional approach to traffic management has been via end-to-end congestion control implemented in such transports as TCP, by which the transport dynamically senses the amount of bandwidth available and adjusts its transmission rate to match. A large portion of Transport Area work both within and outside of the IETF is devoted to development of new congestion control algorithms or to refinement of existing ones.

Unfortunately, the traditional end-to-end congestion control approach suffers from a serious flaw: since each transport endpoint has a limited view of the network, in which it sees only the results of its transmission attempts on particular end-to-end paths, the whole notion of fairness can be seen by a transport protocol only in terms of fairness between end-to-end *flows*. As Bob discussed in his presentation at the Transport Area Open Meeting, network operators tend to think of fairness not in terms of flow rate equality but in terms of volume accounting—in other words, in terms of how much traffic load a particular application or user is placing on the network regardless of whether that load consists of one end-to-end flow or many and regardless of whether the user causes the flow to be active continuously or intermittently. The problem with the traditional per-flow definition of fairness is exemplified by BitTorrent, an application that routinely uses dozens of concurrent TCP connections to different remote hosts. Each of those connections uses standard TCP congestion control and thus is entirely fair—in the per-flow sense—to other TCP applications. However, because BitTorrent's dozens of flows in aggregate consume dozens of times the bandwidth of a competing application using only one flow and because all of its flows run continuously, it is perceived by users and administrators as extremely unfair to other applications.

On the other hand, the volume-accounting view of fairness fails to take load variability into account. In other words, a given volume of traffic that causes considerable congestion during a time of peak load might cause little or no congestion at other times. Enforcing a simplistic view of fairness in terms of volume accounting can thus prevent applications from opportunistically taking advantage of available network capacity.

ISPs are increasingly deploying traffic-rate control devices in the middle of the network. Those devices sometimes attempt to enforce a sense of fairness among different applications and/or users. Unfortunately, the rate-control policies the ISP can enforce effectively are limited by what an ISP's routers can heuristically discover through deep-packet inspection; even when the ISP's intentions in setting the rate-control policies are honorable, the results often do not really correspond to what either the ISP or its users expect, thereby causing confusion and anger.

Network operators tend to think of fairness not in terms of flow rate equality but in terms of volume accounting—in other words, in terms of how much traffic load a particular application or user is placing on the network regardless of whether that load consists of one end-to-end flow or many and regardless of whether the user causes the flow to be active continuously or intermittently.

There are many difficulties in trying to come up with a truly workable notion of fairness for traffic management on the Internet. Bob proposes that deciding on such a notion is not the IETF's job; instead, he says, the IETF should focus on developing design-time accounting metrics and management mechanisms that enable sensible resource-sharing

*Continued on next page*

2. For details on the proposal, see draft-briscoe-tsvwg-relax-fairness-00.txt and Bob Briscoe's presentation slides at http://www3.ietf.org/proceedings/07dec/slides/tsvarea-3/sld1.htm.

*Directions in Internet Transport Evolution, continued from page 31*

policies to be enforced at run time.[2] Whatever the case, finding a reasonable way to escape both the per-flow fairness mind-set of traditional congestion control and the congestion-insensitive volume-accounting mind-set of administrative traffic control mechanisms and synthesising them into a scheme that enables end-to-end transports to work with middle-of-the-network devices so as to create truly useful fairness policies represent the next big challenges to be addressed in the Transport Area.[3]

## End/Middle Interaction

The IETF already has three working groups—MIDCOM, NSIS, and BEHAVE—that are concerned at least in part with improving the way traditionally end-to-end transports traverse and interact with middleboxes, such as NATs and firewalls. Several non-IETF projects, such as UPnP and NAT-PMP, take similar-but-different approaches to the middlebox interaction/traversal problem. Still another approach—extending STUN into a middlebox control protocol—was discussed at the SAFE BOF. Why did we end up with so many different approaches to solving the same problem? Because, as Lars pointed out in the open meeting, any of us can design a middlebox control protocol, but "nobody has designed one that anyone really wants to deploy."

We cannot simply hope that this problem will go away during the transition to IPv6, because even if NATs eventually disappear, firewalls are clearly here to stay. Many feel that NATs are here to stay, too, because of perceived benefits unrelated to IPv4 address space limitations such as administrative isolation/modularisation of address space and obfuscation of internal addresses from the viewpoint of external hosts. Furthermore, there appears to be renewed interest in using NAT to provide interoperability between the IPv4 and IPv6 universes. Because IPv6 is only now starting to see widespread deployment, it is still somewhat malleable, so now would be a great time to make any changes required to enable IPv6 transport protocols to work well over NATs and firewalls.

Because IPv6 is only now starting to see widespread deployment, it is still somewhat malleable, so now would be a great time to make any changes required to enable IPv6 transport protocols to work well over NATs and firewalls.

One possible explanation for the lack of deployment of current middlebox control/interaction protocols is simply that the need for them is not yet great enough, but perhaps it will be soon. As IPv4 address space pressure increases, the cost of static IP addresses will increase, and higher-profile players will start seeking cost-effective solutions to give their hosts full functionality even from behind NATs. As multilevel NAT scenarios become more prevalent and adjacent network domains increasingly end up using overlapping private address spaces, current traversal solutions make legitimate traffic arrive at unintended destinations, creating both efficiency concerns and security concerns that may increase the pressure for explicit middlebox control.

Mobility may also create pressure for the deployment of control protocols in order to reduce the power-draining keepalive traffic that current ad hoc traversal solutions require to hold their UDP bindings open. On the other hand, simply moving middleboxes and applications away from fixed-rate binding timers and keepalives and toward using binding timers with exponentially increasing periods might address the keepalive problem without explicit middlebox interaction.[4]

Therefore, on one hand, maintaining a wait-and-see attitude toward the current crop of middlebox control mechanisms and avoiding new work in this area until the waters clear up a bit might be an appropriate strategy. On the other hand, there are inherent risks with this strategy, mainly because most of the current mechanisms have limitations that may further increase the Internet's brittleness if those protocols become widely deployed without undergoing more-careful analysis and standardisation. For example, some control protocols, such as UPnP and NAT-PMP, do not address multilevel scenarios at all, whereas others can, but only when adjacent private address spaces do not overlap. This suggests that one potentially worthwhile near-term project in this area is to perform a careful, mechanism-neutral, side-by-side analysis of the currently available middlebox interaction mechanisms, clearly identifying the limitations of each and the potential risks to the Internet's future evolution if a given mechanism were to become the de facto standard for end/middle interaction. If we can't identify the right middlebox control mechanism, at least we can try to consolidate what wisdom we do have on the alternatives. This should provide useful guidance for vendors and customers that may now or in the future be considering deployment of middlebox control mechanisms.

3. For more-complete background and motivation, see "Flow Rate Fairness: Dismantling a Religion," Bob Briscoe. ACM CCR 37(2) 63-74, April 2007. http://www.cs.ucl.ac.uk/staff/B.Briscoe/projects/refb/#rateFairDis.
4. See "A Simpler Way to Reduce Keepalive Traffic." http://www1.ietf.org/mail-archive/web/safe/current/msg00073.html.

Aaron Falk, IRTF Chair

# IRTF Report

***By Aaron Falk***

*What follows are summaries of several upates on the Internet Research Groups (RGs), some of which were reported during the Technical Plenary at IETF 70.*

Since the July 2007 IETF meeting, one new IRTF (Internet Research Task Force) RFC has been published: RFC 5050, "Bundle Protocol Specifications." A document to formalise IRTF RFCs is currently being developed.

The IRTF has some new work items on its agenda. One is a proposal to follow up on the IAB's work on unwanted traffic within the IRTF. Interest has also been expressed in forming a research group (RG) that would focus on network virtualisation and another one focused on developing a Quality of Service policy framework. In the IETF Transport area, an IRTF RG dedicated to cross-layer communication has been suggested.

At IETF 70, three RGs met: the Host Identity Protocol Research Group (hip), the IP Mobility Optimisation Research Group (mobopts), and the Routing Research Group (rrg). Additionally, subsequent to the IETF, the End-Middle-End RG has decided to close based on lack of energy.

We would like to use this opportunity to offer details about the achievements, current work items, and future plans of most of the IRTF RGs as well as updates of discussions from IETF 70, held in December 2007.

The CFRG serves as a bridge between theory and practice, bringing new cryptographic techniques to the Internet community and promoting an understanding of their use and applicability.

## Crypto Forum Research Group (cfrg)

The CFRG serves as a bridge between theory and practice, bringing new cryptographic techniques to the Internet community and promoting an understanding of their use and applicability. It is a forum for discussing and analysing general cryptographic aspects of security protocols. IETF working groups (WGs) that are developing protocols that include cryptographic elements often find it useful to bring questions to the CFRG.

### *Current Work*

Members of the CFRG recently discussed message authentication code (MAC) requirements, in the contexts of both draft-irtf-cfrg-fast-mac-requirements and the TCP-AO (TCP–Authentication Option) currently under design in the TCPM (TCP Maintenance and Minor Extensions) WG. New work such as SHA-1 and MD5 that makes digital signatures less vulnerable to attacks against hash functions was presented in draft-irtf-cfrg-rhash-01.txt.

### *New Work*

New work in the form of draft-dharkins-siv-aes-01 has been reviewed, discussed, and revised. The work presents a new method for authenticated encryption that is more robust against misuse than most other modes are. It is under consideration in the TLS (Transport Layer Security) WG and other areas.

The draft of "An Interface and Algorithms for Authenticated Encryption" was approved as an RFC. The work has been adopted by the TLS WG as the basis for its use of AES GCM (Advanced Encryption Standard Galois Counter Mode). It is being adopted for other IETF uses as well.

### *Future Work*

The RG expects that a discussion of MAC candidates will follow the discussion

*Continued on next page*

IRTF Report, continued from page 33

of MAC requirements. References to some candidates have been provided.

### Delay-Tolerant Networking Research Group (dtnrg)

Members of the DTNRG are concerned with addressing the architectural and protocol design principles that arise from the need to provide interoperable communications with and among extreme and performance-challenged environments where continuous end-to-end connectivity cannot be assumed. In other words, is it possible to interconnect highly heterogeneous networks even if end-to-end connectivity may never be available? Examples of such environments are spacecraft, military and tactical programmes, certain forms of disaster response, underwater, and some forms of ad hoc sensor/actuator networks. Another example is Internet connectivity in places where performance may suffer, such as in parts of the world that are still developing.

In 2007, the RG published two RFCs:

- RFC 4838: Delay Tolerant Networking Architecture (Informational)
- RFC 5050: Bundle Protocol Specification (Experimental)

***Current Work***

There are a few areas of current work that are fairly mature and are likely to be completed as RFCs in 2008:

- LTP (Licklider Transmission Protocol): a transport protocol for high-delay environments
- Security: authentication and privacy for the DTN (delay-tolerant networking) bundle protocol
  - There was agreement at a meeting held in Dublin to favour counter-node crypto because of its length-preserving properties. This becomes important when fragmentation is performed.

***A Few Noteworthy Technical Items***

- Structure of the namespace
  - DTN uses URIs (Uniform Resource Identifiers) to identify endpoints, which include a scheme. Ongoing discussion revolves around the semantics associated with such schemes and how applications make use of them.
- Bit-level reliability
  - The bundle protocol does not currently contain a checksum or CRC (cyclic redundancy check) on the data (or blocks, which are similar to headers). Some of the folks involved in the RG would like to add this capability. [The security protocol uses a mechanism to ensure that bundle contents do not get modified either intentionally or unintentionally in transit, but some feel this approach may be too heavyweight.]
- Multicast
  - Although this issue has received some attention in the past, not much activity has been seen recently. The interaction with multicast and custody transfer can be tricky, and it remains an area of investigation.

***Future Work***

The DTNRG will meet at IETF 71 in March 2008 in Philadelphia.

Members of DTNRG are concerned with addressing the architectural and protocol design principles that arise from the need to provide interoperable communications with and among extreme and performance-challenged environments where continuous end-to-end connectivity cannot be assumed.

Upcoming discussions are likely to include some of the aforementioned technical issues in addition to discussion of the future of one or more "reference implementations" (RIs) of the bundle protocol. At present, we have identified one RI, but there are now multiple implementations that were demonstrated to interoperate during a DTN interop held at IETF 67. Issues revolve around whether there should be more than one RI for different operating environments and what its real purpose or purposes may be, such as education, demonstration, and performance.

For more information, see http://www.dtnrg.org and http://irtf.org/charter?gtype=rg&group=dtnrg.

### End-Middle-End Research Group (emerg)

The goal of the End-Middle-End Research Group is to evaluate the feasibility and desirability of an architectural change to the Internet that allows explicit interactions with middleboxes, such as firewalls. We have considered a higher-level DNS-based naming scheme, in the manner of URIs, coupled with signalling protocols that are used to initiate and modify transport-level connections, such as TCP, UDP, SCTP, or DCCP flows. The aim is to investigate possible designs for a straw man experimental protocol.

A joint meeting with the HIP RG took place at IETF 69 that was very successful. The EME NUTSS draft provides connectivity that keeps stakeholders in mind. NUTSS makes explicit policies of middles and ends. EME could help HIP with middlebox traversal because EME is a mechanism for relaying a policy request to the right box. Some people question whether EME should carry application semantics and whether modifications to the packets should be allowed. The complete minutes are available at http://www3.ietf.org/proceedings/07jul/minutes/HIPRG.txt.

The goal of the End-Middle-End Research Group is to evaluate the feasibility and desirability of an architectural change to the Internet that allows explicit interactions with middleboxes, such as firewalls.

The group published a draft (draft-irtf-eme-francis-nutss-design-00.txt), but recently there has been no activity. Disbanding the RG is under consideration.

### End-to-End Research Group (e2erg)

The E2ERG focuses on issues related to the end-to-end nature of communication. Historically, the group has focused its energies on one or more topics that are of particular interest to its members before moving on to another topic. Although the group is closed, it is sometimes necessary to invite other participants to meetings when there is an area of mutual interest. Two recent and current interests are (1) a review of the current state of congestion control and (2) questions about the provision of end-to-endness in an increasingly heterogeneous networking environment, particularly as it relates to management, routing, and characteristics of delivery services across qualitatively different subregions of the network. Typically, the group meets two or three times a year for two days at a time—at times and locations other than the times and locations of IETF meetings in order to avoid time conflicts. Most of the meetings have been in the United States, but the most recent meeting was in London. The next meeting will be in Cambridge, Massachusetts, in February 2008.

### Host Identity Protocol Research Group (hiprg)

The IRTF HIP research group (HIPRG) complements the IETF HIP working group. Its two main goals are:

- To provide a forum for discussion and development of aspects of the HIP

*Continued on next page*

IRTF Update, continued from page 35

architecture that are still in research phase and not ready for working-group-level standardisation

- To stimulate, coordinate, discuss, and summarise experiments on deploying HIP; to provide feedback at some later date to the IAB and the IESG regarding the consequences and effects of wide-scale adoption of HIP. For the latter goal, the RG had planned to produce an experiment report, which currently exists in draft form (draft-irtf-hip-experiment-03.txt).

To date, most of the energy of the RG has been devoted to the first goal. There have been and continue to be various drafts on such issues as privacy extensions for HIP; basic and advanced NAT traversal; the i3 architecture and HIP; DHT (distributed hash table) as a HIP lookup service; process migration using HIP; SIP (session initiation protocol) and HIP interactions; TCP piggybacking of HIP messages; middlebox interactions; HIP and multicast; and network operator concerns with HIP. The RG has also pushed documents into the rechartered HIP WG (NAT traversal, legacy application support, native API) and has published its own IRTF-track document: "draft-irtf-hiprg-nat-04.txt." There has been clear, ongoing interest on the part of a wide range of individuals and groups in studying how to extend the HIP architecture.

It has been difficult, however, for the RG to make progress on goal number 2. The chairs observe that coordinating and conducting experiments—particularly those oriented toward answering deployment questions—are much more difficult tasks than originally thought, especially compared with extending HIP. Since 2006, the HIP RG chairs have encouraged additional collaborative experimentation and dissemination of results. The chairs believe that encouraging wider-scale experiments and collaborations for the purpose of answering specific deployment questions about HIP is a priority. There are three open-source implementations of HIP that continue to mature, which means that software availability will become less of a barrier over time.

Since 2006, the HIP RG chairs have encouraged additional collaborative experimentation and dissemination of results. The chairs believe that encouraging wider-scale experiments and collaborations for the purpose of answering specific deployment questions about HIP is a priority.

The HIP RG met at IETF 70 in Vancouver, Canada, and the group plans to meet again at IETF 71 in Philadelphia in March 2008. At its last meeting, the session featured a presentation of a few new HIP ideas regarding multicast and Internet connection sharing, updates on HIP projects and deployments, and discussions regarding the potential use of HIP as part of the peer-to-peer SIP overlay solution.

### Internet Congestion Control Research Group (iccrg)

Ever since congestion control got included as part of TCP in 1988, this function has helped the Internet survive by ensuring its stable operation. However, it has long been agreed that this mechanism is not ideal; in fact, it shows deficiencies in the face of heterogeneous link properties, such as high capacities, long delays, or noise.

Now, after almost two decades, it seems that the demand for more better-performing mechanisms has become so strong that people are beginning to use alternatives that do not adhere to the standard anymore. Because this can lead to adverse interactions among different mechanisms, it is a major goal of the ICCRG to move toward consensus regarding (1) which technologies are viable, long-term solutions for the Internet congestion control architecture and (2) what might be an appropriate cost/benefit trade-off.

As a starting point for designers of new mechanisms, the group is currently working on two documents: a survey of congestion-control-related RFCs (which should help avoid reinventing the wheel) and an overview of open issues in the field of congestion control.

A process for evaluating experimental congestion-control proposals has been crafted with the goal of arriving at a recommendation to the IETF regarding publication of such proposals as RFCs (this process is currently applied to two documents describing TCP variants—CUBIC and Compound TCP—and discussions are ongoing).

### Network Management Research Group (nmrg)

The NMRG continues to study the behaviour of management protocols by using traces collected from operational networks. The number of traces available is steadily increasing, although we are still short of traces from enterprise networks where we expect a larger percentage of commercial management applications to be deployed.

On the technical side, the NMRG is working on exchange formats for SNMP (Simple Network Management Protocol) traces (currently in second last call) and a common framework (and associated tools) for data aggregation/separation of SNMP traces (under active development by a small group, to be posted as an ID in January 2008). Based on this common framework, people analyse the periodic behaviour of SNMP traffic or they identify basic data retrieval algorithms used by management stations. Another aspect is the development of suitable online visualisation tools.

After almost two decades [of congestion control], it seems that the demand for more better-performing mechanisms has become so strong that people are beginning to use alternatives that do not adhere to the standard anymore.

The NMRG did hold a one-and-a-half-day-long meeting in November 2007, which was hosted by the University of Twente in the Netherlands. Another meeting was held last year in Prague. In October 2007, an *IEEE Communications Magazine* paper on research directions in network management was published that originated from a joint NMRG/EMANICS workshop held in October 2006. Several regular NMRG attendees also participated in a July 2007 seminar on autonomic management held at Dagstuhl in Germany.

The SNMP measurement format ID is under second last call and will likely enter the IRTF publication process in January. The initial ID on data aggregation/separation will be posted in January 2008. It will then go through the RG process, which may mean that it gets last called in the summer. If things go well, it could be ready for the IRTF publication process in the second half of 2008. Thus, there will likely be a discussion about future work items during summer 2008.

### Scalable Adaptive Multicast Research Group (samrg)

The SAMRG was formed in June 2006 and is cochaired by John Buford of Avaya Labs Research and Jeremy Mineweaser of Massachusetts Institute of Technology's Lincoln Laboratory. The scope of the RG is to research application layer multicast techniques that leverage native multicast and that can adapt to different application requirements.

The SAMRG held two meetings in 2007: one was an interim meeting, which was held in January 2007 in conjunction with the P2P Multicasting Workshop; the other was a meeting held at IETF 69 in Chicago.

*Continued on next page*

*IRTF Report, continued from page 37*

The main result of the RG was development of the following drafts and publications related to work items in the charter of the RG:

- Problem statement and requirements

  draft-irtf-sam-problem-statement-01.txt

  draft-muramoto-irtf-sam-generic-require-01.txt

- Technology survey

  H. Yu, J. Buford. Advanced Topics in Peer-to-Peer Overlay

  "Multicast." *Encyclopaedia of Wireless and Mobile Communications* (ed. B. Fuhrt), CRC Press, forthcoming.

- Framework for the SAM design

  draft-irtf-sam-hybrid-overlay-framework-01.txt

  A key part of the SAM framework involves leveraging the design of Automatic IP Multicast without Explicit Tunnels (AMT) (draft-ietf-mboned-auto-multicast-08) by extending it to permit ALM connection.

- Hybrid ALM protocol proposals, including:

  Waelrich and Schmidt: The Hybrid Shared Tree Architecture

  Lei, Fu, Yang, and Hogrefe: Dynamic Mesh-Based Overlay Multicast Protocol

- Test bed for SAM experimentation and demonstration

  Participants from WIDE have developed and tested an XCAST router on a private PlanetLab. The work is discussed in draft-muramoto-irtf-sam-exp-testbed-00.

  The XCAST (multidestination multicast) router is an element of the SAM framework due to the synergy between overlay routing and multidestination routing in the underlay network, which has demonstrated message savings of 30 percent. The topic will be covered in an article to be published in *Computer Communications Journal* called "Exploiting Parallelism in the Design of Peer-to-Peer Overlays," by John Buford, Alan Brown, and Mario Kolberg.

The next SAMRG meeting is scheduled for IETF 71 in Philadelphia in March 2008. Goals for further work include moving the XCAST router to the public PlanetLab for use by the entire RG and integrating this with the extended version of AMT described in the SAM framework specification.

For more information, see http://www.samrg.org.

A key part of the SAM framework involves leveraging the design of Automatic IP Multicast without Explicit Tunnels (AMT) by extending it to permit ALM connection.

### Routing Research Group (rrg)

The RRG is chartered to research routing and addressing technology that is not yet ready for engineering efforts. For the moment, the RRG has elected to focus on the problem of finding a scalable routing and addressing architecture for the Internet. In the current architecture, a multihomed site either injects one or multiple provider-independent address prefixes into the routing system from multiple locations or otherwise injects one or more prefixes it received from one provider into the routing system through other providers. Either case necessitates global scope for the address prefixes, and it results in a scalability

issue. Locally scoped address prefixes are not sufficient, because legacy transport (TCP and UDP) connections use a specific prefix for connection identification, thereby tying the connection to a specific access link and creating a single point of failure. The primary thrust of most of the proposals currently before the group involves decoupling the location semantics from the identification semantics.

Since rechartering in early 2007, the RRG has met three times and has heard approximately 20 different technical proposals or updates to proposals. The group continues to receive new proposals and to refine a number of existing proposals. The mailing list has been reasonably active.

Currently, the goal is to drive the group to overall rough consensus through debate and comparison of proposals. The group plans to start the explicit windowing process in its next meeting (tentatively planned to coincide with IETF 71). The process is expected to take a year. During that time, new proposals are encouraged, and existing proposals are open for revision, potentially incorporating useful ideas from individual efforts and from the group's feedback. The group is interested in devising by the end of the meeting a single, scalable routing architecture that it can recommend to the IETF for further development.

For the moment, the RRG has elected to focus on the problem of finding a scalable routing and addressing architecture for the Internet.

## Transport Modelling Research Group (tmrg)

The Transport Modelling Research Group is chartered to produce a series of documents on models for the evaluation of transport protocols. The documents will include a survey of models used in simulations, analysis, and experiments for the evaluation of transport protocols. The output of the group will also include a broad set of simulation test suites and a set of recommendations for test suites for experiments in test beds. The group's goal is to improve its methodologies for evaluating transport protocols.

Recent accomplishments include the following:

- The first document from the TMRG, Metrics for the Evaluation of Congestion Control Mechanisms (draft-irtf-tmrg-metrics-11), is in the final stages of review by the Internet Research Steering Group.
- Gang Wang, Yong Xia, and David Harrison have produced an Internet-Draft on an NS2 TCP Evaluation Tool Suite (draft-irtf-tmrg-ns2-tcp-tool-00.txt), along with a Web page with simulation scripts. The document describes a tool for use in the ns-2 simulator for generating scenarios with typical topologies and traffic models and for evaluation of the results via a range of metrics.
- Lachlan Andrew organised a workshop in November at California Institute of Technology called the TCP Evaluation Suite Round Table, where scenarios for evaluating congestion control mechanisms were discussed. As a result of the workshop, a short paper titled "Towards a Common TCP Evaluation Suite" was written and submitted to PFLDnet2008 (the yearly workshop on Protocols for Fast Long-Distance Networks).

### *Future Work*

Future plans include the further development of best-current-practice scenarios for the evaluation of congestion control mechanisms in simulators and test beds and building on the paper "Towards a Common TCP Evaluation Suite." Plans also include completion of the Internet-Draft on Tools for the Evaluation of Simulation and Testbed Scenarios (draft-irtf-tmrg-tools-04). For more information, see http://www.icir.org/tmrg/.

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

# IETF Meeting Calendar

**IETF 71**
9–14 March 2008
Host: Comcast
Location: Philadelphia, PA, USA

**IETF 72**
27 July–1 August 2008
Host: Alcatel-Lucent
Location: Dublin, Ireland

**IETF 73**
16–21 November 2008
Host: Google
Location: Minneapolis, MN, USA

**IETF 74**
22–27 March 2009
Host: TBD
Location: North America (Provisional)

Register now for

## IETF 71

9–14 March 2008

Philadelphia, Pennsylvania, USA

http://ietf.org/meetings/71-IETF.html

Early bird registration: 635 USD (through Friday, 29 February 2008)
Regular registration: 785 USD
Full-time students: 150 USD with on-site proof of ID

**IETF 71 is being hosted by Comcast**

Special thanks to

Cisco Research and Microsoft

for hosting IETF 70

Special thanks to

Comcast

for hosting IETF 71

**The ISOC Fellowship to the IETF is sponsored by**

Afilias Google intel

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

Afilias ARIN Cisco

Comcast Motorola NeuStar RipeNCC


# IETF Journal

**IETF 70**
**Volume 3, Issue 3**
**December 2007**

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
Peter Godwin
Russ Housley
Olaf Kolkman

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


Internet Society