Volume 8, Issue 1 • June 2012

the IETF® Journal

*A report from IETF 83, March 2012, Paris, France. Published by the Internet Society in cooperation with the Internet Engineering Task Force**

Inside this Issue

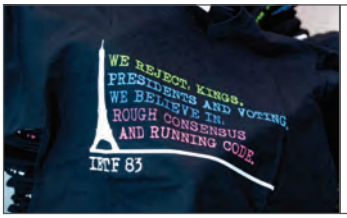

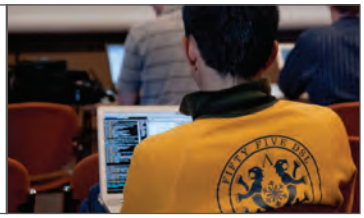

From the Editor's Desk

By Mat Ford

The beautiful city of Paris, France, played host to the 83rd meeting of the IETF, which, as always, contained a vibrant and diverse mix of Internet technology discussions, debates, and proposals from attendees drawn from every region of the globe. In this issue of *the IETF Journal* we have tried to capture a flavour of the proceedings.

Our cover article provides a useful introduction to the important topic of provisioning user identity for online services. As more and more applications and services run 'in the cloud', building efficient and scalable infrastructure to support them and lend them increased flexibility is important.

Another question that received renewed attention during IETF 83 is that of congestion control for real-time communication. On page 18, Randell Jesup provides an interesting introduction to the problem space, some of the relevant preexisting work, and one proposal now being discussed.

Jose Saldana had a baptism by fire as a new attendee when he offered a tutorial on online gaming network traffic and found many interested attendees ready to hear him speak. He writes about his experiences and his work on improving efficiency of gaming traffic on page 20.

Also in this issue are our regular columns by the IAB, IETF, and IRTF chairs; coverage of hot topics discussed during plenary meetings; and a piece on the contributions fellows make to the IETF.

As always, we are grateful to all our contributors. Send comments and suggestions for contributions to ietfjournal@isoc.org. And remember, you can subscribe to the hardcopy version or via email at https://www.internetsociety.org/publications/ietf-journal/ietf-journal-subscription.

An Introduction to Simple Cloud Identity Management

By Chris Phillips

Introduction

Simple Cloud Identity Management, or SCIM for short, made its first foray into the standards process at IETF 83 with a standing room only birds-of-a-feather (BoF) session. Since then SCIM has been working on finalizing its charter, which went to the area directors in late April and has been a topic of interest in numerous identity- and access-management communities, such as the

Continued on page 5

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

Photos: Richard Stonehouse/Internet Society

Message from the IETF Chair

By Russ Housley

Russ Housley, IETF Chair

More than 1,300 people from 56 countries came to Paris, France, and actively engaged in developing the future of the Internet. It was exciting to see so many people collaborating.

IETF 83 was very well attended. Of course, Paris in the springtime is an attractive location, but people came to work with others to make the Internet better for everyone.

Until the last minute, we did not have a host for IETF 83. We held a t-shirt-design contest, and IETF participants selected the winning design. With little time to spare, Cisco stepped forward and sponsored the meeting. Many thanks! Dave Ward was recognized at the plenary meeting on Wednesday for his role in making this happen. He was awarded a "Super Host" cape, which was designed and sewn by Amy Vezza from the IETF Secretariat.

Dave Ward dons a "Super Host" cape, a gift for providing last-minute sponsorship for the Paris meeting.

Our sponsors played an important role in making IETF 83 successful. AFNIC, NBCUniversal, Scality, and Tail-f served as meeting sponsors. Orange sponsored the network connectivity. Huawei sponsored the welcome reception. Thanks to all for your support.

Since IETF 82, three working groups (WGs) have been chartered and five have closed—our count remains steady at 115 WGs. Between meetings, the WGs and their individual contributors produced 576 new Internet-Drafts and updated 1,144 existing Internet-Drafts, some more than once. The Internet Engineering Steering Group (IESG) approved 125 Internet-Drafts for publication as RFCs. The RFC Editor published 115 new RFCs.

IETF tools continue to improve. The second major update to the Datatracker was released shortly before IETF 83, providing a major upgrade to the database schema. The new schema allows volunteers and contractors to more easily provide new capabilities. Support for WG chartering and rechartering activities was released in May. A significant enhancement that will allow anyone in the community to track the status of documents of interest will be released by the time this article is published.

There's been a lot of hype about "big data" on the Internet. Despite the hype, there are some real technical challenges associated with big data, and some people think that the IETF is a good place to tackle them. In my opinion, the IETF is very well suited to address many of them. If you agree, please put forward your ideas for BoFs (birds-of-a-feather meetings) in this area.

IETF 84 will take place in Vancouver, BC, Canada, from 29 July–3 August 2012. Google will be the host. Scheduling information for the upcoming IETF meetings can always be found at http://www.ietf.org/meeting/. I look forward to seeing you in Vancouver.

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

Recent IESG Document and Protocol Actions

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

Dave Ward photos: Richard Stonehouse/Internet Society

Words from the IAB Chair

By Bernard Aboba


Bernard Aboba, IAB Chair

IAB Retreat

The Internet Architecture Board (IAB) held its 2012 retreat 10–11 May, in Washington, DC, U.S.A. As in the past, the focus of the retreat was to set the agenda for IAB work in the next year and to review the programmes and initiatives that have been put in place to manage the work. An update on the retreat will be provided in the next issue.

IETF 83 Technical Plenary

The IETF 83 Technical Plenary featured Leslie Daigle, who provided an update on World IPv6 Launch, as well as a panel introduced by Hannes Tschofenig covering Implementation Challenges for Browser Security.[1] The panel included Chris Weber, who presented "When Good Standards Go Bad," an examination of how new browser features can enable attacks that were not possible before; Eric Rescorla, whose presentation, "How do we get to TLS Everywhere?" looked at TLS (transport layer security) deployment barriers; Ian Fette presented "Lessons Learned from WebSockets (RFC6455);" Tom Lowenthal presented "Cryptography Infrastructure," which addressed issues with the certificate trust model; and Jeff Hodges, whose presentation titled, "It's Not the End of the World," offered an overview of both progress over the last decade and current initiatives.

RFC Series

At IETF 83, Fred Baker, chair of the RFC Series Oversight Committee (RSOC) introduced Heather Flanagan, the newly appointed RFC Series Editor (RSE), to the IETF community. Also, the IAB approved publication of RFC Editor Model (Version 2)[2] and Independent Submission Editor Model[3] as informational RFCs within the IAB stream.

IESG/IAB Statement on IAB Member Roles in Evaluating New Work Proposals

Over the years, questions have been raised about the role of IAB members in evaluating new work proposals. To provide clarification, Spencer Dawkins worked with the IESG and IAB to craft a statement[4] that includes the following paragraph:

> In providing architectural oversight and guidance for a BoF (birds-of-a-feather), IAB members are rarely speaking for the IAB, because the IAB has rarely determined an IAB consensus position on new work being proposed. Individual IAB members do not have a privileged role in determining whether a BoF will result in a chartered working group. IAB members providing architectural oversight must ensure that their role is not misunderstood by BoF proponents or by the larger community of interest.

Privacy Programme

The Privacy Programme, led by Alissa Cooper, released a questionnaire on IPv6 privacy implementations.[5]

IP Evolution Initiative

The IP Evolution Initiative, led by Hannes Tschofenig, published the Smart Object Workshop Report[6] as RFC 6574.

Continued on next page

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

Words from the IAB Chair, continued

IPv6 for IAB Business Initiative

The IPv6 for IAB Business Initiative has brought online an experimental audio and web conferencing service operating over IPv6.

Comments on OMB Circular A–119

The IETF and IAB provided comments on OMB Circular A–119 Federal Participation in the Development and Use of Voluntary Consensus Standards and in Conformity Assessment Activities.[7]

IANA Evolution Programme

On 3 April 2012, Leslie Daigle (Chief Internet Technology Officer, ISOC) and Russ Housley (Chair, IETF) met with U.S. National Telecommunications and Information Administration (NTIA) Administrator Lawrence Strickling relating to the Internet Assigned Numbers Authority (IANA) protocol parameter function. A one-page summary of IETF–IANA oversight,[8] written by IETF administrative director (IAD) Ray Pelletier, was left behind. Subsequently, the IANA webpage on the IETF website[9] was updated to include pointers to all documents relating to IETF–IANA oversight on the IAB, IETF Administrative Oversight Committee (IAOC), and ICANN websites. With the recent reissuance of the IANA RFP by NTIA,[10] the IAB will once again be submitting an ICANN performance evaluation.

ICANN Relations

With respect to ICANN relations, the IAB has replied to ICANN questions[11] relating to the IAB statement on Interpretation of Rules in the ICANN gTLD Applicant Guidebook.[12]

ITU–T Coordination Programme

At the request of the ITU–T Coordination Programme, the IAB has initiated an IETF-wide Call for Comment on IETF and ITU–T Standardization Sector Guidelines.[13] The Call for Comment ended 25 May 2012.

IETF/IEEE–SA Coordination

The IAB appointed Dan Romascanu as IETF liaison to IEEE–SA.[14] One of Dan's responsibilities is to coordinate a meeting between the IAB, the IEEE 802 Executive Committee, and IESG, which will take place in the Bay Area of California in the United States in July 2012.[15] In March 2013, IETF and IEEE 802 will meet in the same hotel in Orlando in adjacent weeks, so a second meeting will probably be scheduled for that week. Topics for discussion will include a recent liaison from IEEE 802.1 to the IESG suggesting potential areas of coordination[16] and discussion of how well the coordination mechanisms described in RFC 4441[17] have been working in practice.

References

1. Technical plenary materials, http://www.ietf.org/proceedings/83/technical-plenary.html
2. RFC Editor Model (Version 2), http://tools.ietf.org/html/draft-iab-rfc-editor-model-v2
3. Independent Submission Editor Model, http://tools.ietf.org/html/draft-iab-ise-model
4. IAB/IESG statement on IAB Member Roles in Evaluating New Work Proposals, http://www.iab.org/?p=5741
5. IAB Privacy Program releases IPv6 Privacy Questionnaire, http://www.iab.org/?p=5842
6. Smart Objects Workshop Report, http://tools.ietf.org/html/rfc6574
7. IETF and IAB comment on OMB Circular A–119, http://www.iab.org/?p=5837
8. IETF Oversight of the IANA Protocol Parameter Function, http://www.iab.org/?p=5779
9. IETF-IANA Oversight, http://www.ietf.org/iana.html
10. IANA RFP, https://www.fbo.gov/index?s=opportunity&mode=form&id=72dc5eb7b831f44f5eadb6c2f44a60ef&tab=core&_cview=0
11. IAB responds to ICANN questions concerning the interpretation of rules in the ICANN gTLD guidebook, http://www.iab.org/?p=5723
12. IAB posts statement: The interpretation of rules in the ICANN gTLD Applicant Guidebook, http://www.iab.org/?p=5631
13. Call for Comment on IETF and ITU–T Standardization Sector Guidelines, http://www.iab.org/?p=5829
14. IAB appoints Dan Romascanu as IETF liaison to IEEE–SA, http://www.iab.org/?p=5681
15. IAB and IESG to meet with IEEE 802 Executive Committee, http://www.iab.org/?p=5793
16. Liaison to IESG from IEEE 802.1, https://datatracker.ietf.org/documents/LIAISON/liaison-2012-03-19-ieee-8021-the-iesg-liaison-to-iesg-from-ieee-8021-attachment-1.pdf
17. The IEEE 802/IETF Relationship, http://tools.ietf.org/html/rfc4441

Corrigenda

The article, "Internet Society Announces Winners of Postel, Itojun Awards and Applied Networking Research Prize," in the last issue of *the IETF Journal* (Volume 7, Issue 3, pp. 15-16) incorrectly stated that the Internet Society granted the 2011 Itojun Service Award. The Itojun Service Award was established by the friends of Itojun and is granted by an award selection committee. *The IETF Journal* sincerely apologises to the award committee and Itojun's family for this error.

The same article failed to recognize the role of the Internet Research Task Force in coordinating and awarding the Applied Networking Research Prize. Again, we sincerely regret this omission.

An Introduction to Simple Cloud Identity Management, continued from page 1

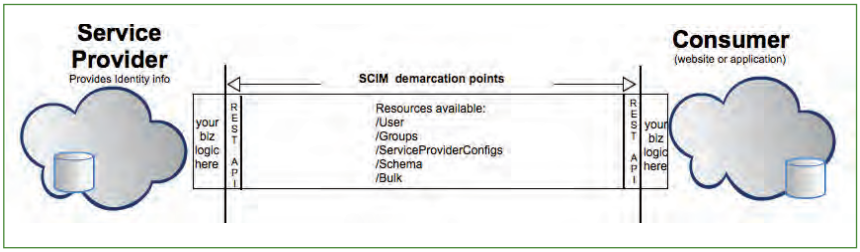

Figure 1. Key elements of SCIM

Internet Identity Workshops (IIW), working groups (WGs) of Internet2 and TERENA, the Kantara Initiative, and several advocacy campaigns of SCIM contributors.

What is SCIM?

The SCIM protocol takes a pragmatic approach to the challenge of provisioning user identity across cloud-based service providers. The *Simple* in Simple Cloud Identity Management is more than just a name; it is a principle participants have used to evolve the concept and hope to continue as it goes through the IETF process to become a formal standard. The SCIM website has a phrase that succinctly sums up the approach: "In essence, make it fast, cheap, and easy to move users in to, out of, and around the cloud."

Clearly articulating what is in scope for SCIM has been a vigorous exercise of how to strike a balance between complexity, simplicity, and the appropriate level of utility to provisioning identity and identity-related information.

Why now?

Without broad adoption of a standard way to do user provisioning, services have had to build custom systems. In turn, anyone doing business with that service must bear the cost of using a custom provisioning interface for each service and a custom schema with little re-use along the way. Provisioning within organizations also suffers this same fate.

While there are standards in this space, adoption is low. There is an unexpected scaling challenge in that instead of worrying about whether it can 'scale up,' the primary question is whether it can scale to just the right size. An organization may only need a fraction of what other protocols offer and need to operate on the corresponding fraction of resources and infrastructure.

SCIM's pragmatic approach is attractive as it is designed to be nimble, less burdensome to adopt than other protocols, and it is more cost effective than to write, staff, and maintain a custom provisioning environment.

The Protocol

The SCIM protocol exposes a common user schema and extension model expressed in JavaScript Object Notation (JSON)[1] format or XML format over HTTP using a Representational State Transfer (RESTful)[2] API. Figure 1 illustrates the following key elements in SCIM:

- the service provider, which holds the identity information being operated on;
- the consumer, which is a website or application using the SCIM protocol to manage identity data maintained by the service provider; and
- resources, which are service provider-managed artifacts containing one or more attributes.

SCIM requests are made via HTTP operations and responses are returned in the body of the HTTP response, formatted as JSON or XML depending on the request, with the request status indicated in both the HTTP status code and the body of the response.

Full details of the protocol responses can be found on the SCIM website[3].

Schema

The SCIM schema was inspired by the approach to schema taken by Portable Contacts[4], along with some extra elements from initial participants. The Portable Contacts model offers more flexibility than other formats to capture data complexity, which translates into the ability to capture complex data relationships at a user level. In

Operations and Expected Actions for HTTP

HTTP Operation	Action
GET	Retrieves a complete or partial resource.
POST	Creates a new resource or bulk modifies resources.
PUT	Modifies a resource with a complete, consumer-specified resource (replace).
PATCH	Modifies a resource with a set of consumer-specified changes (partial update).
DELETE	Deletes a resource.

Continued on next page

An Introduction to Simple Cloud Identity Management, continued from page 5

Resources with Respective Endpoints

Resource	Endpoint	Operations	Description
User	/Users	GET, POST, PUT, PATCH, DELETE	Retrieve/Add/Modify users
Group	/Groups	GET, POST, PUT, PATCH, DELETE	Retrieve/Add/Modify groups
Service Provider Configuration	/ServiceProvider-Configs	GET	Retrieve the service provider's configuration
Schema	/Schemas	GET	Retrieve a resource's schema
Bulk	/Bulk	POST	Bulk modify resources

keeping with the spirit of simple, the core schema for SCIM is intended to meet 80 percent of a user's basic attributes and allow implementers to get up and running as quickly and easily as possible. If mappings between the identity set being used and SCIM do not exist, extensions are available to tailor the schema. SCIM endpoints can be interrogated for schema much like Lightweight Directory Access Protocol (LDAP) servers so schema customizations are discoverable.

A side effect of this schema style is that SCIM can encapsulate more detail about an identity than LDAP's inetOrgPerson or a Security Assertion Markup Language (SAML) profile. The approach to address this is to have recommendations for LDAP and SAML mappings that will provide guidance on how to take a high-fidelity SCIM schema and transform it to a lower-fidelity format. The intended outcome of the mappings is to make it easier for implementers and to help foster consistency for those who need mappings. Other mappings may be published depending on the demand for them.

SCIM, Schema, and Scope

Schema has been one of the most frequently discussed items on both the SCIM and IETF lists with topics like:

- should attribute values be constrained based on other attribute values?
- should applying certain access control methodologies using certain attributes be expected or required? and,
- what is expected of SCIM regarding management of unique identifiers?

Each question is interesting on its own as an identity-management topic. That said, SCIM's guiding principle of simplicity suggests that those requirements should be overlaid/profiled atop the specification to meet the demands of each unique use case. Constraining SCIM to a minimal standard schema with a flexible information model, combined with structured data transport, has broader utility and encourages adoption. At the same time, there is opportunity to leverage SCIM as the lower-level building block for those advanced uses and focus on the application of policy and delegate transport of the data to SCIM.

SCIM's guiding principle of simplicity suggests that those requirements should be overlaid/profiled atop the specification to meet the demands of each unique use case.

Running Code

Specification interoperability sessions are held where possible to exercise implementations and gain running code experience. Three have been held so far with the second held in Paris just before IETF 83[5] with a total of nine participants. The face-to-face sessions have been invaluable for identifying gaps or ambiguities that need to be clarified. The most recent was held at an Internet Identity Workshop from 1–3 May in Mountain View, California, U.S.A.[6] There were nine participants, some of whom have SCIM implementations in production.

The Road Ahead

SCIM has been at the 1.0 state since December 2011 as a result of strong community participation. Vendors have been running code for months and a number have SCIM as features in their products[7] collecting valuable real-world experience living with the technology and going through their product lifecycle with it. Others have stayed on the sidelines, either to see if the protocol gains more traction or because they are not comfortable with the uncertainty surrounding a new protocol and are waiting to see what will happen with SCIM and the IETF.

Draft charter milestones have been proposed to adopt the SCIM core schema, RESTful interface definition, and use cases as a living document by the end of summer 2012 and by summer 2013 to have formalized SAML

bindings and LDAP mappings. This fills some gaps for those who may have been waiting to see what will happen and provides guidance on other areas of the protocol.

The SCIM WG has a number of topics vying for attention that we will introduce here as a glimpse into future discussions.

Defining and Recommending Possible Topologies

Describing possible deployment topologies will help identify where and how SCIM could be applied as a way to encourage adoption. The word *cloud* in SCIM is a bit of a misnomer and shouldn't preclude SCIM from being deployed internally. Internal deployments may be more compelling than connecting to a Software as a Service (SAAS) vendor. There is more customization and more ownership of provisioning processes internally and hence equal if not more utility in deploying SCIM to simplify the internal provisioning environment and consolidate to a common model with SCIM as a foundation element.

Detailing End-Point Security Approaches

The challenge is that one size does not fit all and not all sites are equally equipped to execute one model over another. SCIM requires TLS 1.2 to be used and recommends OAuth Bearer Token as a method to encourage interoperability between SCIM endpoints, but remains flexible to allow other protocols to be used.

One of the bindings being examined with the SCIM protocol is SAML and how it could and should interact with SCIM. From a security perspective, SAML offers SCIM a trust model through SAML federations (private or public) and raises other interesting security conversations. Should all endpoints within a trust set be *equally* trusted?, is just one of many questions in this space. It is likely SAML environments would benefit from the option to augment their experience with formalization around provisioning to SAML end points.

Schema

There are a number of discussion points that could appear under the schema heading. Some have been mentioned already but not the more subtle and slippery slope of a 'just one more thing' approach to schema additions. This is manifested by the appearance of more and more attributes outside the core schema until there are more attributes appearing in the extension than in the core. SCIM is flexible enough to allow this, but should this pattern be encouraged, discouraged, or not something to be concerned about? Is this the antipattern of fortifying the core attribute model? It may very well be an acceptable way to use SCIM if it simplifies the entire infrastructure or is a technique to keep the schema evergreen and allow for maximum flexibility as the years go by. It will be an interesting discussion to dig into.

This brief tour of provisioning topics highlights only some of the areas to be explored. As SCIM heads into the next round of improvements, keeping the balance between simplicity, practicality, and flexibility—even though they may be at odds with each other — will improve the durability of SCIM as a tool in the middleware toolkit. If these topics strike a chord with you, or if you have thoughts and would like to contribute, please join in by trying SCIM[8] and joining the mailing list at scim@ietf.org.

References

1. A lightweight text-based open standard designed for human-readable data interchange
2. http://en.wikipedia.org/wiki/Representational_state_transfer
3. http://www.simplecloud.info/specs/draft-scim-api-00.html
4. http://www.portablecontacts.net/draft-schema.html
5. http://code.google.com/p/scim/wiki/FirstInteropEvent
6. http://code.google.com/p/scim/wiki/SecondInteropEvent
7. http://code.google.com/p/scim/wiki/Implementations
8. http://simplecloud.info

It was standing-room-only at the Simple Cloud Identity Management birds-of-a-feather session.

Photo: Erik Wahlström, Nexussafe.com

Update

The System for Cross-domain Identity Management (scim) working group was chartered by the IESG on 21 June, 2012, in the Applications Area of the IETF. The agreed charter of the working group, including goals and milestones, is available at http://datatracker.ietf.org/wg/scim/charter/.

Implementing Identity Management Solutions

By Carolyn Duffy Marsan

Internet users are clamouring for better and easier ways to ensure the privacy and security of their digital information scattered across websites. Internet standards bodies, including the IETF, are responding to this demand.

The Internet Society held a panel discussion in March 2012 at IETF 83 in Paris, France, about emerging authentication and authorization standards being developed by the IETF, the OpenID Foundation, and the W3C. The panel's aim was to update attendees on the development of these standards and how they will intersect with each other in commercial implementations.

Lucy Lynch, director of Trust and Identity Initiatives for the Internet Society, led the panel discussion, calling it a "timely" topic because the IETF's OAuth working group (WG) will soon have two core specifications—for web authorization and bearer tokens—accepted as proposed standards. "I thought this was a good opportunity to look at IETF-sponsored protocols and see what happens when they actually get implemented in the wild," she added.

Bob Morgan, an IT architect for the University of Washington and an initiator of a trust framework for U.S. colleges and universities called the InCommon Federation, says the goal of all of this standards work is to create a framework for "pervasive usable identity."

"Billions of people have their digital stuff scattered about the Internet. They need to manage it. How do they get to that stuff? How do they control sharing that stuff?" Morgan explained. "That's the infrastructure we are trying to build to enable people to do that securely and with some notion that their privacy is not being totally abused while they do it."

Morgan has been developing authentication technology for 25 years, starting with Kerberos and leading up to Security Assertion Markup Language (SAML). He said the new standards under development by the OAuth WG and the OpenID Foundation are directly competitive with SAML, which is a positive development if they enable new security systems that scale to protect Internet users.

Panel speakers (left to right): Bob Morgan, University of Washington; Michael Jones, Microsoft; Hannes Tschofenig, Nokia Siemens Networks; John Bradley, identity-management expert; Harry Halpin, W3C team member

Photo: Richard Stonehouse/Internet Society

"We're trying to create an attribute economy, where information about everyone and ultimately about everything can be found from authoritative sources and where there is economic motivation to provide that information with reasonable privacy protections," said Morgan. "We need some great technology to make that vision come true."

Hannes Tschofenig, a network engineer with Nokia Siemens Networks and OAuth cochair, said it has been a positive experience for the IETF's OAuth WG to interact with industry players who are rapidly developing common-but-proprietary specifications for access control among data-sharing cloud applications.

"Reaching out to those communities is, I believe, a really important aspect that has to go along with the actual technical work," Tschofenig said, adding that those groups have created interesting designs that solve real-world problems. "We still see technical work happening outside the IETF for good reasons."

Michael Jones, a standards architect at Microsoft and editor of several OAuth and OpenID specifications, explained how Facebook, Google, Microsoft, and others designed a simple specification for security tokens that is in the OpenID specification.

"In order for digital systems to act on trust information, they have to have some notion of where they are getting the information from and what it is," Jones said. "What do I mean by a security token? It's nothing much more than a set of claims that one party makes about a subject. [For example], to do single sign-on, we use a set of particular claims that make assertions about who the party is that's logged into the system."

Jones said several components of a usable digital identity system—including security tokens, cryptographic signing and public key encryption—are

underway within the IETF. "These are all reusable pieces that are each small, easy to build, and easy to use," he said. "Some of this work came out of the OpenID world, where we made a conscious choice that we were building reusable components that we could repurpose. The IETF is a great place to do these general-purpose pieces."

Identity-management expert John Bradley told the panel that the OpenID group tried to integrate its OpenID 2.0 specification with the original OAuth specifications, but that the two specifications were extremely difficult to implement together. Now the OAuth group is reconsidering some issues related to SAML, such as handling structured responses and decoupling the identity provider from the attribute. OAuth also is collaborating with the Kantara Initiative for digital identity, which is building user consent modules allowing separation of protected resources from consent management.

"There are a bunch of different things that have gone into OpenID Connect, with different influences from different sources," Bradley said. "We just got to the 'implementer's draft' stage."

Harry Halpin, a W3C team member, described the group's efforts to create a common JavaScript cross-browser library for cryptographic permits. The idea for this "Web Security API" stems from a May 2011 workshop that W3C held regarding identity in the browser.

"Currently, OAuth is heavily bound to bearer tokens of [Transport Layer Security] in most implementations, but you could imagine greater and higher security flows built with a PKI infrastructure that would be easier to implement if it were available to the browser," Halpin said. "Putting PKI into the browser is a hard problem, and it is not something that we at W3C have solved. But we hope to include as many experts from the IETF to solve it in cooperation with the browser vendors."

Halpin said the W3C wants to encourage the emerging identity ecosystem.

"The world of identity management is very difficult, very fractured, and very contentious. The situation as it stands today, in which the authentication process revolves around user names and passwords, is not going to stand," Halpin said. "It's becoming increasingly unsafe to do high-value transactions over the Web…We think our Web cryptography WG will have a definite impact, and I implore everyone who wants to learn more about it to contact [me.]…We want to work with all major proposed authentication mechanisms."

In a question-and-answer session following the opening remarks from each speaker, panelists were asked how end-users would know when a particular Web site is acting as their identity provider. Panelists explained that end-users with, for example, a Google account that allows them to log into other sites will have Google as their identity provider. Corporate users, on the other hand, will have IT departments provide an internally managed identity-management system. It is difficult, though, to design a user interface that lets end users know what's going on behind the scenes with the identity management process.

"There are certain things that are out-of-scope for the W3C, including user experience and user interaction on the browser," Halpin said. "A lot of the hard problems in getting users to understand privacy are going to be out-of-scope of the kind of standards that we normally do. There is a tremendous amount of regulatory [activity] moving at a speed that may be faster than the technical standards bodies."

Bradley pointed out that the OpenID Foundation has a WG called Account Chooser, which is working on a standard user interface for Web sign-in that will allow end users to preconfigure various identity providers. "It's an important thing that when people log into a Web site with their user ID and password, they have…some sense of the information that [data sharing sites] can get about them," he said.

> "The world of identity management is very difficult, very fractured, and very contentious. The situation as it stands today, in which the authentication process revolves around user names and passwords, is not going to stand."
>
> —Harry Halpin
> W3C Team Member

Overall, the panelists were optimistic that the standards under development by various groups including the IETF would help create the underlying infrastructure needed for identity management in the cloud.

"The landscape is moving really rapidly right now, and it seems to be moving in a positive direction that could lead to mass deployment solving a known problem that's been open for decades," Halpin said.

Browser Security: Many Challenges, Some Progress

By Monika Ermert

The IETF 83 technical plenary looked into browser security challenges, a timely topic as illustrated by ongoing work in several working groups (WGs) and debates in at least two lunch meetings on issues around web authentication. With a number of highly publicized security incidents in 2011, web security has gotten a lot of attention from the media, politicians, and governments, said Hannes Tschofenig, a member of the Internet Architecture Board (IAB).

"Even my favorite gaming platform, Steam, got hacked—it's really bad," said Tschofenig, an engineer who is also a driver of the IAB privacy effort. He also pointed to the fact that many aspects of browser security really require attention outside of the IETF. While the protocol work is done within the IETF and sister organizations like the W3C and its web security- and web authentication-related WGs, there is a longer chain of events that need to happen to ensure secure deployment.

Thomas Lowenthal, panelist, and privacy strategist at Mozilla

Making Browser Communication More Secure

As Eric Rescorla, a developer of transport layer security (TLS) standards, said, "On paper, web security does not look that bad." The TLS WG has delivered more than 27 RFCs, and every major browser vendor supports TLS. Yet Rescorla is not at all satisfied with the level of adoption.

The vast majority of web traffic (90 percent or more) is not encrypted. Only approximately one percent of sites offer secure http (https). "That number should be 100 percent, but the hassle of getting https running on a site has driven people away from it," Rescorla said. He explained that even for the IETF.org site it was impossible to get a certificate working in 48 hours. "If a man with a Ph.D. can't do this, who can?" he asked. Especially for noncommercial offerings, the cost-benefit analysis can result in deciding against making the effort.

Aside from the complexity of the technology, when https was introduced it sat beside http. At the time, there was a risk that mixed content allowed downgrade attacks leveraging the unfamiliarity of users with https. When users, for example, typed the familiar http string into their browsers and a man-in-the-middle blocked the redirect to https, both partners—user and site—thought everything was all right, when in fact it was not.

Looking toward potentially useful alternatives, Rescorla pointed to having certificates stored in the DNS. This work is being done in the IETF DANE WG. Still, Rescorla advised that existing practices and standards would be here for a long time. Shedding legacy technology is slow and difficult, he explained, as any server that wants to have TLS must possess both credentials—the alternative solution doesn't offer adopters any quick benefits.

> "The hassle of getting https running on a site has driven people away from it ... If a man with a Ph.D. can't do this, who can?"
>
> —Eric Rescorla
> Codeveloper of Transport Layer Security

Certificate Mess

Both Rescorla and the Mozilla Foundation's Tom Lowenthal bemoaned the problems of a fragmented and confusing certification market. According to Lowenthal, fixing cryptography is not the problem; fixing implementation is. Major mistakes, Lowenthal said, are that of the roughly 1,500 certificate authorities many, if not a majority, routinely sign things that are wrong. They issue certificates for names that do not exist (or cannot exist), or for names that are unqualified domain names, such as .mars domains.

Moreover, certificate providers often reside in different regions with different legal or cultural norms. For some, Lowenthal said, it has been okay to intercept secured traffic; in general the whole technology just had that "one little dependency in the middle" that made it completely imperfect in some way. Lowenthal pointed out that he had no real solution for dealing with the problem. Currently, Mozilla is working on an alternative solution for web authentication. For its browser ID

Photo: Richard Stonehouse/Internet Society

concept, Mozilla tries to exclude a middleman, similar to the web ID provider from the architecture.

How to Protect Existing Applications from Security Problems Created by your New App

Web Sockets coauthor Ian Fette, developer at Google for Chrome, reported how engineers tried to fix three kinds of problems detected during the development of web sockets: breach of established security boundaries, cross-protocol attacks, and interoperability issues. To protect (sometimes undocumented) security boundaries, for example between the open and the corporate network sitting behind a firewall, the web socket WG added an "origin header" (RFC 6454: The Web Origin Concept) that allows a check before accepting packets from across the security border.

An additional challenge-response established by the web socket client should help avoid cross protocol attacks. "The web-socket client," Fette explained, "would send a key that gets hashed and that signals acceptance of the connection response." Fette also pointed out that despite these measures, many servers passed on requests.

Fette drew several conclusions from the web-socket experience, mainly that developers were "deploying into an environment that had existing security assumptions." Therefore developers need to protect the existing infrastructures against new attacks enabled by the new protocol.

Ian Fette, panelist, and senior product manager at Google

Photos: Richard Stonehouse/Internet Society

Transition to Greener, More Secure Fields?

Chris Weber from Casaba Security described what developers were up to in their attempt to secure Web communication. "You pretty much have to be a rocket scientist to get it right these days," he said. Cryptography, for example, exists on different levels: the transport level, the session level, and the application level.

As Weber explained, if developers miss one tiny thing in complex environments, it can compromise the entire operation. What causes difficulty between the life of a developer and the life of a security evangelist is that all browsers have tiny differences, even in the application of existing standards. "Chrome, Firefox, Internet Explorer, each one implements things a little differently, especially things like URL handling and the way URLs are parsed, the way post messages might be treated or XML is treated," said Weber. "In addition, implementers are not delving as much into new standards as developers might hope. Therefore, it could create pain and vulnerabilities."

Questions from the floor

Despite all the warnings and all of the problems, PayPal's Jeff Hodges said the [Internet] world was better off today than it was a few years ago—due mainly to the rise of a multivendor market against an earlier monopoly market that saw a lot of proprietary development. Thanks to the multivendor efforts and the convergence on open standards, many more eyes are looking at the problems, said Hodges. He also spoke about transition, yet underscored that he saw "attack surfaces arguably shrinking."

World IPv6 Launch, New TLDs, and a New RFC Editor

During the Technical Plenary, ISOC chief Internet technology officer Leslie Daigle described World IPv6 Launch on 6 June, in which participants will make their services available over IPv6 and keep them enabled permanently. "This time it's for real," said Daigle, quoting the Launch's motto.

Newly appointed RFC editor Heather Flanagan provided participants with a look at future work.

With regard to ongoing IAB activities, there was discussion about the recent IAB statement on new generic TLDs and IDN TLDs. According to the IAB, an update of relevant RFCs is needed.

World IPv6 Launch: The Future Is Forever

By Phil Roberts and Mat Ford

Exactly six years since the shutdown of the experimental IPv6 network, the 6bone [RFC 3701], 6 June 2012 marked another step toward progress in the universal deployment of IPv6. World IPv6 Launch saw major Internet service providers (ISPs), home networking equipment manufacturers, and Web companies around the world coming together to enable IPv6 for their products and services permanently. This launch event, organized by the Internet Society, saw regular business operations with IPv6 'on by default' become a reality around the globe. The ultimate goal is commercially supported IPv6 products and services at Internet scale by the end of 2012.

The need for IPv6 deployment and the advantages of an open, globally addressable network should be very familiar to readers of *the IETF Journal*. In a nutshell, the healthy future of the Internet depends on the rapid deployment of IPv6. World IPv6 Launch served to accelerate planning for some, encourage adoption for others, and helped define IPv6 as the 'new normal' for global internetworking.

World IPv6 Launch by the Numbers

- **2,696:** Website participants now available over IPv6
- **60:** Network operators who are delivering over 1% IPv6 traffic
- **5:** Home-router vendor participants

http://ams-ix.net/sflow-stats/ipv6/: Evidence of doubled IPv6 traffic is seen in this monthly graph from the AMS-IX Internet Exchange.

http://eggert.org/meter/ipv6: The combined impact of World IPv6 Day (2011) and World IPv6 Launch is illustrated by the number of IPv6-enabled websites in this global ranking by Lars Eggert.

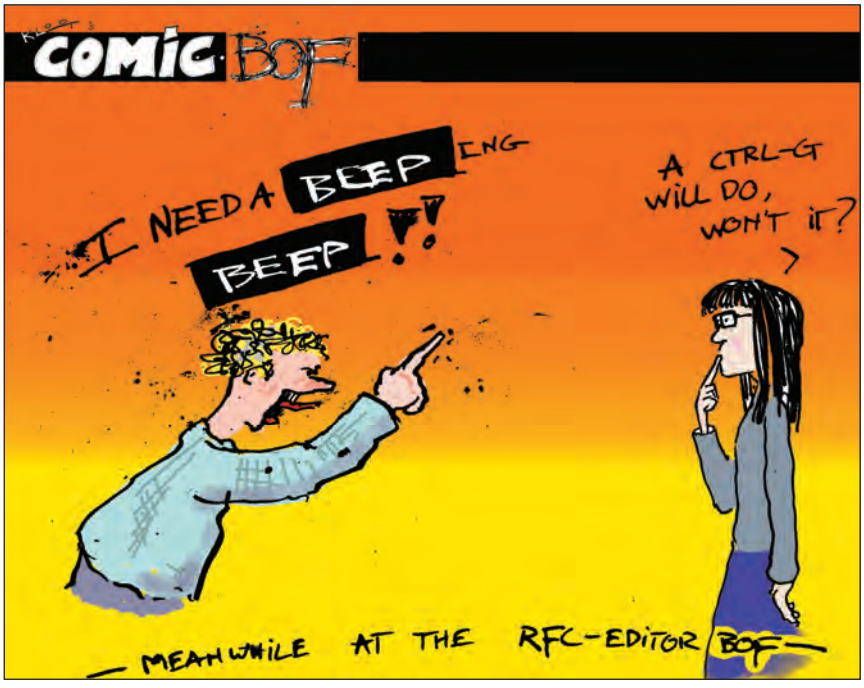

In order to ensure real and lasting impacts from the initiative, strict criteria were established for the three participant categories. For access networks, a commitment to make IPv6 part of their regular business was required. This meant new subscribers getting IPv6 on by default from 6 June 2012 and no special user configuration required to access the IPv6 Internet. To enable verification of progress towards this goal, access networks were also required to be delivering 1 percent of visits to major IPv6-enabled websites over IPv6 by 6 June. The level of IPv6 usage depends on the IPv6 capabilities and configuration of home networks, so the enabled customer base needed to be greater than 1 percent to reach this target. Measurements were made by the major website participants to verify meaningful participation.

The second category of participant was home router vendors. Again, the requirement was that IPv6 become part of regular business. This meant a majority of products shipping with IPv6 on by default, meaning no user configuration is required to use IPv6. The University of New Hampshire Interoperability Lab provided independent verification of IPv6 interoperability.

Finally, websites wishing to participate were required to enable IPv6 access to their main website *permanently*, meaning IPv6-enabled users can now access their content over IPv6 without any additional configuration.

Industry leaders in all three categories of participation made early commitments to the Launch and welcomed additional commitments from others.

For more information, see http://www.worldipv6launch.org/.

Exploring IPv6 Deployment in the Enterprise

Experiences of the IT Department of Futurewei Technologies

By Tina Tsou, Diego R. Lopez, John Jason Brzozowski, Ciprian Popoviciu, Charlie Perkins, and Dean Cheng

In most of today's enterprise environments, information technology (IT) has changed roles from that of a pure cost center to that of a strategic business enabler. With this transition has come an ever-growing focus on technologies measured by innovation, operational excellence, security, and business continuity. Except for in the case of service providers, IT is not a core competency of most enterprises. Nevertheless, new technology adoption is viewed as critical to an organization's competitiveness. Given the inevitable IPv6 transition, enterprise IT departments now find themselves at the forefront of IPv6 deployment, faced with specific challenges and opportunities.

Except in the case of service providers, IT is not a core competence of most enterprises. Nevertheless, new technology adoption by IT is viewed as critical to an organization's competitiveness.

The IT department of Futurewei Technologies, which may be considered the U.S. research and development (R&D) arm of Huawei Technologies, is a good example of a midsize enterprise committed to adopting IPv6 and learning from the transition process. Futurewei is located in Santa Clara, California, U.S.A., and occupies several office buildings with several hundred engineers and staff members. Futurewei manages and operates its own campus networks for daily business needs, and, as a registered U.S. company, operates relatively independently of its parent company, which is headquartered in China.

As part of an ambitious IPv6-enabled next generation campus infrastructure project, the IPv6 transition process began by acquiring dual-stack Internet access and services from Comcast. The IPv6 infrastructure is an overlay logically separated from the existing IPv4 network and is targeted to enable IPv6 transport and services and to effectively explore new concepts in deploying IPv6 in an enterprise environment.

Over the past two years, industry perception about the best practice for IPv6 deployment in the enterprise environment has changed significantly. The core-to-edge approach is viewed as very practical. Tunnelling has fallen out of grace with many IPv6 practitioners, who believe the operational headache is not worth the level of access tunnels provide. More recently, those who have already deployed IPv6 have begun to explore ways to eliminate IPv4 from their environments as soon as possible in order to avoid the costs of simultaneously managing two stacks. In this context, the team at Futurewei is focussed on leveraging recent IETF work to explore deployment options that would improve some of the current best practices.

Environment Setup

The environment architecture reflects the goals and design guidelines of the initiative:

1) enable all employees to have IPv6 access to Futurewei and

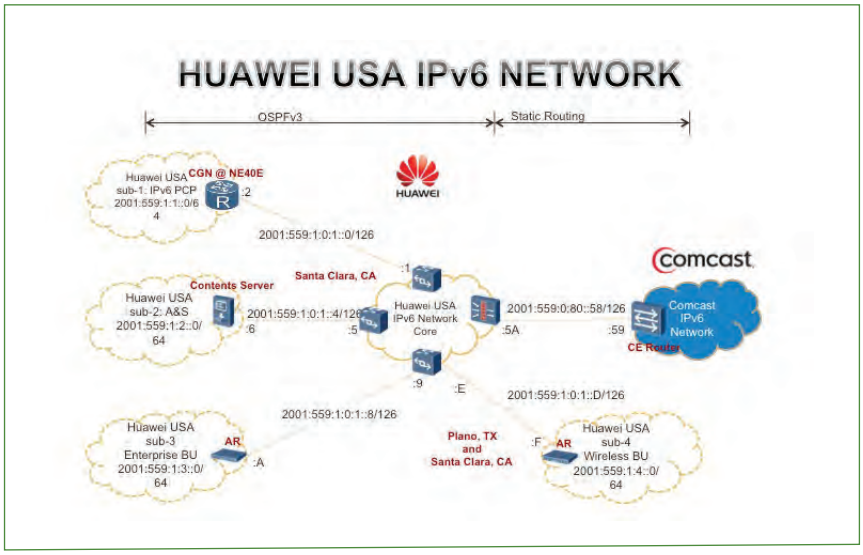

Figure 1: Futurewei IPv6 Environment Layout

Continued on next page

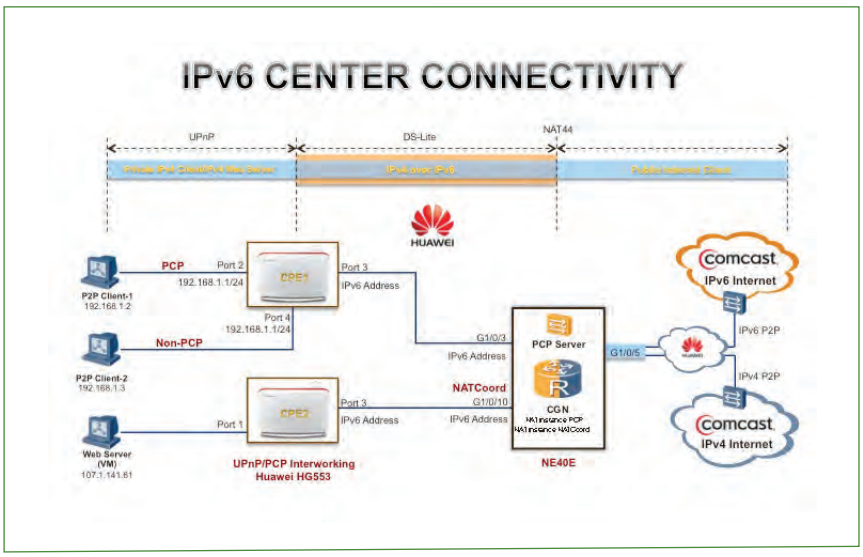

Figure 2: CGN Plus PCP Test Layout

Huawei Intranet- and Internet-based services;

2) explore practical IPv6 deployment and transition options for medium-sized enterprises;

The campus routing topology is organized as core-to-edge (hub-and-spoke), with emphasis on maintaining an IPv6-only stack for as much of the infrastructure as possible. Such an approach reduces operational costs more than a pervasive dual-stack infrastructure and accelerates the transition to the targeted IPv6-only architecture.

3) enable employees to innovate and collaborate with external partners;

4) and, enable Huawei product teams to test the implementation of their IPv6 innovations and standards.

The first steps were to acquire IPv6 address space and to enable IPv6 Internet access. The R&D center decided to use provider-assigned address space for this medium-sized organization, as it provides an incentive to explore multi-homing constraints and implications. A native IPv6 link was acquired from Comcast Business Services along with a /48 globally routable IPv6 allocation. The address space was distributed along topological and functional boundaries covering several participating teams and test environments (see figure 1).

The campus routing topology is organized as core-to-edge (hub-and-spoke), with emphasis on maintaining an IPv6-only stack for as much of the infrastructure as possible. Such an approach reduces operational costs more than a pervasive dual-stack infrastructure and accelerates the transition to the targeted IPv6-only architecture.

All elements supporting end-to-end services were updated to support IPv6. IPv6 has been enabled on all networking infrastructure elements (switches, routers, wifi gateways, service platforms, such as firewalls, and Integrative Directory Services, or IDS), data center resources (compute, storage, load-balancers), and infrastructure services (Dynamic Host Configuration Protocol—or DHCP—and DNS). Key business and productivity applications have also been IPv6-enabled, and employees are provided incentives to use them over IPv6: email, eSpace, file transfer, messenger, VoIP, and voice conferencing, to name a few. Employees are also provided with native IPv6 Internet access.

The project team is helping employees enable their devices—from laptops to iPads—to use the IPv6 infrastructure. The team conducts regular seminars to provide users with the awareness and knowledge necessary to conduct business in an IPv6 environment, and to stimulate and enable innovation.

The environment is instrumented to monitor utilization of the IPv6 infrastructure and services, and to evaluate user experience. Testing new deployment solutions for IPv6 goes beyond functionality; it collects statistics that enable quantitative evaluations. A virtual team consisting of corporate IT staff and IPv6 project resources operate the environment.

Exploring Carrier-Grade Network Address Translators (CGNAT) as an IPv6 Transition Mechanism in Enterprise Environments

In addition to dual-stack, there are other IPv6 transition mechanisms and technologies that have been or are currently being developed within the IETF, which may be required for

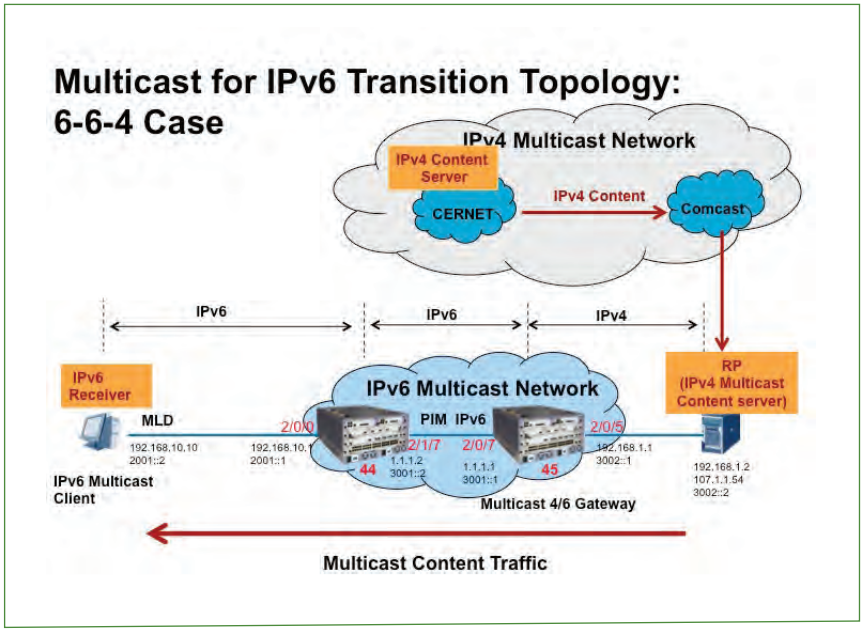

Figure 3: Multicast Test Infrastructure

After two years of development, the general problem statement for this work was recently accepted as a formal mboned working group draft. The primary focus of this problem statement, and the detailed work programme it lays out, is the distribution of broadcast content by service providers, but the results will be equally applicable to the distribution of such content in enterprise networks.

business in the near future. With that in mind, the IT engineers developed a practical plan to test running code implemented in compliance with relevant RFCs and IETF drafts, including NAT44, NAT64/DNS64, PCP, DS-Lite, 6rd, multicast transition, extensions to DHCP/DHCPv6, and RADIUS, among others. These tests have contributed to progress within the IETF. In fact, a significant amount of work on and preparation for the Port Control Protocol (PCP) demo at IETF 81 and the IP multicast transition demo at IETF 82 were accomplished by Futurewei IT engineers. In addition, the IT department plans to extend and apply these new technologies to the business. For example, the IPv4-based customer data centre used by the marketing department today will be accessible by IPv6 customers around the world via NAT64/DNS64 installed on the site. Video applications and eSpace will be supported by two or more IP multicast transition technologies to accommodate external IP access from customers, business partners, and Huawei offices in China, as well as other parts of the world.

The CGN plus PCP solution first demonstrated at IETF 81 was tested more extensively in our IPv6 lab environment (see figure 2).

In this transition scenario the core of the network is IPv6 only, enabling native forwarding for IPv6 end-to-end. At the same time, IPv4 access is provided over tunnels. The solution can be considered the opposite of 6PE, in which instead of a Multiprotocol Label Switching (MPLS) core with an IPv4 control plane where IPv6 traffic is tunnelled via MPLS, we have an IPv6 core where IPv4 is tunneled via DS-Lite. This solution can be very appealing to large enterprises having to mitigate a shortage of IPv4 (RFC 1918) addresses. At the same time it reduces the footprint of the infrastructure, no longer having to support two IP stacks per node with the inherent additional operational costs.

This solution was used to demonstrate ubiquitous accessibility of resources, such as FTP services, behind provider translating devices. Vodafone partnered with the Futurewei engineering team to run tests over the infrastructure services that could be enabled by its customers or by client applications that need to receive gratuitous flows.

Use Case: Content distribution over an IPv6-only core.

To support a useful IT environment on an IPv6-only core, the Futurewei IT engineers had to devise a solution for distributing content over multicast. In partnership with the Huawei product teams and associated universities, a translation solution for both control and forwarding planes has been developed and tested in the Futurewei IT environment (see figure 3).

This solution is related to an emerging body of work on multicast transition in the IETF to which Huawei has given strong support. After two years of development, the general problem statement for this work was recently accepted as a formal mboned working group (WG) draft. The primary focus of this problem statement, and the detailed work programme it lays out, is the distribution of broadcast content by service providers (such as Internet Protocol

Continued on next page

television, or IPTV), but the results will be equally applicable to the distribution of such content in enterprise networks. The basic issue is that content sources and/or receivers may still support only IPv4. Running the core in dual-stack mode and carrying the content in both IP versions is wasteful of bandwidth. The work in hand will specify the mechanisms required to support IPv6-only multicast distribution in the core network interfacing with IPv4 source networks and/or IPv4 networks or subnetworks at the receiving end.

The work programme needed to meet the objectives is laid out in the mboned document and consists of several parts, as described here.*

- A discussion in the mboned WG about how receivers can acquire the multicast group addresses (and unicast source addresses, in the case of Specific-Source Multicast, or SSM) of the channels in which they are interested in the IP version they can use. This problem is much simpler if the receivers are dual-stack.
- A specification of an IPv6 multicast address format that embeds IPv4 multicast addresses, allowing stateless translation between IPv4 and IPv6 addresses. This eliminates the need to provide coordinated translation tables at multiple nodes if the control or data plane crosses multiple IP version transition points. This is already a chartered mboned work item, but the details are still under discussion.
- A related draft on DHCP provisioning of the IPv6 multicast prefixes used in the stateless translation mechanism. This is a work item in the Softwires WG, supporting operation of multicast consistent with the DS-Lite transition mechanism (RFC 6333) that is being done in that WG.

On the control-plane, listeners use MLDv2, PIM for IPv6 is used for routing over the infrastructure, and PIM for IPv4 is used towards the IPv4-only content sources.

- A draft on translation between the Internet Group Management Protocol (IGMP) and Multilistener Discovery (MLD), needed in Customer Edge equipment (the "B4") in DS-Lite environments. This has been submitted to the Protocol Independent Multicast (PIM) WG.
- A draft on operation of PIM routers in dual-stack environments, also submitted to PIM.

**References for drafts related to this work follow this article. They still have the status of individual drafts at the time of writing.*

Deploy360: Helping You Deploy New Technologies Quickly and Efficiently

The IETF creates protocols based on open standards, but some are not widely known or deployed as quickly as we would like. To help bridge this gap between completed standards and real-world deployment, the Internet Society launched its Deploy360 Programme in January 2012.

Deploy360 collects and creates technical resources from industry experts and makes them available in a free and open environment so other organizations can deploy technologies like IPv6 and Domain Name System Security Extensions (DNSSEC) quickly and efficiently. Deploy360 provides case studies, videos, technical documents, tutorials, and more, and continuously spreads this information through social media, speaking engagements, and ION Conferences.

Have you already deployed IPv6 or DNSSEC?
We can work with you on a case study, tutorial, whitepaper, or other resource to help spread your knowledge to those following your lead with their own deployments.

Are you looking to deploy IPv6 or DNSSEC but you're not sure where to start?
Visit and explore http://www.internetsociety.org/deploy360 to see what we already have. If you are looking for something specific and it's not there, let us know by emailing us at deploy360@isoc.org.

The IETF community works hard to finalize standards to help the Internet run better. Now help the Internet Society Deploy360 Programme get real-world deployment information into the hands of the professionals who need to deploy those standards.

To get more involved, email us at deploy360@isoc.org or visit us at http://www.internetsociety.org/deploy360.

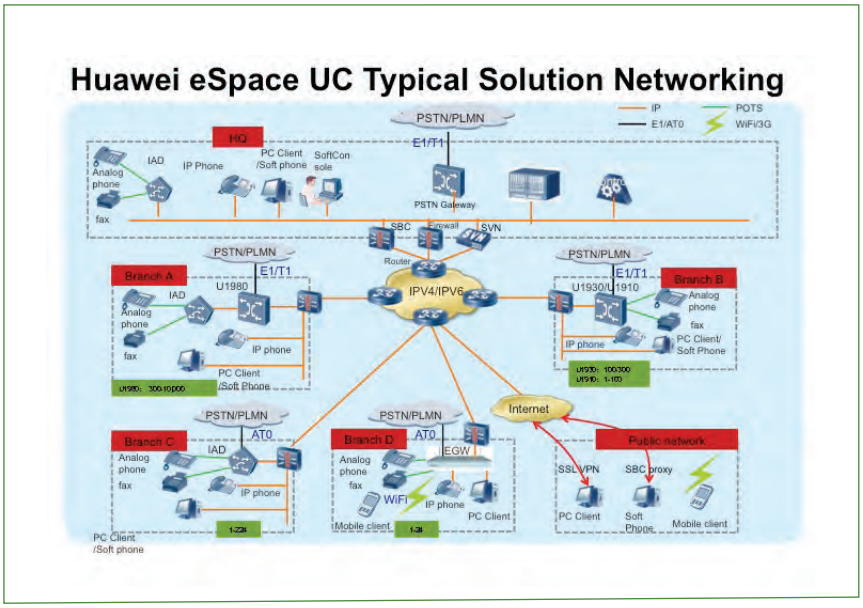

Figure 4: eSpace Unified Communication Target Architecture

The multicast solution described in combination with the CGN-based IPv6 integration option provides an innovative new approach to IPv6 transition that takes into account recent operational experience (in particular, the idea of operating as much of the infrastructure as possible in IPv6-only mode), as well as standardization developments.

The results of the tests referred to at the beginning of this section were demonstrated at IETF 82.

To ensure that as many services as possible are accessed over IPv6, multicast content is delivered to end users (i.e., multicast listeners) over IPv6. The traffic is transported over the IPv6 infrastructure while translation is used at the data centre edge to enable access to IPv4-only content sources. On the control-plane, listeners use MLDv2, PIM for IPv6 is used for routing over the infrastructure, and PIM for IPv4 is used towards the IPv4-only content sources.

The multicast solution described in combination with the CGN-based IPv6 integration option provides an innovative new approach to IPv6 transition that takes into account recent operational experience (in particular, the idea of operating as much of the infrastructure as possible in IPv6-only mode), as well as standardization developments.

IPv6 Exploration Continues: Future work at Futurewei

The IT department at the Futurewei R&D centre plans to expand upon the solutions evaluated thus far and apply them to larger scale deployments. For example, a project is ongoing to enable IPv6-only hosts to access the IPv4-only data center supporting corporate marketing applications. This solution will test NAT64/DNS64 technologies. Another Futurewei project will enable IPv6 operation for eSpace, Futurewei's unified collaboration platform that facilitates off-campus, off-corporate network collaboration (see figure 4).

The Futurewei team is sharing its lessons learned with the community at large to help both technologists and customers with their IPv6 transition planning. Many of the tests performed continue to support the work done in the IETF by applying new ideas in a real environment.

We thank the following people for their comments and support.

Comcast: Jason Livingood, Kent Porter, David Chai, Frank Rudnick, Yiu Lee, Tim Ruelas

Nephos6: Mo Khalid

Viagenie: Marc Blanchet

Vodafone: Loris Cardullo

SI6Networks: Fernando Gont

Huawei: Frank Zhong, James Huang, Yunshan Zhu, Frankie Zhang, Steve Anderson, Wendell Rios, Kenneth Durazzo, Susan Hares, Daniel Lo

References

draft-ietf-mboned-v4v6-mcast-ps, IPv4-IPv6 Multicast: Problem Statement and Use Cases.

draft-tsou-pcp-natcoord, Using PCP To Coordinate Between the CGN and Home Gateway Via Port Allocation.

draft-ietf-softwire-multicast-prefix-option, DHCPv6 Option for IPv4-Embedded Multicast and Unicast IPv6 Prefixes.

draft-tsou-mboned-multrans-addr-acquisition, Address Acquisition For Multicast Content When Source and Receiver Support Differing IP Versions.

draft-perreault-mboned-igmp-mld-translation, Internet Group Management Protocol (IGMP) / Multicast Listener Discovery (MLD)-Based Multicast Translation ("IGMP/MLD Translation").

draft-taylor-pim-v4v6-translation, A Translator For Protocol Independent Multicast (PIM) Interworking Between IPv4 and IPv6.

draft-lopez-v6ops-dc-ipv6-01, DC Migration to IPv6.

Problems in Low-Delay Internet Communication: Congestion Management

By Randell Jesup

The Internet has a problem: there's no agreed-upon reasonable standard for low-delay communication. In particular, congestion control for low-delay communication, typically running over User Datagram Protocol (UDP), is lacking. Real-time communication, especially interactive communication, is at best unreliable and currently depends on over-provisioning of any critical links, especially the last-mile hops (access links and/or local wireless links). Typically we are talking about voice-over-IP (VoIP), but all sorts of other communications are or will want to be real-time, including telerobotics, hardware control, and others.

When standard congestion-control mechanisms are used for communications that need low delay, they build up queues in routers and hosts, which may delay a data stream by seconds with little warning; this can be disruptive for real-time media, and can be fatal for low-delay control mechanisms.

There have been some experiments and work in lower-delay congestion control protocols, though none have directly or fully solved the problem. There have been experiments in the Transmission Control Protocol (TCP) space (such as TCP Vegas and Cx-TCP), in User Datagram Protocol (TFRC, DCCP, LEDBAT), and some in Stream Controlled Transmission Protocol (SCTP) (related to the Cx-TCP work).

Due to the increasing need for real-time communication, a number of people are trying to develop standards that will allow appropriate sharing of bandwidth and avoid congestion collapse. This is especially relevant since video and adaptive audio are now often part of media streams, and unlike "classic" fixed-rate VoIP protocols, they can adapt to changing network congestion.

The problem is simple: existing congestion control methods, in particular TCP, are loss-based and, in determining link bandwidth, must force the intermediate routers into a loss state (congestion). For a tail-drop router (the most common), this means maximum delay. If this is a small number of milliseconds (ms), this may not be a problem, but combined with even a hint of BufferBloat[1], the delays can quickly make real-time interaction impossible. Delays of 100, 200, 500ms, or even multiple seconds, are possible.

Even algorithms such as LEDBAT[2], which was designed as a 'scavenger' protocol to make use of 'extra' bandwidth while getting out of the way of user-initiated transfers, will engender a typical 100ms of delay in the bottleneck node, which is problematically high for many real-time applications. For example, the ITU recommends in G.114 (and elsewhere) that one-way-delay (mouth-to-ear) be kept below 150ms for best subjective audio quality[3], especially if echo isn't perfectly controlled. In just one bottleneck node, 100ms is a killer when added to the other delays in a VoIP call. This is especially problematic as LEDBAT flows might be assumed not to interfere with user-initiated VoIP calls; so applications might use them indiscriminately and without user knowledge (e.g., background updates or backups).

An added complication is the need to work correctly in an environment with effective Explicit Congestion Notification (ECN) or Active Queue Management (AQM), which is a focus of efforts to combat BufferBloat.

Current Efforts

There is an effort underway as a derivative of the rtcweb effort (part of the W3C/IETF joint WebRTC project) to develop and standardize congestion control protocols to deal with those problems as best as possible given the need to compete with TCP flows. With tail-drop routers, large TCP flows may always 'win' and force the buffers to expand, but there's little we can do about that unless or until active queue management (AQM) is the norm. About all we can do is mark real-time packets properly so routers have the option of handling them separately from TCP and other loss-based flows. (A number of access routers/modems do this now for many classic VoIP flows). It would be nice if a solution (or solutions) could be found for generic UDP flows, Real-Time Transfer Protocol (RTP) media flows over UDP (which carry some inherent timing information), and TCP/SCTP/etc. flows, though the initial focus is RTP flows.

When standard congestion control mechanisms are used for communications that need low delay, they build up queues in routers and hosts, which may delay a data stream by seconds with little warning; this can be disruptive for real-time media, and can be fatal for low-delay control mechanisms.

Other considerations will include interactions with the TCP slow-start algorithm, the impact of probing for bandwidth availability on other flows, and the advantages (or disadvantages) experienced by relatively late-arriving flows.

This effort is underway on the rtp-congestion list[4], and a birds-of-a-feather (BoF) meeting is proposed for IETF 84 in Vancouver, British Columbia, Canada, with the goal of chartering a working group (WG) to address this problem. There will also be a one-day IAB/IRTF Workshop on Congestion Control on July 28 in Vancouver. See the notice at http://www.iab.org/cc-workshop/.

There are several proposed ways to attack this problem, though more research and simulation is needed. One approach is from Google for a delay-sensing algorithm that infers the state of the bottleneck router from packet arrival delay deltas. Algorithms in this class are known to work, but their fairness (both with themselves and TCP) and ability to adapt to AQM have not been explored yet. Other options would be to use or leverage Cx-TCP or DCCP. LEDBAT in its current form is probably not an option, but a congestion-control algorithm based on the same principles may be a viable candidate.

One important consideration is to develop a coherent strategy for managing the different classes of flows on the Internet, with smooth-as-possible fallbacks when the preferred solution isn't available (such as when the bottleneck is a tail-drop router).

A Delay-Sensing Congestion-Control Algorithm

The following is based on draft-alvestrand-rtcweb-congestion[5], and is only a high-level description of the actual algorithm. This and other similar algorithms have been in use on the Internet in small amounts since at least 2004.

The basic idea is that we monitor the drift between when a packet was sent and when it was received (one-way-delay). We don't need to know the actual one-way-delay value (which can be challenging to measure), only changes up or down in the delay value.

An increasing delay (after filtering) implies that the bottleneck node's queues are increasing in depth. If the flows on the link are relatively static in bandwidth, such as a bandwidth-constrained access link largely occupied by the flow itself, then the signal from the filter will be very clear. More complex environments should show a signal, but the filter may take longer to converge.

Conversely, when the filter shows decreasing one-way-delay, then the assumption is that the bottleneck queues are draining.

This information can be used to estimate the amount of available bandwidth on the bottleneck link, and thereby whether the congestion control algorithm should allow the flow rate to increase, decrease, or remain the same. Other inputs are necessary, such as packet loss, Explicit Congestion Notification (ECN) markings, but these do not yet appear in this draft. Part of the research needed is to determine what the best response to such inputs is.

Unlike LEDBAT and Cx-TCP, there is no explicit nonzero queuing target; this algorithm attempts to use as much bandwidth as possible while keeping the queuing delay at or close to zero. This implies that it can't be 100 percent efficient, but in a relatively static situation it can come very close. Another research consideration is determining algorithm efficiency in different scenarios, especially nonstatic scenarios and scenarios with larger aggregations of delay-sensing flows on the bottleneck link.

There are many areas for useful research and innovation within this area. In addition to adapting existing algorithms to fulfill the need and comparing them, there is significant work to be done in improving these proposals.

Future Work

There are many areas for useful research and innovation in this area. In addition to adapting existing algorithms to fulfill the need and comparing them, there is significant work to be done in improving these proposals, such as adding loss and ECN support to the Google proposal, developing appropriate startup-time heuristics, methods to minimize the impact on the current and other flows when probing for additional bandwidth, avoiding "swings" in fairness that can cause a major loss of utility (such as for interactive video calls), and many others.

As mentioned, we hope to charter a WG in Vancouver, British Columbia, Canada, and work in it to develop one or more RFCs to address these issues and move existing proprietary and ad-hoc congestion methods into a standardized framework.

References

1. http://www.bufferbloat.net/
2. http://tools.ietf.org/wg/ledbat/
3. http://www.itu.int/rec/T-REC-G.114-200305-I
4. https://sites.google.com/a/alvestrand.com/rtp-congestion/
5. http://tools.ietf.org/html/draft-alvestrand-rtcweb-congestion

Shooting Around the Corner: The Problem of Real-time Services

By Jose Saldana, Dan Wing, Julián Fernández-Navajas, José Ruiz-Mas, Muthu A.M. Perumal, Gonzalo Camarillo, Michael Ramalho

A new draft has been submitted to the Transport Area working group (WG): draft-saldana-tsvwg-tcmtf. The primary goal of this draft is optimization of real-time flows independently from the use of real-time protocol (RTP) (RFC 3550: A Transport Protocol for Real-Time Applications, RFC 3551: RTP Profile for Audio and Video Conferences with Minimal Control). This draft proposes that a number of small packets be compressed, multiplexed, and bundled into one packet, with the resulting packet being forwarded using a tunnelling scheme. This proposal has demonstrated its ability to save bandwidth while concurrently reducing the packets' per-second rate. Preliminary tests have shown that significant bandwidth savings can be achieved while not introducing delays that could degrade the real-time user experience for gamers and the like.

It's 22:37 in San Jose, California. After a hard day of work, Jack is at home playing a game he has just bought: a first-person shooter game. He is exploring the scenarios and shooting virtually everything that appears in his screen. He spots a running enemy named "Wang" and he shoots. His screen shows Wang falling down.

It's 14:37 in Tokyo, Japan. Wang is spending his lunch break at school playing the same game. He is running through the same scenario and notices an enemy named "Jack" nearby. He quickly goes around a corner in order to avoid getting shot. But shortly after he has already turned the corner, a message that reads "you are dead" appears in the screen. He gets angry and tells his girlfriend: "Have you seen this? This player is cheating! He's shot me around the corner!"

The latency problem

What's the problem? What Jack, Wang, and Wang's girlfriend are probably ignoring is the real cause of the problem: network latency. The server where they are playing the game is in California, only 80 km from Jack's house, but 8,200 km from Wang's school. So the network latency for Jack is very small and mainly due to equipment, whereas Wang's packets must travel through a submarine cable along the bottom of the Pacific Ocean. Figure 1 illustrates how when Jack shoots, that information quickly arrives at the server, which calculates the result and decides that Wang is dead. But the packet telling Wang's computer this information takes 150 milliseconds to arrive, due to a number of routers, the access network, and even the speed of light based delay. By the time this information arrives at Wang's application, he has already hidden around the corner.

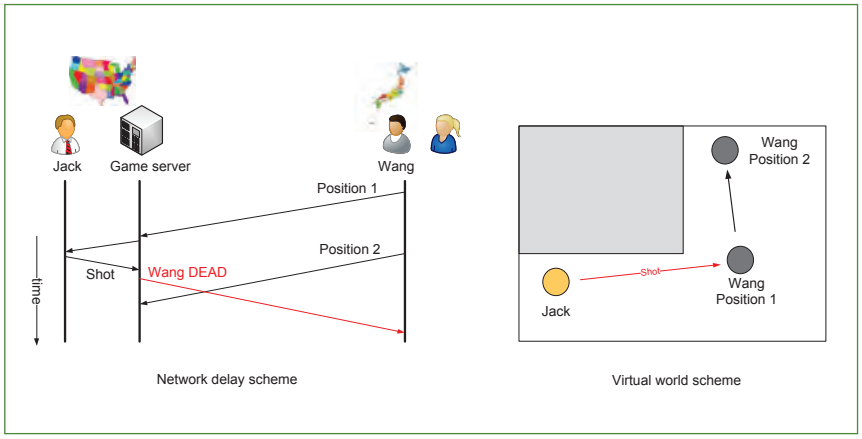

Figure 1. Schema of the latency problem for Wang

This is only one example of the Internet's limitations when dealing with real-time services. The same problem appears in VoIP (voice over IP), remote desktop solutions, video conferencing, and database access, among others.

The Internet was designed as a *best-effort* network, which does not warrant a maximum delay. This was enough for traditional delay-insensitive applications, such as web browsing, e-mail, or FTP. However, new real-time services are raising the question: Is the Internet adequate for them? Although many quality-of-service (QoS) mechanisms have been specified within the IETF, and used in many network scenarios (such as enterprise intranets), the Internet remains mostly *best-effort*-(non QoS)-based transport.

> New real-time services are raising the question: Is the Internet adequate for them?

The efficiency problem

Another interesting question is related to packet size, which usually ranges between 40—1,500 bytes. Since every packet must include the IP and the TCP or UDP (User Datagram Protocol) headers, we have a simple rule: the bigger the packet, the better the efficiency. Traditional services tend to

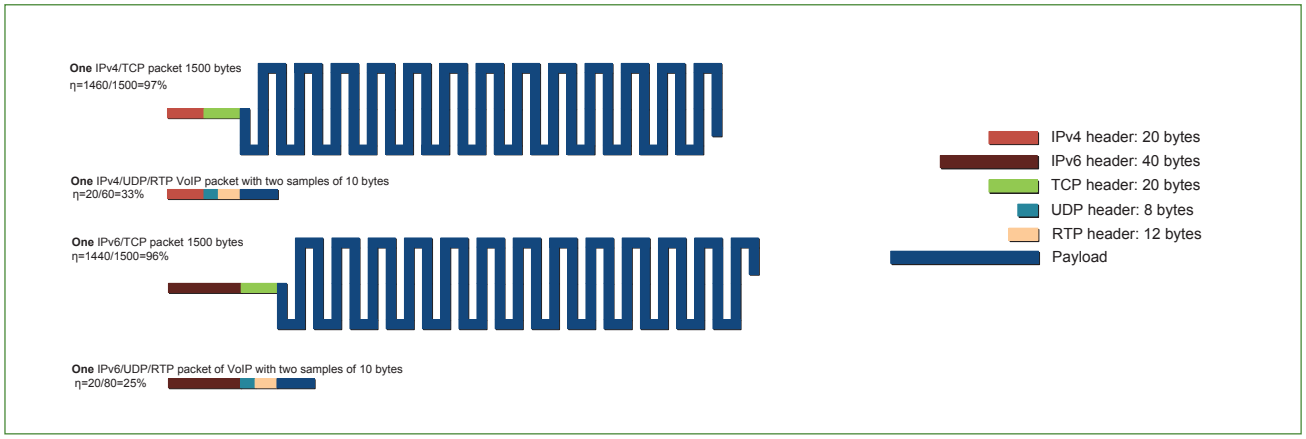

Figure 2. The efficiency problem. *Note: Packet sizes are to scale.*

maximize packet size; since delay is not critical, they can wait until they have almost 1,500 bytes to send. But the problem with real-time services is that the information has to arrive quickly, sometimes with many applications, and sometimes with a fixed cadence. As a consequence, tiny information chunks are sent using a high frequency, which implies very small packets. Figure 2 illustrates how a packet of 1,500 bytes has an efficiency of 97 percent, whereas a VoIP packet carrying two voice samples of 10 bytes each has its efficiency reduced to 33 percent; in other words, only one out of three bytes is voice information while the others are headers. If IPv6 is used, the problem becomes even worse—20 more bytes are included in the IPv6 header and VoIP small packets' efficiency is reduced to 25 percent. Logically, the problem is almost negligible for big packets.

There is another issue that real-time services have highlighted: The packets per-second limit. Traditional network-equipment-processing capacity was designed to manage a mix of big and small packets (referred to as Internet Mix or IMIX). If the average packet size diminishes dramatically, and the number of packets per second is high, the route lookup function may become a bottleneck in addition to link speed.

Description of some real-time services

In this section we will focus primarily on describing VoIP and online games, as they are the services for which tests have been deployed. Other services, such as remote desktop or video conferencing, could also be considered.

VoIP

VoIP was one of the first widely deployed real-time services. It uses Real-time Transport Protocol (RTP), the IETF protocol designed to deliver real-time content, which uses 12 bytes for protocol header and works over UDP, so the total overhead is 40 bytes for IPv4 and 60 bytes for IPv6. One RTP packet includes a small number of voice samples, which use different voice codecs, some of them having a high robustness to packet loss. Since interactivity is high, retransmission of lost packets is not deployed; the destination application has to deal with packet loss, trying to conceal it. The maximum recommended round-trip delay is around 150 milliseconds.

Online games and virtual environments

Online games are a real-time service in which avoiding delay is critical. Some research has shown that the maximum round-trip delay tolerated by gamers is about 150 to 200 milliseconds. In addition, gamers are extremely fickle: if a game does not work properly, they may leave the server and never return. Games use client-server architectures for many reasons: the convenience of maintaining the consistency of the game, synchronization, cheating prevention, and recording high scores. But the main reasons are commercial: The providers can charge for use or sell the server software.

Gamers are extremely fickle: if a game does not work properly, they may leave the server and never return.

Different game genres can be distinguished:

- *First Person Shooters (FPS):* A virtual scenario shared by some tens of players is created. Every player controls an avatar that has to accomplish a mission or kill all the enemies. The weapon can be improved according to the score. The action is fast, and the aim of the player is crucial. Although the sessions may last up to two hours, the game is divided into short rounds, which may be of some minutes. The client application generates high rates of tiny UDP (non-

Continued on next page

Real-time Transport Protocol) packets, which go to the server. When the new game state has been calculated, it is sent to every player. Depending on the number of surviving players, these packets can be bigger.

- ***Massively Multiplayer Online Role-Playing Games (MMORPGs):*** This genre has become very popular, especially in Asia. Some titles have achieved more than 10 million subscribers (e.g., *World of Warcraft*). They are normally set in magical or historical environments and the player controls an avatar with a long-term life. He obtains better abilities, weapons, and curses by means of missions. Trading is also permitted. The player can connect to different servers, called *shards,* which are shared by some thousands of gamers. The action is less interactive than in FPSs. For example, in the case of MMORPGs, for purposes of fighting, first you select the enemy with your mouse, and then you may choose the weapon or curse to use against them. Nevertheless, these games can still be considered a real-time service: The speed and the ability of the player are decisive in determining who wins and who loses. The vast majority of these games use small TCP packets, which is surprising, since they are "real-time services" using TCP.
- ***Real Time Strategy (RTS):*** In these games, the player creates a city or a civilization, manages the resources, and amasses armies to attack other players and dominate the virtual world. The real-time requirements are looser.
- ***Sports games:*** There are a great variety of these types of games. The differences between their network behaviours are big, making it very difficult to establish general rules.
- ***Some other virtual environments,*** such as *Second Life,* also have become popular in recent years. A single server is usually shared by a large number of users.

Although games are real-time services, they do not use RTP to deliver information to the server. Instead, they use bare UDP or sometimes TCP. Also, unlike VoIP, they do not present a fixed-packet cadence; in other words, inter-packet time may vary depending on a player's actions and changes to the game's environment.

Traffic optimization

Now that we have recognized the two problems, and the kind of traffic generated by real-time services, the question becomes, *can we optimize the traffic in order to improve packet efficiency as well as to reduce the packets per second?*

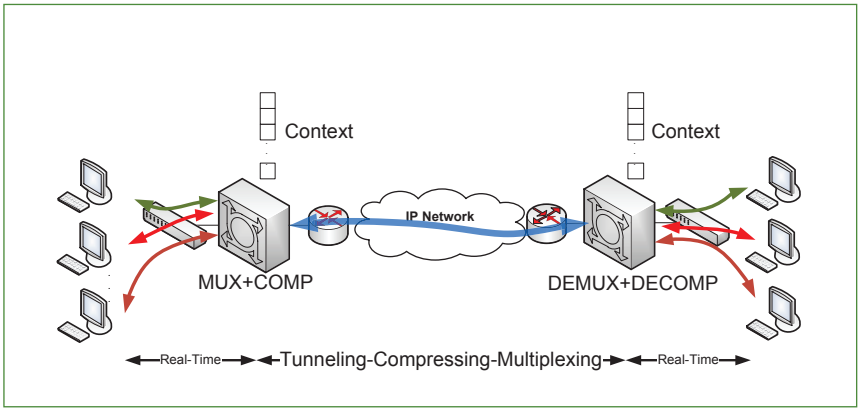

Figure 3. Tunnelling, compressing, and multiplexing traffic flows schema

The idea of Tunnelling Compressed Multiplexed Traffic Flows (TCMTF)

Multiplexing a number of payloads into a single packet can be seen as a solution for improving network efficiency. If only one flow is present, the number of samples included in a packet can be increased, but at the cost of adding new packetization delays. However, if a number of flows share the same path from an origin to a common destination, then a multiplexer can build a packet in which a number of payloads share a common header. A demultiplexer will then be necessary at the end of the common path in order to rebuild the packets as they were originally sent, thus making multiplexing a transparent process for the two hosts that are exchanging data.

The headers of the original packets can also be compressed in order to save more bandwidth, using one of the header compression schemes defined by the IETF. The headers rely on the fact that many header fields are the same for every packet in a flow. Additionally, they use *delta* compression in order to reduce the number of bits required by a field; they only transmit the difference between the value of a field in a packet and in the previous one. Bandwidth is saved, but at the cost of the need for storing a series of values, the *context*, in the compressor and decompressor These store the header fields that are the same for every packet of the flow, and can thus be avoided. There is another counterpart: The possibility of context desynchronization, which would imply bursts of corrupted packets. Nevertheless, modern header compression schemes, such as Robust Header Compression (ROHC), are indeed robust against this.

Logically, a tunnelling scheme will be necessary in order to send the bundle via the public Internet. So we have a global scheme, including tunnelling, header Compressing and Multiplexing Traffic Flows (TCMTF), as shown in figure 3.

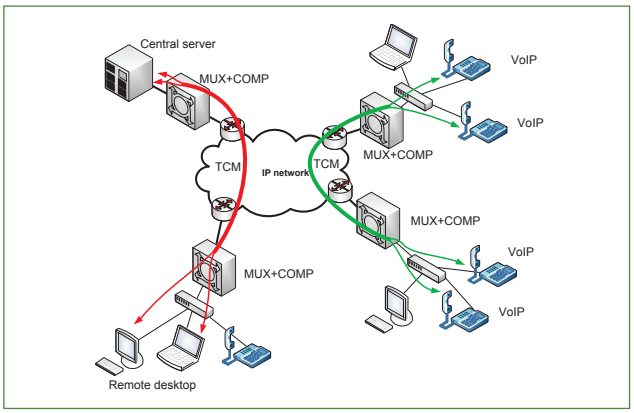

Figure 4a. Scenario in which optimization can be deployed.

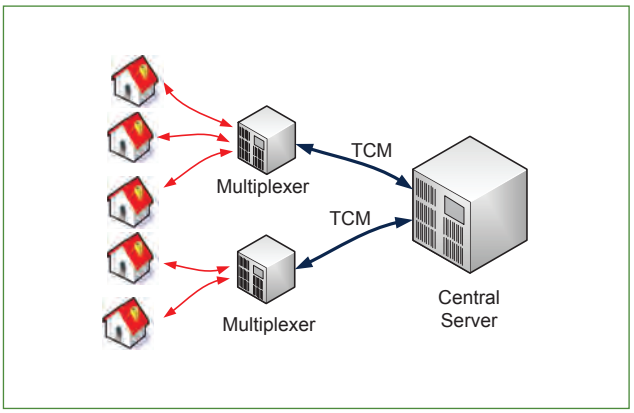

Figure 4b. Scenario in which optimization can be deployed.

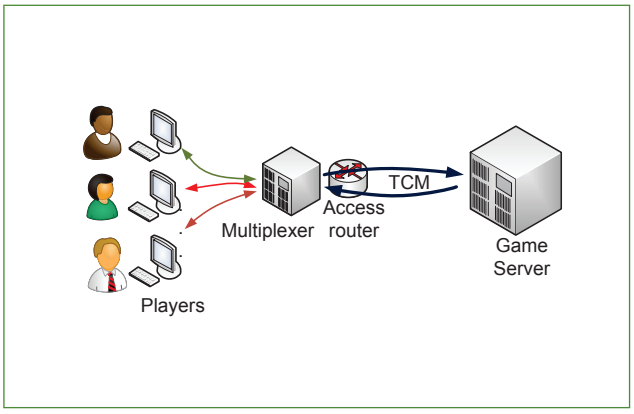

Figure 4c. Scenario in which optimization can be deployed.

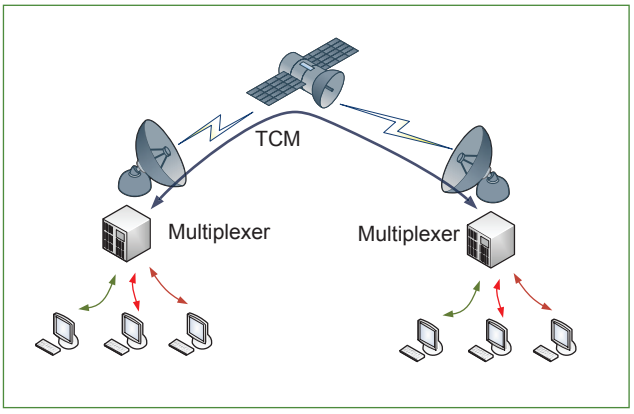

Figure 4d. Scenario in which optimization can be deployed.

Internet cafés, a very popular way to use the Internet in developing countries, are also places where people may simultaneously play the same game using the same server.

Scenarios in which optimization can be deployed

We can find different network scenarios in which the traffic of a number of flows shares the same path. Figure 4a illustrates an enterprise with a number of offices, in which tunnels can be established for merging VoIP flows (green lines), and also for multiplexing the traffic of a remote desktop application (red lines). For example, the traffic of the users of a game in a town or a district can be multiplexed by the ISP and sent to the central game server (figure 4b). Internet cafés, a very popular way to use the Internet in developing countries, are also places where people may simultaneously play the same game using the same server. So a tunnel could also be established from the router, or even from the computer of one of the players, to the game server (figure 4c). Another scenario is a satellite link (figure 4d), which may manage the bandwidth by limiting the transmission rate, measured in packets per second to and from the satellite. If small packets are used, this will result in poor utilization of the bandwidth, establishing an upper bound for the number of calls that can utilize the link simultaneously. Multiplexing small packets into a bigger one would improve the efficiency. This may be especially interesting for compressing TCP acknowledgements.

Current status and new proposal

The IETF has already tackled this problem for RTP, first with cRTP in 1999 (RFC 2508: Compressing IP/UDP/RTP Headers for Low-Speed Serial Links), which compressed headers across links; then enhanced to work over networks with loss or reordering with ECRTP in 2003 (RFC 3545: Enhanced Compressed RTP (CRTP) for Links with High Delay, Packet Loss and Reordering); and finally with RFC 4170: Tunnelling Multiplexed Compressed RTP (TCRTP) in 2005, which had the main aim of improving the efficiency of multiple RTP streams across a network. This is useful in the scenarios in which

Continued on next page

Shooting Around the Corner: The Problem with Real-time Services, continued

many VoIP conversations share the same path: we can do "voice trunking" between two offices of an enterprise, or group a number of conversations of a network provider. TCRTP, approved as a "best current practice," merged three layers (figure 5): First, RTP/UDP/IP headers were compressed using ECRTP; next, a number of header-compressed packets were multiplexed with PPP Multiplexing (PPPMux). Finally, the bundle was sent using an L2TP tunnel (RFC 2661: Layer Two Tunnelling Protocol).

But many things have happened since 2005:

- The outbreak of wireless access networks, which enable people to access the Internet from almost anywhere. These wireless scenarios are prone to packet loss, and may add bigger delays than the ones we can find in wired environments.
- The approval of ROHC (RObust Header Compression) in 2007 as RFC 4995: The RObust Header Compression (ROHC) Framework, updated by RFC 5795: The RObust Header Compression (ROHC) Framework in 2010. This header-compression standard was specifically designed for links with high loss and high round-trip times. It not only compresses RTP, but also IP and UDP. In addition, ROHC for TCP was defined by RFC 4996: RObust Header Compression (ROHC): A Profile for TCP/IP (ROHC-TCP)
- The approval of RFC 5856: The Integration of Robust Header Compression over IPsec Security Associations in 2010 as a framework for integrating ROHC over IPsec (ROHCoIPsec), which targets the application of ROHC to tunnel mode Security Associations (SAs). It reduces the protocol overhead associated with packets traversing between IPsec SA endpoints. This is achieved by compressing the transport layer header and inner IP

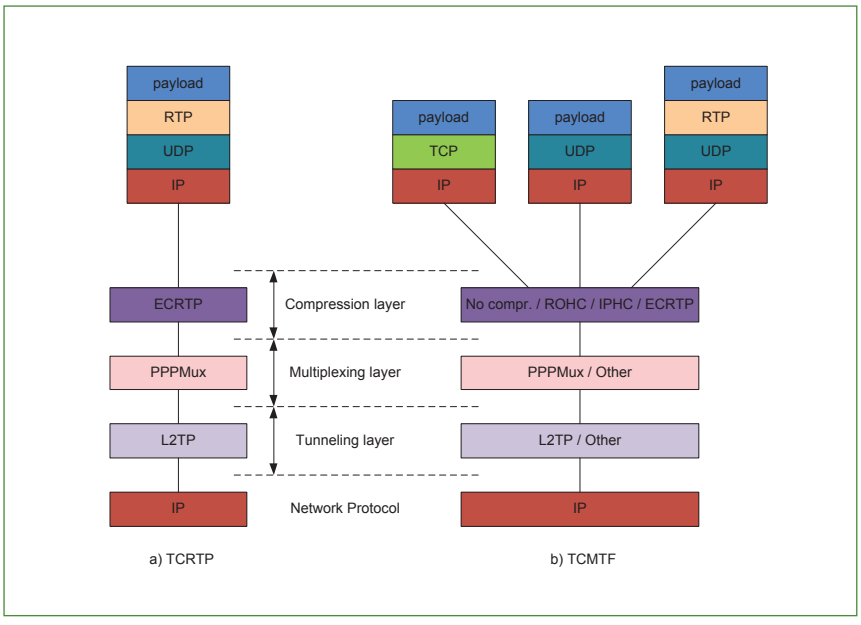

Figure 5. Protocol stack of TCRTP (left) and protocol stack of TCMTF (right).

Newcomer Experience: Arranging a Last-Minute Online Gaming Tutorial

By Jose Saldana, University of Zaragoza

IETF 83 in Paris, France, was my first IETF meeting. I enjoyed the Newcomers' Orientation, where I met many interesting people, and the meeting mailing list, where people asked all sorts of questions from running routes to directions to the Louvre. The variety of these discussions and the wide range of people I met gave me an idea: I offered to host an informal tutorial to discuss the kinds of network traffic generated by online gaming.

I was amazed by the positive response. I worked with the IETF Secretariat to book a room for my tutorial on Tuesday morning just after my presentation to the Transport Area working group (WG). About 25 people attended the 50-minute tutorial that I arranged via email list and coordinated with the Secretariat in only about three hours.

I demonstrated three games and presented some of the traffic optimizations that we've studied and are now trying to standardize, including tunnelling, header compression, and multiplexing. We used Wireshark to capture the traffic and saw the different options that game developers use for each genre: UDP for first-person shooter and real-time strategy games, and TCP for massive-multiplayer, online, role-playing games.

I learned from the audience's questions and comments, and several more requests for tutorials rolled in over the rest of the week. All in all, I was surprised by the speed with which things happen at IETF meetings, and I left the meeting very glad that I decided to stay the whole week.

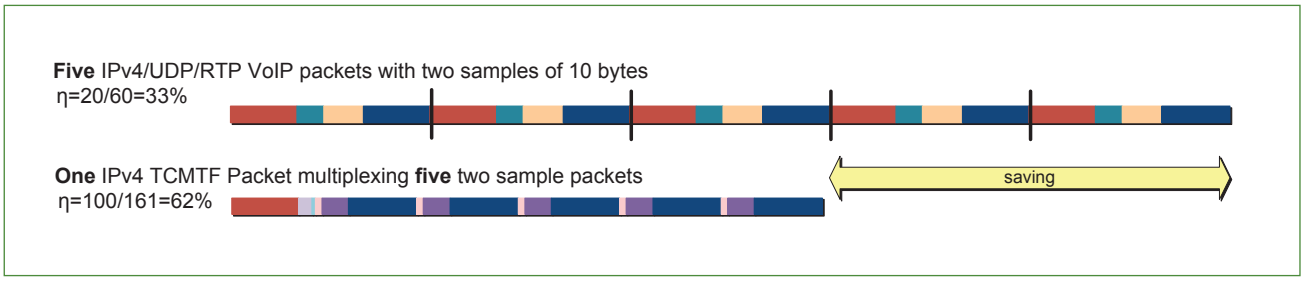

Figure 6a. Header compression results: VoIP

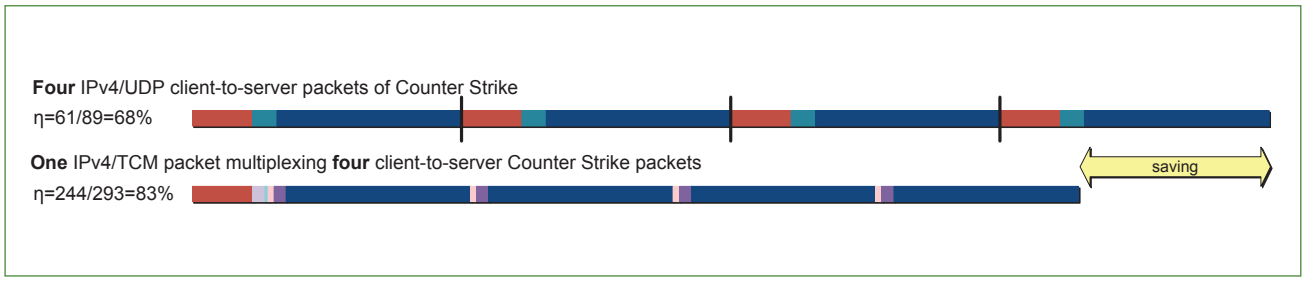

Figure 6b. Header compression results: FPS

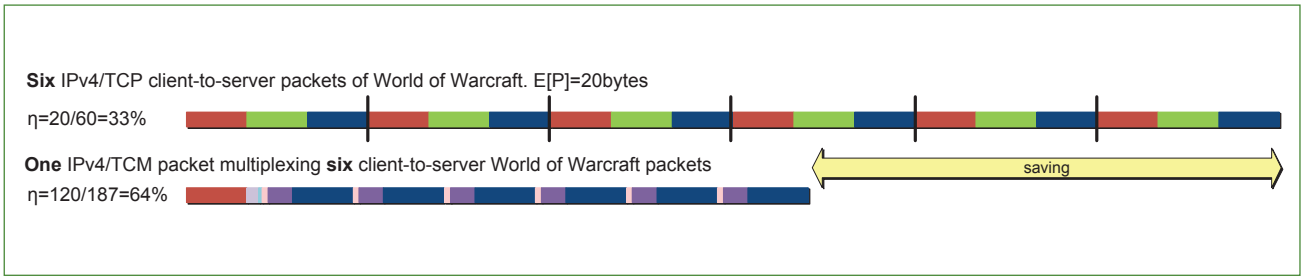

Figure 6c. Header compression results: MMORPG

header of packets at the ingress of the IPsec tunnel, and decompressing these headers at the egress.

- The popularity of many real-time services—online games have become very popular applications. As we have seen, they do not use RTP protocol.

As a consequence, we have considered doing the same thing as TCRTP, but also in cases when real-time flows do not use RTP. We can use ROHC to compress their headers, and do something similar in order to include a number of packets into a PPP Multiplexing (PPPMux) bundle.

Figure 5 illustrates this new proposal, which includes the three layers, but also considers more options: Different traffic types can be used and compressed—TCP, UDP, and also RTP—in the same way it was compressed by TCRTP. The compressing protocol will have to be selected depending on many factors: the scenario, the availability of processing and memory resources, etc. In addition, a null header compression is considered, taking into account that in some cases there may be many context synchronization problems. With respect to multiplexing and tunnelling, other options different from PPPMux and L2TP may also be considered.

Finally, as mentioned previously, in some services, interpacket time is not fixed. So we must define a policy in order to determine which packets are multiplexed in each bundle. We can do that either by defining a fixed number of packets or defining a maximum packet size. Another option is to define a period or a timeout, which may be more adequate for setting an upper bound for the added delay.

Preliminary tests

Figures 6a-c illustrate the savings that could be achieved by TCMTF. *(Note: The colors of the headers correspond to the layers in figure 5. Headers and payloads are to scale.)* Figures 7a-c show the bandwidth savings that have been obtained for the same services, by means of simulations based on ECRTP (Enhanced Compressed RTP) or IP header compression (IPHC) over PPP, Layer 2 Tunnelling Protocol, Version 3 (L2TPv3) and PPPMux.

Continued on next page

In Figures 7a-c we can see the bandwidth saving, measured as the quotient of TCMTF bandwidth divided by the native one. In figure 7a it can be seen that significant savings can be achieved when multiplexing different numbers of G.729a voice flows, depending on the number of flows and on the number of samples per packet (1, 2, or 3 samples, which means 10, 20, or 30 bytes of payload). The savings present an asymptote, which implies that when the number of multiplexed flows is high, the difference will be small in terms of bandwidth.

Tests have also been deployed for an FPS (*Counter Strike*), which sends UDP packets (figure 7b). The graph shows the bandwidth savings depending on the period and the number of players. If the period is small, the added delays can be kept in the order of 10 or 20 milliseconds, in order not to annoy the players. It must be taken into account that the average added delay is half the period.

Finally, figure 7c shows an example of the gains achieved for an MMORPG—the bandwidth savings are higher than the ones obtained for the FPS, however the number of players and the multiplexing period must be higher. This is not a problem, since the interactivity of these games is not as critical as in FPSs.

Conclusion

Summing up, the TCMTF proposal is able to mitigate the efficiency problem by sharing a common header across multiple payloads. Additional delays will be incurred, but they are small enough that they will not harm subjective quality. As we have seen, being able to both optimize RTP flows and bare UDP or even TCP can save bandwidth and reduce the number of packets per second generated—compelling advantages in the scenarios we've illustrated here.

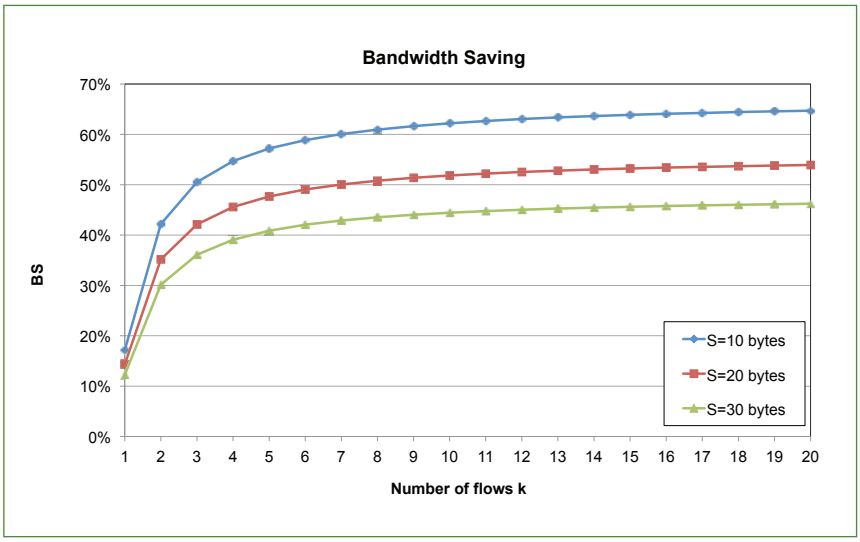

Figure 7a. Bandwidth savings for VoIP: G.729a codec with two samples per packet.

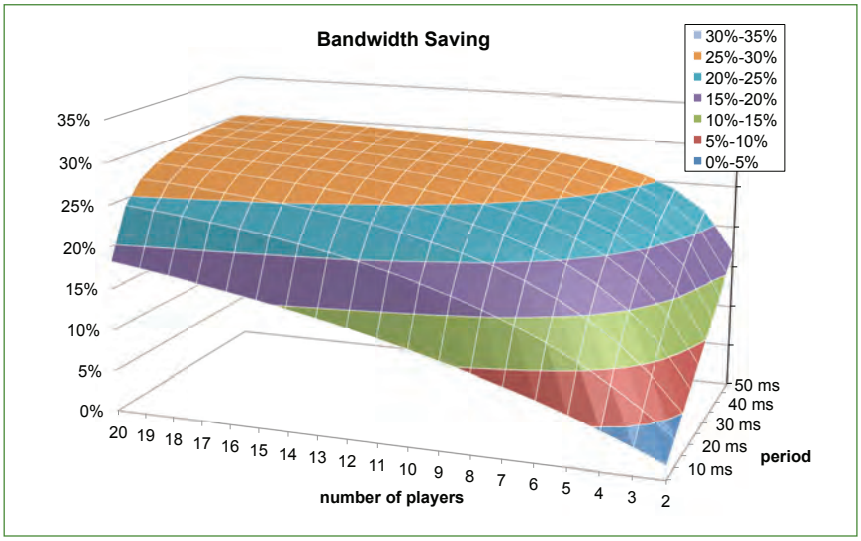

Figure 7b. Bandwidth savings for FPS: *Counter Strike*

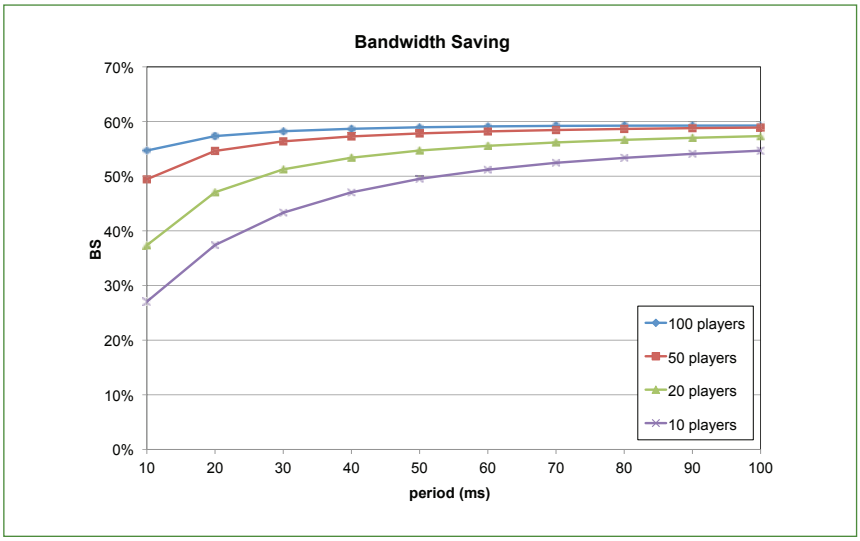

Figure 7c. Bandwidth savings for MMORPG: *World of Warcraft*

IRTF Update

By Lars Eggert

During IETF 83 in Paris, France, six Internet Research Task Force (IRTF) research groups (RGs) held meetings:

- NCRG—Network Complexity
- MOBOPTS—IP Mobility Optimizations
- DTNRG—Delay-Tolerant Networking
- ICCRG—Internet Congestion Control
- SAMRG—Scalable Adaptive Multicast
- CFRG—Crypto Forum

In addition, the Network Management research group (NMRG) met at a one-day workshop on the "Usage of Netflow/IPFIX in Network Management" on the Saturday following the IETF meeting.

Since IETF 82, the Network Complexity research group (NCRG) was chartered. Its first meeting was held in Paris during IETF 83. The NCRG aims to define and analyze the complexity of IP-based networks. Areas of interest include defining "network complexity" and relevant metrics, methods, and ideas to contain, control, or reduce complexity in IP-based networks, as well as collecting use cases regarding specific network designs or failure cases where complexity played a role. The group's charter is available at http://irtf.org/ncrg.

Also, the Network Virtualization (VNRG) research group has decided to close, due to lack of energy and participation.

On the IRTF RFC Stream, three new RFCs were published since IETF 82, one from the Anti-Spam research group (ASRG), and two from the Host Identity Protocol research group (HIPRG).

No IRTF open meeting was held at IETF 83. The IRTF open meeting is the venue where awardees of the Applied Networking Research Prize (ANRP) deliver their invited talks. Because no ANRPs were awarded for IETF 83, no award talks needed to be scheduled. As announced in the previous issue of this column, the ANRP selection committee has decided to move from a nomination/selection cycle per IETF meeting to a yearly nomination/selection cycle covering all the year's IETF meetings. Consequently, 2012 is a transition year for the ANRP, and one effect of the transition was the decision to skip the award for IETF 83.

ANRPs *will* be awarded for IETF 84 in Vancouver, Canada, and IETF 85 in Atlanta, Georgia, U.S.A. The combined nomination/selection cycle for these two IETF meetings in 2012 has begun!

The ANRP is awarded for recent results in applied networking research that are relevant for transitioning into shipping Internet products and related standardization efforts. It is supported by the Internet Society in coordination with the IRTF. See http://irtf.org/anrp for details on the award and instructions for nominating researchers for the prize.

Finally, the IRTF has adopted guidelines for handling intellectual property rights in contributions. In a nutshell, the IRTF follows the IETF Intellectual Property Rights (IPR) disclosure rules. Please see http://irtf.org/ipr. In summary, the IPR rules for IRTF participants are:

- If you include your own or your employer's IPR in a contribution to an IRTF research group, then you must file an IPR disclosure with the IETF.
- If you recognize your own or your employer's IPR in someone else's contribution, and you are participating in the discussions in the research group relating to that contribution, then you must file an IPR disclosure with the IETF. Even if you are not participating in the discussion, the IRTF still requests that you file an IPR disclosure with the IETF.
- Finally, the IRTF requests that you file an IPR disclosure with the IETF if you recognize IPR owned by others in any IRTF contribution.

In closing, please join the IRTF discussion list (http://www.irtf.org/mailman/listinfo/irtf-discuss) to stay informed about these and other happenings.

Pilot IETF Fellowship Programme for Regulators Sets Solid Groundwork for Future

By Markus Kummer

As part of a pilot IETF Fellowship Programme for Regulators—a collaboration between the Internet Society's Public Policy and Internet Leadership departments—five representatives from telecommunication regulatory authorities in Cameroon, Côte d'Ivoire, and Senegal were invited to IETF 83 in Paris, France, where they had the opportunity to learn about the IETF and to experience first hand its unique model of standards development.

The programme was a great success for all parties involved. It enabled the participants, who are already experienced in standards development processes, to exchange views with IETF leadership and the Internet Society Board of Trustees, and to network with experts from the Internet technical community in the spirit of collaboration. Their main interests included finding concrete solutions to reduce the price of Internet connectivity in their regions and expanding network access in rural areas. In addition, the programme enabled the IETF leadership to gain a better understanding of the concerns and priorities in developing countries.

The Internet Society plans to build on this pilot by granting fellowships to participants from regulatory authorities, enabling them to attend future IETF meetings. To start, fellowships to IETF 84 in Vancouver, British Columbia, Canada, will be provided to engineers from the Africa, Asia-Pacific, and Latin American regions.

[The programme] enabled the beneficiaries, who are already experienced in standards development processes, to exchange views with IETF leadership and the Internet Society Board of Trustees and to network with experts from the Internet technical community in the spirit of collaboration ... [It] enabled the IETF leadership to gain a better understanding of the concerns and priorities in developing countries.

From left to right: Lambert Bouo, Agence des Télécommunications de Côte d'Ivoire; Sally Wentworth, Internet Society; Josephine Adou, Agence des Télécommunications de Côte d'Ivoire; Markus Kummer, Internet Society; Aline N'Dakon, Agence des Télécommunications de Côte d'Ivoire; Pauline Tsafak, Ministère des Postes et Télécoms du Cameroun; Seyni Faty, Agence de Régulation des Télécommunications et des Postes, Senegal; Nicolas Seidler, Internet Society.

Photo: Richard Stonehouse/Internet Society

ISOC Fellows Make Ongoing Contributions to IETF

By Carolyn Duffy Marsan

Internet Society fellows are giving back to the IETF community in numerous ways, from acting as scribes in working group (WG) meetings, to submitting Internet drafts and engaging in meaningful online dialogue–all long after their financial support has ended.

The ISOC Fellowship to the IETF Programme provides grants to network engineers from emerging economies to pay for their meeting fees, travel, and other expenses so they can attend IETF meetings. The goal of the programme is to increase the diversity of input to the IETF and to increase global awareness of the IETF.

Increasingly, IETF leaders and WG chairs are seeing a major return on ISOC's investment in terms of former ISOC fellows contributing their time and effort towards standards development work.

"We are asking returning IETF fellows to scribe at least two meetings," said Steve Conte, senior manager of Internet leadership at ISOC. "They are contributing to the overall notes for that day, but it also ensures that they get a lot of out of the meeting by being active and engaged participants."

Conte said ISOC hopes to see its former fellows make long-term contributions to IETF WGs.

"This is another benefit of ISOC's Fellowship to the IETF programme," Conte said. "It is helping to foster ongoing engagement in the IETF by folks from emerging and developing countries."

IETF 83 Fellow Sandra Cespedes

For example, Fernando Gont is an ISOC fellow who is making significant contributions to the TCP Maintenance and Minor Extensions (TCPM) WG. Gont is a network engineer from Argentina who has worked on projects for the U.K.'s National Infrastructure Security Co-ordination Centre (NISCC) and the U.K. Centre for the Protection of National Infrastructure (CPNI) where he has reviewed the IP and TCP protocol specifications from a security perspective.

Michael Scharf, cochair of the TCPM WG and an engineer with Alcatel-Lucent, said Gont has presented this work at three TCPM meetings. "Fernando's work is very important for improving TCP specifications," Scharf said. "Fernando authored a quite significant share of the recent

Continued on next page

IETF 83 Fellow Paventhan Arumugam

Increasingly, IETF leaders and WG chairs are seeing a major return on ISOC's investment in terms of former ISOC fellows contributing their time and effort towards standards development work.

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

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

Photos: Richard Stonehouse/Internet Society

ISOC Fellows Make Ongoing Contributions to IETF, continued

RFCs published through TCPM. He is one of only a few TCPM working group contributors who checks the consistency between the TCP standards and what is actually implemented in real-world protocol stacks, in particular regarding security mechanisms. This is important for the IETF since some security mechanisms are not well documented in the RFC series."

Scharf said Gont has contributed up to seven RFCs through TCPM, all related to TCP security, as well as two individual drafts that fall in the scope of the TCPM WG.

"The ISOC Fellowship to the IETF Programme helped a lot to bring his work to the IETF and to move it forward in the standardization process," Scharf says. "Fernando is an important contributor to the TCPM working group….As far as I can tell, this fellowship programme significantly supported the recent work performed by the TCPM working group."

Vinayak Hegde, an ISOC fellow who is a lead architect at Indian mobile advertising firm Inmobi, was an ISOC fellow at three meetings. He regularly contributes to the IP Performance Metrics (IPPM), Content Delivery Networks Interconnection (CDNI) and Benchmarking Methodology (BMWG) WGs.

"I am also part of the recently formed Performance Metrics for Other Layers (PMOL) Directorate and help review drafts which are related to performance measurements from other WGs," Hegde said. "I also have contributed one draft on HTTP latency measurement. In various IETF meetings, I have scribed meetings of IPPM [IP Performance Metrics], PMOL, CDNi [Content Delivery Networks Interconnection], and JOSE [Javascript Object Signing and Encryption] WGs…I have also contributed to the IETF by talking about the IETF standardization process at SANOG," the South Asian Network Operators Group."

Hegde said that ISOC's fellowship programme has allowed him to connect with many Internet engineering experts from around the world.

IETF 83 Fellow Brenda Nyangweso

"The conversations with people who come to IETF [meetings] have given me new ideas, insights and problem-solving techniques in my regular work," Hegde said. "Due to my regular participation in the working groups, I have come to know about past and current work in the area of Internet performance. I also realized that you need to have a strong base in statistics and analytics if you [want] to consistently contribute to this area."

Zartash Uzmi, a professor of computer science and electrical engineering at Lahore University of Management Sciences in Pakistan, said he has benefitted from interacting with the Internet industry at IETF meetings. He has attended three IETF meetings as an ISOC fellow, presenting on an Internet draft at one meeting. He also acted as a scribe for the Forwarding and Control Element Separation (ForCES) WG.

"I am an academic, and there is no greater pleasure than getting your research endorsed by the industry," Uzmi said. "Being an ISOC fellow provided me the opportunity to discuss and present my work with the practitioners at IETF. I feel enthusiastic in contributing to IETF, and I draw personal satisfaction by doing so."

IETF 83 Fellow Thiago Marinello

Photos: Richard Stonehouse/Internet Society

IETF Ornithology: Recent Sightings

Compiled by Mat Ford

Common Swift (Apus apus)

Getting new work started in the IETF usually requires a birds-of-a-feather (BoF) meeting to discuss goals for the work, the suitability of the IETF as a venue for pursuing the work, and to assess the level of interest in and support for the work. In this article, we'll review the BoFs that took place during IETF 83, including their intentions and outcomes. If you're inspired to arrange a BoF meeting, please be sure to read RFC 5434: Considerations for Having a Successful Birds-of-a-Feather (BoF) Session.

SCIM—Simple Cloud Identity Management

Description: The Simple Cloud Identity Management (SCIM) specification is designed to make managing user identity in cloud-based applications and services easier. The specification suite seeks to build on experience with existing schemas and deployments, placing specific emphasis on simplicity of development and integration, while applying existing authentication, authorization, and privacy models. Its intent is to reduce the cost and complexity of user management operations by providing a common user schema and extension model, as well as binding documents to provide patterns for exchanging this schema using standard protocols. In essence, make it fast, cheap, and easy to move users into, out of, and around the cloud. For more details, see http://www.simple-cloud.info/.

Proceedings: http://www.ietf.org/proceedings/83/minutes/minutes-83-scim.txt

The Simple Cloud Identity Management specification is designed to make managing user identity in cloud-based applications and services easier.

Outcome: A very productive discussion that indicated lots of interest in the work and support for doing the work in the IETF. Ongoing discussion will address the scope of the work and how to label it (not everyone is happy to see the word *cloud* used in this context). See our cover article by Chris Phillips for more discussion of this developing area.

WEIRDS—WHOIS-based Extensible Internet Registration Data Service

Description: This work is aimed at designing a replacement for WHOIS that can be delivered as a RESTful service, with an eye toward avoiding a number of the issues that have prevented IRIS (Internationalized Resource Identifiers) deployment as a WHOIS replacement. The impetus for this work is the existence of three already deployed experimental services similar to the approach being proposed for IETF work, as well as the burgeoning number of internationalized domain name (IDN) TLDs in the domain name system root zone. See the cover article from the last issue for more detail: http://www.internetsociety.org/articles/something-weirds-way-comes.

Proceedings: http://www.ietf.org/proceedings/83/minutes/minutes-83-weirds.txt

Outcome: Following the BoF meeting held during IETF 82 in Taipei, Taiwan, this meeting focussed on how to scope the charter for a working group (WG) on this topic. If working on a solution that includes domain names starts to slow things down, the focus will shift to numbering resources only.

Continued on next page

Barn Swallow (Hirundo rustica)

IETF Ornithology: Recent Sightings, continued

RPSREQS—Remote Participation Services Requirements

Description: For many years, the IETF has provided tools for remote participation in a variety of activities. Some IETF participants also used their own tools when they felt the need. The IETF now wishes to support enhanced remote participation that is as seamless as possible, approaching the quality of direct physical attendance for the various roles—including chair, presenter, and simple attendee. Before deploying the new tools and services needed for this enhanced remote participation, the requirements for such tools and services must be defined. This meeting allowed for discussion of the draft requirements and participants' experiences of remote participation during the IETF 83 meeting.

Proceedings: http://www.ietf.org/proceedings/83/minutes/minutes-83-rpsreqs.txt

Outcome: Good discussion, more to come as this work progresses.

The IETF now wishes to support enhanced remote participation that is as seamless as possible, approaching the quality of direct physical attendance for the various roles—including chair, presenter, and simple attendee.

ANTITRUST—Does the IETF Need an Antitrust Policy?

Description: Standards development organizations have done a range of things in response to antitrust legislation over the years. This meeting addressed the question of whether there is anything that the IETF could or should usefully do in this regard.

Proceedings: http://www.ietf.org/proceedings/83/minutes/minutes-83-antitrust.txt

Outcome: The meeting identified a need for two types of materials to be produced: (1) educational materials aimed at chairs that clarify the laws of various countries, and (2) materials that put this information within the context of the IETF.

RFCFORM—RFC Format

Description: The newly appointed RFC series editor, Heather Flanagan, is soliciting input on the question of formatting the RFC series. This topic has a long history of passionate debate within the IETF community and this meeting was intended to capture as many of the different requirements and concerns about moving to a new format as possible before any firm proposals are tabled.

Proceedings: http://www.ietf.org/proceedings/83/minutes/minutes-83-rfcform.txt

Outcome: Very well organized session, no conclusions reached yet.

Blackbird (Turdus merula)

NVO3—Overlay Networking

Description: Support for multi-tenancy has become a core requirement of data centers (DCs), especially in the context of data centers supporting virtualized servers known as virtual machines (VMs). A key multi-tenancy requirement is traffic isolation, so that a tenant's traffic (and internal address usage) is not visible to any other tenant and does not collide with addresses used within the data center itself. Another key requirement is to support the placement and live migration of VMs anywhere within a data center, without being limited by DC network constraints, such as the IP subnet boundaries of the underlying DC network.

This work proposal will develop an approach to multi-tenancy that does not rely on any underlying L2 isolation mechanisms to support multi-tenancy. In particular, the proposal will develop an approach where multi-tenancy is provided at the IP layer using an encapsulation header that resides above IP. This approach should provide an Ethernet service. It may provide an IP service; an important goal is to develop a layer-agnostic framework and architecture that meets data center requirements.

This was a packed meeting with lots of support for adopting this work within the IETF. The chairs presented a well-developed draft charter and there was considerable support for it.

Proceedings: http://www.ietf.org/proceedings/83/minutes/minutes-83-nvo3.txt

Outcome: This was a packed meeting with lots of support for adopting this work within the IETF. The chairs presented a well-developed draft charter and there was considerable support for it. Many people indicated they would be interested in actively participating in this work by writing and reviewing documents and there was no dissent. A proposed charter for a WG is now in review.

I2AEX—Infrastructure-to-Application Information Exposure

Description: A non-WG-forming BoF to investigate infrastructure-to-application information exposure and communications requirements in fully controlled (such as data centers) or partially controlled environments (such as content delivery networks (CDNs)). Application-layer traffic optimization (ALTO) was designed to provide network infrastructure information to applications that may not be under the same administrative control as the network. ALTO only reveals limited information about the network infrastructure. By comparison, other protocols that monitor and manage infrastructure may reveal much more information, but are typically only accessible to the operators of the network infrastructure. CDNs and data center applications have some requirements to operate over the Internet, possibly between administrative domains. ALTO was initially designed to address peer-to-peer application requirements, but it was designed to be extensible. This BoF seeks input from the IETF community about whether ALTO (with possible extensions) or other protocols might be most appropriate to satisfy the information access requirements of CDN and data center applications. As this is not a WG-forming BoF, but an input-gathering exercise, discussion of proposed work in this BoF should not lead to any expectation that any specific proposed work will be authorized.

European Magpie (Pica pica)

Proceedings: http://www.ietf.org/proceedings/83/minutes/minutes-83-i2aex.html

Outcome: Operator input indicated that if trust and security issues can be resolved, then they are willing to expose the information to applications. Application developers indicated they are willing to make use of this information. The ALTO model provides a useful abstraction but it is critically missing a publish-subscribe mechanism. There is a choice of protocols that may apply here. Work may proceed in the ALTO WG, or maybe in another venue.

IETF 83 At–A–Glance

Registered attendees: 1318

Newcomers: 230

Number of countries: 56

IETF Activity since IETF 82 (Nov 2011–Feb 2012)

New WGs: 3

WGs closed: 5

WG currently chartered: 115

New Internet-Drafts: 576

- 200 updated
- 46 updated more than once

Updated Internet-Drafts: 1114

IETF Last Calls: 133

Internet-Drafts approved for publication: 125

RFCs published: 115

- 64 Standards Track and 5 BCP
- 40 Informational and 6 Experimental

IANA Activity since IETF 82 (Nov 2011–Feb 2012)

Processed 1250 IETF-related requests, including:

Reviewed 110 I-Ds in Last Call and reviewed 102 I-Ds in IESG Evaluation

Reviewed 98 I-Ds prior to becoming RFCs and 60 of them contained actions for IANA

Processing goal average for IETF-related requests: 92%

Updated RFC Editor Queue: http://www.rfc-editor.org/current_queue.php

More-complete RFC Editor Report: http://www.rfc-editor.org/ietf.html

Good News for Time Zone Database

- Astrolabe dropped copyright infringment suit against David Olson and Paul Eggert noting that it was “based on a flawed understanding of the law.”

RFC Editor Activity since IETF 82 (Nov 2011–Feb 2012)

Published RFCs: 124

Internet-Drafts submitted for publication: 103

- 82 IETF WGs

Thanks for the Code

Code sprint was very successful

- Incremental improvements to datatracker
- Deployed datatracker

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

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

http://www.internetsociety.org/ietfjournal

Photos: Richard Stonehouse/Internet Society

IETF Meeting Calendar

IETF 84
29 July–3 August 2012
Host: Google
Location: Vancouver, BC, CA

IETF 85
4–9 November 2012
Host: North American Cable Industry
Location: Atlanta, GA, USA

IETF 86
10–15 March 2013
Hosts: NBC Universal, Comcast
Location: Orlando, FL, USA

IETF 87
28 July–2 August 2013
Host: TBD
Location: Berlin, Germany

For more information about past and upcoming

IETF Meetings

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

Special thanks to

CISCO

for hosting IETF 83

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

Afilias afnic Google intel Microsoft Backing from nominet trust www.nominettrust.org.uk SIDN

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

Afilias ARIN American Registry for Internet Numbers CISCO comcast Juniper NETWORKS RIM TWNIC www.twnic.net.tw RIPE NCC SIDN VERISIGN

Photos: Richard Stonehouse/Internet Society

This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 3.0 Unported License.

pub-IETFJv8.1-20120702-en

Published three times a year
by the Internet Society
Galerie Jean-Malbuisson 15
1204 Geneva, Switzerland

Editor
Mat Ford

Associate Editors
Megan Kruse • Wendy Rickard

Contributing Writer
Carolyn Marsan

Editorial and Design
The Rickard Group, Inc.

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

Email
ietfjournal@isoc.org

Find us on the Web
www.internetsociety.org/ietfjournal

Editor's Note
The IETF Journal adheres to the *Oxford English Dictionary*, 2nd Edition.

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

IETF Journal

IETF 83 • Volume 8, Issue 1 • June 2012

IETF Journal

Internet Society
Galerie Jean-Malbuisson 15
1204 Geneva, Switzerland