Report of the

IETF 80

Prague, Czech Republic
March 27 – April 1, 2011

Prepared by Monika Ermert
For the CENTR secretariat

Table of Contents

Highlights 3

DNS-TLS marriage still faces obstacles – „there are a lot of sharks“ 3

Browser Cabal or future of the net: Real-time Communication on the Web 4

Privacy Issues in Web RTC 6

The Fight over MPLS between IETF and ITU 6

Working Groups, BoFs 8

DNSOP 8

DNSEXT: IETF and ICANN discussing variants/aliases/bundles 10

SIDR recharterd 11

Paws 11

Plenaries 12

IETF Budget 12

RFC editor 12

IANA - statistics 13

Future IANA contract 13

New IAB Chair 14

Meeting Venues selected until Jan, 13 14

Highlights

DNS-TLS marriage still faces obstacles – „there are a lot of sharks“

A full-fledged IETF working group has been created since IETF 79 to work on the „marriage“ of the Domain Name System and Transport Layer Security (DNS and TLS). The main goal of the Working Group on „DNS based Authentication of Named Entities“ (DANE, with Dane being the Czech word for *tax)* is to allow the use of DNSSEC enabled DNS to associate a TLS server's certificate with an intended domain name. As Paul Hoffman (Internet Mail Consortium) put it during the presentation in Prague, the core idea is; if one can trust DNSSEC for the address one is using, one can trust it for the certificate as well.

The idea that the DNS (or the DNS administrators) could take on the role of a trust anchor instead of third-party certification authorities has been around for a while, but clearly became virulent with the roll-out of DNSSEC. Some see DANE as the killer-application for DNSSEC in general, not the least some DNS registries. Jim Galvin (Afilias) said the mid- or long-term vision was to allow every DNS server to provide https by default by self-certifying when being set-up.

Problems identified, though, are lack of last-mile DNSSEC security – in case users do not operate their own DNS resolver – and also the question of the value of the DNS-stored self-certification. DNS administrators currently do not check the „content“ of what they put into the DNS, which could also be true for the signatures added to a DNS entry, Peter Koch (DENIC) pointed to for some time. If consequently deployed - and with all security considerations take into account – DANE could be a good solution, one long-time DNS expert said. Yet he was afraid that the standard development in the IETF DANE working group would face opposition from commercial certification providers. The closer to finalizing the documents, the tougher opposition would be, he said. During the session one participant spoke about „a lot of sharks in the room“. Phillip Hallam-Baker (from CA provider Comodo) raised several concerns and promoted his own draft proposal for fixing CA problems (especially mis-issue of certificates recently experienced by Comodo itself).

According to the time plan of the WG, DANE is already close to delivering for IESG review a protocol for using DNS to associate domain names with keys for TLS and DTLS, plus an additional protocol to associate them with keys for IPsec. The date mentioned in the WG milestones is September 2011. A first call for more caution and consideration – and slower path? - was given during the Prague DANE session, with substantial support for getting a requirements document. The WG settled on quickly finishing a „use case type“-document with an initial list of use cases distributed and heavily discussed on the DANE mailing list by now (see below).

In general the current version 06 of the DANE document considers a new resource record (TLSA RRType). The format of this record contains three values: a certificate type, a reference type and the certificate for association. The document lists two possible certificate **types (**1, end entity certificate and, 2, certification authority's certificate) and **three reference types** (full certificate, SHA 256 hash of the certificate, SHA 512 hash of the certificate). With regard to crypto algorithms a proposal was made by Steve Kent (BBN) and supported by IETF Chair Russ Housley to put these in a separate document.

Issues discussed and still open issues in the document can be followed on a issue tracker list. It is unclear, for example, if bare keys will be an option. For the time being -according to Hofmann- they are not foreseen in TLS and therefore not a possibility in DANE. Yet lack of clarity in the TLS specification has been acknowledged during the Prague session, for example, by Eric Rescorla. A discussion about the relation between DANE and PKIX took place with Hallam-Baker warning that PKIX would be overridden by DANE, while Richard Barnes in a presentation explained that for harmonizing PKIX with DANE in some instances second certificates could become necessary for end entity (domain issued) certificates. For domain issued bare keys there still was a need to generate and store certificates, Barnes explained.

Potential use cases/requirements for the requested new requirements document written by Richard Barnes (BBN) are:

Use Case 1 "CA Lock":
- The certificate my server presents in TLS will be chain to this CA.
- Clients should accept a TLS server certificate only if it chains to this CA, but may also require that it chain to an existing trust anchor.

Use Case 2 "Cert Lock":
- The certificate my server presents in TLS will be this specific certificate.
- Clients should accept a TLS certificate only if it matches this certificate, but may also require that it chain to an existing trust anchor.

Use Case 3 "New TA":
- The certificate my server presents in TLS will chain to this CA, which should be treated as a TA.
- Clients should accept a TLS server certificate if and only if it chains to this CA.

Use Case 4 "Certificate as Bare Key":
- The certificate my server presents in TLS will be this specific certificate.
- Clients should accept a TLS certificate if and only if it matches this certificate.

Another lengthy use case document was just posted Hallam-Baker. It includes assumptions about the lack of deployment of DNSSEC, see here.

Browser Cabal or future of the net: Real-time Communication on the Web

Communication seems to be on a steady move to port 80, the web, a trend the IAB decided to be worth to be discussed in a technical plenary about the future of applications in Prague, featuring experts of the IETF and the web standardization body W3C. Are IETF standards and standards in general still able to shape the net in an ever faster moving world of applications, apps?

Jonathan Rosenberg, well-known developer of SIP, a major protocol for Voice over IP, and working for Skype said a root cause for the trend to unstandardized applications was the difference in innovation cycles. Due to the elimination of dependencies, the faster innovation cycle of apps was taking over from the older, much slower telecom innovation cycle.

Rosenberg said he had checked his phone to prepare for the talk to illustrate the trend to find 83 little apps icons on his phone of which 53 could be categorized in some way as cloud apps. Most of these were proprietary. The web model and applications using the web were heading in this same direction, said Rosenberg. „Both trends are new techniques for easy distribution of software to the client that allows an entity to build both the server and client pieces of this. Back in the original days of original Internet and SIP, desktop and hardware devices were a problem. But through app stores and web, software on people's client is not a problem like it used to be."

While Henry Thompson from the W3C spoke about a “real rise of port 80” and invited the IETF to cooperate with the W3C on new applications, Mark Nottingham, who acts as a liaison for the IETF to the W3C, added that the capabilities of the next generation of browsers were „truly astounding“ with browsers allowing access to local file systems, web sockets, audio, video. In some cases the browsers were acting as operating systems. Nottingham said: „This architecture is not being pushed by the IETF. It is not pushed by the W3C. It is pushed by a group of browser vendors, a cabal.“ This group, Nottingham said, was pushing a vision where standards would not apply anymore.

A clear contradictory point was made by Leslie Daigle, Chief Technologist at the Internet Society and ex-IAB chair, while speaking on a personal basis. With regard to the http and application hype, she said, some of the discussion was focused on „we don't really need to do […] new standardization of application protocols because we can migrate everything over http or over the http infrastructure.” But the sad fact of the matter was that these overlaid protocols would be constrained at some level by the underlying http semantics. Answering the question what should still be standardized, she said „interoperability is of course the reason why you want to standardize something or specify it.“ Only standardization could allow for further building blocks for future innovation.

IETF and W3C obviously have joined to give their answer to the trend: A BoF on Real-Time Web Communication (RTC) started work that shall allow future applications to easily hook up to browsers by finding standardized protocols for audio, video, gaming and collaboration supporting protocols. Instead of adding these functionalities in packaged plugins or browser extensions, standard interfaces should allow the use of a set of standard protocols for real time communications available for browsers – and possibly other Internet platforms.

Harald Alvestrand, co-chair of the BoF, (and working for Google, but reiterating the mantra that he was speaking personally in the IETF) confirmed that it was browsers the potential RTC WG was targeting first and foremost. From a browser vendor perspective the ease for application developers to bring their applications to the browser was attractive. Currently the only way to embed interactive features was by using Flash or RTSP. Instead of having to develop a new version of an application to support various browsers, the planned set of standards and corresponding standard APIs (these shall be developed by the W3C) would lower the barrier for new applications. Alvestrand's BoF Co-Chair, Rosenberg, when asked for the interest a provider like Skype had in this work, said: “We go where the user goes. If he goes to the desktop, we're there. If he goes mobile, we're there. If he goes to the Web, we're there.”

Several things are not fully clear after the BoF: Will the future IETF WG fulfill its task by just listing a set of standards? An initial list of standards to be considered was

1) RTP/ RTCP
2) a baseline audio codec for high quality interactive audio. Opus will be considered as one of the candidates
3) a baseline audio codec for PSTN interoperability. G.711 and iLBC will be considered
4) a baseline video codec. H.264 and VP8 will be considered
5) Diffserv based QoS
6) NAT traversal using ICE
7) RFC 4833 based DTMF transport
8) RFC 4574 based Label support for identifying streams purpose
9) Secure RTP and keying
10) support for IPv4, IPv6 and dual stack browsers

Yet during the Prague session ICE (for NAT traversal), for example, was said to have been taken from the list. Another issue discussed was in what way APIs, which shall be developed by the W3C (only?), had to be treated as "protocols".

While many observers see the Web RTC work as a considerable step of the IETF to make up lost ground in the application area, others think the work comes much too late. Opinions about the very ambitious time plan diverge in the same way. The IETF milestones include a finalized draft reflecting what the protocol set should be (August 2011, final document for IESG in December 2011) and a documentation specifying the mapping of protocol functionality to W3C-specified API (November 2011, final document to IESG April 2012).

Privacy Issues in Web RTC

Special consideration had to be given to security and privacy issues, Alvestrand said in Prague. The real-time communication applications "will not only be used by your friends, but also by your enemies". An issue discussed was, for example, the control of the web camera in order to avoid new types of surveillance. Also the control of incoming voice communication by the user was necessary with allowance to start the respective communication flow only after consent from the user's side. A short discussion about what "consent" meant here – machine/application consent or end user consent, something requested in privacy legislation like, for example, the currently reviewed EU privacy directive – was inconclusive.

Eric Rescorla in a presentation about security issues said it was, for example, not ok to let browsers send TCP and UDP to arbitrary locations, so before sending traffic from a sender to a recipient it had to be verified that the recipient wanted to receive it from this sender. Beside such consent to communication, access to local devices and communications security (key storage) had to be considered. Privacy issues -according to the agreement with the W3C- will be covered by W3C work.

Privacy meanwhile is part of ongoing work of the IAB, that has tabled a draft proposal about a potential "privacy consideration" extension to RFCs. The privacy considerations would be added in the same way as the security or IANA considerations. Additional work on a potential do-not-track-tag was discussed rather controversially during the Websec WG.

The Fight over MPLS between IETF and ITU

While cooperation with the W3C seems to work smoothly, the IETF is in a struggle again with the International Telecommunication Union (ITU) over follow-up operational mechanisms for Multiprotocol Label Switching (MPLS). MPLS, originally standardized by the IETF (since 1996, with around 100 documents on MPLS and CCAMP), allows to establish "virtual links" between distant nodes by attaching labels to packets (IP, ATM or others). It is for example essential for quality of service. On February, 25, the ITU-T Study Group 15 "determined" their OAM recommendation for MPLS – the determination was made by a completely unusual vote as consensus could not be reached in the study group.

IETF Chair Russ Housley reacted promptly making a strong statement via an ISOC press release (and via statements to the press) warning that the ITU move "takes us off the path of global interoperability for this technology.” The issue of who and how could speak for the IETF came up several times as an item of discussion during the IETF plenaries (see below). Both IETF/ISOC and ITU in March released more press and public statements, putting the fight between the two standardization bodies in front of a much wider audience than usual.

The fight over MPLS has been smoldering ever since the ITU began to develop “T-MPLS” in 2006/2007. T-MPLS according to the ITU was a “sub-set of MPLS that was specifical for application in the transport network”. The IETF in 2007 indicated to the ITU that T-MPLS (in sum five draft recommendations and one OAM recommendation by 2008) was in clear conflict with IP/MPLS, for which standardization in the IETF had began some time in the mid nineties. IETF and ITU in 2008/2009 finally agreed to establish a Joint Working Team (JWT, documented in RFC 5317) to “bring transport requirements into the IETF and extend IETF MPLS forwarding, OAM (Operations, Administration and Management), survivability, network management and control plan protocols to meet those requirements through the IETF Standards process.”

After the February, 25th vote the ITU argued that development of MPLS OAM had been stalled and specific requirements by the ITU had been rejected in the IETF; an ITU report mentions for example a waiver of “rate negotiation” as one controversial issue.

During the Prague meeting Malcolm Betts, from the Chinese hardware vendor ZTE, presented the view of the ITU Study Group 15, which mainly centers around the IETF development not making up to it at the IETF meeting, and also tabled a request for a IANA code point allocation allowing the ITU MPLS OAM to be identified on the wire. The IETF leadership in Prague on the other hand clearly rejected the code point request. IETF Chair Russ Housley said: „The normal process is to develop one solution for one problem.“ The IETF did see no reason for a second one and would continue to develop MPLS OAM in the IETF “in the spirit of the JWT agreement”, as IETF leaders underline.

MPLS – a technically or politically motivated fight?

How dangerous would two standard-variants be for interoperability in the Internet? Betts once more rejected the IETF warning about a path to non-interoperability. He said differences between the two variants were nearly “invisible” provoking tense questions from several IETF participants, like Nurit Sprecher from Nokia Siemens Networks, for a reasoning for an ITU standard in the first place.

Leaving IAB- Chair Olaf Kolkman explained to this reporter that with divergent OAM solutions it would become more complex and thus also more costly to run a “transport network” and an MPLS/IP network. “While the ITU-TI experts that do not agree with the technical argument that two solutions threatens the integrity of the Internet, the “one protocol for one job” was the reason that the ITU-T and the IETF agreed that there should be one standard and set out to develop one a few years ago.” Another IETF participant (from Ericsson) complained in Prague that the ITU was granting itself the right to trample on the IETF's Intellectual Property to the MPLS standard.

From various comments it seems somewhat obvious that IETF participants are concerned that the MPLS cause might serve as a precedent for future ITU “re-use” of existing IETF standards, with some pointing to base protocols including TCP/IP. The issue by now seems to be highly politicized with governments taking opposite positions in the still ongoing ITU discussion. The US Government according to IETF information has requested the MPLS OAM ITU standard to be processed under the so called “traditional approval process” (TAP) as opposed to the “alternative approval process” (AAP). The TAP includes a three months consultation period for member states (not yet started due to the draft text for G.8113.1 not yet officially available and a subsequent statement by the Director of the ITU Standardization Bureau still to be made).

According to the TAP procedure, the recommendation needs 70 percent support out of member states answers for the consultation to be up to approval by the study group. Upon request the ITU secretariat explained that “true opposition from one Member State is sufficient to prevent approval (except at WTSA)”. The relevant provision says “Should any Member States be of the opinion that consideration for approval shall not proceed, they should advise their reasons for disapproving and indicate the possible changes that would facilitate further consideration and approval of the draft new or revised Recommendation.”

The ITU following its Plenipotentiary mandate currently works on “cooperation and collaboration with standardization organizations”, but this seems to become more difficult instead of easier.

Housley gave another very detailed report on the issue from the IETF's perspective to the ISOC Board during the IETF meeting in Prague.

Working Groups, BoFs

DNSOP

DNSSEC, IPv6 and DNS work in the Behave and MIF WGs were touched during the Prague meeting. Participants mainly agreed to send the document on DNSSEC Operational Practices (Version 2) into last call with changes presented by Matthijs Mekking (NLnet Labs) during the WG (most recent changes were an additional section on the motivation for algorithm roll over, and another one on non-cooperating operators). The document represents the current status, said co-author Olaf Kolkman, with the option to develop a new version in 2-3 years. There is a general feeling that DNSSEC is just taking off (with a lot of large TLDs supporting DNSSEC right now, including .com) and documents related to it might still need further refinement.

WG participants, for example, said while trust anchor bootstrapping was a problem that needed attention, the solution presented in a new document by Joe Abley (ICANN) might need a closer look. The issue Abley is addressing is how validators might determine an appropriate trust anchor for the root zone to use at start-up or when other mechanisms for a graceful key rollover are not available.“ The approach proposes that a validator would identify the trust anchors valid for current use (http://data.iana.org/root-anchors/root-anchors.xml), then retrieve the corresponding X509 identity certificates for the key identified and then finally start to validate. Critical remarks made during the session regarding the approach were that it could be viewed as somehow „circular in where you go to get your trust anchor“ (Russ Mundy, Sparta).

More DNSSEC-related work was proposed by Mekking, who shortly presented a follow-up to the key-timing document, which is close to last call. Mekking listed several issues he wants to address in the new document: rollover considerations, key types, key goals, unraveled key states, rollover centric logic and new rollover scenarios. WG chairs requested WG review of both (existing and new key timing document). Other DNSSEC drafts to watch are the DNSSEC Policy and Practice Statement Framework (version 04, which was expected to include views from an ongoing CENTR survey) and „Changing DNS Operators for DNSSEC signed Zones“ (currently not a DNSOP WG document).

New work was taken up with regard to an expected increase in the query load on the DNS root servers and the IP6.arpa authoritative servers due to „stupid queries“ in IPv4. For non-delegated subdomains servers have to send NXDOMAIN responses with a high probability that these answers would be repeated further increasing the query load. George Michaelson from APNIC presented figures to demonstrate what he said could develop into a potentially huge problem: With a v6/v4 transport rate of 1,78 percent and a in-addr/ip6.arpa rate of 7,56 percent the problem was still manageable, yet with growing IPv6 traffic this can change dramatically, Michaelson said. Sources for bad traffic were not only queries for un-delegated (private) addresses as in IPv4, but also queries for link local, site local and multicast addresses, unique local addresses (ULA) and attempts to tunnel (6RD, 6to4, Teredo). Michaelson asked to address the issue now despite the fact that the AS112 documents (to deal with the issue in IPv4) still were in last call. His proposal is to get a v6 prefix assigned to AS112. The WG in its majority favored to start work on the issue right away.

Finally DNSOP chair, Peter Koch, and DNSEXT chair, Andrew Sullivan, pointed to several work items of other WGs looking for the DNS solutions to, for example, v4-v6 translation. Sullivan made reference to work in the Behave WG that would allow registering a „well know (DNS) name“ to allow a AAAA query for it.

A nice overview over the various proposed mechanisms to learn about NAT64 prefixes is given here.

It includes the following five DNS-assisted solutions (and four more solutions using DHCP and other protocols). The five DNS-based proposals are:

EDNS0 option indicating AAAA Record synthesis and format

The document korhonen-edns0-synthesis-flag *defines a new EDNS0 option [RFC2671], which contains 3 flag bits (called SY-bits). The EDNS0 option serves as an implicit indication of the presence of DNS64 server and the NAT64 device.*

EDNS0 flags indicating AAAA Record synthesis and format

The document EDNS0-Flag *defines 3 new flag bits (called SY-bits) into EDNS0 OPT [*RFC2671*] header which serve as an* implicit *indication of the presence of DNS64 server and a NAT64 device.*

DNS Query for a Well-Known Name

I-D.savolainen-heuristic-nat64-discovery *describes how a host requiring information for local IPv6 address synthesis or for NAT64 avoidance sends a DNS query for an AAAA record of a Well-Known IPv4-only Fully Qualified Domain Name (FQDN). If a host receives a negative reply, it knows there are no DNS64 and NAT64 in the network.*

DNS Resource Record for IPv4-Embedded IPv6 address

I-D.boucadair-behave-dns-a64 *defines a new DNS Resource Record (A64) that is a record specific to store a single IPv4-Embedded IPv6 address [*RFC6052*]. Using a dedicated Resource Record allows a host to distinguish between real IPv6 addresses and synthesized IPv6* *addresses.*

Learning the IPv6 Prefix of a Network's NAT64 using DNS

I-D.wing-behave-learn-prefix *proposes two DNS-based methods for discovering the presence of a DNS64 server and a NAT64 device, and then a mechanism for discovering the used NSP. First, a host may learn the presence of a DNS64 server and a NAT64 device, by receiving a TXT Resource Record with a well-known (TBD IANA registered?) string followed by the NAT64 unicast IPv6 address and the prefix length. The second method proposed is to specify a new U-NAPTR [*RFC4848*] application to discover the NAT64's IPv6 prefix and length.*

DNSEXT: IETF and ICANN discussing variants/aliases/bundles

The DNSEXT Working Group deliberated how to proceed with the document on DNS aliasing (variants, bundling), especially against the background of ongoing work in ICANN on IDN variants. The co-author of the IETF document on „DNS Resolution of Aliased Names“ (version 01 see here), Suzanne Woolf, proposed to send a status message to ICANN and allow for input from the policy debates before finishing up the document.

Others warned about waiting for ICANN, with Ted Hardie arguing the technical committee should feed its observations right away instead of confronting the „policy community“ with technical considerations – and possibly feasibility – afterwards. Jaap Akkerhuis from NLnet Labs warned about waiting for ICANN since the latter was preparing several studies (an issue report based on case studies of variants in Chinese, Arabic, Cyrillic, Indic and Latin, see here) resulting in time delays for the IETF requirement document. Akkerhuis also said he was pessimistic „when it comes to what the rather high-level ICANN studies could add“ to the IETF document.

Paul Hofmann pointed to such studies coauthored by himself in 2000/2001. After an inconclusive discussion Thomas Narten, who is acting as the liaison to ICANN for the IETF, recommended to allow both bodies to proceed in parallel with the IETF sending a clear message when there was one. Participants in Prague, while not seeing deeper problems with the requirement statement as it stands, did not see it ready for last call yet.

Woolf announced an overhaul of terminology for version 02 and asked for text on use cases with regard to IDN and non-IDN aliasing (the Arabic variant section is still empty in the document). Also some more discussion on the solutions (currently the draft lists CNAME, DNAME, BNAME and Zone Cloning) was necessary. Regarding CNAME there was a quick discussion about potential limitations because of its current deployment levels and its original character as a temporary, light-weight measure for migration. Another possible quick fix to „aliasing“ mentioned by Phillip Hallam-Baker (Comodo) were DNS redirects which might lead to negative side effects.

Woolf's main question was if the five requirements listed in the problem statement were the „right requirements“. The five requirements in the document are:

- DNSSEC support of any solution,
- backwards compatibility,
- no overhead for registries, authoritative servers, clients in comparison to existing mechanisms (for example provisioning or existing resource record solutions),

- new Resource Record Types as a possibility
- no simple shift of costs from DNS authoritative service providers to DNS users

A question of substance discussed at the meeting was if expectations for the IETF work on aliases went as far as asking for „equivalence“ of sets of names or if it was enough to define one out of a set of names as canonical with the others to be considered as variants. Hardie said he had a „niggling fear“ that in the end there could be demands for the IETF to create another layer of indirection, to allow several names to be pointers to some kind of „metarecord“. According to the draft protocol statement „no requirement for complete interchangeability or identity“ had been articulated so far. Such equivalence would be extremely difficult to define in the DNS, he said.

A last agenda item was discussed at the Prague meeting: resolver improvements with regard to the handling of NXDOMAINs proposed by Paul Vixie (ISC). The measures proposed for optimization are:

- Re-validating a delegation when a parent NS RRset TTL expires
- Stopping a downward cache search when an NXDOMAIN is encountered
- Upgrading the credibility of NS RRsets upon delegation event

The discussion in Prague focused on the handling of NXDOMAIN answers. Peter Koch (DENIC) asked for clarifying side effects of potential „aggressive negative caching“.

SIDR recharterd

The Secure Inter-domain routing WG has stepped up to the next level in securing the routing system. As the protocol suite for securing the origin of a route is nearly completed – documents are either in IETF last call or in the IESG review process – the WG started work on securing the AS path in Prague. From its outset according to WG Chair Sandra Murphy it has been noted that protecting the origin was not enough to secure the routing system, because a valid origin could be appended to a bogus path. During the Prague meeting the new SIDR Charter that takes up securing the path was accepted by the IESG. Four documents were presented during the Prague meeting, including a basic threat analysis by Steve Kent (BBN), an operational requirement document by Randy Bush (IIJ) and two documents on the core idea for securing the AS path by what the authors call “BGPSEC” (overview and protocol) by Matt Lepinski (BBN). The basic idea of BGPSEC is to add a new type of certificate, the BGPSEC router certificate, “that binds an AS number to a public signature verification key, the corresponding private key of which is held by one or more BGP speakers within this AS.” A deployed BGPSEC would provide an “attribute called BGPSEC_Path_Signatures” consisting of a “sequence of digital signatures, one for each AS in the AS path of a BGPSEC update message”. The interest of operators seems to be high; participants from Level3, Cable&Wireless and Deutsche Telekom were participating in the discussion. Yet there are some concerns with regard to the prerequisite for more mighty hardware to process the additional certificates.

Paws

Interesting new work has been started with regard to more effective work with still unused TV frequencies, so called white spaces. The PAWS BoF presented a concept driven by the US National Broadband Plan and by companies that have been awarded FCC licenses to set-up white space databases (Comsearch, Frequency Finder Inc., Google Inc., KB Enterprises LLC and LS Telcom, Key Bridge Global LLC, Neustar Inc., Spectrum Bridge Inc., Telcordia Technologies, and WSdb LLC).

Standardization at the IETF of a query method to the various data bases and a basic data model for the data bases will allow to make unlicensed (secondary) of “under-used” TV band frequencies while protecting the primary (licensed) users from interference. Hardware vendors have a big interest in a standard solution because it will allow them to develop the necessary additional capabilities of end devices.

For the time being according to Basavaraj Patil (Nokia) it would be too expensive to put intelligence for a mere sensing solution to allow end devices to query for free white space. The solution envisaged for now was to put intelligence about free TV spectrum capacity at a given time in a given location in national/regional databases that will be queried by end devices. Other regulatory bodies already preparing for the unlicensed use of white spaces, according to Basavaraj Patil, are the British regulator Ofcom and the Finnish government.

How far requirements of other regulatory authorities will be included seems to be an open question. The ECC report on “Technical and Operational Requirements for the possible operation of cognitive radio systems in the White Spaces of the frequency band 470-590 MHz” was referenced. With regard to other standardization bodies the IEEE Wireless standards are mentioned.

Plenaries

IETF Budget

The IETF revenue for 2010 was 101.000 USD lower than expected (unaudited result for 2010 3,033 Million USD), yet lower expenses (especially savings in tools' development, 243.000 instead of 575.000, and 125.000 less in meeting expenses) resulted in positive net result. ISOCs contribution to the IETF budget therefore could be reduced from 2,1 to 1,8 Million USD.

For 2011 the IETF will raise meeting fee from 635 USD to 650 USD. Day passes will continue to be available for 350 USD. Meeting expenses are expected to increase by 2 percent, secretariat costs by 5 percent. IT development is expected to rise from 243.000 to 430.000 (disregarding the fact that IT development came in under budget in 2010, see above). Revenues budgeted for 2011 are 3,317 Million USD, budgeted expenses are 5,003 Million (plus an additional 429.000 for IT tools' development). ISOC's contribution according to this budget would once more be 2,115 Million USD.

RFC editor

The cost for the RFC editor expenses are also expected to rise in 2011. Glen Kowack, after having worked as a transitional RFC editor – and being heavily criticized for his proposals on the future tailoring of the RFC editor series organizational structure during the meeting in Beijing - was succeeded by leaving IAB Chair Olaf Kolkman. Kolkman will work as a transitional RFC editor until the planned new structure is fully in place. Kolkman has tabled a draft to document current consensus positions on the new structure (meant to be a temporary placeholder until rfc 5620 on the future RFC editor model has been updated, according to consensus on the RFC-interest list). Consensus, according to Kolkman's documents, is sharing the workload of the classical RFC editor into the RFC Series Editor, the Independent Submission Editor, the RFC Production Center, the RFC Publisher and the RFC Series Advisory Group (RSAG).

IANA - statistics

Russ Housley presented figures for IANA's work. IANA has processed 1468 IETF related requests since the IETF 79 in Beijing (680 private enterprise numbers, 79 port numbers, 47 TRIP (Telephony Routing over IP) ITAD numbers, 15 language subtag requests, 46 media type requests). IANA reviewed 136 I-Ds in last call, 134 I-Ds in IESG evaluation and 117 I-Ds prior to becoming RFCs (with 57 including IANA action).

Future IANA contract

By the end of the IETF week the IAB filed its comments in reaction to the Notice of Inquiry by the US National Telecommunications and Information Association (NTIA) on the future IANA contract. The core aspect of the IAB's position is that it is not favoring a split of the IANA functions - protocol, DNS and IP address allocation – as the functions were technically related, linked historically, and also shared the principle of bottom-up consensus governance processes.

The IAB in its position paper reminds the NTIA about the IETF/IAB role and mandate "to approve the appointment of an organization to act as IANA on behalf of the IETF" (RFC 2850) and requests to be involved in any selection of a new operator: "Should any changes to the existing IANA Functions operator be proposed, the successor will have to meet the requirements of the IETF as documented in RFC 6220 and stability and security of the continued operation must be assured." Changes to the current operation, which the IAB does not see as necessary, "would inevitably be disruptive".

Other contributions to the NTIA Notice of Inquiry-consultation favor a split of the IANA functions. A very strong statement arguing for a split is the joint statement of SWITCH, the Swiss Registry, and Swiss Telecom Regulator OFCOM. Several Government or official statements (e.g. Arabic countries, China) emphasize the need to allow for independence from US government oversight, see:

CNNIC*: "Moreover, it is also suggested the approval process by DOC be removed in the IANA process when change of request is submitted by ccTLDs to reflect the fact that the ccTLDs are to serve local community of the nation. The change of root zones should be maintained by IANA itself completely without involving DOC and Verisign."*

United Arab Emirates Telecommunications Regulatory Authority*: "We believe the current structure which is based on a procurement contract from a single government is not an appropriate model to maintain a resource that is being used and owned by the entire world. We believe that this structure must be enhanced to fulfill the above objectives, especially given that the Internet has reached a mature stage of development."*

Kenyan Government*: "However, we would like to propose transition over to an arrangement similar to the Affirmation of Commitments that replaced the U.S. Government's MoU with ICANN. (..) The requirement for approval of Internet DNS root zone by the US Department of Commerce's NTIA should transition into a multi-stakeholder relationship, where various stakeholders (Root Server Operators, IETF/IAB representing the protocol developers; RIRs the IP address functions, ICANN the gTLDs, ccTLDs and GAC would manage and oversee the functions."*

Supreme Council of Information and Communication Technology "ictQATAR"*: "Automation of IANA's Root DNS management services, transparency related to Root DNS changes and proper auditing on IANA functions performance are required changes and well be welcomed by the Internet technical community. The Time frames required to complete Root DNS changes and other IANA functions should be clear and documented, Root DNS changes requests status should be visible to the requester in a real-time mode."*

Egyptian Government: "Egypt believes that the IANA functions could be enhanced through more transparency and through accountability to the whole community which could be significantly improved by removing (or at least narrowing the scope of) a unilateral contractual oversight; consequently providing more flexibility and responsiveness of ICANN in accordance with a constantly evolving Internet."

(for a list of all contributions, see here)

Generally speaking, there seems to be overwhelming support for ICANN as future IANA manager and the DoC/NTIA role as steward for the root zone, and the IANA functions in general, is accepted by the many contributors.

The split of functions put on the table as one option by the NTIA is questioned by most – because of the lack of proved advantages. But one might very well speak of a "rough consensus" in the international community with regard to the need for more transparency for the IANA: Egypt said IANA was a black box, for example, and many others request live access to the status of change requests. This "rough consensus" on better transparency (something ICANN wrote should also apply to the NTIA part) falls pretty much in line with requests by CENTR in its joint and additional individual member submissions. AFNIC in its submission cautions against ICANN's request for structural change of the contract (MoU instead of procurement contract) as long as accountability and transparency mechanisms have not be enhanced.

Another position to note is, for example, the ITU request that any future IANA contract should oblige the manager to implement the provisions of Recommendation ITU-T E.910 for the management of .int. The ITU complains that ICANN did not react to related requests.

It should also be mentioned that the EU Commission, which had been eager to have a broader consultation of the new IANA contract, did not file any submission.

New IAB Chair

Bernard Adoba (Microsoft) took over the IAB Chairmanship from Olaf Kolkman (NLnet Labs) during the Prague meeting. While Kolkman had been heavily engaged in DNS issues, Adoba has chaired the Radext WG, worked on the Radius protocol suite (Authentication, Authorization, and Accounting (AAA) management) and has been active in the Ecrit (Emergency calls on the net) and Geopriv (geolocation and privacy) WGs. He recently is co-authoring several drafts regarding focus issues the IETF/IAB seems to look into – on privacy and on the role of classical standardization in a web application driven environment (see Browser Cabal, above, draft text here). Adoba in a short interview said one main target he had for his term was to reduce the workload of the IAB by establishing new committees and share the workload of the IAB with them.

Lars Eggert, Nokia Labs, has been selected as new Internet Research Task Force (IRTF) Chair. Eggert announced open meetings of the IRTF to bring more visibility to ongoing research work. He also announced to cooperate closely with the ISOC to support a travel grant program for scientists who want to present their work at the IRTF.

Meeting Venues selected until Jan, 13

Meeting venues for 2011/12 have been selected (81 Quebec, 82 Taipei, 83 Paris, 84 Vancouver, 85 Atlanta, 86 Orlando, 87 Europe)

The next IETF meeting will take place in Quebec, Canada, June 24-29