Report of IETF 83 Paris

25 – 30 March, 2012

by **Monika Ermert**

for the CENTR secretariat

Table of Contents

Highlights **3**

Just put it into the DNS – Secure Routing via the reverse DNS tree 3

Another attempt at the WHOIS protocol – reloaded 4

Choosing the best Web ID concepts - and discussions on browser security and privacy 5

Ipv6 – From Test flight to Take-Off

Ipv4 - A new Working group 8

Working Groups **8**

DNSOP WG 8

Simple Cloud Identity Management (SCIM) BoF, it is not what you think 9

DNSEXT 9

Security Area Meeting: New crypto-algorithm not that urgent? 10

IETF News **11**

Highlights

Just put it into the DNS – Secure Routing via the reverse DNS tree

Securing BGP routes has been a topic for years with IETF standards for Route Origin Authorization (ROA) being ready and the securing of the routing paths by BGPSEC being pushed forward in the Secure Inter-Domain Working Group (SIDR). SIDR participants gathered at the IETF Paris were (in their majority) not amused by two drafts presented by Joe Gersch, Founder of DNSSEC provider Secure64 and DNSSEC co-author Dan Massey from Colorado State University. Gersch and Massey are proposing to build on the DNSSEC secured DNS, in particular the reverse in-addr.arpa tree, to check for routing announcement verification (ROVER).

Gersch and Massey when presenting in GROW and SIDR underlined ease o use with the DNS structure in place. They propose to create two new DNS records, the Route Lock (RLOCK) Resource Record which shall indicate if an administrator has opted in and a Secure Route Origin (SRO) Resource Record that shall hold the ASN authorized for a BGP route announcement.

> „The RLOCK record allows prefix owners to indicate whether the DNS is being used to publish routing data. The SRO record allows operators to indicate whether an IPv4 or IPv6 prefix ought to appear in global routing tables and identifies authorized origin Autonomous System Number(s) for that prefix. The published data can be used in a variety of contexts and can be extended to include additional information. This work is part of an on-going effort and is accessible in an active testbed."

After receiving a BGP announcement a query of in-addr.arpa could result a „valid", (RLOCK set and DNS query matching the BGP announcement), a „bogus" (RLOCK set, but no DNS query match) or „viable" (no RLOCK set, so administrator does not use ROVER).
One trick necessary to perform is to introduce a new naming convention for CIDR blocks in in-addr.arpa, as the current reverse DNS naming method covers complete IP addresses, but no address blocks. The naming convention has to adapt prefixes that are not on a octet boundary to the octet structure of the DNS tree. The operation for prefixes not matching the octet boundary according to the naming convention is

> A. Truncate the name to remove the least significant octet. Add a "m" label to this domain name to indicate "mask".
>
> B. Convert the least significant octet to binary, separating each bit into its own label (with a "." character).
>
> C. Truncate the binary labels to the N significant labels that correspond to the given prefix_length.
>
> D. Reverse the string and add ".in-addr.arpa."

The resulting in-addr.arpa-strings look like 82.129.in-addr.arpa for the prefix 129.82.0.0/16 (at octet boundary) or 1.0.m.82.129.in-addr.arpa for 129.82.64.0/18. An Ipv6 address like 2607:fa88:e000::/35 is represented by 1.1.1.m.8.8.a.f.7.0.6.2.ip6.arpa. A potential problem was seen by several experts in the DNS WG in the fact that there were potentially two versions for prefixes on octet boundaries

The DNSOP WG hummed in favor of further discussing the naming convention draft, but fundamental objections to ROVER, from both DNS and routing experts were made during or in between sessions. "Confusion" and a luring away of "focus" from the RPKI system that is just about to take off, was an issue addressed by routing experts like Rüdiger Volk from Deutsche Telekom. Volk questioned the possibility of

effectively “filtering routes” using ROVER while an event like the Pakistan Telecom You tube-hijacking already could be prevented by now through the ROA concept developed in SIDR. Some SIDR participants, including Co-Chair Sandra Murphy, warned that with proposing an alternative to the RPKI protocol suite the system of secure routing might be “split” and administrators might have to somehow do sequential checking of the different systems in order to check for the validity of BGP routes. It is unclear so far if the DNS-based version will also be extended in any way to allow securing BGP paths. Another fundamental problem mentioned by DNSOP Co-Chair Peter Koch (DENIC) was circular statements, with trust being lost and a big bootstrapping problem in case of an attack.

On the other hand some DNS experts did first experiments with the ROVER test bed and applauded its ease of use, see the evaluation of Stephan Bortzmeyer (AFNIC). It remains to be seen if ROVER will, after DANE, become the next DNS approach to secure the net. The DANE WG which is seen as a top challenger to certificates and TLS (and was not welcome to parts of the TLS community when it started) did not meet in Paris. According to DANE Co-Chair Ondřej Surý the documents were close to IETF Last Call. During the Vancouver meeting re-chartering might be one issue on the agenda.

Another attempt at the WHOIS protocol – reloaded

No consensus could be reached on starting a Working Group on a new Whois protocol after a first BoF meeting for a Worthwile Extensible Internet Registration Data Service BoF in Taipei last year (see CENTR report on the first WEIRDS BoF). But the push to start an IETF WG originally started very much by the numbers registry organizations (RIRs, especially ARIN and RIPE) is increasing. The pressure is not only coming from the RIR community, but there are also commitments made by the Internet Corporation for Assigned Names and Numbers (ICANN) toward governments (“we work on a new Whois”) and by large name registries like VeriSign and Afilias to assist ICANN (see for example a clause in VeriSigns proposed follow-up contract for the .com-registry, advertised by Scott Hollenbeck from VeriSign).

Also new gTLD-registry hopefuls like Google (represented at the BoF session by Warren Kumari) seem to be very interested in a new Whois-protocol. He would use it, Kumari said during the BoF session.

The requirement document tabled by Murray Kucherawy (Cloudmark) in its fourth edition adds additional requirements: beside the support for international values, machine readability, a “well-established” and “extensible” data format and support for a minimum set of fields (with an overlapping core set for names and numbers), authentication either for the protocol or the underlying transport mechanism is mandated, and cache values have to be set (to prevent re-iterated requests). No change has been made with regard to “classes” of users of the new Whois (limited anonymous requests, requests for infrastructural/security reasons, privileged law enforcement requests).

The updated requirement document also makes a step to entangle names and numbers and concedes that names and numbers might end up to develop different specifications. It cautiously says:
“It is hoped that this work will also influence the development of requirements and specifications for domain name registries at some point in the future.”
Also the Post-BoF draft WG Charter (Version 7, by Andy Newton from ARIN) allows for a future split of work between the name and the number sphere.

> “Should the Working Group reach a point where it determines that the problem of producing a grand unified specification for both numbers and names appears to be intractable, it will be permitted to divide the problem into separate tasks and amend its milestones accordingly.”

Yet discussion on the need of the new Whois has not died down completely with some well-known arguments including the failure of IRIS, the lack of agreement about policy issues (including privacy issues) and the lack of incentives to deploy.

The latter might well be solved by ICANN mandating a potential new standard (see recommendations from the Whois Review Team and comments to them here) at least for their domain (gTLD-registries and registrars). Some pressure potentially could trickle down to ccTLDs, gTLD registries for example were asking if "ccTLD registries" should be completely excluded from the better Whois process. ICANN's Government Advisory Committee at ICANN for years has pushed for cleaner Whois Data (with the US supported by LEAs taking a center role) and the Whois Reviews have been set alongside other reviews ICANN is obliged to according to the Affirmation of Commitments.

During the public comment period of the review team there was a lot of support for a unified, new Whois policy (and potentially new data models), but there were also warnings against centralization of WHOIS data at ICANN. This would mean that ICANN would provide "a lookup services for name registration data it does not control", Andrew Sullivan (Shinkuro), WEIRDS BoF Co-Chair wrote to the Review Team. Also user groups like the ICANN At-large Advisory Committee warned against a clash of jurisdictions. The privacy concerns have meanwhile been added to the security section of the WEIRDS requirement document (beside denial of service attacks and redirections loops).
The WEIRDS milestones look ambitious with a a protocol for registry queries shipped to the IESG next year in June.

Choosing the best Web ID concepts - and discussions on browser security and privacy

WebID and federated access got a lot of attention at the IETF 83. Not only is there work ongoing in the Web RTC and OAuth, and the WebSEC working groups, the technical plenary looked into browser security issues. Additionally the competing concepts of OpenID (and OpenID Connect) and BrowserID were compared during a lunch panel attached to the Oauth WG meeting. The Internet Society which has been working for several years on the issue digital identity and privacy (check about ISOC participation in related W3C Wgs and Kantara initiative) in its traditional lunch covered the topic "digital identity".

Beauty Contest of Single Sign On Solutions

With regard to the different Single Sign On-concepts, OpenID currently in use at a number of large platform providers like Google, Yahoo or Microsoft and smaller providers more and more rely on identification via these services. For example newspaper providers in the US according to US lawyer Wendy Seltzer often just ask for Facebook Connect sign-in. Content providers or smaller services leverage the authentication of Facebook for their own service instead of setting up their own web authentication system. Users on the other hand benefit from the possibility of these single sign on concepts as the make getting on to a new platform faster. Once they have registered with Facebook, they can use their "Facebook identity" for a variety of services.

Given that Facebook Connect, and OpenID (and the various spin-offs developed by the various companies) and Oauth-variants are in place why did the Mozilla Foundation decide to not join forces, but develop their own alternative, Browser ID?

Ben Adiba from the Mozilla Foundation said Mozilla wanted to allow more user control and transparency, and had looked for a system not only large US platforms could act as Identity provider. Mozilla has chosen the users' email addresses as the unique identifier[1], making every email provider an identity provider. Users already did understand that an email address was a "slice of the identity" of users.

[1] While Other Web ID (including Online ID) concepts also used email, with security vulnerabilities detected last year by a group of researchers. See paper published soon at the IEEE conference, and related story.

When using BrowserID a user is asked for an email only, a notification request then is sent to the respective mailbox and after the user sent back the notification, in the future the browser can check with the browser-id enabled email provider before setting long-term or short-term certificates for the user.

The main architectural conceptual difference between the OnlineID and BrowserID pushed forward by Mozilla is to allow the user to not disclose his whereabouts on the Web to their identity provider. Instead of refering a user's request arriving at a third party site (relying party) back to the Identity provider (OpenID provider) the Browser authenticates the users sign in via his authenticated email with him. (check on the flow diagrams of OpenID 2.0 and BrowserID below).

Adiba said, people would be shocked if they would experience that the reception desk of their next hotel would call back to the responsible authorities to ask if it was ok that the passport holder was there before processing the check-in. Yet that was what happened with OnlineID, he said, and in order to allow the user to prevent the ID providers to amass a lot o information about the users, Mozilla had decided to build the alternative.

It was not about only shifting trust to another place, as identity providers would not be approached for the credentials before the sign on onto another platform. Instead the BrowserID used the short-time credentials stored in the browser before. Depending on how long the lifetime for the certificates was set, there certainly was a problem with compromises, for example when a device was stolen or lost or when an account was hacked.

Because the certificates could be used from the browser, the concept certainly was adopted to real time communication. Eric Rescorla at the Web RTC workshop presented a generic model to rely on various ID providers, in order to authenticate a caller before the online call is put through.

Another feature that shall in the future help the users to control the sharing of data (also with the browser) is a button that will allow users to leave any services the logged on via the browser with one click. Users were regulary confused by the fact that they were still logged in to a third party services.

What is still lacking with regard to BrowserID is a solution for the “delegation” (or third party access to user resources) case that is addressed by Oauth (for example, granting a foto services to fetch private fotos from the user's site to perform a printing job). Adiba said Mozilla was looking into it and certainly wanted to reuse existing standards as much as it could.

Facebook which started out with their own Facebook Connect for the Single Sign on meanwhile has switched from proprietary-only to including OAuth (and using it for authentication, too, which Also competitors see as problematic) OpenID, too, in OpenID Connect, has married third party resource access and easy authorization (SSO) by integrating Oauth functionality to OpenID.

Challenges for browser security and the call for more transparency and privacy

With the race for better authentication ongoing, security in browsers still face a lot of security problems, as the participants at the technical plenary agreed. Transport Layer Security was only provided by one percent of all sites, a result of the complexity to use it, TLS author Eric Rescorla, acknowledged. Rescorla pointed to DANE and “storing the keys or hashings of the keys or some such thing in the DNS” as a potential alternative to X509.

Still https was not use enough and mixed content sites (putting http and https-content side by side) created problems. A rather bleak picture was drawn by Tom Lowenthal who said the certification system that relied on several hundred certification companies was broken and with a future move to only allow clean sites (with the protocol for HSTS nearly completed in the WebSEC Working Group) a lot of sites could go dark at once, especially if one large certificate authority controls the vast majority of certificates.

Several developers explained how vulnerabilities were created with the browser not in control by application engineers, but instead huge, diversified machines and new applications – often written without

considering new forms of attacks made possible by those. Ian Fette from Google explained this for the websocket standardization, where violations of network boundaries (internal network, outside network) had to be prevented by origin checks for requests and the mandatory exclusion for websocket requests to talk to something else than a websocket server. Also the embedding of the websocket concept into an environment where existing proxies did not implement necessary new security standards could result in attacks to “suck data out of a corporate network” or in “poisoning caches”.

While Fette recommended that the effects of bringing new applications into the wild had to be checked really closely (and existing infrastructure had to be protected against new attacks enabled by the new things) Jeff Hodges from Paypal said during the panel, that things were better than a few years ago. Because of competition in a “multi-browser” world and also because of the movement of browsers to open standards instead of proprietary plug-ins there was much more attention and sensibility for the developments. The attack surface, he said, was “arguably shrinking”.

Data minimization and transparency as a principle from the regulator was laid out to the participants of the Security Area Working Group meeting in Paris. British law professor Ian Walden (Queen Mary College) explained the core aspects of the EU data protection reform that, for example, pushes for an opt-in regime or cookies, something that some developers heavily criticized. Eric Rescorla (who had presented the RTC security concept involving cookies to allow easy authorization for cross-platform phone calls, for example) warned, such regulation was not implementable. Walden acknowledged that the draft proposal was not perfect, but recommended that standardization bodies should not only “click the box” when considering privacy issues of new standards. Instead to make privacy by design real, more guidelines should be given in the specs, he said, pointing or example to Cookie specs (which are W3C specs).

The IETF has been considering to put privacy into the security consideration part of RFCs, and in fact, the privacy board has been checking on draft standards they found to have potential effects with regard to privacy and data protection. Walden also explained the differences between privacy and data protection, the latter very much being a central European concept.

For the flow of the two main competitors see:
Browser ID

1. Client prepares an Authorization Request containing the desired request parameters.

   Client sends a request to the Authorization Server.

   Authorization Server authenticates the End-User.

   Authorization Server obtains the End-User Consent/Authorization.

   Authorization Server sends the End-User back to the Client with an Access Token and ID Token.

OpenID 2.0 (note that OpenID Connect takes on board Oauth for resource access)

The end user initiates authentication by presenting a User-Supplied Identifier to the Relying Party via their User-Agent.

1. After normalizing the User-Supplied Identifier, the Relying Party performs discovery on it and establishes the OP Endpoint URL that the end user uses for authentication. It should be noted that the User-Supplied Identifier may be an OP Identifier, as discussed in Section 7.3.1, which allows selection of a Claimed Identifier at the OP or for the protocol to proceed without a Claimed Identifier if something else useful is being done via an extension.

(optional) The Relying Party and the OP establish an association -- a shared secret established using Diffie-Hellman Key Exchange [RFC2631]. The OP uses an association to sign subsequent messages and the Relying Party to verify those messages; this removes the need for subsequent direct requests to verify the signature after each authentication request/response.

The Relying Party redirects the end user's User-Agent to the OP with an OpenID Authentication request.

The OP establishes whether the end user is authorized to perform OpenID Authentication and wishes to do so. The manner in which the end user authenticates to their OP and any policies surrounding such authentication is out of scope for this document.
The OP redirects the end user's User-Agent back to the Relying Party with either an assertion that authentication is approved or a message that authentication failed.

2. The Relying Party verifies the information received from the OP including checking the Return URL, verifying the discovered information, checking the nonce, and verifying the signature by using either the shared key established during the association or by sending a direct request to the OP.

Ipv6 – From Test flight to Take-Off
Ipv4 - A new Working group

During the technical plenary and also during a dialogue with local press – IETF press conferences seem to become rather regular exercises – the second edition of World Ipv6 Day was promoted by Leslie Daigle, CTO at the Internet Society (ISOC). After the test flight 2012, when over thousand companies had turned on Ipv6 for a day, 2012 Worldv6 Day (June, 6) was about a full launch. Content providers like Google, Yahoo and Microsoft have committed to not turn back from full Ipv4-IPv6 dual stack operation. Network providers (including Comcast, Time Warner Cable, Free France and Dutch provider XS4All) participating in the v6 launch day commit to offering every new customer v6-service and also to bring v6 traffic in their network to 2 percent.

“This time it's for real” is the catch-word for the event that will be further promoted by ISOC. The organization hoped to get additional network and content providers on board. For example there is no network provider currently involved. Observers also pointed to the fact that so far there is a focus on US-based providers participating. From Europe beside Free (with its 2,5 Million Ipv6 customers) and the smaller Dutch XS4All there are only national research networks from several EU countries listed. ISOC's US base might be one reason, one operator said to this reporter. In fact large EU telcos and also mobile operators currently are absent from the list.

While ISOC beats the drum for v6launch day, there is growing displeasure of parts of the IETF community about v6Ops or Softwire WG agendas crammed with Ipv4 life-saving or life-prolonging specifications. Operators still needed to support Ipv4 content and still had to support existing devices, IAB Chair Bernard Aboba (Microsoft) said during the press conference, the IETF had to support these efforts as well. There is an expectation that the IESG might establish a v4exit WG to bundle the respective efforts and clean up the agendas of v6 oriented Working Groups.

With regard to IESG WG creations an idea to establish a email-related WG, because of related work in various places at the IETF.

Working Groups

DNSOP WG

The Domain Name Operation WG discussed the the reverse naming convention for CIDR blocks (see “Just put in in the DNS” above), the naming convention was said to potentially be useful not only for ROVER, but also for other applications. The issue was accepted to need IETF work in the room as was the issue on long or short Time to live (TTL) values for DNS servers.

Vasileios Pappas (IBM Research), Eric Osterweil (VeriSign) together with a group of authors revived a draft proposal to allow longer TTLs for servers higher up in the DNS tree (TLD servers) in order to rise availability during DDOS attacks. DDOS-related effects could be reduced by 70 percent if using a TTL of seven days for NS+A/AAAA, DNSKEY, DS records of the root zone (for more numbers from a lab experiment see the presentation). The proposal by Pappas and Osterweil had already been presented at the recent OARC meeting (see here), together with other options to secure higher availability. During the OARC Duane Wessels pointed also to the known downside of long TTLs:

> "The downside to long TTLs, of course, is the need for better planning and longer schedules when changing NS records or their IP addresses. In particular, if your zone itself is the target of an attack, short TTLs may give you useful flexibility in responding to it."

Another argument for short TTLs for NS and DS records was DNSSEC and potential failures in key management. "Especially, impact of KSK rollover failure is huge and its recovery requires cooperation of child/parent zone and full resolver operators", said Yoshiro Yoneya from the Japanese ccTLD registry. Short TTLs allowed to mitigate failures caused by inconsistencies between parent and child zone. Yoneya presented several options for discussion, including having short TTLs only on DS records, or on both DS and NS records. Shortening the TTLs only during changes (Key rollovers) would be another option, but needed a lot of attention. Yoneya recommends having "DNSSEC KSK rollover failure recovery practices".

While some participants were hesitant with regard to the revival of the long-TTL draft discussion, the TTL issue will be discussed further in the WG.

Other drafts pointed to by the Co-Chairs (needing DNS review) were:

IANA Reserved IPv4 Prefix for Shared Address Space
Special-Use Domain Names
Update to SPF (SPFBIS) (because of potential for amplification attacks)

Simple Cloud Identity Management (SCIM) BoF, it is not what you think

During the BoF on Simple Clould Identity Management (SCIM) engineers from Cisco, UnboundID (a cloud-service company from Austin that just raised 12.5 Mio venture capital) and Phil Hunt (blogging on independentid.com) presented ideas about how to provision identity data into new (Cloud) services in bulk form. The core target of a potential IETF standard (obviously there has been related work at OASIS) according to the presenters are companies or organizations that want to populate new services (in house or in the cloud) with existing user (employee, member) data. Existing protocols (including SAML, LDAP, OAUTH or SPML (OASIS) were said to be not sufficient, but potential hook-ups. Quite astonishing for the observer, the draft charter does not include the limitation of the work for companies and organizations owning a large set o user data. The draft charter on the non-WG mailing list states:

> "The Simple Cloud Identity Management (SCIM) specification is designed to make managing user identity lifecycle in cloud based applications and services easier."

Ondřej Surý from czNic proposed to think about EPP for provising, the Area Director, Pete Resnick, concluded by describing the problem to be solved as "movement and management of identities", he did not see relation to the "cloud" and he recommended to get down to what the entities and the data model were.

DNSEXT

The DNSEXT WG is about to be closed, both Co-Chairs said, Paris would be the last f2f meeting. Before closure several documents are expected to be shipped, including:

1. A new query type (and general update) for DNS Incremental Zone Transfer Protocol (IXFR), IXFR-ONLY. The document would replace the existing RFC 1995.

2. An update to RFC6195 (RFC6195BIS)for streamlining of IANA resource record type allocation procedure.

Yet, just close to closure, there also was a new document, again on resource records, placed on the DNSEXT table, by Ed Lewis (NeuStar) which shall document IANA registered resource code types that have not been well documented. Some, according to Lewis are abandoned efforts. The list of Rrs in the Lewis proposal which was welcomed on the mailing list as highly useful is below:

TYPE	Value and meaning	Reference
EID	31 Endpoint Identifier	[Patton]
NIMLOC	32 Nimrod Locator	[Patton]
SINK	40 SINK	[Eastlake]
NINFO	56 NINFO	[Reid]
RKEY	57 RKEY	[Reid]
TALINK	58 Trust Anchor LINK	[Wijngaards]
CDS	59 Child DS	[Barwood]
URI	256 URI	[Faltstrom]
CAA	257 Certification Authority Authorization	[Hallam-Baker]
TA	32768 DNSSEC Trust Authorities	[Weiler]

Yet another draft proposal for a new resource record was presented to the WG by Ted Hardie. "The Reachability Method (RM) DNS Resource Record" shall establish a record for "providing adjunct reachability methods for network hosts or resource which are only accessible within limited reachability domains" (like enterprise domains). There is a IPR declaration from Google for this draft. The reachability method was a mechanism that provided access to a limited reachability domain from outside the usual administrative boundary, according to the draft.

Smaller reply size with denial of existence answer;

- Introduce Wildcard Flag bit;
- Opt-Out for unhashed names;
- One type of denial of existence

Security Area Meeting: New crypto-algorithm not that urgent?

~~THE NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY (NIST) WILL SELECT A NEW CRYPTO ALGORITHM LATER IN SUMMER, TIM POLK FROM NIST, SAID AT THE 83. INTERNET ENGINEERING TASK FORCE (IETF) IN PARIS LAST WEEK. THE SO CALLED SECURE HASH ALGORITHM (SHA) IS USED TO COMPUTE SMALL NUMBERS FROM WHOLE DATA SETS – A SO CALLED FINGERPRINT - TO ALLOW CHECKING IF THE DATA HAS BEEN HAMPERED WITH. AS THE CURRENT SHA ALGORITHMS, SHA1 AND SHA2 WERE THOUGHT BE BE VULNERABLE TO ATTACKS DOCUMENTED STARTING IN 2005, NIST IN 2008 STARTED A COMPETITION THAT RESULTED IN 64 PROPOSALS BEING SENT, OF WHICH FIVE HAVE BEEN SELECTED AS FINALISTS. WHILE POLK SAID THE RESULTS OF THE CONTEST WERE VERY GOOD, NOT ONLY BECAUSE THE NIST NOW COULD CHOSE BETWEEN FIVE EXCELLENT, AND STRONG ALGORITHMS, THE COMPETITION ALSO HAD AN UNEXPECTED RESULT. „WE HAD THOUGHT THAT SHA2 WOULD BE ON ITS LAST LEGS BY NOW, BUT INSTEAD IT HAS PROVEN TO BE STRONG," HE SAID. INSTEAD OF LOOKING OR THE~~

~~SELECTED CANDIDATE AS A SUCCESSOR TO SHA2 NIST THEREFORE WAS CONSIDERING TO PROPAGATE IT AS COMPLEMENT, NOT THE LEAST BECAUSE SHA2 IN SOME FEATURES STILL WAS BETTER THAN THE NEW CANDIDATES. RUSS HOUSLEY, CHAIR OF THE IETF, WARNED NOT TO OVERLOOK THAT MANY ADMINISTRATORS HAD NOT EVEN CHANGED FROM SHA1 TO SHA2, TO INTRODUCE AN ADDITIONAL CRYPTO STANDARD THEREFORE COULD BE PREMATURE AS MIGRATION OR SUPPORT FOR DIFFERENT STANDARDS AT THE SAME TIME WAS COSTLY FOR THE USERS.~~

IETF News

Anti-Trust Education and Anti-Trust pointers in IETF "Note Well"

The IETF will task a design team to draft educational material on anti-trust regulations, participants of the 83 meeting of the organization decided during an Anti-Trust-BoF. The Chairs of the currently 120 working groups which develop new Internet protocol standards will have to be made aware about how to prevent potential Anti-Trust-like behavior (including talking about prices), the IETF legal counsel, Jorge Contreras, and Scott Bradner (Harvard University) explained during a well attended BoF meeting. The discussion on an Anti-Trust strategy was triggered by a court case brought against the European Telecommunications Standards Institute (ETSI) last year. ETSI had been dragged into a case including several hardware providers (Ericsson, Alcatel-Lucent) which were sued by Trueposition before a US Court for conspiracy in anti-competitive behavior.

The fact that engineers participate at the IETF as individuals and not as representatives of their companies would not effectively shield the SDO from cases brought against it if participants would talk on pricing during the sessions, Contreras said. The participants nevertheless decided against a full-fledged Anti-Trust policy and chose additional wording in its „note well" announcements that are used as pointers to the IPR policy and other guidelines and norms of the organization at every meeting. Even a policy would be no guarantee against law suits as the True Position vs ETSI case had shown, because ETSI had had an Anti-Trust Policy. In the US one could be sued for anything, Contreras said, and a policy would result only in being sued differently.

Subpoenas, Blue sheets, IRTF Logo, Time zone database

Other news from the "legal front" include the start to publish subpoenas received by the IETF Trust, Marshal Eubanks included one subpoena just received in his report. The Trust had after consulting Legal Counsel decided that it could in the future publish all subpoenas received (if there will be a dedicated space for it, is unclear). Also in order to spare some digging into "old blue sheet"-records (which allow to check who attended the meetings of the IETF WGs) the IETF will publish the blue sheets collected. In order to prevent Spam, participants will in the future only be asked for their name (not for their email addresses). With regard to IPR, there were two issues pending, one included the IRTF logo which after being "lost" in the transition from CNRI to the Trust and current secretariat function. Eubanks said that the Logo had been recreated by IRTF Chair Lars Eggert, yet Eggert did not yet sign a contract on the recreated Logo with the IETF Trust. IETF Chair Russ Housley reported during the Administrative Plenary also that the trial about the time zone database came to a lucky end, after astrology software company "Astrolabe" dropped the suit and apologized that it had a wrong understanding of the law. US Civil Rights Organization Electronic Frontier Foundation had stepped in to represent the IETF pro bono in the suit.

Correspondence with ICANN on new gTLDs

The IAB has sent letter to ICANN warning against potential conflicts from different interpretations about the syntax of DNS names (RFC 1123). After ICANN sent back questions pointing to existing IDN ccTLDs

that did not match the proposed IAB “most conservative”-rule, the IAB reacted with an answer to ICANN recommending to retain what TLDs are there. It was up to the ICANN evaluation panels to decide whether applications could potentially result in instabilities. The IAB while preparing for clarifications in the RFCs according to the experts, appealed to ICANN to be conservative with regard to accepting respective applications.

Budget and future meetings

Bob Hinden, Chair of the The IETF Administrative Oversight Committee (IAOC), reported about a balanced 2011 budget. The organization despite a less-well attended meeting in Taipei (only 923 participants) had revenue of 3,318 million US Dollar in 2011 (against budgeted 3,317). With lower than budgeted expenditures (4,872 instead of 5,004 Million US Dollar) the contribution of ISOC could be smaller than planned (1,923 instead of 2,325 Million US Dollar). For 2012 with a new RFC editor hired now expenditures are expected to climb to 5,408 Million US Dollar (against smaller revenues of 3,246 Million US Dollar). After Cisco stepped in as a last-minute host for Paris, all meetings in 2012 will be hosted (Vancouver by Google, Atlanta by the North American Cable Industry), but hosts are still looked for for 2013 (Orlando, Berlin, Vancouver again), and 2014 (London, Toronto, Honolulu). Obviously IETF meetings in Europe while drawing larger crowds (the Paris meeting came close to 1400 attendees) are more difficult with regard to finding hosts. Meetings in Asia suffer from difficulties to find less expensive venues according to the IETF leadership, so the next meeting planned for Asia is only in 2015 (Yokohama).

The IETF got some nice press attention during the Paris meeting for “tuning the network of the Concorde conference hotel” by – beside a lot of other “tricks or black magic” - re-channeling 321 access points.

IETF 84 is meeting in Vancouver, July