Report of IETF 84 Vancouver

30 July – 3 Aug, 2012

by **Monika Ermert**

for the CENTR secretariat

Table of Contents

Highlights

Is WEIRDs weird?

Progress at the „Web Extensible Internet Registration Data Service“ (weirds) working group, which had their first full meeting in Vancouver, might be slower than the number registries community or ICANN staff and large gTLD registry members are hoping. The meeting in Vancouver under the chairmanship of Olaf Kolkman (NLnet Labs) and Murray Kucherawy (Cloudmark) discussed the issues of service differentiation/authentication, versioning, bootstrapping, URI templates and the data object model without coming to conclusions or even first hums on the issues (see minutes).

Andy Newton from the North American Number Registry ARIN and author of the core WEIRDS candidate proposal reminded the group about the commitment to not allow the more complicated domain names stuff to slow down the new WHOIS standardization advanced by the number registries community. But this might just be what will happen. For the five RIRs, consensus on the future protocol might be easier than for the rather diverse domain name community that integrates:

- gTLD versus ccTLD registries,
- registries with strong authentication, validation and access requests from their authorities versus registries with strong data protection legislation,
- registries with a strong interest to solve internationalization of the Whois versus registries that see no big need to change port 43 Whois at all.

A study prepared by a group of developers from CNNIC, Nominet, .br and ICANN staff illustrated the diversity in data objects nomenclatures in 124 registries (raw data of the inventory here). According to the presentation given by Ning Kong (CNNIC) domain whois data collected contained 68 data elements using a total of 550 different labels. Also, there are a total of 392 not classified extensions for domain whois data, with the largest number of extension in use for registrants (41 TLDs).

One of the (several) hot issues in WEIRDS is service differentiation and the authentication (and authorization?) thought conditional for it. Sean Shen (CNNIC) in his presentation listed four basic levels for authentication: public (open to all), group (membership open), authorized (open to particular people with permit) and private (open to non than registrant).

During the debate Peter Koch, DENIC, warned to not to get involved in what participants saw more as implementation issues, like access policy. Koch pointed to obligations for some registries to lie about data they could provide because of privacy issues. With regard to the potential protocols available for authentication – either pure client based or involving a third party – were „digest authentication“ (RFC 2617), Oauth 2.0 or Transport Layer Security (TLS). Cookie-based authorization, which Alex Mayrhofer from nic.at said might be considered, was said by Marcos Sanz from DENIC to violate the rest-ful paradigm. Other attendants moaned that the rest-ful paradigm had been started to be abandoned some time ago. Making a protocol like Oauth mandatory to implement might only lead registries to ignore the protocol, participants warned. URI templates, presented by Kucherawy, were seen as unnecessarily adding complexity.

For bootstrapping John Levine, president of the Coalition Against Unsolicited Commercial Email (CAUCE) proposed inter alia to use TLD.weirds.arpa in the DNS for the location of the new services. Former IAB member John Klensin warned against „crossing the ICANN-IANA border several times“ from pushing the ICANN protocol down to the IETF to having the IETF requesting delegation of the new dedicated (and policy related) zone to IANA/ICANN.

Crossing IETF – ICANN boundaries / a decade of Whois accuracy and enforcement discussions

The WEIRDS work might be seen as a case of the shifted boundaries between the I-organizations (IETF, ICANN, IANA), or at least as a case where the different organisations seem to be at odds over their respective competency in guarding the DNS.

ICANN and the IETF have both started to work on a Whois successor protocol. For the IETF a problem statement first arrived from the numbers community, which seems to be driven by the growing access requests to their Whois data.

ICANN on the other hand has discussed over Whois accuracy and enforcement for more than ten years and has been called for urgent action many times in the meetings of the Governmental Advisory Committee, with especially US- UK and Australian (less often Dutch) law enforcement being brought in as re-enforcement by the respective governments. Whois was seen as important enough by the US administration to dedicate a special review to it according to the Affirmation of Commitment Review system (AoC).

The first report is just out and its recommendations have been welcomed by the Intellectual Property Constituency, the Business Constituency trademark groups (INTA), but was questioned by privacy and technical experts for flaws and neglectance of challenges, for example, on different privacy legislation and the fallacies of a potential central Whois registry services provided for by ICANN.

What has led to some criticism from DNS experts is that ICANN's Security and Stability Advisory Committee (SSAC) now has presented its own work on a data model in report No 54 already, ahead of the ongoing work in the IETF. Also ICANN has published a request for comment for a new Whois testbed service, with the IETF WEIRDS group only just starting to work. NLnet Labs is one of those bidding: Olaf Kolkman (NLnet Labs), co-chair of the WEIRDS working group, said testbed providers would have to work on the basis of the specification delivered by the IETF WG.

The ICANN data model will in any case be fed into the upcoming work of the new IETF design team for WEIRDS for the data objects, with so far three of eleven members participating in both the IETF and ICANN data object efforts (ICANN staff members Francisco Arias and Steve Sheng and SSAC Vice-Chair Jim Galvin from Afilias).

The composition of the IETF data object design team has resulted in some questions over balance. So far out of its members there are

- 5 RIR
- 2 ICANN gTLDs
- 2 ICANN staff
- 1 ccTLD
- 1 from ISC.

The Chairs will reconsider this distribution, according to one of the recent messages on the mailing list.

What complicates the relation of ICANN with the IETF is the new IANA contract awarded by the NTIA to ICANN. The contract has additional “touch points“ allowing interventions by the US administration also with regard to protocols, as one expert source said. While administrators are assuring that the IANA contract was “business as usual“ the contract would allow stricter oversight over protocol work, especially where policy aspects are involved. According to Keith Davidson, just elected new ISOC Board of Trustee member from ISOC NZ, there are renewed deliberations about a potential split of IANA functions, with each functions being served by staff reporting to the relevant community. This would not result in outside oversight change, though.

A proposal to potentially use weirds.arpa in DNS for a bootstrapping mechanism for information about the whois concept in use and/or a trust anchor might also involve IANA, which on the other hand is requested to stay away from policy decision, according to the new IANA contract. ICANN staff dedicated to the IANA work according to the new contract has to keep away from policy work, on the other hand, hence the warnings about caution where the boundaries between the organisations are crossed.

Current WEIRDS candidates are:
https://datatracker.ietf.org/doc/draft-designteam-weirds-using-http/
https://datatracker.ietf.org/doc/draft-kucherawy-weirds-requirements/
https://datatracker.ietf.org/doc/draft-sheng-weirds-icann-rws-dnrd/
https://datatracker.ietf.org/doc/draft-hollenbeck-dnrd-ap-query/
https://datatracker.ietf.org/doc/draft-newton-et-al-weirds-rir-query/

Measuring and routing around filtering – politics@IETF?

A research team at the Cooperative Association for Internet Data Analysis (CAIDA) in cooperation with universities in Italy and the European IP Address Registry provider, RIPE, creatively used the measurement of malware background noise in the networks of Egypt and Libya to get more insights about filtering methods. Libya's approach had been much more sophisticated, said CAIDA-researcher Alberto Dainotti during the presentation of his findings at the Internet Research Task Force session in Vancouver. The Internet Research Task Force has awarded one of the annual Applied Networking Research Prizes to Dainotti.

The malware traffic analysis had revealed that in addition to taking down routes, Libyan authorities had also more focused packet filtering and jamming of satellite traffic. Dainotti said that ongoing work was aiming to establish an automatic alarm system that would send an alert notice when malware traffic background noise signals went down, be it as a result from censoring attacks or natural disasters.

In a Bar BoF session potential remedies for filtering or for kill-switch attacks were proposed by Dutch academic Johan Pouwelse (TU Delft) asking for standardization of several components for public interest microblogging services. The alternative to “Twitter” (yet without business interest and therefore less to be affected by gag orders) would allow for anonymous, secured and potentially non-Internet-connectivity dependent posts. The topic was welcomed by IETF participants, but a gap analysis was necessary to define what protocol work was still necessary for “censor-free media“, the BoF agreed.

Pouwelse, who teaches at the University of Technology in Delft and leads a p2p research group that has developed the tribler p2p platform, had been highly political in his call to the IETF. “Freedom to spread information“, he wrote in the draft presented, “is under active attack in various corners of the Internet.“ Control mechanisms developed were more and more sophisticated and complex. “The age of cyber supression is upon us and we need to act,“ he warned the engineers, and requested that the “forces favoring freedom“ should re-group under a single initiative in order to “impact the live of millions“.

The Power of Netistan

A high-level view on the state of “Netistan“ finally was given by Google's Chief Technology Advocate Michael Jones. Jones said that being the state with the largest population on earth, its netizens still had to become aware of their power. “Governments need us more than we need them“, Jones said. Outages and the reversal of the SOPA, PIPA legislative efforts (and potentially the ACTA failure in Europe) could nevertheless be seen as one instance in which regions of Netistan used their power somehow.

Acting globally as other Internet users Google is exposed to all sorts of national legislation and sometimes had to make choices about how to adapt. With regard to privacy, Jones, who is one of the creators of Google Earth, described the influence national legislation could have on its services.

With regard to Google Street View, for example, German data protection concerns were taken up for the services in general. In some instances services, on the other hand, were not offered, for example, there was no official Gmail offer in China in order to avoid being legally obliged to hand over sensitive data. When asked with regard to a listing of companies as enemies and friends of “Netistan“ Jones put the US, Australia and France to the countries “on probation“ because of attempts to establish domain name blocking and three-strikes regimes, respectively.

Wanted: new IETF Chair

IETF Chair Russ Housley has announced he will not run for another term of office. So for next year the IETF has to look for a new Chair. Housley is founder of his own security consultancy, Vigil Security LLC and has worked, for example, on security standards for the NSA. During his IETF tenure the organization has – due to visa issues – started to hold more meetings outside of the US than ever.

Working Groups

DNSOP

The DNS Operations working group had another discussion on the publication of the DNSSEC key timing document. With a follow-up version sitting on the shelf and waiting to be advanced, some recommended to kill the older version while several participants pushed for the document to be shipped quickly, which finally the WG decided to go for.

New proposals presented were related to omniscient AS112 servers, as well as a potential standard for monitoring statistics of DNS servers and a the potential use of DNSSEC for automated updates between child and parent servers. With Omniscient AS112 Warren Kumari (Google) proposes “mechanisms by which AS112 name servers could answer authoritatively for all possible zones, reducing the add/drop problem to one of delegation within the DNS without operational impact on the servers themselves“. There was no agreement if the concept could work in that form at all, discussion will be continued on the list.

The DNS Server Statistics Management Information Base (MIB), proposed inter alia by members of the Canadian ccTLD CIRA, put forward standard management protocol to receive statistical (read only) data from anycast nodes of various DNS operators about one's zone. Alongside the CIRA proposal there is another one from ISC that wants to include the function in BIND 10.

Another new document presented in the WG is an attempt to promote leveraging DNSSEC for NS, DS record and glue updates in the parent zone. The idea is to avoid out-of-band authentication necessary for updates so far. While “fixing the broken child-parent-relationships“ was something welcomed by many, there were also concerns that registrars or resellers sitting between registrants and registries must not be bypassed. Despite some participants pushing for a hum at the session to avoid delay, the WG chairs decided that more discussion had to take place first on the list.

Dane

More work leveraging DNSSEC was presented in DANE. With the core specification being just published as a proposed standards track document and the milestones of DANE reached, the working group discussed closure, hiatus and potential re-chartering. The group could go to hiatus, Paul Hoffman, co-author of the just published core RFC and of several new proposals said. Yet given the new work presented by Hoffman and others at the Vancouver meeting, DNSEXT ex-chair Andrew Sullivan (Dyn) warned to let the ongoing work piling up instead of having the WG to re-charter and address the issues.

New work presented covers the use of the DANE concept for S/MIME and for SMTP and Mail User Agents.

DNSSEC on S/MIME, according to the draft document of Paul Hoffman (VPN Consortium) and Jakob Schlyter (Kirei), provides for signing of a certificate for a user in his zone. The draft so far just heavily references the DANE specification, in general it allows “to associate an S/MIME user's certificate with the intended domain name“ by relying to secure DNS. Effects of DNSSEC on SRV-based auto-configuration and TLS certificate verification in IMAP, POP3 and message submission is considered by Tom Finch's (University of Cambridge) draft. DANE according to Finch could be used for stronger authentication of server TLS certificates. Finally, Finch in a separate draft proposes how TLSA records in the secure DNS can be used for SMTP server authentication.

Discussion was about questions like what would be secured in the DANE/SMTP proposal, communication or identity of the sender (for the MUA draft), how security semantics should look like (SMTP draft) and if normalization for the left side of mail addresses would be a must for the S/MIME concept plus how would privacy issues be dealt with. Instead of putting personal contact information in the DNS, hash of left-hand-side to avoid tree-walking was an option, according to Richard Barnes (BBN).

Once more challenged was the assumption of another draft under consideration, that presents DNSSEC as a potential tool to decide “am I connecting to the right server“ (see draft Miller/St. Andre, both Cisco). While they argue that for virtual hosting DNSSEC would help with secure delegation and Dane with identity verification, several participants disagreed.

Opinions considerably vary about the nature of security that secure DNS can provide. For some DNSSEC only ensures that something is coming from the zone it purports to be coming from, but not about the legitimacy of it being in the zone. There would always be a certain gap with regard to the security of the provisioning, Richard Barnes said. The problem could be worse with regard to the mail case.

Barnes also committed to work on the more general problem of indirection present for all DANE-based applications, the WG agreed. Once this would be solved, DANE steps for the various use cases would be straightforward.

After the discussion there was consensus that the WG should re-charter and possibly start with the indirection document.

RTC Web and RMCAT (RTP Media Congestion Avoidance Techniques) BOF

With the RTC Web WG pretty advanced - and two mandatory open source, royalty free audio codecs being picked by the group at the Vancouver meeting, Skype/Microsoft played the evil empire when, just after the IETF meeting, published their own RTC API not compatible with the work done. Skype engineers explained their move by pointing to a lack in ubiquitous deployability and also a lack to honour “key web tenets“, like stateless interactions.

Skype's/Microsoft's judgement is pretty sharp, saying that the RTC Web proposal **“***shows no signs of offering real world interoperability with existing VoIP phones, and mobile phones, from behind firewalls and across routers and instead focuses on video communication between web browsers under ideal conditions. It does not allow an application to control how media is transmitted on the network*.“

Equally sharp answers were fired back from participants of the WG (and the WG Chairs) claiming that there had in fact been consensus “to stick to a higher-level API rather than trying to move the level of the API downwards towards a 'low level API'.“ (Harald Alvestrand, Google). A somewhat more detailed analysis of the claims made by Microsoft about the issues they saw with the prepared standard was quickly delivered by Eric Rescorla (who is authoring several security related RTCWeb drafts).

Rescorla challenged, for example, Microsoft's assertions that statelessness was still a basic web tenet or that the RTC Web standard (implemented by Chrome and Mozilla) would not interoperate with VoIP applications. In Rescorla's view “while it's clear that the Microsoft proposal is a lot more work for the application developer; it's a lot less clear that it's sufficiently more powerful to justify that additional complexity.“

Martin Dürst from the W3C also questioned that the Microsoft proposal was a W3C submission, as it had not been listed by the W3C. Yet on Microsoft's html5labs page CU-RTC *is edited as a submission to the W3C, with IAB-Chair Bernard Aboba from Microsoft being noted as the core author*.

Work on one issue mentioned as essential by Microsoft in its blog post on CU-RTC, namely congestion management, just started at the IETF in Vancouver in the RTP Media Congestion Avoidance Techniques BoF. One concrete proposal presented by Harald Alvestrand (Google) devises a concept allowing the receiving end to estimate and control bandwidth rate and then transmit estimates to the sender. Having the estimates, the sender can choose to transmit at any rate that is smaller than bandwidth rate.

HOMENET and addressing issues (just a glimpse)

The homenet Wg has been touching DNS issues for some time and in Vancouver they discussed what approach to take with regard to naming inside the home network, while at the same time allowing addressability of the single nodes in the home network from the outside world. The DNS community has been called to pay attention to the development and to check on the documents.

In general, there are heavy discussions ongoing for some time about how differentiation between local and global names might be made and also how to avoid user confusion about these differences (e.g. potentially different resolving for one and the same name under different settings).
Now, at least three different proposals on a naming architecture (beside the homenet numbering drafts) that will even allow addressing the local names from outside are on the homenet table (no WG documents yet, though).

France Telecom authors Philippe Lemordant and Daniel Migault and Wouter Cloetens (SoftAtHome) state that home networks have grown (and will continue to grow) into a large set of applications, objects or devices managed by the CPE. As IP numbers while with IPv6 available in abundance are just not easy to use for addressing these applications and devices, naming conventions were necessary.

Lemordan/Migault/Cloetens filed two documents containing two potential solutions for a naming architecture. In IPv6 Home Network Naming Delegation Architecture they propose that the CPE “acts as the DNS authoritative server of the home network also called the delegated DNS server. The Naming Delegation is configured between the Delegated DNS Server and the Delegating DNS Server managed by the ISP”. The use case considered here was the End User who subscribes for a specific delegated domain with his ISP to establish his home network. The ISP's infrastructure would protect the CPE from heavy load.

In the other draft document the three authors describe a Front End delegating DNS server as the target for all queries. The user's CPE on the other hand would act as „Back End Network of delegated DNS servers“.

> *„All DNS queries for any Home Network are addressed to the Delegating Front End Server. The response is expected to be stored on a CPE, and the Front End Delegating DNS Server sends a DNS Query to that CPE before answering to the initial DNS query. The negotiation between the CPE and the ISP is using DHCP options. This document provides options so Front End Delegating and the Delegated DNS Servers configure their respective zone files and so that CPEs restrict access and protect themselves from unauthorized DNS queries.”*

„The Naming Delegation is requested by the CPE. The CPE DHCP Client and the ISP DHCP Server exchange DHCP Options to properly set the Naming Delegation. More specifically, the CPE DHCP Client (resp. the ISP DHCP Server) configures the DNS(SEC) Zones of the Delegated DNS Server (resp. Delegating DNS Server). For the Delegating DNS Server, the necessary pieces of information required to set the Naming Delegation are the IP address of the Delegated DNS Server, and if DNSSEC is used, the Delegation of Signing Information. For the Delegated DNS Server, the necessary information is the Delegated Domain associated to the Home Network."

The US answer (Kerry Lynn, IEEE, and Dan Sturek, Grid2home) to the naming architecture problem for homenet is "extended multicast DNS" (xmDNS). XmDNS would, according to the authors, "support multi-hop LANs that forward multicast packets but do not provide a unicast DNS service." Like mDNS, xmDNS wants to designate "a portion of the DNS name space to apply to the site-local network and specifies rules for its use".

In the ongoing discussion on the homenet mailing list, Andrew Sullivan (Dyn) stated "If we're serious about systems inside our target networks being somehow addressable from the global Internet, then it seems to me a mechanism to integrate easily to the global DNS without requiring major lock-in (or major upgrade of DHCP) is going to be something we'll have to tackle."

HTTPBIS re-chartering and taking off for HTTP 2.0

HTTPBis is about to re-charter to work – after having finalized HTTPBIS- on HTTP 2.0. A lengthy discussion took place in the second of two sessions of the HTTPBis WG about the method to advance. There seems to be sound consensus to start from the Speedy proposal (instead of clean slate) and pick from other proposed documents or ideas as the work advances. The documents to be drafted as listed by the Chair Mark Nottingham are

- a document (or set of documents) that is suitable to supersede RFC 2616 as the definition of HTTP/1.1 and move RFC 2817 to Historic status
- a document cataloguing the security properties of HTTP/1.1
- a document (or set of documents) that specifies HTTP/2.0, an improved binding of HTTP's semantics to an underlying transport.

Expectations for HTTP 2.0 are that it will improve perceived latency in most cases over HTTP 1.1, that it will not require multiple connections to a sever to enable parallelism, that it will retain the semantics of HTTP 1.1 (HTTP methods, status codes, URIs, and where appropriate, header fields), that the interacting of HTTP2.0 with HTTP 1.1 will be clearly defined and that new extensibility points an policy also will be clearly identified. Congestion control again is something the WG will discuss.

A very long debate occurred during the HTTP 1.1 and the HTTP 2.0 sessions on security considerations, especially the question if TLS should become mandatory. There seems to be no consensus on that, alternative proposals mentioned "optimistic TLS" (no checking of certificates).

Finally another topic under review, but with no consensus reached yet, is a new HTTP status code for when access to resources are denied "due to legal demands". This would allow for greater transparency on things like DMCA take down requests. Potentially geographic granularity might be put in place, explanations on where the legal demand came from are discussed for the code. The WG has no consensus on going forward or killing the proposal.

IETF/IAB News

Open Internet endowment

ISOC president Lynn St. Amour and IETF Chair Russ Housley announced a "family launch" of the Open Internet Endowment.
With contributions to be collected first from IETF members and later in the year the general public, the endowment is planned to become a new additional funding resource to secure the IETF budget which still is heavily subsidized by the ISOC (ISOC paying around 2 out of the 5 Million US Dollar IETF budget).
The endowment, while being primarily used for the IETF, also allows its board to sponsor other "open Internet" initiatives if the sponsor had not dedicated his contributions exclusively to the IETF. While there were fifty-some contributions over the first week, participants pointed to technical problems with credit card payment on the website.
There were also questions about the choice of an endowment (as opposed to a foundation) and potential consequences for sponsors from countries outside the US (tax deduction).

In the attempt to establish yet another new source of funding, the IETF for the first time has arranged its mini-exhibition space "Bits & Bites". According to the IAOC/IASA report, the first Bits & Bits income for the IETF was 58.000 US Dollar.

The IETF has not been able to get a bid for an IETF meeting agenda and scheduling tool and has extended its request for proposals on this to August, 31.

Open Internet Standards

With the declaration on a "Modern Global Standard Paradigm" the IETF, jointly with the with the World Wide Web Consortium (W3C) and IEEE, addressed to the International Telecommunication Union (ITU). The statement will be presented at the ITU World Telecom Standardization Assembly (WTSA) in November. The three organizations by signing this document in the next weeks agreed to commit to the basic principles for the development of open standards.

The high-level principles included are due process and appeals process, consensus and balance, transparency, openness and access to everybody. Yet the text still has to be worked on, Housley acknowledged during the plenary meeting in Vancouver.

In the version presented by Housley the standardization bodies will commit to *Fair, Reasonable and Non-Discriminatory (FRAND) license model. Yet after several complaints that in the IETF currently there was no obligation to license IETF standards according to FRAND, a* slightly changed version *was published that (possibly) allows for a little more flexibility in what licenses can be used.*

The non-involvement of the IETF community in developing the paradigm was also criticized by some.

IAB preparing for a Security Program

The Internet Architecture Board works on a series of long term programs. With the privacy work, for example, now rather developed – not only are the respective documents close to finalizing, but privacy considerations have in fact been raised in many working groups at the Vancouver meeting – the IAB seems to look into establishing an additional security program. The IAB will only decide on this over its next sessionsk, according to Hannes Tschofennig (Nokia Siemens Networks), co-author of the privacy and now also of the security documents.

The next IETF meeting will take place November 4-9, 2012, in Atlanta, USA.