Report of IETF 82 Taipei

13 – 18 Nov, 2011

by **Monika Ermert**

for the CENTR secretariat

Table of Contents

Highlights

Do we need DNSSEC for DANE and other DNS, do we need DANE for WEBSEC?

The DANE Working Group which started working to put transport layer security (TLS) information into a DNSSEC secured DNS over a year ago is having considerable discussion now about use cases without DNSSEC.

The overall goal of „DNS-based Authentication of Named Entities“ (DANE) according to the specification is to use the DNS and DNSSEC to provide additional information about the cryptographic credentials associated with a domain, so that clients can use this information to increase the level of assurance they receive from the TLS handshake process. While originally claims were made that DANE would be the killer application to help pushing for DNSSEC adoption – providing an add-on in security by delivering a single trust-anchor concept TLS security could be based on – there are now those who warn that mandatory use of DNSSEC could slow down DANE (preventing cases where PKIX information is just put in the DNS).

WG Chair Warren Kumari pointed out that not all clients were DNSSEC-aware and not all zones could do DNSSEC-chaining up to the root zone. Moreover, using non-DNSSEC secured DANE only in restrictive cases (one option according to document author Richard Barnes) would not create new security risks, Kumari said. But others warned that analysis of answers would become hard if the WG would move away from fully DNSSEC validated and potentially even try to hand out differentiated answers for the match/non-match-cases that were “DNSSEC insecure” or “DNSSEC indeterminate”. The more complex decision tree was sketched by framework Author Richard Barnes (BBN Technologies) in his presentation.

The responsible Security Area Directors quite obviously lean towards a clear-cut solution including DNSSEC publication and validation. Tim Polk, Security Area Director and NIST engineer, warned that a discussion about such a layered security use case model could cause the WG „to rat hole and not get the document done“. Polk said it was DANE task was „to specify what happens when information is delivered securely“. Clients would figure out individually what they want to do in other cases.

Since the meeting there was a hint on the WG's mailing list by Kumari that the Internet Engineering Steering Committee, which will do the standard peer review on the documents, had „strongly expressed the view that 'if we don't have DNSSEC as a must, we are outside our charter'“. The WG now seems to be trying to come to some consensus to proceed as if a TLSA record did not exist in all cases where DNSSEC validation was „insecure“ or „indeterminate“.

After about a week of intensive discussion Barnes now has posted a new ASCII graph for the straightforward DANE cases to the mailing list, answering the co-chairs request to come to a quick closure for a KISS version of the DANE-protocol.

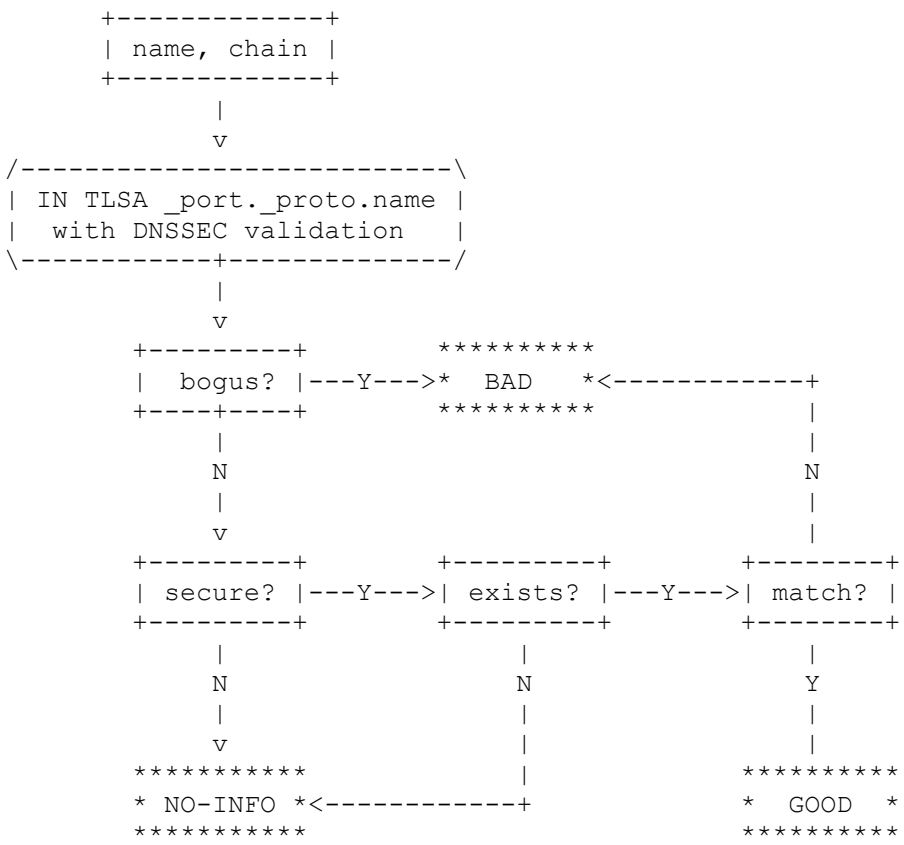

DNSSEC	Exists	Matches	Outcome
bogus	*	*	BAD
secure	yes	no	BAD
secure	yes	yes	GOOD
secure	no	*	NO-INFO
other	*	*	NO-INFO

The Working Group in Taipei furthermore tried to resolve a list of open issues in order to keep to the still ambitious time plan of sending the last called framework and use case documents to IESG review before the next IETF meeting.

Last Mile Security issue

For several issues there was consensus at the Taipei meeting to defer them and not to include them in the active drafts, but to potentially come up with additional documents, again in an effort to get DANE in place as fast as possible against the DNSSEC-, and especially last-mile-unfriendly trend of local rewriting (favored by BINDs repute policy zones, see below).

On last mile security the WG agreed to consider additional text in the security considerations section to make DANE users aware of the issue. If validation is done at the ISP and not in the local stub resolver, man in the middle attacks are possible on the last mile. Yet the WG should not try to give operational advice for the DNS, DNSEXT co-chair Andrew Sullivan said, while his co-chair pointed to RFC 3655 for reference. Another long discussion circled around the question on how to include assertion of self-signed

server certificates, which would make it easier for people using them, according to Alexander Mayrhofer from NIC.at. Discussion is ongoing on how bare public keys should be handled.

A question for the time being, according to Barnes, was if software currently could handle a one-key only usage scenario. Another question to be dealt with in the security considerations are cases of conflicting information between PKIX and DANE, where according to the working group DANE should override PKIX.

Websec – protect against future DigiNotars

While the DANE WG is trying to finalize its standard track documents, a first draft addressing an issue covered by DANE was presented by Ian Fette from Google during the session of the WEBSEC WG in Taipei. As Google had protected Chrome users against the recent DigiNotar-attack by hard-coding hashes of Google's VeriSign certificates, the company had been approached by several parties with requests about how to use the idea for other sites than Google. The DigiNotar attack was made possible by the fact that compromised certification authorities could start to sign certificates for sites at will.

As hard-coding certificates or CA hashes for many sites into the browser would not scale, Fette now came up with the concept of „dynamic pinning of public keys“. When first connecting to a site, the hashes of the correct certificate or of the legitimate root certificate (of the CA) would be received and for the time set the browsers then would check if the respective Subject Public Key Info was there before TLS sessions are opened. The hash values and the max time according to the draft document will be stored by the user agent together with metadata for „Http Strict Transport Security“ (HSTS). HSTS is the core draft of the Websec WG and shall enable web sites to declare themselves accessibly only via secure connection (Https).

Depending on how stable the certification policy was, Fette explained, the site would choose a longer or shorter running time for the pins. Shorter times would allow faster changes in the policy, longer time could lead to stale data in case the policy was changed before the run-out time. Yet the longer times would also provide less points to attack, as the core problem of the dynamic pinning was the bootstrap problem. When first fetching the hashes, a man in the middle could send over wrong key material resulting in compromise and non-reachability of the correct site. Pinning also does not help in cases your own certification authority has been hacked. The WG did discuss the relation to DANE that addresses the case of false CAs, too. The advantage of the pinning spec is that it can be used for „some more security“ as a stand-alone feature, while DANE according to current plans is based on DNSSEC. Fette considered the possibility that DNSSEC might help with the bootstrapping problem.

Another attempt to reform Whois

Good old Whois-services are still lacking a standard data format and and also support for internationalized contact data. These reasons together with the need for service access differentiation were presented at a BoF in Taipei by the proponents for a new Working Group on „WHOIS-based Extensible Internet Registration Data Service“ (weirds) to push for a new attempt to reform the Whois. A draft of the the proposed Working Group Charter for WEIRDS lists the following three goals:

> 1. Complete support for internationalization of queries and responses, including a standard representation for IDNs and non-ASCII text in queries and responses, and ways to indicate language or character set preferences for responses. The Working Group will need to determine whether language or character set (or both) is an appropriate selector.
> 2. machine-friendly data model that allows for the (current) developments of innovations in TLDs and in number registries.
> 3. Support for differential service levels, including bulk access, according to different classes of user.

BoF Co-Chair Murray Kucherawy (Cloudmark, Carrier Grade Network Infrastructure Provider) prepared a requirements document and underlined that his company needed WEIRDS output because he did not have enough reliable information to make security related decisions. Kucharawy also said that if registry providers or registrars were not giving him the needed information he would had to assume bad things and would „put pressure on them to come to the table“. Quite obviously, law enforcement is one of the potential customers, as law enforcement, especially in the US, but also in the UK (SOCA) have been complaining for years about gaps in the Whois. During the debate in Taipei the proponents referred to recommendations by ICANN's security and stability advisory committee calling for

the adoption of common terminology with regard to Whois data (Domain Name Registration Data, DNRD, DNRD Access Protocol, DNRD-AP and DNRD Registry Services, DNRD RS)
the adoption of a replacement domain name registration data access protocol that supports the query and display of Internationalized DNRD as well as addressing the relevant recommendations in SAC 003, SAC 027 and SAC 033.
the development of a a uniform and standard framework for accessing DNRD that would provide mechanisms to define and implement a range of verification methods, credential services, and access control capabilities.

One question discussed intensively during the BoF was if a common protocol for IP address registries and domain name registries made sense. It is especially the RIRs that have started to implement an alternative to the existing port 43 Whois (and also to the IETF standardized IRIS protocol, an earlier attempt to reform the Whois and to allow layered access, for example).

During the BoF meeting chaired by Kucherawy and Andrew Sullivan (Shinkuro), three out of four existing approaches to implement new Whois services were presented by representatives of the RIRs, namely ARIN, RIPE and LACNIC. For a domain name registry variant, Steve Sheng from ICANN presented a „pilot RESTful Web Service for querying Domain Name Registration Data (aka WHOIS data).“

All four implementations use „Representational State Transfers-based Web services“ (REST) based on HTTP (allowing to use caching, referrals, authentication, version control, https and accommodating Unicode). Andrew Newton from the North-American IP-Registry ARIN, which has operationally implemented restful services in 2010 and meanwhile was serving 40 percent of the requests via this channel according to ARIN statistics, described delivery of metadata about registration data and facilitated search mechanisms (for example “/rest/pocs;first=John&last=Doe” to search for all entrances for John Doe) as the benefits of the standard Whois data conventions.

Despite the common interest of the IP-address registries and ICANN they have implemented slightly different variants (see below). All of the variants use XML as primary output format, and also support XHTML, JSON or plain text.

ARIN	Our Whois-RWS mirrors the major structured data types used by ARIN in URL patterns: /rest/poc/XXXX for points of contact /rest/net/XXXX for networks (IP prefixes) /rest/org/XXXX for organizations /rest/asn/XXXX for autonomous system numbers /rest/rdns/XXXX for reverse DNS delegations
ICANN	ICANN’s prototye service uses the HTTP protocol and conforms to the REST architecture. The client sends its request with the following URL structure: /rest/domain/XXXX for domain name request /rest/contact/XXXX for contact request (by contact ID only) /rest/host/XXXX for host request.

RIPE	Single-resource lookup service: Given a primary key, a type, and a registry, it always returns one and only one object. Can be used to identify resources by URL bookmark. Resolution of referenced resources: all attribute values that represent references to other resources contain an xlink anchor that can be followed to navigate and browse networks of resources. Client can navigate through any network of resources via xlinks without requiring any stateful information to be stored on the servers. This comes as a benefit of the two previous features. Normalization of continuation lines, end-of-line comments, and other RPSL intricacies, and normalization of comma-separated values when they represent references to multiple resources.

Despite the pilot and even operational implementations (at ARIN, e.g.) there was no consensus about starting a WG. A majority (given the hums asked for by the BoF Chairs) seemed to be accepting the RIRs initiatives, but there were questions during the session about why there was a need for a WG instead of a quick standardization of the RIRs joint concept.

Several experts questioned the attempt to standardize yet another new alternative for the Whois, because of earlier failures to deploy existing alternatives to port 43. IRIS was not widely deployed because of its complexity, Newton and others said, making another attempt necessary. Interestingly, Marcos Sanz, IRIS co-author from DENIC, reminded that for IRIS it were the domain name registries that pushed for the Whois reform, while the RIRs initially did not buy in. For the new attempt there seems to be some reluctance now on the site of the domain name registries.

The failure of earlier Whois-reforms was not a protocol-related problem, but more a political problem, Peter Koch from DENIC and ex-IAB member John Klensin argued. Another new protocol driven by ICANN and the RIRs would not solve the „layer nine"-issues. Observers now expect another BoF before a decision about a new WG will be made

IPv6 – Conflicting concerns on the way to transition

Despite assurances by the experts that there is no change for traceability of IPv6 addresses, there is a feverish production of new draft proposals for source address validation, many of them being tabled by Chinese experts, and there is a haste to get IP address logging recommendations in place for so called Carrier Grade NAT (CGN) by large US cable providers. NATs and transition mechanisms in general are seen as a problem in China, for example, but also by law enforcement more generally.
Because NATting and also tunneling will make it more difficult to trace traffic back to the source/end user based on current mechanisms, more logging and, when it comes to the Chinese operators more centralized Source Address Validation (SAVI) switches, were discussed in the BEHAVE and SAVI WGs respectively.

Source Address Validation – at edge or core?

Source address validation (SAVI) has been developed as a method to better prevent the forging of IP addresses (spoofing) in IPv4 and IPv6. Developed originally to be deployed in the local LAN – close to the user – there is now one version operational, for v6 DHCP. In a joint draft, researchers from Cernet/Tsinghua University and China Telecom in Taipei propose to expand source address validation to various scenarios, including

> 1: Dual-stack with stateful (for this SAVI switch should snoop DHCPv6/PCP protocols and bind the relationship of <IPv6, MAC, Switch-Port>

2: Dual-stack with stateless (SAVI switch saves the relationship of <IPv6, MAC, Switch-Port > or <IPv4, IPv6, MAC, Switch-Port>
3: CPE-behind with stateful (SAVI switch should snoop the DHCPv4/PCP protocols interaction and bind <IPv4, MAC, Switch-Port> relationship
4:CPE-behind with stateless (SAVI switch does the same thing with scenario C, with tunnel initiated by CPE)
For Lightweight 4over6, similar mechanism, with SAVI switch having to listen to these address allocation protocols and bind relationship of <IPv4, MAC, Switch-Port, Port-range>.

Another proposal wants to shift the SAVI Switch from the edge network to far more centralized positions. Jun Bi, Director of Network Architecture & IPv6 Research Division, Network Research Center of Tsinghua University, which is closely cooperating with Cisco, said at the SAVI session of the Taipei IETF that it was difficult to deploy the necessary Savi-Switches locally, as this would mean to deploy this for the one million students on CERNET, China's academic and research network.

Yet Jun Bi's question if the IETF would go along with a proposal to “deploy something on layer 3” (away from the user and closer to the network center) was not well received at the IETF. The constraint to keep the SAVI switch and filtering function at the local level contradicts still standing comments on the SAVI framework document (currently in IESG last call) about potential privacy risks. SAVI solutions that were shifted to the provider would increase the propagation of potentially private host identity information, IETF experts who did a privacy review of these drafts warned.

CGN Logging

But it is not only the Chinese operators that have their problems with transition. Traceability will also become an issue for all operators (including those in China) that deploy a CGN. As many users will share addresses behind such NATs, logging becomes essential. Cisco Engineer Senthil Sivakumar, said that NAT logging was necessary for legal requirements, traceability and data retention. Operators still had to come up with better solutions for the logging, because otherwise they will collect over a petabyte of data in a year for one CGN only, Chris Donely from Cable Labs explained.

From the point of view of data protection officials, the logging and even more the SAVI proposals underline that precaution needs to be put in place with IPv6. Data officers from Germany, Mexico, Canada, Belgium, the UK, and Ontario recently published a declaration requesting that end users should be enabled to decide to have dynamic addresses, instead of the possible static addresses with IPv6 and that privacy extensions protocol would be on by default. “You have to ensure this, for example, for routed Android devices”, Tahar Schaa, German IPv6 expert said.

Privacy Concerns

There has been quite a lot of noise from the privacy community with regard to potential privacy tussles of IPv6 always-on-addresses, with the declaration in IPv6 privacy adopted.

The IETF is looking into the privacy issues, with a new privacy directorate in the making. IETF Chair Russ Housley explained that IPv6 offers more alternatives in the use of the lower portion of the 128-bit address: “One choice is to use your 48-bit MAC address. This choice means that you have a unique identifier for your laptop, and it stays the same regardless of the network or hot spot being used. This has the worst privacy implications, but it has a very simple network configuration.“ The other choice, he said, „ is to use a cryptographically generated address. This choice means that you get a new identifier each time you connect. This has the best privacy implications, but there is some overhead associated with generating the address and convincing the nearest router that it is okay.“

Working Groups

A word about the DNS WGs

Neither the DNS Operations Working Group nor the DNSEXT Working Group were meeting in Taipei, and the DNSEXT meeting group is expected to close down in Paris officially, according to its chair Olafur Gudmundson. The main DNS related issues currently under way in the IETF is DANE (see highlights above), other issues touching the DNS are the ongoing discussion about potential new Whois work and internationalization issues in PRECIS and Email Address Internationalization WGs.

Following up to the controversial discussions at the RIPE meeting on the reputation policy zones in BIND and domain name blocking by registries on injunctions by US authorities, debates on the trend to more filtering continued. In an interview with this reporter, former root server manager Bill Manning heavily criticized the Internet Software Consortium (ISC), which is providing BIND: „ISC has turned into an arms dealer“, Manning said. After years promoting DNSSEC ISC now offered rewriting and the blocking of bad reputed zones. Manning warned that failures in DNSSEC validation would result in people turning away from DNSSEC before it would really take off – he compared the effect to the „IPv6 has a problem, nothing to gain, so let's turn it off-situation“. But questions to several registries if they would consider turning away from BIND's products were answered negatively.

While there is considerable resentment about the „high security“-path ISC is taking, including its closer involvement in law enforcement actions (for example the Operation „Ghost Click“ by the FBI and several non-US law enforcement) registry experts asked by this reporter said they would stick to BIND for its stability. In the Ghost Click operation ISC had been appointed by a Manhattan Federal Court judge to replace the malware-distributing DNS servers of the defendants for 120 days. ISC also might have annoyed some customers by the company's expansion to the new gTLD back-end service market, one expert said.

PAWS

The Paws Working Group is pushing for rapid adoption of its basic use case and framework drafts as White Space pilot scenarios, and some proprietary operational projects in the US are already underway. Access to White Space frequencies shall open up additional spectrum for wireless connectivity for users roaming or fixed. White Spaces are underused frequencies allocated to Broadcasting companies and organizations. The PAWs WG started at IETF 81 (after a BoF at IETF 80) to standardize the a method for white space data base discovery, a method to access the database and ensure security mechanisms for discovery, database access methods and query/response formats. Devices that want to use White Space frequencies have to query a database in order to check what channel is available at a given location.

The Use Cases presented at the Working Group in Taipei include Wifi-like services in urban areas, where additional space is welcome, broadband internet access in rural areas (where PAWS is a real alternative already developed in some parts of the US) or the case of offloading data traffic from a 3G or 4G network to the White Space channel available. The main idea of the WG is to reuse existing protocols for querying the data base (for example running the White Space application over Https and IP).

Despite postponing the shipping off of the base documents to the IESG to next year (after an originally very ambitious goal to ship first documents by the end of the year) there is a clear rush to finalize the PAWS specs. Presentations on the data model (by Virgina-based Key Bridge Global LLC) and on a White Space protocol (by Telcordia) were sent to the WG long after the regular cut-off data and the protocol documents reached WG participants the evening before the meeting. Nevertheless one of the Chairs even tried to push the WG for adoption of the Telcordia document, despite being criticized by one participant (Orit Levin from Microsoft) and admonished sharply about procedure by the Area Director.

Obviously there is a race to the TV White Space market, with the FCC having licensed nine companies as White Space data base providers: Comsearch, Frequency Finder, Google, KB Enterprises, LS Telecom,

Key Bridge Global, Neustar, Spectrum Bridge, Telcordia and Wsdb. Next year devices that will be able to connect via TV White Space frequencies would be already available, say experts in Europe and the US. But while the Federal Communications Commission after years of deliberations has decided to open up the frequencies for free, in Europe regulators are hesitating and, according Alexander Kholod (Ofcom Switzerland and currently Chair European Conference of Postal and Telecommunications Administration, CEPT), are waiting for companies „to come up with business plans“ before they want to decide how they will allocate the frequencies.

For some time broadcasters in Europe had been rather conservative fearing interferences, but in the EU White Space Research project CogEU, broadcasters and their research institutes have been participating in developing, for example, a first version of a EU geolocation White Space database (recently tested for Bavaria).

One White Space pilot project is currently underway in Cambridge, UK, involving BBC, BskyB
BBC, BSkyB, BT, Cambridge Consultants, Microsoft, Neul, Nokia, Samsung, Spectrum Bridge and TTP. TV White Space standardization work in Europe (driven for example at the ETSI) and the US are currently ongoing in parallel.

SIDR

The SIDR working group is continuing its path to securing the routing protocol BGP, discussing the WG documents on BGPSEC without any mentioning of the policy debates at the recent RIPE meeting. While the documents for the Routing PKI architecture and specs are far advanced with seven documents sitting the RFC editor queue for up to 200 and some days (potentially a result of the fact that 2011 will be the IETF's second busiest year for RFC production) already a heated debate has started about route leaks potentially non-controllable by the path-to-path securing of BGPSEC.
As discussion in the WG about the issue was somehow cut short, Danny McPherson from VeriSign and Shane Adamante from Level 3 now posted a little draft describing the problem they see. The core problem is described in the quote below:

> As currently defined, BGPSEC only provides two functions:
>
> 1. Is an Autonomous System authorized to originate an IP prefix?
> 2. Is the AS_PATH represented in the route the same as the list of ASes through which the NLRI traveled?
>
> In order for an attacker (AS 1) to divert traffic from ISP1 for prefix P through their AS they simply fail to scope the propagation of the target prefix P (received from ISP2) by announcing a (syntactically correct) BGPSEC update for prefix P to ISP1. This vulnerability is what the authors refer to as a 'route leak'. It is important to note that the default behavior in BGP [RFC4271] is to announce all best paths to external BGP peers, unless explicitly scoped by a BGP speaker through configuration. Because ISP1 prefers prefixes learned from customers (AS 1) over prefixes learned from peers (ISP2), they begin forwarding traffic for prefix P destinations through the attacker's AS (AS 1). Voila!

Two questions discussed during the Working Group session in Taipei were how to deal with mergers of AS, with an example made be Level 3 who bought Global Crossing including all their AS. Potentially this issue could be solved by counting these AS as zero value in the path chain, Randy Bush, Internet Initiative Japan and one of the core RPKI developers proposed. This zero value calculation might also help for IX systems that do not have AS numbers or only use private numbers, Bush explained during the session.

IETF News

The IETF will be meeting in Paris next time, but has not yet found a sponsor for this meeting, which is happening in March 2012. IAOC Chair Bob Hinden said, as it had been expected that 75 percent of the meeting costs would be covered by sponsoring, that the gap had to be filled either by asking the ISOC – who is supporting the IETF budget with its org-revenues – for additional support, or to have IETF participants pay for it with higher meeting fees, potentially spread across the three meetings next year.

It is interesting to note that the IETF seems to more and more cut the deals with hotels before having settled on a main sponsor for the venue. This is also the case, for example, for the next upcoming meeting in Berlin in 2013. While venues like Paris or Berlin might attract more participants than Taipei (which did attract 300 less participants than the IETF a year ago in Beijing), the lack of sponsorship still does constitute a financial problem for the IETF, who calculates with 75 percent of the meeting costs to be borne by sponsors.

In the administrative plenary IETF Chair Russ Housley reported about a legal complaint filed by Astrolabe, a company providing an astrological Almanach, against the volunteers who managed the time zone database since the mid-eighties. Astrolabe held that the time zone database has violated its copyright. But while the time zone database relies on the Almanach for some historical data, Housley said, after talking to several lawyers, he had been assured that the case was „frivolous“. Housley said that Electronic Frontier Foundation has taken up the case „pro bono“ to defend Arthur David Olsen and Paul Eggert in the case. Meanwhile the database that had been shut down due to the law suit for some time was moved as planned to IANA, with ICANN restarting the database on October, 14.

The IETF is still working to complete work to put the new RFC editor in place. Interviews have been held with 36 people according to Fred Baker. The list meanwhile is down to around 10 people with a decision to be expected before the IETF in Paris.

With IETF work done under the current Intellectual Property regime reaching the end of the legally foreseen time span of 35 years, the IETF has to think about how to organize „re-assignment“ of the rights from the authors. According to Marshal Eubanks, Chair of the IETF Trust the IETF cannot „just send around a form saying I agree my RFCs are going to stay with the IETF Trust.“ The issue was a concern to every standardization body, Eubanks said, but he had not an immediate solution, The first IETF RFCs will need re-assignment starting January 2013.

During the technical plenary there was a short discussion about the effect of IAB Best Current Practice Documents (BCPs) or architectural documents, like for example the architectural implications of IP anycast, or of application features in the DNS, or the IAB considerations on privacy. Such documents were open to comment and discussion from the community, IAB members said. Dave Crocker, who had brought up the topic, on the other side complained that there was not enough discussion given the nature of the documents as guidance for the larger community.

Thomas Narten was re-elected to serve as the IETF-contact to ICANN; Peter Koch will serve as IAB-contact to ICANNs Root Server System Advisory Committee (RSSAC).
Kilnam Chon received the Jon Postel Award 2011. The ItoJun Award for young IPv6 developers was given to Alexandre Cassens (Free France) and 6RD developer Remi Déspres, who in fact is not as young as the earlier Itojun awardees. But as obviously the ISOC Selection Committee wanted to honor the strong IPv6 engagement of Free, Déspres as the brain behind the rapid deployment-idea, had to be given due credit. Déspres in a short speech said that it took the Free engineers just five weeks from his first presentation of the 6RD idea to implement IPv6 access for all of their customers. Currently Free has 1,5 million IPv6 users.

IETF 83 meeting is being held in Paris, March 25-30, 2012