Report of the

IETF 79

Beijing, China,
November 8 - 12, 2010

Prepared by Monika Ermert
For the CENTR secretariat

Table of Contents

Highlights

IETF goes to China

For the first time the Internet Engineering Task Force (IETF) had a meeting in the People's Republic of China, co-hosted by Tsinghua University, the Internet Society of China (ISC) and the China Internet Network Information Center (CNNIC), – and IETF and the host obviously had agreed to “behave well”.

The IETF had asked for and received – according to the participants – their own unfiltered network, a network managed by the IETF itself in cooperation with the host. The unfiltered and unmonitored network had been a prerequisite for the IETF coming to China. One of the trouble tickets of the Network Operating Centers, ticket No 279[1], illustrated the kind of privileges granted, or better what a possible lack of these might cause – the trouble ticket made clear that addressing the global DNS route servers without going through a proxy was only open to privileged networks. As a privileged network the IETF meeting network also provided access to sites like Twitter, You Tube or Facebook , sites that were blocked for users of the Shangrila hotel network (Internet access in non-listed hotels close by seemed even tighter in their blocking policy, this reporter even faced access problems to parts of the IETF website).

China on the other hand had received a quasi-acknowledgement of its “one-country-policy”: the Taiwan Chapter of the Internet Society for this had to be “degraded” to the Taipei Chapter. The ISOC according to ISOC president Lynn St. Amour was hoping to “deepen” the relationship between ISC and ISOC. Granting the Jon Postel Award to Wu Jianping from Tsinghua University earlier this year had been completely coincidental, Lynn said at the first ever press conference at an IETF meeting, with around 20 Chinese journalists attending.

From a Chinese point of view the IETF coming to China was a “big thing”, as Wu put it during the press conference. Wu even said a “dream is becoming true”. Mao Wei, Director of CNNIC and CNNIC's representative in many international for a like ICANN or the IGF, underlined the importance from the Chinese point of view during the 4. US-China Internet Industry Forum (UCIIF) organized by Microsoft that took place on Monday, Nov 8 to Tuesday, Nov 9 in Beijing. Mao Wei said the IETF's first visit to China showed that „the influence of China's Internet for the global Internet is growing“.

Mao Wei said despite growth in participation rates (Chinese developers were in the majority with 366 IETF 79 participants and US participants only 340) China still had a long way to go in order to play its part in standardization with roughly one percent of RFCs being proposed by Chinese developers. CNNIC is especially active in internationalization efforts, Chinese companies like Huawei, but also

---

1 *With regard to the request to check on the problem to access the f.root-server, see*
dig -4 @f.root-servers.net . soa
; <<>> DiG 9.4.3-P3 <<>> -4 @f.root-servers.net . soa
; (1 server found)
;; global options: printcmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: REFUSED, id: 12759
;; flags: qr rd; QUERY: 1, ANSWER: 0, AUTHORITY: 0, ADDITIONAL: 0
…
the NOC team gave the answer:
One of our upstreams (CERNET) maintains a root cache for internal use, and we were accidentally redirected to use that cache. Unfortunately, we weren't added to the list of netblocks authorized to use the cache. Hence, for IPv4, we got redirected to a server that refused to answer us.
We've fixed things so queries now go out as expected. Could you test and verify that it's working now for you? All my tests come up clean.

network operators like China Telecom and China Mobile are especially interested in transition technologies for Ipv6. In the BoF meeting on a “Lightweight IP Protocol Design” Chinese participants clearly were in the majority.

In what way the IETF Bejing meeting will push for further attention for the IETF in China still remains to be seen, yet several Chinese participants told this reporter their intention to attend future meetings (especially the Quebec one was mentioned). The IETF will come back to the other China, Taiwan, in one year, obviously in an effort to demonstrate some neutrality.

As the IETF in the future intends to switch to a one-one-one geographical meeting schedule (one meeting in North-America, Europe and Asia each year) Chinese participation will be easier, visa- and travel-expense-wise. During the IETF press conference, Hu Qiheng, Chair of the ISC, said other cities in China were looking forward to host IETF meetings in the future. Other than for IETF meetings in the US according to Russ Housley there had been no Visa issues for the Beijing meeting. On the other hand there was one bitter complaint during the administrative plenary about the relatively tight checks of badges at the entrance to the conference area, which the complainant, Sam Weiler of Sparta called a „new policy“.

The hosts, Hu and Wu, on the other hand both applauded the “IETF culture of openness” and “selfishness contribution” to the development of the Internet. Hu said the IETF did have “strict regulations” with regard to Intellectual Property. “If you want to apply for patents or royalty fees, do not come to the IETF”, Hu warned. At least participants had to publish patents and share the knowledge. From the IETF standpoint the promotion of “one network” was one of the important points to deliver in China.

DNS – prepare for delivery of more secure services?

With the „DNSSEC thing“ around there have been proposals to use the secure DNS as a trust anchor for certificates instead of using traditional PKIX or OpenPGP from third parties. After a first Bar BoF in Maastricht a formal BoF was held in Beijing, with discussions about the various proposals peaking before on the mailing list and in the hallways. Ondrej Sury, CEO from cz.nic who co-chaired the BoF he expects the „Keys in the DNS“ (KIDNS) to be formally set up in the next two weeks.

Generally speaking there are two schools of thought on the idea to use the DNSSEC-secured DNS as a trust anchor. The proponents of the working group and authors of the so far five drafts all point to the relative ease to allow a check on a domain owner by using the existing DNSSEC key infrastructure. The basic idea, as Paul Hofmann from VPN Consortium presented it during the BoF was just “use the DNS to allow the host to declare what public key he uses” (with various methods to identify the key by the key itself or a hash of it, the key in a self signed EE cert or a hash of it, a hash of a CA cert that expected to be the user's trust anchor store or the CA Cert itself). Also Hofman explained that “the key can be associated with all secure services running on a host with a single domain name”.

The „DNS purists“ point to the „insecure qualities“ of the secure DNS, namely there is no verification for domain name owners during registration in the first place and, secondly, the last hop could still be attacked, according to the warning.

One draft author, Philip Hallam-Baker, points explicitly to the verification/identification problem, writing that

"A signed CERT record does not and cannot express any assertion concerning the existence, trustworthiness or accountability of either the key holder or the domain name holder."

Hallam-Baker in one of his two draft proposals also lists a number of possible attacks. His second proposal explicitly covers behavior of a "Certification Authority Authorization DNS Resource Record" to "allow a public Certification Authority to implement additional controls to reduce the risk of unintended certificate misssue." There was a lot of criticism against the ideas presented in this draft by the DNS purists.

The concerns about the insecurity of the last hop according to Sury will be addressed in the working group. Sury said during the RIPE meeting in Rome that the charter of the KIDNS WG would be extended to include „securing the last mile".

With regard to the proposals so far there are now five drafts, of which only two were presented during the Beijing BoF. The general concept was explained by Paul Hofmann. Performing a DNSSEC lookup followed by a Transport Layer Security (TLS) negotiation and certificate validation would be, as Hofman put it, have a long latency. Typical PKX certificate validation (DV validation) more over had known security issues, because of the convention that if „you trust one certification authority you trust them all". According to Hofmann and other proponents DNSSEC provided a solution in allowing the host to declare what public key it was using.

For TLS Hofmann explains in his coauthored draft (together with Jakob Schlyter from Kirei and Warren Kumari and Adam Langley from Google):

TLS and DTLS use certificates for authenticating the server. Users want their applications to verify that the certificate provided by the TLS server is in fact associated with the domain name they expect. Instead of trusting a certificate authority to have made this association correctly, the user might instead trust the authoritative DNS server for the domain name to make that association.

An easy and straightforward explanation of the idea also is presented in the draft of Simon Josefsson, an Open Source Software developer, on key assurance via the DNS:

TLS supports X.509 and OpenPGP certificate based mechanisms to authenticate a server. Users want their applications to verify that the certificate provided by the TLS server is in fact associated with the domain name they expect. Instead of trusting a certificate authority to have made this association correctly, and an X.509/OpenPGP implementation to validate that properly, the user might instead trust the authoritative DNS server for the domain name to make that association. This document describes how to use secure DNS to associate the certificate chain transferred by TLS with the intended domain name.

In a draft now supported by two „heavyweights", IETF Chair Russ Housley (Vigil Security) and Security Area Director Tim Polk (NIST), an extended use of DNSSEC in connection with SMIME and IPSEC is elaborated. For SMIME Turner, Housley and Polk explain:

To encrypt the message, the originator needs the recipient's key agreement or key transport certificate. To obtain the recipients certificate, the originator composes the email, selects sign and encrypt, and hit send. The mail client/DNSSEC client reviews the local store and determines that no

certificate is available. The mail client then queries the DNS to determine whether certificates are available for that domain. If a CERT resource record (RR) [RFC4398] is available, the mail client examines the certificate to determine if it is a CA certificate or end certificate. For domains with multiple users, the certificate would be a CA certificate and would include a SIA extension [RFC5280]. (…) If an appropriate certificate is available (and validates according to local policy), the client can encrypt the message. The originator includes their own certificates in the message, so this process is not required to validate or decrypt the original message or for a response.

For IPSec the process would be similar, the authors wrote.

Hofmann in his presentation touched questions about the expected or to be standardized behavior for cases where DNSSEC was not available for a domain or where Keys or Certificates were not available. Hofmann pointed to earlier work for signatures for email stored in the DNS in DKIM.

During the BoF session there was not much discussion, presumably partly a result of Hofmann's very cautious explanation of the scope. Support for the concept in general by the security area AD and the IETF chair who both participated in the session also might have led to caution on the side of the DNS purists. A lot of problems with the concept would become apparent during the work on the drafts, said one participant, and would be controversially discussed then.

Ipv6 – get prepared to panic

The number of Ipv6 transition technology proposals before the IETF is breath-taking. There were 66 Ipv6-related proposals tabled for the IETF in Beijing in various working groups, not only the classical Ipv6 WGs (v6ops, v6man) and those looking into transition technologies like behave, software, but also in many other groups and open area meetings like apparea, intarea, mext, dhc and so on. The time necessary for those people that had to decide which of the drafts were important enough to implement, and which were crappy enough to invest time in to try to get off the table, was enormous, Suzanne Woolf from ISC said to this reporter.

There are now attempts to give orientation through meta-documents, namely an "Annotated Bibliographyof Ipv4/IPv6 Transition and Coexistence" by Ed Jankiewicz. Jankiewicz said, people were starting to realize that they needed a variety of coping mechanism and therefore "we went from a situation where nobody paid attention to many things going on at the same times right now." He saw an "explosion" of new drafts – with often older, but recycled ideas - after the last IETF in Maastricht, he said. Jankiewicz also proposed three fundamental rules for deploying the various sets of transition technologies:

- First, do no harm
- Keep it simple
- Keep moving towards more native Ipv6 (eg. Date after which there will be no more deployment that will not support Ipv6)

Moreover implementors should not try to reinvent the wheel and write new draft proposals before checking on the existing ones. Jankiewicz lists pretty much what has been put on the table over recent years, categorizing transition and coexistence scenarios and architectures, and the various tools for address mapping (address translation, NATs in applications, dual stack lite-approach), tunneling mechanisms (Teredo, 6rd, tunnel support protocol, residual Ipv4 over Ipv6 infrastructure, Dress Plus

Port, Iron-Ranger and ISATAP solutions, softwires, L2TP), various translation approaches and connectivity checking and delay avoidance.

There was a fierce debate on the reiterated request from several network operators (Rogers Comm., Cox Comm. Telstra, Frontier Comm.) to the IETF to get a “re-useable /10-address block” to allow the CPE to CGN network to be deconflicted with the customers' networks. The problem the operators say they have is that they have to reuse their existing v4-addresses to an extent that causes problems. The use of NAT444 while the operators concede leads to breakage was the “least disruptive way of managing those heritage devices during transition”; transition to Ipv6 would be pushed for at the same time, the operators promised. But they did not at all meet with approval from the community.

While some just told the presenters to go away and push for Ipv6, others pointed to the risk of leakage of those public, but only privately used addresses to the net. In conclusion not even the proposal to allow the operators to follow up with statistics to document and thereby substantiate their need. An interesting question is if the operators can – as the said they would have to – request addresses with the RIRs now (and thereby fasten Ipv4 depletion). Some of the opponents argued that if the operators would be able to substantiate the need to the RIRs they wouldn't have come to the IETF in the first place. A representative from China Telecom said it would just not be fair to allow the use of the rare Ipv4 public addresses to solve the problems of a handful of private companies.

Another Ipv6-transition problem brought up by Jason Livingood (Comcast) during many Wgs was the issue of AAAA “whitelisting” of domains by content operators like Google to avoid possible breakage in requests to dual stack deployments. Google several times had argued that they could not afford to cut of even the measured 0,05 percent of customers, who faced very long delays or eventually time-outs because their systems did not work well with dual-stack sites (obviously this is for example a problem for MAC OS machines). Therefore operators like Google provided AAAA records only to those domains that had announced to them that they were capable of receiving the AAAA answers and then were taken on to an Access Control List (ACL). Everybody else who in principle could very well be able to talk Ipv6 was excluded.

The big problem with the whitelisting to Livingood is that once broader deployed by various content providers it was very difficult for users to know if they did not receive Ipv6 answers because of such ACLs. Also the rules for being taken on the lists (and de-whitelisting events) were opaque and possibly followed different criteria everywhere. While Livingood still laid out the option to introduce whitelisting on a broad (and somewhat standardized) way, it seems pretty clear that his document much more favors to ban White-Listing. Lorenzo Colitti said to this reporter that whitelisting certainly didn't scale in the long run, yet it had been used by Google to avoid to cut-off its users.
Just for fun, here is a video about the day of Ipv4 exhaustion (from Cisco).

Working Groups, BoFs, Plenaries

DNSOPS

The working group started a new discussion on “Nameserver Control Protocols” without a conclusive decision if the WG will take this up as a work item. The topic according to the experts had come up several times and been ruled out of scope by the DNSOPS Chairs earlier. Still a group of people have prepared a problem statement that sums up that a possible WG should address the “perceived need for an interoperable way to manage (monitor, control and configure) name servers”. The goals of a possible WG would be “review of existing work to date, such as draft-dickinson-dnsop-nameserver-control-00, which proposes a solution based on Netconf and Yang)” and the production of “documents specifying a name server control protocol that addresses the requirements of interoperable management of name servers.” A requirements document can be found here

The prepared BoF was withdrawn from the Beijing Agenda, but may take place at a later stage. While some participants were not convinced there was a problem, others said time might be ripe for the work. Stephane Bortzmeyer from Afnic said, with the advent of cloud computing it might become much more common to have servers in different domains and therefore there was a possibility that somebody could create a "very general system" for the management of these servers.

Two proposals were discussed shortly during the DNSOPS session. Ning Kong from CNNIC presented a draft on "an automated synchronization mechanism for configuration data of DNS name servers." The draft document by Ning Kong "proposes an in-band synchronization mechanism to automatically synchronize configuration data among multiple DNS name servers." Any part of the config data of a name server could be "constructed as a similar DNS zone file, and be automatically synchronized by DNS messaging and notifying mechanisms. The proposal envisages to extend synchronization to the servers of managed DNS service customers.

The second proposal presented by DNSOP Co-Chair Steven Morris (ISC) develops a "common data model for describing the configuration and operation of a basic, but usable, generic name server", it could be used as a basis of a set of NETCONF (RFC 5277) operations and capabilities, the authors write. Morris pointed to basic functionalities: zone manipulation like add or remove zone, ACL creation, control name server behavior, statistics to obtain information from name server and manipulation of small amounts of zone data.

Two older drafts were discussed in the first part of the session, Johan Ihren's (Netnod) draft on DNSSEC-Key-Timing and the rfc4641bis document prepared by Olaf Kolkman who could not be there for the presentation. On the DNSSEC key-timing document Ihren said that algorithm roll-over and piecemeal approaches in zone updating still needed work. He also asked if the document should be split into several. But the WG finally tended to a publication of the current version with a bis-version to be started right away.

The DNSOP WG also briefly discussed the Whitelisting Implication document by Jason Livingood, see Highlights above and the DNSSEC history project, presented by Steve Crocker, see the "Short Update on DNSSEC" below.

Other documents in the pipeline are (matureness indicated in colors by the WG Chairs)

draft-ietf-dnsop-name-server-management-reqs-04.txt
draft-ietf-dnsop-default-local-zones-13.txt
draft-ietf-dnsop-as112-ops-05.txt
draft-ietf-dnsop-as112-under-attack-help-help-04.txt
draft-ietf-dnsop-dnssec-dps-framework-03.txt
draft-ietf-dnsop-resolver-priming-02.txt

DNSEXT

The Working Group discussed four major items and split up in the second part of the meeting to collect input to documents that are on the WG's plate in small break-out sessions.

The first document discussed is a proposal from Patrik Fälström for a new DNS resource record (Uniform Resource Identifier RR) "for publishing mappings from host names to URIs". The new record

according to the draft text shall allow to overcome a limitation of the current URI-lookup that does not allow for a subset of URIs connected with a domain. The URI RR on the other hand shall "enable the querying party to select which ones of the NAPTR records one is interested in." The mechanism was complementary to existing queries for NAPTR Rrs (or S-NAPTR). URI RR would reuse SRV and Enumservice registrations, said Ted Hardie, who presented the draft in Beijing. Questions to the WG were if this was the right mechanism to bind service to the URA at which the service was delivered, Hardie explained, and also if there should be a limitation to already registered services. Another question was if the reuse of "weighting mechanisms from the DDDS" was ok?

The second draft discussed was the still to be finalized problem statement on name equivalence (variants and aliases - CNAME, BNAME), for which DNSEXT Chair Andrew Sullivan pushed for a quick finalization, otherwise the document had to be abandoned. Co-Author Suzanne Woolf from ISC committed to another version at the beginning of December.

A presentation by Hiroshi Kitamura (NEC Corporation) promoted the combination of Ipv4 and Ipv6 DNS queries to reduce latency and load. Current two-Queries method urged two answers, and if either was lost, one had to rely on complicated recover procedures. Also twice as much traffic was produced. The biggest risk of parallel or consecutive double queries could lead to problems unsolvable for normal users who would then turn away from Ipv6. The proposal is to simplify the queries by sending both A and AAAA queries in one packet, combining A+AAAA to a meta type query, or just query for AAAA and map the address to v4 in case needed. Several participants asked for statistical evidence of the problem before going on.

On the last proposal, an Internet draft on optimizing resolver behavior prepared by ISC there was a short, but fierce discussion. The draft that proposes three practices for interested parties:

- Revalidating a delegation when a parent NS RRset TTL expires.
- Stopping a downward cache search when an NXDOMAIN is encountered.
- Upgrading the credibility of NS RRsets upon delegation events

The document was only informational, said Frederico Neves when presenting the draft. Peter Koch from Denic, Co-Chair of the DNSOPS WG warned against the proposed introduction of NSR queries and aggressive negative caching that the IETF decided to be out of scope for DNSEC. Both were attempts to "change the fundamental underpinnings of how the DNS works", said Koch. Discussion on the draft is ongoing. For a presentation of Steve Crocker, see the "Short update on DNSSEC" below.

Name-Based Sockets

With IP addresses often changing during individual sessions on the net, or problems arising from the Ipv4-IPv6 transition and from NATs the proponents of the Name-Based Socket (NBS) BoF see a need to push locator management away from applications, possibly to the operating systems.[2] Applications

2 *A good list of problems for applications was presented by Brian Carpenter in the NBS BoF and in the Application Area meeting. Carpenter in his proposal on referrals wrote that applications*

"- cannot assume that an address by which you reach a host from location A also works from location B.
− IP addresses no longer all have global scope, they often have limited reachability, and may have a limited lifetime.
− Can no longer assume that a host with a fixed location has a single fixed IP address, or even a stable IP address.
− A public IPv4 address often no longer identifies a single customer/user/host, without knowing the port number.
− A private IPv4 address is meaningless out of the private network.
− Addresses and port numbers may be different on either side of a NAT, and firewalls may block them.
− The Internet has two address formats (IPv4 and IPv6)."

according to the general concept discussed in the BoF should not deal with IP-addresses anymore, they should use addresses. Fully qualified domain names according to Javier Ubillos, from the Swedish Institute of Computer Science, who presented a draft architecture document on name-based sockets, would bring additional benefits compared to other names. The proposed charter is here.

BoF Co-Chair Christian Vogt, co-author (together with Mingwei Xu and Zhongxing Ming of Tsinghua University), said one motivation was to allow application developers to focus on innovation without trying to solve the various problems described by several presenters during the BoF. Stuart Cheshire from Apple pointed for example to the rejection of IPv6 address queries leading application servers to retry until time-out was reached.

Vogt has written a easy to read first overview about the concept, also comparing the nbs design to other solutions that try to address the problem (provider independent addresses, surrogate IP addresses, socket interface abstraction) plus the shortcomings of these as he perceives it (routing table scalability, duplication of efforts, extra-administrative overhead, no support for address changes). The Name-based socket architecture envisaged by Vogt and the other authors shall use an IP header extension using names (either fully qualified domain names or other names) to allow smooth implementation without major changes to applications, APIs, middleware or network. With the need of shorter TTLs to ease name resolution possible effects – additional query load - to the DNS were discussed. Salem Bhatti did tests on “zero TTL values for edge-site DNS records, concluding that they were possible, and the load did seem manageable. A first test implementation for the Ubillos/Vogt/et alii draft was done at Tshinghua University.

With Apple and Microsoft two large OS-representatives made presentations in the BoF showing interest. But Thaler also pointed out that opinions varied about how far the Name-Based sockets work should reach out: for a new API, for the behavior of an API or for an additional protocol (that would IETF-codify the exchanges between the two servers). Thaler said, he was not convinced that a new API was necessary, as there were APIs out there. Thaler favors to work on using and optimizing on several trends already visible in application development. Most new applications anyway were using higher layer APIs or frameworks and not classical sockets, and were generally not using address, but names. Cheshire pointed to Apple APIs, but also Java API's.

Both Cheshire and Thaler's said these APIs should just be used more. Thaler proposed a to-do-list to be considered including a relationship with dynamic DNS providers for the host, applications that use names and not IP addresses, application or session layer reconnects, and optimized reconnect time for DNS and TCP. Moreover DNS servers and API frameworks should respect small TTLs and an ability should be provided for hosts to communicate predicted name-to-dress changes.

At the end of the BoF those who asked that work should go ahead were in the majority, despite some skepticism by Transport Area Director Lars Eggert who asked for the chances of adoption. During the Application Area Open Meeting another pitch to dress the problem by Brian Carpenter was questioned. While Carpenter's analysis was welcomed his proposal to develop a new „general referral mechanism“ was not discussed further.
For the mailing list go to keyassure@ietf.org.

Internationalization

The work on internationalization proceeds slowly. To adapt string prep in various protocols that use it, a matrix has been created by the Precis WG to allow reviewers to give their input. The declared goal is “to assess whether a new method based on the new IDNA2008 algorithmic approach is the appropriate path forward for existing stringprep protocols as well as for other application protocols requiring internationalized strings. Issues to be contemplated for the stringprep reviews included case folding, case sensitivity or preserve case, user input, user interaction, normalization, classes (U-, A-

Label, domain name, email, restricted identifier, less restrictive identifier) what is published/seen, security/authentication issues, impacts of false positives and false negatives, tolerance of changes in the community, deliminators such as ".". A first check on stringprep in XMPP/Jabberclients was done by Peter St. Andre,

A new problem presented by John Klensin during the application open area session in Beijing was that the new version of Unicode disallowed characters that had been allowed before, something that had been expected to be very rare. Possible solutions would be either to freeze the IDNA 2008 standard at Unicode version 5.2, to adapt it to the new Unicode 6 or to do nothing. The decision, the experts think, will set a precedent for future changes in Unicode.

During the WG session on Email Address Internationalization the documents SMTP-bis (draft-ietf-eai-rfc5336bis-04), header-bis (draft-ietf-eai-rfc5335bis-03) and DNS-bis (draft-ietf-eai-rfc5337bis-dsn-01) were pushed to last call, with November, 26, being the dealine. More last calls were started at the WG session and ending on December, 3, are the POP-bis document, the Post-delivery Message Downgrading for Internationalized Email Messages (draft-ietf-eai-popimap-downgrade-00) and IMAP-bis (rfc5738bis).

Besides this the WG took time to discuss the i18n Mailto plan and the preference of either "comments" or "group syntax". The final sense of the room summarized by Pete Resnick was

- for from field continue to use group syntax as downgrade for addresses
- for destination fields to use group syntax, while documenting this was a syntactical break
- for nested groups there should be a proposal advising to use the entire set of groups within and make them the outermost groups and "2047" the whole thing
- sending of group syntax is fragile

The purpose is to allow legacy clients "not to choke on the new i18n-addresses" while not expecting them to be able to answer to the internationalized, but still only to the ASCII address. A Mailto-document according to the WG time plan is expected to be finalized by April, as was a document on EAI and mailing lists.

John Klensin said he had announced to the IESG the WG would provide a series of "advice document" on "forming addresses", "advice for MUA" and "advice about EAI deployment". Drafts on this still had to be written, said Klensin. Jiankang Yao (CNNIC) announced that he had prepared an advice document for non-ASCII addresses and eai deployment. The MailtoURI document was ready for a bis-version. An IRI-document also had to be developed.

Open questions also include if there was a need for a EAI-specific document for submission servers.

A very short DNSSEC update

There was no more extra DNSSEC Update at the Beijingg IETF, but several TLDs reported their signed zone going to the root (.nl) and being signed (.asia) over the week. Steve Crocker who promoted the DNSSEC history project (asking for contributions from DNSSEC developers, but also from governments about the process to get DNSSEC started) said according to an older count from October, 10 2010, were 53 TLDs signed and 45 in the root.

One interesting piece of news on the pending adding of the .arpa keys to the root zone came up during the week after the IETF. According to observers on the DNSSEC-deployment mailing list .arpa's DS records still have to be pushed to the root by the NTIA. With the Interim Trust Anchor Repository (ITAR) of IANA being put to rest .arpa validation was hampered, at least until ISC took the up the records in the DLV. Some experts seem to be somewhat upset at the delay. A request by this reporter to the NTIA resulted in a “no comment at this time” answer.

Big additional zones that will be signed in the near future are .com, .net (see VeriSigns published time plan). Other big TLDs, like .de, still did not commit to go forward. Denic runs a test to the end of the year and meets this week Wednesday for another DNSSEC-Testbed meeting in Frankfurt.

IAB Plenary: China preparing for end of Ipv4 addresses

Preparing for the IPv6 roll-out was a very pressing issue in China, different Chinese telecom operators and a representative from the China Internet Network Information Center said at the first day of the 79. Meeting of the Internet Engineering Task Force, which started on Monday in Beijing. Ipv4 addresses were already scarce, said Zhao Huiling, Research Vice President of China Telecom. By the end of 2010 China Telecom was „facing a gap of 20 million Ipv4 addresses“, Zhao said in what she said were personal comments.

The depletion of Ipv4 addresses expected to happen by the end of next year was, Zhao said, the biggest technical challenge for the time being. The Chinese incumbent has 64 Million fixed network customers and according to Zhao expects ongoing growth in the fixed and mobile network. „For a lot of new services and new applications, we additionally need billions of new addresses over the next five years“, Zhao said.

China Telecom was participating in the national IPv6 China Next Generation Internet (CNGI) project and had together with Tsinghua University started a lab, said Zhao. The company was also in the middle of intensive testing of various transition technologies in several regions of the country. For transition Zhao said, „we think there are four options. One is Ipv4 address optimization to raise use efficiency“, she said. The second was using so called private Ipv4 addresses, this allows operators to hide a lot of users behind Network Address Translation Devices (NAT). „We do not think this is a good solution, it is only a temporary solution“ said Zhao. Zhao also ruled out the purchase of Ipv4 addresses on the „market“, something many experts have been speculating about. In the end, Zhao said, IPv6 was the only way to go. She said she expected that there would be a „cocktail“ of technologies for transition. v6ops-Chair Fred Baker later the week said China Telecom was on a very aggressive schedule for Ipv6 adoption.

Bill Huang, General Manager of the China Mobile Research Institute agreed that IPv6 deploymnt was urgent, but he said. the pressing issue of Ipv6 was not adequately addressed. „Our view is, based on the result of what we have seen, the fact that many people are talking about Ipv6 does not really mean that we are ready“, Huang said. „Unfortunately, there has not been a very wide deployment in general, so, a lot of problems are only starting to be uncovered.“ Huang pointed to several technical issues he saw not yet solved by the IETF. He mentioned transition and translation mechanisms, one big problem was also the compatibility of applications.

Service platforms, said Huang, had been developed for years based on IPv4. Migration here was not easy. -Also packet filtering still was a problem with IPv6 devices. Huang asked the IETF to do more to really fix still existing problems. At the IETF there are several groups working on IPv6 and Ipv4-IPv6 transition mechanisms.

Lee Xiaoping, CTO of the CNNIC confirmed that real deployment of Ipv6 in China is still very low. CNNIC is not only domain name registry, but also national IP-address registry. The ratio of Ipv4 to Ipv6 queries was 286-1, according to CNNIC. Only 102 /32 address blocks are in use, and many things still were test activities, many Ipv6-queries came from inside China, said Lee. Lee said, the core problem was lack of applications and content. „We need to more IPv6 applications. If we don't get more IPv6 application, there will be no IPv6“, he said. The lack of IPv6 support from a service like Skye was heavily criticized by IETF participants.

IAB Program on privacy and data protection in standardization

IAB Chair Kolkman on Monday gave an update on the IAB ongoing privacy work that includes a highly political debate about how Internet developers should address policy issues like privacy in their work. One of the ideas is to add a section on “privacy considerations” to every RFC document (analogue to security considerations, iana considerations). The Center for Democracy and Technology has proposed to adopt two documents that look into policy considerations in general like „how much could new standards become a tool for censoring or third party control over users. In the privacy documents CDT references also the European Data Protection Directive as a „political“ standard. The IAB together with the W3C is preparing a workshop on data protection and privacy (Dec, 8-9) and received over 50 contributions in response to its call for papers.

IETF Administrative Plenary

Fierce Debate about the future RFC Editor

There was a fierce debate about the future of the RFC Editor function after a presentation of Glen Kowack who was hired as a consultant to look into the issue, document the process and develop recommendations for the future RFC Series Editor function. Kowack's main three recommendations are:

- The RFC Series Editor has to have technical writing expertise (he does not need to be a computer scientist, but have a lot of knowledge about the technology)
- The RFC Editor must listen to the Community (IAB hires and fires the Editor)
- The RFC Editor function must be managed by one person (while the independent submission stream is independent)

The editor should be selected by a search and selection committee, not only by IETF And IAB leadership. Kowack was attacked by Ted Hardie who questioned the recommendations (“the worst piece of shit”) mainly because Kowack, according to him, was proposing a system of “adult supervision” over the technical community and had not tried to be more visionary about the development of the series. Leslie Daigle (ISOC) said Kowack had not described the reasoning behind the recommendations at all and therefore had missed to fulfill the tasks given to him. The core question obviously seems to be how much authority is attached to the Editor function to shape the

development of the series. Discussion is ongoing, and Kowack was asked to produce a new version of the current document by November, 22.

Budget

According to IAOC Chair Bob Hinden's presentation the budget will increase incrementally only over the coming years. Expenses in 2010 according to a Year end forecast are 4,66 Million Dollar and slightly lower than the budgeted 4,697 Million Dollar. For 2011 the IAOC expected many expenses going up about two percent, said Hinden, Secretariat cost would go up 5 percent after steady costs for two years. The total budget plan notes expenses of 5,003 Million US Dollar, revenues also are expected to rise slightly, mainly from raised meeting fees (up to 650 US Dollar from 635 US Dollar), totaling. A decision about the day pass experiment is to be taken in December by the IAOC, with community input still welcome. With additional revenue from meeting fees the IAOC hopes to lower future contribution by the ISOC who pays for nearly a half of the expenses of the organization. (Hinden said the hope was to bring ISOC contribution needed down to 1,76 Mio in 2012, while IETF still needed 2,215 in 2011).

The next IETF meeting will be held in Prague, March 23.3. - 1.4. 2011 (just right for an April Fools' day RFC)