Council of European National Top Level Domain Registries

Report on IETF 92 Dallas

22 - 27 March 2015

Table of Contents

Highlights

DNS Privacy work ongoing in both DNSOPs and DPRIVE

Privacy has become a major work item in many areas (for transport area, see below), good old DNS only being one of them. With proposals being picked (DNS over TLS looks like the winner in the DPRIVE WG for now, implementation still is far out. While some operators point to the need to keep a grip on DNS traffic for security monitoring reasons, choices are also influenced by the question of how a solution influences the respective business model. Data minimization keeps data away from the root zone operators, for example.

DNS Privacy Work - DNS over TLS gets a nod

DNS over TLS seems to be the preferred variant for more private DNS queries and answers. The Dprive WG in Dallas after an undecisive round of hums on the potential bundling of different solutions expressed a clear preference for DNS over TLS (and DNS over TCP at the same time). This „rough consensus" was meanwhile confirmed over the DNSOP mailing list. DNS over TLS is the draft concept proposed by VeriSign lab researchers, VeriSign Inc. (plus Paul Hofmann who presented the draft in Dallas).

Of three drafts for a more private DNS presented and discussed in Dallas, Private DNS (Philip Hallam-Baker, Commodo) got the least support, while the Confidential DNS (Wouter Wijngaards, NInet Labs&Glen Wiley, VeriSign Inc) did rank second according to the „hums" of the Dallas Dprive participants.

Confidential DNS

Wiley presented changes in Confidential DNS as being simplified compared to earlier versions. The core concept of the draft remains the new „ENCRYPT" resource record (RR), structured as flags, algo, id, (in decimal) – and data (in base-64). The domain name of the ENCRYPT record is '.' (the root label) for hop-by-hop exchanges. Clients fetch the ENCRYPT RR from the server they want to contact, and use the the public key retrieved as a result from ENCRYPT RR to encrypt a shared secret or public key that the client uses to encrypt the sections of the DNS query. The key is refetched after the TTL expires. If the key fetch fails or the encrypted query fails, a fall-back to non-TLS DNS is performed.

Confidential DNS offers both opportunistic and authenticated, the latter using DNSSEC. The authenticated version has the key included in an extra DS record in the parent's delegation for the authoritative server. For recursive servers the key is at the reverse IP address location.
The major shortcoming is that Confidential DNS is that it has not yet been implemented, contrary to DNS over TLS. „Only in our dreams", Wiley said, it was implemented.

DNS over TLS

The DNS over TLS also saw some changes in the new version. Notably, the DNS over TLS draft does now include a two step mechanism. First, a client will try to start a TLS connection over a special, newly assigned port. If the port is available, the connection is opened, if it it blocked, the connection should be set up by sending a real or dummy query over port 53. The request should be flagged with a new EDNSO flag „TLS Ok" (TO). In case the client gets back the TO bit the TLS

secured DNS connection can be established. The order of using port-based or start-TLS based encryption can be made locally.

If TLS fails, a fall-back to plain DNS is performed. According to Paul Hoffman who has joined the DNS over TLS author group (and announced that three alternative proposals would be withdrawn) the group's choice was to not have people think about how DNS had to be changed. Instead clients should know how to do this. Blocking by middle boxes can create problems for DNS-over-TLS (see RFC3207). The draft describes

a) DNS client sends T0=1 receives T0=0	→ fallback to normal encryption
b) DNS client sends T0=1 receives T0=1 middlebox does not understand TLS negotiation	→ if cleared go on, otherwise fall back to normal, or retry
c) DNS client sends T0=1 receives no response	→ fall back to normal, retry later

In general, clients that attempt TLS and fail can either fall back on unencrypted DNS, or wait and retry later, depending on their privacy requirements.

Issues brought up during the Dallas session on DNS over TLS was how multiple applications wanting to open TCP/TLS connections at the same time should be dealt with. There number of TCP queries might be bigger than envisioned, said Peter Koch (Denic). The initial synchronization therefore still needed discussion. Paul Wouters (Apache) said, there was no added value from DNS over TLS, when encryption of connections through VPN was used anyway. Plus people would not be safe from resolvers when using the concept, unless they would decide to use their own resolver.

The DNS-over-TLS concept mainly targeted people using public DNS and no VPN – it was about „privacy for everybody", and §"a form of privacy with a limited scope", John Heidemann told the sceptics. What made DNS over TLS attractive, a representative from Microsoft said during the Hums, was that it needed the „least innovation".

Yet another proposal: Do not focus on TCP alone

In the meantime another potential solution has been added based on using DTLS/UDP instead of TLS/TCP. The authors, Dan Wing and a group of Cisco authors do argue against focusing on DNS over TCP only because of TCP „head-of-line blocking" and the need for complete TCP handshakes to resume sessions (more round trips). So far there was not enough support to adopt the DTLS document as working group document. The discussion nevertheless is going on.

Measurements: How private do you get with DNS Privacy mechanisms

VeriSign Labs researchers this time did p
resent a draft on the evaluation of privacy mechamisms – according to Allisen Mankin and co-authors, there is a need for much more differentiation on what level of privacy could be provided by the proposed DNS privacy enhancing mechanims. The group came up with a draft on the „Evaluation of Privacy for DNS Private Exchange".
Privacy gains should be checked for the various links:

Stub -> Recursive
Stub -> Proxy
Proxy -> Recursive
Recursive -> Authoritative

DPRIVE is mainly looking into the Stub-Recursive link, but evaluation in the ongoing work for more privacy concerned a broader spectrum.

Another aspect to differentiate/classify solutions concerns the nature of the attacker. The draft lists

pervasive monitor (Type 1A)
direct monitor (Type 1B, specific target)
malicious monitor (Type 2).

Mechanisms listed include Tor-like mixing networks, hiding requests in dummy traffic, private information retrieval techniques and encryption. Protocols used for these include IPSEC, TLS (DNS over TLS) or special confidential TLS (which is described in the text as special purpose encryption or private DNS („special purpose encryption with a privacy broker).

The various concepts offer different privacy properties, with the potentially best effect available from a mix of several mechanisms:

„Consider a hypothetical system in which mixing networks (for unlinkability) and randomized encryption (for undetectability) can both be applied, thus providing for unobservability, a stronger property than either of the two along."

To classify the various mechanims in this way, while academic to some extent, might help to clarify potential effects and efforts necessary. The document still is under consideration by the DPRIVE WG.

More Privacy Work in DNSOP

DNSOP also touched briefly one another privacy enhancing proposal, the qname minimization. Document author Stephane Bortzmeyer pointed to questions over the potential negative effects of the minimization of data. VeriSign in particular has called to consider the potential „trade-off" qname minimization was triggering: „While enhancing privacy, it may also reduce visibility into security threats."

Name collisions or a recently announced Microsoft remote code execution vulnerability would not have been detected with hiding full domain name queries from root servers. VerSign calls for further analysis and the potential development of „new methods for sharing information within the DNS", while eagerly underlining it was fully in favor of qname minimization (VeriSign promised RAND licensing option for its qname related patents – how much that offer attracts or rather deters operators has to be seen).

Other questions concern the hiding of the qtype and two alternative versions of qname minimization, „aggressive" and „lazy" and offering a different amount of data. While several interventions underlined with privacy as the ulterior motive, losing of data was of no concern to the WG. Bortzmeyer said, it was clear that there were players with different or even competing interests.

DNSOP WG in the Hot Seat: What to do about special name applications

One of the big issues the DNSOP WG finally comes around to deal with are alternative names and a potential second application stream, beside the one of ICANN. The list of special name applications (RFC6761) has recently been growing considerably. Two applications were touched briefly in Dallas during the DNSOP session:

.alt as an alternative space for names that are not supposed to be rooted in the DNS, and where „normal registration and lookup rules do not apply", the proposal is presented by Warren Kumari (Google) and Andrew Sullivan (Dyn). On the question if there was interest for such a space, Kumari said yes.

.onion as a label used by the TOR network, with around 30.000 onion names out (.onion" names label Tor hidden services and are resolved via TOR servers, hashofpublickey.onion). The reason why Tor is now urgently seeking to get accepted as a special use TLD is a decision by the CA/Browserforum to no longer grant certificates under an exception. If .onion gets not approved by the IESG (who is in charge of deciding about the special use TLDs) certificates needed for https-connections https://www.facebookcorewwwi.onion/, https://blockchainbdgppzk.onion/ or the Intercept (https://y6xjgkgwj47us5ca.onion/) will fail. Richard Baines (Mozilla) who presented the .onion case, beside the Certificate-deadline noted also reasons more general (and true for other special use TLDs), namely the continued leakage of potentially private information in the DNS and the increasing load of bogus queries on DNS resolvers.

Another group of special TLD applicants results from the ICANN commissioned study on name collisions, yet the proponents meanwhile have cut their requests from more than half a dozen to two for this point it time: home and corp.

This is not the end of the list, though. There is also a request for an experimental, non-DNS gnu name system TLD, .gns. All in all a number of 41 requests was mentioned by the Chairs which has led to concerns that the IETF might attract requests of applications that want for some reason to circumvent the more tedious (and expensive) TLD application procedure at ICANN.

The Chairs wanted to defer most of the discussion mostly to a dedicated intersessional meeting, now set for May, 12, with an f2f-option for DNSOP members present at the RIPE meeting in Amsterdam. Currently announced time is 1600-1800 UTC.

Innovating Transport, the Spud Bof and TCPINC:

„A series of tubes" - a fierce debate about SPUD

A fundamental change to the Internet could be made, if Spud (which stands for Substrate Protocol for User Datagrams) would not only become an IETF Working Group, but also be widely deployed. But not so fast. Spud is a product of the IAB's stack evolution program, and a January Internet Architecture Board workshop in Zurich on „stack evolution in a middlebox internet" (SEMI). It was discussed highly controversial, some see it a s a „peace offer to middle boxes", others warn that it could deal a blow to a neutral net.

The basic idea is to create a mechanism for grouping UDP packets together into a „tube" with a defined beginning and end in time. State will be easer to be kept. Devices on the network path between the endpoints speaking SPUD may communicate explicitly with the endpoints outside the context of the end-to-end conversation.

Those presenting Spud said the motive was to get rid of deep packet inspection, but allow for network management and conveying necessary information for that. You might give up some information to the middleboxes, in order to get transport, Ted Hardie, Google explained.

To differentiate UDP when used for SPUD from regular text based usage a magic Bit (frist 32 bits) will be used .)

The List in the Spud header fields according to Joe Hildebrand (Cisco) could be:

- 32-bit constant magic number (d80000d8 (hex), or 1101 1000 0000 0000 1101 1000 (binary))
- 64 bits defining the id of this tube
- 2 bits of command
- 1 bit marking this packet as an application declaration (adec)

- o 1 bit marking this packet as a path declaration (pdec)
- o 4 reserved bits that MUST be set to 0 for this version of the protocol
- o If more bytes are present, they contain a CBOR map

Several reasons why the IETF should tackle the work were brought up during the very well attended Spud BoF session in Dallas:

- an attempt to innovate transport (in the more and more ossified stack and transport layer)
- allowing for network management while protecting privacy by making dpi which breaks end-to-end encryption unnecessary
- enabling to reduce complexity in the Real Time Communication Web (RTC Web); Spud should help against the complexity illustrated in designs like „encapsulation of SCTP over DTLS over ICE/UDP provides a NAT traversal solution"

The new concept, while solving some issues of a middlebox internet, creates new problems on the other hand. One is that suddenly a new layer of meta information about data packets is injected in the net. It had to be made sure that people would not stick their name on to „Mr. Potato Head", Ted Hardie said. Because the meta information might suddenly leak the very information that people tried to secure/hide in encrypted streams before. Mark Nottingham, Akamai, warned not to make such changes without broad (IETF external) debate and transparency about the newly created meta information.

While encryption of communication on that meta-layer is under consideration, as Christian Huitema from Microsoft presented in Dallas (using DTLS). The prototype now promoted for experimental use, has been stripped of privacy mechanisms to keep it easy.
Finally Spud was also prone to specific new attacks, as Christian Huitema acknowledged later. For example it was well possible that attackers sent close Spud packets to the middle boxes, breaking open media streams.

A WG was not formed this time, but work can be expected to continue. One of the BoF Chairs told this reporter, that a draft charter could be drawn up and brought back to another BoF at a later point in time.
Some attendees warned to give up the end to end principle by allowing for the new meta layer. An interesting question to discuss is how much the Spud effort equals offering a peace treaty to middleboxes (as Huitema described the initiative) and how much on the contrary it is giving in to the middlebox internet. Instead of making traffic talking to the middleboxes, they could also just shut up and encrypt everything, on participant said.

TCPINC

Encryption of transport as well as encryption throughout the stack as promoted by the IETF recent statement is in fact on the table at the IETF. The TLS WG saw a fierce debate over the choice of TCPINC or TLS.

Not only the DNS is looking for added protection against passive eavesdropping, the TCPINC WG is doing the same for TCP, the transport protocol. At the Dallas meeting the WG Chairs, and even more so some Area Directors, wanted to get a final decision which of two options to chose to develop further.

The protocol that received some interest over recent month is TCPCrypt, proposed by a group of Stanford (and other) researchers. TCPCrypt could, according ot the draft, provide „unauthenticated encryption and integrity protection at the TCP layer". The idea is that applications would not have to be changed. When hosts agree to perform TCPcrypt, cryptographic keys will be exchanged using the data portion of TCP segments. After that encryption would secure confidentiality and integrity of transmitted applications. Downgrade attacks remain

possible, as well as MITM attacks where an attacker controls a part of the network.
The second proposal is reusing TLS für TCP, a TLS is negotiated at the start of a TCP session. The idea according to author and TLS guru Eric Rescorla (Mozilla) was simple and also geared toward allowing applications to go unchanged:

„The SYN and SYN/ACK messages carry TCP options indicating the willingness to do TLS and some basic information about the expected TLS modes. If both sides want to do TLS and have compatible modes, then the application data is automatically TLS protected prior to being sent over TCP. Otherwise, the application data is sent as usual."

In Dallas there was no clear decision by the WG participants on which one to chose. While some participants recommended to let both parties add to their drafts (and work on running code,for example Steve Kent, BBN), there was a strong call for a quick decision and deployment of one option instead (for example Ted Hardie, Google).

Meanwhile the academic group has started implementing and reacted to requests to use TLV as a framing protocol. A good overview is here. Both proposals can be checked out and implemented on Github here (TLS) and here (TCPCrypt). Implementing the latter can bring you into the „TCPCRYPT Hall of Fame" ;-).

Email privacy and a reopening of Open PGP in the making

More confidentiality throughout the stack has been declared to be on top of the IETF agenda since the IAB declaration in Hawai. There is also a bunch of how to further enhance email privacy, by minimizing email meta data for example.

Minimizing Meta Data - a memory hole

Minimizing meta-data could be done easily right away, according to a proposal presented by Daniel Kahn-Gillmore, representative of the American Civil Liberty Union who has become a permanent attendee for the US civil rights organisations. The idea is based on the „stone age" RFC 822, as used in RFC 6533. For mail bounces or mail forwards the header does not contain the subject line anymore. Instead the subject line is integrated in the body of the email. A subject line not carefully chosen can sometimes reveal the main message of the email (even if the mail body is encrypted, for example „Contract negotiations - it is a go"). With the header put into the encrypted body and a „dummy subject" included in the open header, less information might be leaking.

The mechanism eventually could be extended for additional header fields, Gillmore thinks. The concept was currently discussed in the open source community, a potential solution for MIME was also under discussion, Gilmore said. Phased introduction at MUAs was possible. For the next IETF preparations are underway to reopen the OpenPGP working group to cover either openPGP maintenance work or include additional work, as for example the Kahn Gillmore's „meta-data memory hole".

Ladar Levison: Darkmail

A non-IETF concept to enhance email privacy was presented during the SAAG working group by Ladar Levision, alleged provider of email services for Edward Snowden, who closed down his company when subpoena-ed by the US authorities into handing over his SSL keys to get access to the email stored in encrypted version on the lavabit servers. Levison who is joining forces with Phil Zimmerman's Silent Circle company wants to create a completely new email system for security/privacy sensible customers.

The „Darkmail" suite will include alternatives to IMAP (DMAP), SMTP (DMTP), DIME (MIME) plus new encryption processes. In essence Darkmail is based on encrypting the different parts of the email one by one in an onion-like way and in the end even hide sender and receiver. For not as paranoid users, mail would be sent to a mail provider and the individual receiver would only be decrypted there. In addition users can integrate alternative, and more secure key material.

Wgs and BoFs

DANE: Case-folding, normalization remains an issue

With many documents making their way up through the RFC process (done DANE-SRV, DANE SMTP advanced to IESG, Update to and Operational Guidance for the DANE Protocol and OPENPGPKey in Working Group Last Call) the Working Group focused on the problem of case folding and/or normalization steps for the local part of email addresses.

The well-known issue is seen as potentially causing trouble for OPENPGPKEY (as well as for MIME). When fetching or verifying keys (DNSSEC-)securely stored PGP key in the DNS a potential mismatch can be caused through varying ways to deal with case-folding for the local part of the address. The current approach in the draft by Paul Wouters (Apache) is to refer to the path taken by RFC5321 and its predecessors which is that only the recipient MTA was allowed to interpret the local-part of an address.

„A client supporting OPENPGPKEY therefore MUST NOT perform any kind of mapping rules based on the email address. As the local-part is converted to lowercase before hashing, case sensitivity will not cause problems for the OPENPGPKEY lookup."

While discussion has been going on for some time after the meeting on the mailing list, most mail experts think it is not an issue to be solved in the DANE WG.

The WG also got a presentation by Eric Osterweil on VeriSign Lab's work on S/MIME library status update, tools can be checked out on:
https://github.com/verisign/smaug,
https://github.com/verisign/smaug-tbird-plugin
Osterweil also invited use of DANE provisioning portal,
https://dane-provisioning.verisignlabs.com/ (which was not available from my machine).

Co-Chair, Olafur Gudmundson, finally requested participants to consider the updating of milestones and potential recharter. It was a good time to bring new work. One item presented in Dallas, the use of DANE for association of payment information, could be a candidate for the WG. The proposal written by VeriSign and Bitcoin wallet provider Armory Technologies is, that „a payment association record associates an Internet service identifier such as an email address with payment information such as an account number or Bitcoin address".

With all of this additional information, keying and potentially payment information, stored in the DNS the need to make DNS exchanges more confidential (and more secure through DNSSEC) becomes only more urgent, some experts say.

EPP Ext - VeriSign registers first big bunch of extensions

The EPPext working group has come close to finalizing its work items, especially, it has established a formal registration process for new EPP extensions, based on the finalized RFC 7451. The first 18 EPP extensions registered with IANA came from VeriSign, with some concerns being raised about IPR notices attached ot the extensions. Will the new registration process indeed fulfill its aim to better manage and coordinate the development of EPP extensions?

A list of designated experts has been selected by the IESG to review the extensions filed by various registries/parties. Reviewers are: Scott Hollenbeck (VeriSign) – primary, Alex Mayrhofer (nic.at), Ning Kong (Cnnic), Roger Carney (GoDaddy) and Jim Galvin (Afilias) – all secondary.

The WG has been according to its charter discussed a short list of extensions as candidates for the new registry
Internationalized Domain Name Mapping Extension for the Extensible Provisioning Protocol (EPP) (by Uniregistry) – needs review section and a decision if it will be standards track or not, according to Co-Chair Jim Galvin
Key Relay Mapping for the Extensible Provisioning Protocol (by SIDN) – ready for last call, if not intended to be standards track according to WG Co-Chair
Launch Phase Mapping for the Extensible Provisioning Protocol (EPP) (by VeriSign, Cloud Registry,Centralnic) – ready for last call according to WG Co-Chair Jim Galvin, there is again the question if it will be standards track
Mark and Signed Mark Objects Mapping (by ICANN), ready for last call according to the Co-Chair, intended for standards track

Procedurally the WG intends to make a difference between extensions that are intended to become standards track RFCs, for which EPPExt after rechartering might become a home, Galvin said in Dallas, and those documents that are only informational. The latter will only be reviewed by the designated experts and, given the requests fullfill the formal obligations sent on to the IANA EPP registry.

While the WG still is pondering about the initial document list, and especially on their status of either standards track or informational, and new documents are added via the mailing list, a first big bunch of 18 extensions, all tagged as informational, has been put forward by VeriSign and has already been passed and added to the new IANA EPP extension registry without discussion by the WG. The VeriSign „dump" registration seems to anticipate the future process for informational documents. After a very short exchange over some editorial pointers, plus a slight concern about the IP boilerplate notice on the VeriSign documents the documents were sent on to the EPP registry by Scott Hollenbeck. VeriSigns IPR boilerplate forbids, Alex Mayrhofer from nic.at wrote in his review, to "copy or communicate" the documentation without "written prior consent of Verisign" which in fact on the web would make even viewing problematic.

Also listed in the brand new registry are four older standard track extensions, namely RFC3915 (E.164 Number Mapping for the EPP), RFC5076 (ENUM Validation Information for EPP) RFC4114 (Domain Name Grace Period Mapping for EPP) and RFC5910 (DNSSEC Mapping for EPP).

One interesting issue with regard to the EPP registry is the question how much it will help „to manage and coordinate" extension development, and more specifically to avoid duplicate efforts. According to the RFC 7451 „designated experts should be permissive in their evaluation of requests to register extensions that have been implemented and deployed by at least one registry/registrar pair. This implies that it may indeed be possible to register multiple extensions that provide the same functionality. Requests to register extensions that have not been deployed should be evaluated with a goal of reducing functional duplication." A registrant wanting to sumbit

an un-deployed extension that is similar in functionality to a registered one will be asked to reconsider.

How much IPR notices like the one on VeriSigns registered extensions will turn out to be a barrier for harmonisation, is an open question.

The next step for the WG after finalizing off its initial set of documents is a potential recharter. Galvin noted the WG could recharter to become the official place for extensions that want to become standards track (as long there are documents in the pipeline), additional topics of interested he mentioned are the Whois developments around ICANN and the Registration Operations Workshops (ROW) that have now been held two times in conjunction with the IETF and are on the way to become an institutionalized event linking standardization and operational work. The next ROW will take place on July,19 in Prag, alongside IETF 19.

Modern BoF - Taking up where ENUM left?

The IETF starts talks voice numbering again, following proposals by Henning Schulzrinne (Federal Communications Commission) and Jon Peterson (Neustar, which in fact is provider for the North American Number Registry). Will the ITU like it?

Basic idea is that MODERN will „define a set of Internet-based mechanisms for the purposes of managing and resolving telephone numbers (TNs) in an IP environment". With voice slowly moving to all IP, the proponents argue, there is a need for a new system to manage Tns, because

- the model that Tns have an association to one single service provider is gone
- network locator feature vanishes (instead TN is more individual or organization identifier)
- devices, applications, and network tools increasingly have to request and acquire TN delegations from authorities.

Both a hierarchical or p2p tree for number management would be possible. Privacy of number management is said to be of prime interest.
While in Dallas there was considerable interest to have a WG chartered, questions now are raised as to would the new framework change the number allocation model - from a two-tier model to direct allocation with no need for porting? Would that lead to a more centralized model (with national number registries eliminated)? As discussion on the list involving representative from carriers shows there is a considerable push-back against an „overreach" of a new WG with people pointing to the failure of ENUM.
Jon Peterson who presented the problem statement for the WG, explained that limitations of the DNS (rigid syntax, with security added getting complex) motivated the idea to develop an „independent framework and information model for querying and responding to requests concerning telephone numbers and call routing that allows a richer expression of both questions and answers".
Work items listed in the draft charter for now:

- an architecture overview, including high level requirements and security/privacy considerations
- a description of the enrollment processes for existing and new TNs including any modifications to meta data related to those TNs
- a description of protocol mechanisms for accessing contact information associated with enrollments
- a description of mechanisms for resolving information related to Tns
- a protocol mechanism for resolving TNs which will allow entities such as service providers, devices, and applications to access data related to TNs, possibly including caller name data (CNAM).

ENUM, SPEERMINT, and DRINKS work would be considered.

CFRG selecting new curves for TLS - messy process

CFRG Co-Chairs decide for Curve 25519 (Dan Bernstein) and Goldilocks (Mike Hamburg, MIT) after a quite messy process. There is more work to do when the IETF wants to create its own Crypto standard stream.

The Crypto Forum Research Group has selected two new curves to be used for encryption in Transport Layer Security (TLS). One is curve 25519 (y^2 = x^3 modulo p = 2^255 – 19), developed by Dan Bernstein (University of Chicago and Eindhoven), the other is the much newer Goldilocks (x2 + y2 ≡ 1 - 39081x2y2 mod 2448 - 2224 – 1), designed by Mike Hamburg. Hamburg is a fellow researcher who has cooperated with Bernstein.

The curves are similar in the way that both are elliptic curves, acknowledged by cryptographers to be faster and more secure to side channel attacks. Both curves have been aced as safe by the safe curve project. Obvious differences are, Bernstein's 25519, has been around since 2006 with no successful attack known so far. Goldilocks on the other hand is a newby, and was created only in 2014.
-
At the same time 25519 allows for a 128 bit security level, Goldilocks goes up to 223-bit. The TLS WG of the IETF had asked for curves of 128bit and above, CFRG Chair Paterson said. Both curves are faster than the NIST curves in use so far. CFRG participants, especially Microsoft had asked to include a larger one in the proposal to the TLS. Microsoft had hoped that the curves of its team could be selected, but finally failed after fierce debates on the mailing list.

The CFRG in the end was unable to reach consensus, so that the Chairs finally decided on the selections, based on answers to a list of polls. In the IRTF, contrary to the IETF, Chairs can make such a decision. Paterson during the WG meeting in Dallas hinted at the animosities during the debate – the climate on the mailing list now, after the decision had been made by the three chairs, had improved a lot, he said.

Two crypto experts talking to this reporter noted what they saw as deficiencies of the selection process. One argued that the process envisioned had been to decide on a requirement document – and only then chose for the proposals put forward. The other pointed to attempts to put pressure on the Chairs with regard to accept or reject certain curves.

The IETF selecting crypto standards for their own protocols and protocol suites results from revelations in the documents released by US whistle-blower Edward Snowden. After NIST, the National Institute for Standards and Technology, had to confirm that the NSA had tampered with its algorithm selection at least in one instance (decreasing the level of arbitrariness in arbitrary numbers) the IETF started a process to select its own algorithms – or better have potential candidates checked and recommended by the IRTF Crypto Research Forum(IRT).

While most participants can agree that the process so far has been messy, many nevertheless say that the results are ok. He could assure the group, that the process had not been manipulated, Paterson joked during the WG session in Dallas, „at least not by me". A systematic problem of the Crypto standard selection is that the mathematics involved is only for hard core mathematicians. Even of the WG participants, many do say, that the cannot check on the curves themselves.

While NIST officials had warned IETF did not have enough Crypto excellence on board to select their own curves – an attempt to dis-encourage the IETF/IRTF to strive for crypto independence? - TLS WG Co-Chair

The selection now made by the starting point, TLS WG Chair Sean Turner said, he expected the TLS WG (and other WGs) to make more requests to the CFRG to provide for Crypto for IETF protocols. Acknowledgement for the IETF standard by NIST would be nice to have, but not necessary, most CFRG members agreed in Dallas. The WG did however push the Chairs to

announce and present the selected curves during an upcoming NIST conference that will address the selection of new curves. The CFRG work for the TLS WG group's call will be completed when signature standards will have been chosen. The CFRG will present the complete package of curves and signature algorithms to the TLS WG.

IETF News

IAOC Trust will hold IANA related IPR (and domains)
New IANA Chair Andrew Sullivan (DYN)
IETF going to Latin America for the first time in 2016

IANA transition and IANA IPR

IAOC Trust Chair Tobias Gondrom informed IETF community that the Trust was prepared to hold the IPR for IANA. The answer was elicited by a request of the ICG (IANA Stewardship Transfers Coordination Group). A meeting of the IANA Plan WG, official place for the IANA transfers proposal by the IETF, was canceled for Dallas. IETF Chair Jari Arkko pointed out during the plenary that the community was waiting for the ICANN community to finalize their proposal for oversight for the domain name related services for IANA in the future.

New IAB Chair

The IAB has a new Chair. Andrew Sullivan, Dyn DNS, follows Russ Housely (Vigil Security). Housely had been IETF Chair for the maximum possible time (two terms) and been kept as a leading figure once his time as IAB Chair came to an end. As IAB Chair again he served the maximum time (two terms). His openly declared sponsoring from the NSA obviously seemed to be no problem for members. Sullivan taking over means also a change from a security expert to a DNS expert. Sullivan has served for DNSEXT CoChair when DNSSEC was standardized. Sullivan has also participated in a number of ICANN and IGF meetings.

Meetings and Visa

For the first time in its history the IETF is gathering in Latin America in 2016. The meeting in Buenos Aires will be prepared by a series of workshops in Latin American countries organized by ISOC.

Tobias Gondrom IAOC, reported that the much debated VISA problems for non-US citizens could be resolved for the Dallas meeting with more ease than in years before. In an effort to give much time for the preparation for VISA issues and travel preparation in general meetings will now open registrations sooner. Registration for the Prag meeting has been opened right away after the Dallas meeting.

Next meeting will take place in Prague,19 - 24 July 2015