Session 1: Monday

- Administrivia (5min)
- Document Status (5min)
- Exported Authenticators
  https://datatracker.ietf.org/doc/draft-ietf-tls-exported-authenticator/
- Record Size Limit Extension for TLS (15min)
  https://datatracker.ietf.org/doc/draft-thomson-tls-record-limit/
- SNI Encryption (25min)
  https://datatracker.ietf.org/doc/draft-huitema-tls-sni-encryption/

- Administrivia:
  -
  - Wednesday's agenda is not yet solid due to [reasons]
- Document status
  -
  - ECC cipher suites in RC editor's queue
  - ECDHE_PSK is waiting for EKR's no objection
  - Adopted: DTLS 1.3, Exported auth, TLS cert compression
  - 2nd WGLC for TLS 1.3
  - DANE record and DNSSEC auth chain just completed WGLC
  - in progress: handshake traces (mt), and applying GREASE
  - needs work: delegated credentials
- Nick Sullivan, Exported Authenticators
  -
  - given TLS connection, bind a given cert to the connection.
  - mt: do we know of any other open issues? besides we don't have formal analysis that this is safe to use.
    -
    - Nick: no other open issues. Intuition of a number of cryptographers is that this is ok.
    - EKR: probably would be worth making another pass at this to get formal analysis.
    - SPT: how many read it? 10
    - SPT: **if we don't hear anything by singapore, we'll push to WGLC.**
- MT, TLS record size limit
  -
  - discussions inspired by max fragment length.
  - plan is to deprecate max_fragment_length, define an alternative
  - EKR: this is an individual draft? yes. support bringing it into the WG.
    -
    - what would happen if there was a version of TLS that encrypted the certs?
    - mt: they'd have to respect the limit.
    - ekr: having an exception for plane text records doesn't make sense if we're moving to encrypted certs. Specifically, client knows on server_hello what the server will accept.
  - SPT: How many people red the draft? 10
  - SPT: **Hum for working group adoption?** Obvious, we will adopt it (no hums against.)
- Christian: Encrypted SNI
  -
  - Many of us want to encrypt SNI and now with HTTPS deployment and DNS over TLS, it's starting to make sense.
  - two solutions: 1) HTTP fronting, TLS to "fronting.example.com", inner HTTP is to "hidden.example.com", and 2) a "combined ticket" solution.

- Kazuho Oku: interested in implementing encrypted SNI. bit sad that RTTs increase. we should consider transferring more data over TLS to allow 0RTT.
- bret jordan: question/concern: Google talk at RSA on Beyond Corp strategy and extensive use of middlebox technology to protect network. Concern with Encrypted SNI means that enterprise orgs, will have to decrypt everything that traverses to identify anything.
    -
        - CH: google deploys fronting.
        - KM (AD): do want to make sure that this convo is followed up on ops area. Good that you have attacks enumerated but do you have uses enumerated.
- PHB: you can also view TLS as two separate protocols: 1) key agreement protocol that results in a ticket; and, 2) a session protocol that transfers content… keys don't have to be the same. Also, a "new type of cert" terminology is unhelpful… what you will have is nothing that will look like a PKIX cert. E.g., you wouldn't want these things in CT, either, since this needs to be secret.
- Subodh: some concerns about fronting and HTTP delegation. Where are the application layer solutions?
- EKR: we should be working on this. We have a few too many mechanisms and we should bash them into one. There are some inspection endpoints which examine the SNI and sometimes server cert to determine whether to MITM the connection. There are a bunch of natural extension for co-tennanted solutions where this could be useful and we'll need to enhance some HTTP mechanisms. We don't have solution for cases where you wish the fronting server should not have access to the plain text. Would be great to streamline the discovery of who the fronting servers are.
- KM: just want to emphasize that ops needs to be notified and have a chance to speak up.
- Nick Sullivan: echo EKR's points… separating fronting server that does not have access to plaintext vs. one that does. Maybe there is a more general solution to content tunneling.
- Bryan Ford: we should think of this as does it nest… negotiate a tunnel within a tunnel, does it work? Think about that carefully. Discovery problem is very similar to a routing problem in a higher-level overlay network.
- Joe: how many people read the draft? a lot.
- Joe : **Who thinks that this is something the WG should work on?** Very loud for yes, something like 3-4 hums against.
    -
        - We definitely want to work on this.
- Bryan Ford: to what extent do we need to think about traffic analysis issues?
- Yoav Nir: if I started filtering based on SNI and sometimes the client doesn't send SNI so filter in the server cert sent in response. But now 1.3 encrypts certs so we can't do this and you're taking SNI away from us.
- Daniel Frenke: don't think we should adopt this yet, we should have a more streamlined draft worked on by a design team. SPT: this is essentially that.
- EKR: too early to pick any given solution. We definitely want to work on this… between a coalition of the willing and a design team with solutions.

Session 2: Wednesday

- Agenda/Administrivia
- Exported Authenticators (EKR)
    - draft 21, hopefully close
    - WGLC #2 ended yesterday

- Changes since -19
  - shorten HKDF labels
  - make post-handshake auth imp option
  - add per-ticket nonce, each ticket is assoc. w/ new PSK
  - new section 0-RTT anti-replay
- Mandatory anti-replay (PR# 1059)
  - requires some bounded mechanism, but no specific technique
  - **Should we adopt this? Any objections?**
  - MT: every instance has to ensure that it only accepts the same 0-RTT once... which means an unbounded state problem
    - EKR: imp in NSS would guarantee "as 0-RTT"
    - ... "you must accept 0-RTT data once"
    - MT: We've got a window, only accept in that window, no guarantee either.
    - (No objections)
- PR# 1053: Hashes that aren't hashes
  - HKDF-Expand-Label included a hash function that occasionally is not a hash.
  - essentially, SHA156(empty string) passed frequently to something else.
  - Probably worth saying "you can pass a null value, and not pass a hash"
  - EKR: any opinions?
    - MT: noticed while doing vectors draft, we do this once every handshake. I don't care.
    - EKR: there are two places that it can happen.
    - MT: still, I could care less.
    - Hannes: doesn't make sense to change since people have implemented like this.
    - RLB: I agree with MT and Hannes. Like the current mechanism, can opt with table of hash values.
- Placeholder: NAT/Middleboxes
  - TLS 1.3 starting to show increased connection failure rates.
    - hard to meausre but 1-10%
    - Problem seems to be middleboxen
    - proposals are either make connection look more or less like TLS 1.2
  - Joe S: when will experiments complete?
  - EKR: depends on what we see. Will have data relatively soon, 4-8 weeks. Takes a while to get into a release... but nightlies and betas give some indication of if it will work.

- draft-ietf-tls-dnssec-chain-exstension (Melinda Shore)
  - completed WGLC on this draft. (https://www.ietf.org/proceedings/99/slides/slides-99-tls-sessb-dnssec-chain-validation-00.pdf)
  - excellent feedback so far.
  - (melinda summarizes changes)
  - record ordering (server canonicalization, yes or no?) No one came to mic.
  - use of _udp label for QUIC
    - Ted Hardie: reading of the draft is that UDP label used for DTLS and QUIC.
    - ...: you might have two different TLSA records, one for DTLS and one for QUIC. Maybe call it _quic?
    - Paul Woters: want to make sure we don't create a new plaintext reference
  - tell client imps how to handle unexpected/irrelevant/extraneous records?
  - Joe: we'll close on these remaining issues before reving the draft.
- DTLS 1.3 (EKR)
  - first mentions something about exported authenticators:
    - certificate type extension goes in server [something] message.
    - odd thing is can have EE X.509 extension [somewhere] which is nuts.
    - Suggest we maintain the property where I change the entry in the table and [something]
    - Trying to keep 1.2 functionality.
  - Reminder: ACKs

        - implicit ACKs historically.
        - interacts badly with some TLS 1.3 features (like NST)
        - Solution: intro an explicit ACK
    - current proposal: SACK
        - kind of ripping off the QUIC structure.
        - MT: other thing with it being a handshake message is that it adds to the transcript record and that gets weird.
    - When should receivers ACK
        - supposed to ACK when you're not moving the state machine forward, when messages might have gotten lost, not for non-handshake messages
        - **If anyone thinks this is a bad strategy, please speak up.**
        - Joe S: how many people have had a chance to look at this? Not too many.
        - Janardhan Igengar: would be nice if this is not too complicated.
    - Reduced Header Format
        - MT: we currently have range between 20-64 reserved for us in this demux thing. We only use the lower half of the 20s... this uses the upper half of that range from 32 onwards. Can use that entire space and allows good distinguishing. Don't see us using too much more content types.
            - EKR: essentially the point of doing something different would be to have much longer sequence bits.
            - MT: not sure I've convinced Ian Swette [sp?]
        - SPT: thing about this is that the IoT will think we ned to make this smaller... this seems about as small as you can get to.
            - MT: one optimization we could make in addition, would be to remove the length.
            - EKR: but that would make the ACK'ing problematic (?)
            - MT: other way to do that is to do some internal framing... "I've got an ACK and some other stuff in here"
            - ... real challenge here are the cases when you're changing keys. would need internal lengths for those content types.
    - Connection IDs
        - have spent an enormous amount of time on this.
        - things behind NATs have problems rebinding.
        - also a serious privacy problem, none of the proposals I've seen are adequate let alone completely baked.
        - huge problem in the browser context, not so much in the mobile context.
        - proposal for DTLS is to kind of punt: have an optional but fixed Connection ID. Doesn't change across the connection.
        - We can add a new extension to negotiate functionality. Best scaling involves passing a token for each connection, yucky.
        - (EKR shows a proposal for a Connection ID extension)
        - IDs are used if a client offers and a server answers
        - Each side *sends* with the other's ID.
        - Happy to hear objections to this strategy.
        - Tobias G.: Connection ID is currently a very big problem in IoT space.
            - lake of entropy space in connection ID... 100k entropy doesn't work for IoT. need 10^6.
            - Need this ASAP.
            - Privacy concern is absolutely correct. Need to be able to reneg the client ID.
            - EKR: I've proposed reciever sets.
            - TG: want to avoid collision in the space... if the server controls that ID, you avoid collisions.
            - EKR: this design avoids that.
            - TG: can we do this in 1.3 please?
        - Hannes: data flows in two directions, similar to IPSEC.
        - Ted H: how does this impact RTCWeb?
            - EKR: wouldn't anticipate being able to use this for RTP.

        - ... unaware NAT rebinding typically assumes that one side has an open connection.
    - EKR: clarify, server picks ID for packets that come to him, client picks the IDs that come to them.
    - DKG: two questions: 1) IoT devices are unlikely to be mobile, do we have evidence for that? Seems like it's also active in things like vehicles; and 2) looks to me like a field for arbitrary metadata insertion in each packet and with long lengths, that looks like SPUD but we're not going to call it SPUD.
        - EKR: let me make my threat of violence clear, you need to solve it.
    - Yoav channeling Victor: why is this an assymetric connection ID scheme?
        - EKR: any symmetric scheme people want to pack the target into the connection ID. might not want to have a random ID.
        - MT: whole point of this is to mark packets so the reciever can get them to the right place.
    - MT: I think this is enough of an attractive target that I don't want to see this in DTLS 1.3, want to see that in a separate document... will address 1.2 as we can hit them at the same time.
        - EKR: structure of this is that we can easily add as an extension.
        - MT: if we just do this we don't get the ability to change over time, important for mobility. makes the arb metadata insertion better to deal with.
    - Christian Huitema: we need to do this right and not fast. Quick and dirty stuff is not going to cut it.
        - ... want to be able to reneg. probably want to make it optional so it's not in every packet. Want to have a constraint so that we don't get huge privacy holes.
        - EKR: what about this? Feel like QUIC got bogged down... proposal that got the most attention was an unencrypted connection ID. Should we build a fixed connection ID or something that constantly changes?
        - EKR can prepare a separate draft for this... may not change 1.2 as that's hard.
    - Jan I: [could not understand]
    - Ben Schwarts: in addition to privacy within a connection, if youre a client trying to keep track of a number of servers, it can be essentially a counter. May encourage clients to create a cross-connection... IDs are in a sequential range, these connections are all the same client. Would be nice to not do that.
        - EKR: this might require it being longer!
    - Hannes T: we wouldn't have a problem in IoT case if the NATs wouldn't exist or do rebinding or if the devices would more frequently send traffic.
        - ... vehicles will probably use cellular keeping the connex open. Always the chance to restart from scratch... connection ID wouldn't make a difference, need to start DTLS connection over again.
        - ... so sending a big ID is not going to help at all.
        - DKG (off mic): why have it then?
        - Hannes: we don't have it!
    - Tobias G: for these connections, mobile devices are in the use cases that I see. When you consider reneg, consider that main use case is devices with power constraints. So every RT etc. is very costly. Not reneg every time would be good. Would like things that we can tailor, customize.
        - Would rather have this really soon... problem is out there.
        - Would strongly urge the chairs to consider for 1.2
    - Jan I: could we constrain this to a smaller thing...
        - EKR: any number large enough to be useful will make DKG sad.
        - MT: any size that is useful, is useful (making the point that it's useful for good and bad)
- SPT: will send an email for when DTLS will drop compared to TLS. Now is your chance to get to the mic.

- Chair interrupt:
  - First thing: presenters are keeping it short and to the point. Hold points until after presentation.
  - Plenty of time for discussion.
  - Want to address both political and technical topics
  - **The main question: Is this subject something that the WG should consider? This = "passive decryption of traffic"**
  - **What technical solutions are available, becasue the WG gets change control if adopted.**
- Impacts of TLS 1.3 on Eneterprise network operation (Steve Fenter, Matt Green, Russ Housely)
  - use cases:
    - wireshark PCAP decrypt
    - Fraud Monitoring
    - IDS/IPS
    - Malware Detection
    - Security incident response
    - Regulatory requirements
    - Layer 7 DDoS Protection
    - NPM/APM
  - When problem hits, no one knows where in this universe of 400 boxes the problem is.
    - I need packet level visibility in everywhere across these 400 boxes.
    - a month ago, got a problem in login failures and slowdowns.
    - sniffer guys called in to swoop in and save the day.
    - Many other guys getting called with severity 1 problems and need to decrypt
    - No way to identify the user bc CDN, decrypt the packets to find user_id or other elements.
    - one particular URL was giving 10s response time
    - Tier 2 load balancer, found the same symptom.... etc.
    - would need 5 proxies here and that doesn't scale... millions of dollars.
    - endpoint monitoring doesn't work, as you can't do full-scale pcap.. because of NAT etc.
      - often need decrypted PCAP where there is no endpoint (e.g., firewalls don't often allow you to terminate)
    - tl;dr: a particular db call to a small single-threaded access table was slowing everything down.
  - If TLS 1.3 rolls out without static DH, severity 1 problems will drag out for weeks. Severity 2 will take months.
    - level of pain that enterprises aren't willing or able to handle.
    - if I have a problem that TLS causes, it's basically a DoS attack. **TLS 1.3 is a DoS attack for us.**
- Matt Green
  - This is not technically a problem: if you control the endpoints, you control the secrets.
  - How do you do this that doesn't harm the protocol?
  - possible solns:
    - endpoints being the servers, deliver session keys or MSs through an OOB channel. But # of keys can be very large.
      - have to deliver keys in a very timely manner, can't cache over much time
    - encode keys in band in the TLS protocol.
      - one option is to use a full extension and then include an in-bad inclusion.
      - unfortuantely, legacy systems don't include this functionality.
      - lots of different variants, some of very hard to detec.
      - Hovav: use DUAL-EC
    - Endpoints use (semi)-static keys
      - don't change the protocol, let's do something others can recognize.
      - No changes to 1.3

        - Easy to detect
        - Reduces FS, mitigated by key rotation.
    - Static DH draft
        - (Matt describes draft-green)
    - Security of Static DH
        - leave aside FS, it is cryptographically secure
        - FIPS 800-56A talks about using static DH. TLS 1.2 has DHS
    - concerns
        - easy to have imp flaws.
        - but easy to not affect most users.
    - Harm reduction
        - enterprises don't adopt 1.3, today they're using 1.2 with static RSA.
        - make some dramatic changes to endpoints to deliver session keys.
        - some really really bad ideas (won't go into)
        - Extensions (open to this)
        - pros: no significant protocol changes, well-understaood crypto, detectability.
- Natl Cybersecurity Center of Excellence (Tim Polk)
    - Sponsor of a related draft.
    - NCCOE is all about implementation and adoption.
    - Have been hearing issues with meeting operation reqs for 1.3
    - Objective is to collab with industry, solve problems and get better security than we have today.
    - NCCoE will produce a proof of concept imp and a number of documents... want to prove that we can tighten up the life span of those keys that we are sharing in the enterprise.
    - Would produce an 1800 series practice guide that would say, "if you want to do what we did in the imp, do this."
    - would submit an IETF draft that showed what we did, what worked/did not, here are the key liftetimes that we think we can manage.
- Proposal DOES NOT violate the IETF policy on wiretapping (Russ Housely)
    - RFC 2804 defines wiretapping and this is not that.
    - Want to get as much TLS nowledge from this WG as possible to produce as secure a thing as possible.
- TINFOIL (Stephen Farrell)
    - This whole thing is a terrible idea, we shouldn't do it.
    - Stephen goes through the various arguments listed here: https://github.com/sftcd/tinfoil
        - not in scope for charter
        - could put TLS/DTLS 1.3 at risk
        - TLS is hard
        - 1.3 has undergone significant analysis so far, this has not
        - Static DH is not implementation robust
        - we have a case where law enforcement has tried to coerce a server operator to tap at TLS level
        - should in no way be a standards track document.
        - breaking TLS is not part of the WG charter.
        - ...
    - Question to group: should we document these arguments about breaking TLS?
- Q&A:
    - PHB: agree with both sides. Problem here is coming from the PKI world, saw what happen to bluecoat and co. We didn't make WebPKI holes for them, so they blasted holes into it.
        - on the techincal side, don't like how you're doing it. I'd use a different DH share every time.
        - would like to have this such that if this makes it into the wild, it is not compat with legacy stuff
    - Paul Woters: want to quote RFC-someting-bis

    - talking about DH groups 22, 23, 24. 22 is must not. 23 and 24 will be must not soon, should not now.
- Dan Harkins: there is IPR here out here from a past employer.
- DKG: want to express disappointment with this draft.
    - export cipher suites was brought up. As recently as last year we've had problems with the fact that export cipher suites were standardized 20 years ago.
    - The first time we see a problem with this might be soon... the last time we see a problem will be way way in the future.
- Rich Salz: (applause)
    - I am torn between: prof. and personally I think this is not a good thing. Remember Dave Clarke's talk that we need to tilt the playing field for things that people use.
    - would like to see us wait two years for deployment exp. It's pre-mature. Let's revisit.
- Darin Pettis:
    - We've been here before, part of a large financial organization. We've ditched RSA, we understand that.
    - We've explored technical options, and have not find a better way.
- Roman from CMU:
    - didn't see discussion for security uses, esp. DFIR (incident response).
    - very key to do ad hoc instrumentation and this would help.
- [Cisco business security group]:
    - Don't think security is served by seeing this as a black or white approach.
    - Reminds me of the old discussion of NATs.
    - Community is better served by ack'ing this problem and find a way to solve it.
- Nalini Elkins:
    - There are very real problems from real people doing real things.
    - If you are hearing from real people that there are problems, behooves us to listen.
- mnot:
    - this is not the first time we've seen this thing.
    - 2 years ago, proposals strongly made in HTTP.
    - We chose not to accept that work; reason is that HTTPS is explicitly a two-party protocol. Did not have a way to get the informed consent necessary to change the protocol.
    - You are changing the nature of the protocol pretty fundamentally.
- Ted Hardie:
    - two points: FS is a feature of this protocol. This turns it off in certain contexts, not obvious to the end users that FS has been removed. Can't tell in the first connection if a key will be re-used. Need to have a way with features like this about 1) signal that it's being used; and 2) get agreements to use it from those communicating.
    - if it comes back with those changes, what's the domain analysis? Russ read us section 3 of RFC 2048, but not section 4, that says, "we're not taking a moral stand, but a technical stand". You MUST expect a technology to be used in places you might not expect. Analysis must take into account all of the domains of us.
- Max Pala (cable labs):
    - agree with Stephen.
    - Most of the time this is a problem if your arch is outdated. No one will force you to do this... if you do deploy 1.3, should be conformant.
- Ralph Droms:
    - Living the dream (laughter)
    - Want to emphasize keeping separate the fundamental abstract questions that are being discussed and the particular proposal that is on the table in this document.
- Roland Dobbins:
    - Want to emphasize being able to troubleshoot and need visibility.
    - Often this needs to be on the wire.

        - They may face potential fines and liability if they don't take care of our information.
        - We don't want crypto that is proprietary or that is developed in smoke-filled rooms.
    - Jeff Hodges:
        - want to echo ralph and roland and a few other people.
        - There are enterprise needs here to do this thing.
        - We want to get to FS in the data center and we need a migration path that has been scrutinized by experts.
    - Christian Huitema:
        - want to support Stephen and Ted: take the Lavabit scenario.
        - A provider of a service is being compelled to turn on a feature so that someone can get their traffic.
        - This approach is very dangerous... you assume that you are using the same software in both DCs and on a server.
        - Don't like the feature that this is keeping the wire format unchanged.
        - We need a "big red flag" requirement so that this is only used in the DC.
    - Daniel Franke (Akamai):
        - like draft-00, not draft-01
        - draft-01 is standards track, not ok
    - Kenny Patterson:
        - There is nothing in the current draft that would force the rotation of keys.
        - suggestion: adopt the draft and force key rotation on each connection.
    - Sharon Goldberg (BU)
        - Want to support Stephen.
        - Not confied to DCs, do not support at all
    - Deb Cooley (NSA)
        - I believe if you take the draft here, you control the draft.
        - if you let this go underground, it will happen silently like what happened in the past with Static RSA.
    - Tara Tarikyee (OTF)
        - add voice to those disappointed in this draft.
        - there are a lot of people that depend on TLS for the practice of those rights.
- **Questions the charis want answered (policy)**
    - **The main question: Is this subject something that the WG should consider?**
        - **this = "passive decryption of data center traffic"**
        - **subquestion: Is this wiretapping?**
    - Clarifying questions:
        - Stephen Farrell: don't believe that passive is correct here, draft-green allows active attacks. Allows the attacker to be active.
            - Ralph Droms: separate the solution in the draft from the mechanism.
            - Stephen: A, your saying your draft is broke.
            - ... not clear to me that there is any solution that allows passive that doesn't allow active.
        - DKG: we have considered this question for many weeks.
        - Lucy Lynch: take a hum first on whether or not the group should accept the draft and then take a more general hum.
        - Joe S: "decryption of data in the data center" ommiting the word "passive"
        - Hums: No clarity whatsoever. Seemed pretty even.
    - Stephen: want to take it to the list as to if the WG is interested in documenting reasons to not break TLS.