AL: Adam Langley
AP: Andrei Popov
BM: Bodo Moeller
CJ: Cullen Jennings
DB: Daniel Bernstein
DKG: Daniel Kahn Gillmor
DMG: Dave McGrew
EKR: Eric Rescorla
EN: Erik Nygren
HT: Hannes Tschofenig
JS: Joe Salowey
KI: Kevin Igoe
KP: Kenny Patterson
MSJ: Mike St. Johns
MT: Martin Thomson
PA: Paul August
PH: Paul Hoffman
RH: Russ Housley
RS: Rich Salz
SF: Stephen Farrell
ST: Sean Turner
TA: Tolga Acar
YN: Yoav Nir
YS: Yaron Sheffer

TLS Interim Meeting
Sunday 20 July 2014, 10 am - 4 pm
366 Adelaide St W, Suite 500, Toronto, ON
EKR: logistics for the location
ST: kicking meeting off, note, agenda

EKR: presentation on changes since -01
https://www.ietf.org/proceedings/interim/2014/07/20/tls/slides/slides-interim-2014-tls-2-2.pdf
Removed support for compression.
Removed support for static RSA and DH key exchange.
Removed support for non-AEAD ciphers.

KP: Will there be a requirement on the servers being able to reuse values.
AL: No reason to do that…

HT: Going to add AES-ccm cipher suite.
EKR: Going to pull cipher suites in that are on standards track and conform to the requirements.

Remove custom DHE groups.*
KP: Support for slightly larger groups might extend the life of DH in TLS should quantum computers technologies appear that would be the case otherwise. This might be a reason to allow for custom diffie-

hellman groups as a new named group would not need to be standardized via the IETF.
DB: Current largest DH group is 8192.
PH: We discussed this in IPsec (named groups vs custom groups), and determined it would be faster to push an RFC through the IETF.
AL: Biggest group.
EKR: Is this consensus
ST: Consensus call, remove custom DHE group and point to Dan's draft, no objections… This will be confirmed on the list.
EKR: Porting ECC into TLS as standards track, would import groups.

Reworked handshake to provide 1-RTT mode.*
1RTT A

???: When the server key is exchanged is not signed.
EKR: You are still going to sign it, but you are signing the entire set of messages.

If the server doesn't ask for the certificate from the client, it can start sending application data sooner.

The exact rules of when you allow this in have not been written down, they are a bit complicated, need a smaller group to think about this problem.
PH: If this is going to be allowed, (sending application data with the server finish message) then it needs to be explicitly documented.

New ClientKeyExchange
AL: Can you specify that they must be in strict order numbering? (check)
EKR: File that as an issue.
Issue being added to github issue tracker.

Slide New ClientKeyExchange Syntax

Should we be renaming this message (WTC)

The new message needs a new code point in any event - someone editing issues #58.

Slide #14 – early data syntax

Overal 1RTT Flow
AL: comment I missed,
Don't allow the flexibility of all those lengths and types.
??: like that idea , handshakes and then data.
works if you get rid of CCS,
hold off to post NT

If you are going to do a design team on that, need someone from F5 and

large CDNs.
EKR: going to work with chairs to define a small group to study this.

Extension handling slide –
The keys should be dependent on the transcript of the exchange so far
EKR: Have a list is better than having logic record
Define the same semantics
Enumerate where the extensions go…

Slide – new ServerKeyExchange syntax
Subset of previous syntax
Rename this one

What about the server's signature
Why isn't certificate verify the first thing?
That cuts against the argument for finished.

Slide
Overall 1RTT Flow (slide 5)

Slide What if the client guesses wrong? (slide 15)
Don't make this look like a
For resetting – makes the operation stateless, allows you to use the same for the DTLS stateless challenge
Against resetting – if you has the
Strong, strong weak,
If the client preference order, server preference order, list of what the
Require in the
Violates the monotonic ordering
… long discussion …
why are we not keeping the running
don't know
we should do supersetting in any case,
issue #73

How does client distinguish these two handshakes? (slide 16)

Interaction with Triple Handshake Fix (slide 17)
Need to think about this further to ensure it is ok.

Renegotiation Presentation
Martin Thomson
https://github.com/tlswg/wg-materials/blob/master/20140720_interim/tls13-renego-mt.pdf
Slide b) hiccup
EKR: Not clear where the breaks in the session are identified from the server application.
slide c) rekey
AL:

has some of the same problems (not as bad) as
some services will never do renegotiation
slide d) start over
AL: Suggestion to just rekey after a fixed number of records by hashing to chain the key forward
MT: Problem with having a fixed number if things need to change in the future
EKR: Two reasons - 1) exhaust key usage space 2) protect backwards usage of key
AL: May want to have two different key streams derived from the MSK so that the server and client sides have different key streams
MT: Raise issue to track this?
Issue # (to be filled in when done)
AL: re-key solves the attack backwards in time, but not the forward in time attack. Re-handshake solves both
EKR: Ok with tear-down to solve the going forward attack
MT: In DTLS can use the epoch to replace the CCS
HT: Does not address any issue where the DH key exchange is compromised
MT: Would rather force a new connection to address this issue
YN: Should there be new randoms to the roll over?
EKR: How does this help?
-- no real response
DB: What about do a rekey after every packet - what is the performance hit?
AL: Would be good for some ciphers - kills AES (key schedule setup) okay for ChaCha
EKR: Does not work for DTLS
EKR: Need to resolve in order to move forward on some other issues.
ST: Are we happy to remove renegotiation on the assumption we will provide a rekeying facility of some form and initial client initiated client authentication.
EKR: would have no problem with fatal alert that says - please reconnect with a certificate
AP: problem with alert is don't have a way to give the data back on what cert is needed
MT: issue raised in Denver is that until know what is asked for, there is not enough context to ask for a correct certificate
HUM: Yes - loud - no - a couple - don't know - a few (more than no)
EKR: Order of operations - pull request on renegoation then pull requests on other items that follow
Encrypted Content Type (DKG)
https://www.ietf.org/proceedings/interim/2014/07/20/tls/slides/slides-interim-2014-tls-2-5.pdf
CJ: (Addressing middle boxes) some firewalls will have rules on keep-alives based on seeing something that is TLS data -
AL: will already be freaking out middle ware boxes with TLS 1.3 header anyway
EKR: Also some issues with multiplexing DTLS with other protocols - uses the first byte for doing dispatching correctly

DKG: looking at trying to drop the version as well - so two octets of length will start packet
MT: short fragments will be ok
CJ: add in one octet byte with identifies the packet as data.
EKR: Decide to adopt encrypted content types independent of adding first octet.
CJ: DTLS will needed it - so should have it in TLS also
AL: Why do this given that handshake will be gone
DKG: separate alerts and data - also may have padded vs non-padded application data - add new record in the future
ST: one of the premises of 1.3 - encrypt as much as possible.
DKG: Use a consistent (and current) record format for non-encrypted records - just change the encrypted records going forward
CJ: Should look at the DTLS collision problems first
MT: structure of DTLS - the epoch and sequence number come first - have epoch first
MT: STUN can use a larger range (RFC 5764) than setup
EKR: seems to remember that there were other places to look at for doing the demux
ST: Any objects to the way forward doing encrypted content type?
room - none
Encrypted SNI (DKG)
https://www.ietf.org/proceedings/interim/2014/07/20/tls/slides/slides-interim-2014-tls-2-6.pdf
CJ: When looking at the blocking list, people are who don't want to be blocked get off the ip address of somebody who is going to be blocked. Encrypted SNI relies on the fact that there are enough different people on the same IP address that IP address and DNS name are not the same.
PH: Example of tumbler that is doing this today for third level domains
AL: Issue of timing if you are doing additional network lookups to get pre-handshake key to do this. You are either going to slow down or sometimes leak this information
EKR: There exist history sniffing attacks on cached data
PH: Does not except that there is going to be any DNS privacy in the near future - Needs to make clear that this is going to be a strong dependency on this issue.
PH: DNS privacy is dependent on the DNS community and not specifications - this makes the problem worse.
DB: Can use OpenDNS today to encrypt requests to their servers and get a degree of privacy today.
EKR: pre-key would be known at the CDN but the handshake key is known at the correct server. Allows for distribution of session
AL: This does require the padding extension as well.
AL: This does imply issues on the question of when the hash restarts - if transcript includes this then an active attacker would be detectable even though the SNI has been leaked
CJ: Want to go back to first principles on what happens
EKR: 1.3 servers that do this, 1.3 servers that don't do this and

pre-1.3 servers
DKG: 1.2 servers will return with a certificate that does not match (on request w/o SNI)
AL: Can't switch this on until all front ends support this 1.3 extension
DKG: Bound to a name that identifies a cluster via a service record
MT: Look at IP load balance case - needs to do the upgrade on all of the servers at the same time to prevent issues with different capabilities
CJ: Naive clients always send SNI.
DKG: This is a phase in thing, won't be able to turn on over night
MT: If there was a glue record in the DNS, then even the naive client could do this. (not necessarily - according to others).
Encrypted SNI: Threat Model Analysis (Rich Salz)
http://www.ietf.org/proceedings/interim/2014/07/20/tls/slides/slides-interim-2014-tls-2-9.pptx (fix when pdf is available)

RS: Cost is not going down because as crypto gets faster - key sizes need to get bigger
AL: Looking at cost vs benefit analysis
RS: No the RSA cost -
AL: certificate verify message is the dominate cost
EN: bullet on making passive attackers become active attackers
EKR: Cost is much more than is helpful. Cost on the server side of remember key for length of DNS advertisement is high.

Assumes hosting provider of lots of innocuous noise to cover the small number of people that need to be hidden.

EKR: if google had to go to a bunch of effort to encrypt SNI to allow people to hide things would they do it.
AL: There are instances were we do things based on hard coded things - could hard code in a constant SNI for a small number of sites
AL: Really wants to avoid doing twice as many DNS look ups (for race return)
CJ: Schools block sites that contain malware - currently block on domain name. If it is hidden then need to block on IP address
PH: So expensive to block based on TLS - the block is done either on DNS or the IP address itself for all traffic. Don't do it based on the SNI inside of TLS
YN: do it based on passive rather than active for legal and speed requirements.
Is it important to encrypt the SNI?
Hum - most people hummed SNI-encryption was not important
Extended Master Secret (K. Bhargavan)
EKR: Disscussion in Denver - Agree this should be done - open question - currently TLS PRF only involves randoms. This adds lot of data digested. Is there a weakness in the underlying hash function, can an attacker mange to manipulate the operation to re-allow for this attack.

AL: If a hash function is broken then much larger problems than this.
CJ: Just adopt it
ST: Any objections to adopting this issue?
PH: Can we prioritize this over doing 1.3? so all of this can be pushed forward into 1.3
EKR: Can we assign a code point now?
EKR: Remove fallback from current draft. based on discussion
ST: If stable and reasonbly then can do.
JS: Will come up with something based on the HUM for encrypted SNI.
Monday Session 15:20-17:20
Administrative - Agenda Items
ECC to Standards Track/MTI (Sean Turner)
EKR: Pull into 1.3 - triage the list of curves to be included.

Have a revised draft if there is excitement - otherwise just do the normative DOWNREF process

ST: Any volunteers - Yoav does
EKR: Should the 4492bis list be trimmed down? Suggest not doing it.
SF: Will the bis be for 1.3 only - or also for 1.2

If CFRG comes back with lots of curves - what are you planning to do

EKR: Need one for each purpose @ each security level

the new algorithms should go into the bis document

ST: Should we have one curve for each purpose and level?
HUM: - large hum yes - small hum no - small (but larger) don't know
EKR: Assuming we are not punting the NIST curves from the document.
AL: Not aware of any advantages for the EDH groups - just larger and slower. But not rabid.
ChaCha20-Poly1305 (AL)
AL: Can improve the AES-GCM performance if you increase the side-channel attack range
MSJ: Question on setup times
AL: ChaCha20 has no key schedule setup time
EKR: Kent raised the issue last time on the nonce being explicit rather than a counter
AL: AES uses counter for it's nonce. Used because it is reasonable. Allows for prevention of duplicates

Thinks that this is still OK

DMG: Applauds the use of AEAD algorithms

Thinks nonce is worth getting input on.

FIPS-140 requires that input of nonce allows for arbitrary value

AL: Make the crypto-module not be an input but make it a counter internally
DMG: Possible reason to have nonce is IPSEC use of multiple enrypters
AL: This is TLS not IPsec - will not speak to it here.
DMG: Use a consistent way of using nonces in TLS for all AEAD algorithms
HT: Should be using hardware for doing comparisons rather than software. This is going into all IoT type devices today.
JS: Question to CFRG chairs for comments and concerns on this algorithm
KP: On the agenda for Wed.

General issue is lesser amount of review for ChaCha.

Hesitant to give strong recommendation.

EKR: See if we can get partial review from CRFG and continue our processing and if no problems come up by ready to progress then do so
DB: No vast difference in having hardware implementations of the two algorithms
Downgrade SCSV (Bodo Moller)
PA: Why is this not symmetric
BM: Client chooses and the downgrade dance is always the on the client side

Keep the extension as simple as possible. Could public all versions and play games but much harder to get right.

YN: Problem skipping intermediate versions during the downgrade as server may have a floor on implementation levels (i.e. 1.2 -> 3.0)
BM: Don't do that or don't use this extension
EKR: Please tell us about the type of data you are seeing
AL: Had two changes at the same time, so difficult to tease apart the results. Number of errors has gone too zero. Venders have fixed the bugs since all of the users got broken
JS: Look for a WG last call next week to have time to comment.
Negotiated Discrete Log DHE Groups (DKG)
YS: this is a fine doc, but the reasoning while we're not sharing groups with IPsec is unconvincing. Also (a nit), the reference to IKE groups should point to IANA
AL: If you are getting updated - why not just switch to EC?
DKG: Most will do so, but many will also allow for EDH as well. Named groups improve selection criteria
EKR: Would separate mechanism from the selection of algorithms

Issue with possible confusion during negotiation if the new group types hit a system that does not recognize them

What do we do about strange crypto sizes

AL: Don't standardize bad crypto - these lengths are good.
PH: the numbers used in IKE have been studied much longer
DKG: Used e rather than pi so that the numbers are different and breaking one would allow for a breakage of the other in the event that a break was found
PH: The numbers from IKE have been widely researched

Would not recommend if going to use EC is going to be the MTI as no good traction

EKR: No intent to make a change to MTIs for 1.2
Session #2 - Thursday 24 July
Report from CRFG (KP = Kenny Patterson)
http://www.ietf.org/proceedings/90/slides/slides-90-tls-6.pdf
EKR: discuss randomized curves vs rigid curves please
KP: Hard to produce randomized curve generation process where there are not necessarily hidden properties
Rene: Don't agree with the current consensus, e.g., while "twist security" is a "nice to have", this is certainly not "required".
Ben: useful to indicate length of use of keys
EKR: Sense of universe of things that might be done?
KP: Personal sense is between 25519 family and the NUMS family

Weierstrass off the table
brainpool never really in the running

TLS Crypto Constructs (MSJ)
http://www.ietf.org/proceedings/90/slides/slides-90-tls-5.pdf
EKR: Why do I care about parallellization
MSJ: May have multiple channels - small difference but difference
EKR: Why does it matter if compute over the handshake or the finish
MSJ: If use CMAC than half of AES and no hash function
EKR: Common complaint for 1.2 is need multiple hashes over the handshake.

Not a trivial change to fix this

MSJ: if hold onto the data until you ready to hash - then don't need multiple hashes
TA: Triple handshake changes this
MSJ: Does not change the output lengths of the keys. Just the suite number is bound in.

HSM does not know anything about the TLS suites

PH: What is the motivation for this document
MSJ: How do I write policy language for the machine so that the HSM can prevent some problems
EKR and MSJ argue if the hash is applied to the master secret before

computing the HMAC
DKG: Feedback from the Crypto community is that keys encrypting handshake are also used for data.

This would make analysis simpler as it is better decomposed

RH: Should decompose so HSM either invokes material for the inside or for the outside. Not use the same data stream for the two different operations
EKR: Sympathetic to the hygiene issue, but not to the hash questions

Making CMAC work is not a good think if I have to buffer the entire handshake upfront

New piece of the puzzle, evolving the master secret by modifying the master secret as part of the rollover issues on keys

MSJ: Reason to push on CMAC is the IoT world

On rolling - create a new Next Master secret as part of the derivation process to do the roll overs

KP: +1 on the hygiene for doing analysis work
MT: +1 on hygiene - View of the world is a generation of multiple keys or data streams to produce read key, write key, unique key, iv streams, ...
KI: - CMAC = replacement for CBC-MAC - SP 800-36 - block cipher mode

May have a problem with no block cipher if we are doing ChaCha20

Multiplexing Scheme Updates (MPH = Marc Petit-Hugeunin)
http://www.ietf.org/proceedings/90/slides/slides-90-tls-7.pdf
EKR: Not sure how sympathetic to the solution of this problem

three categories of data

MPH: Removed content type - tell IANA not to allocate in these branches
EKR: If doing TURN then only small confusion point of time. Hard to believe what to resevere so much space to make problem harder
MPH: Upgrade the algorithm rather than create a registry
JS: Not good design for protocol - (EKR claims fault for the original)
EKR: suggest sit down to fix the algorithm
MT: don't need to multiple in one space - because each would toss the packet if it does not process well (DTLS, STUN,...)
ST: If you don't like this - speak up now or forever hold your peace
ST: Discussing having an interim meeting - not in North America - looking at Europe