Complaint regarding a declaration of consensus to adopt a non-hybrid draft

Daniel J. Bernstein, 2025-06-05

This is a complaint to the "Area Directors" (ADs) for the "Security Area" of the "Internet Engineering Task Force" (IETF). The complaint is that Joseph Salowey and Sean Turner, in their roles as chairs of an IETF "Working Group" (WG) named "Transport Layer Security" (TLS), erred—both procedurally and in their conclusion—when they declared that the WG had consensus to adopt a draft named "draft-connolly-tls-mlkem-key-agreement". (The WG has a third chair, Deirdre Connolly, but the other chairs later said she was "recusing as she is an author".)

BCP 9 (RFC 2026), Section 6.5.1, includes the following two paragraphs:

> A person who disagrees with a Working Group recommendation shall always first discuss the matter with the Working Group's chair(s), who may involve other members of the Working Group (or the Working Group as a whole) in the discussion.
>
> If the disagreement cannot be resolved in this way, any of the parties involved may bring it to the attention of the Area Director(s) for the area in which the Working Group is chartered. The Area Director(s) shall attempt to resolve the dispute.

This complaint is hereby invoking the second of these BCP 9 paragraphs.

To review preconditions: I explicitly invoked the first of these BCP 9 paragraphs by email to the TLS mailing list dated 18 Apr 2025 14:02:55 -0000, asking for the "discussion to be on-list for transparency". The chairs sent very few public email messages about this, concluding with email dated 25 Apr 2025 15:04:39 -0400 saying that they stood by their consensus call and that I "can appeal". (The full procedural history is reviewed below.)

The second BCP 9 paragraph has a prerequisite that "the disagreement cannot be resolved in this way". It's not clear that this is triggered merely by the chairs saying that I "can appeal". Can one conclude that the dispute *cannot* be resolved by discussion with the chairs, merely from seeing that the chairs aren't trying to resolve the dispute? The answer has to be yes: using an impossible-to-prove condition to reject appeals would violate the requirement of due process and would violate the requirement to provide an appeals process.

See Section 13 for further discussion of the procedures that should now be followed.

1 Context, part 1: NSA

BCP 188 (RFC 7258), "Pervasive Monitoring Is an Attack", says (among other things) "The IETF Will Work to Mitigate Pervasive Monitoring". This RFC was triggered by news articles in 2013 regarding mass surveillance by NSA and GCHQ.

For example, https://www.theguardian.com/world/2013/sep/05/nsa-gchq-encryption-codes-security reported that NSA was budgeting a quarter billion dollars a year for a project that "actively engages US and foreign IT industries to covertly influence and/or overtly leverage their commercial products' designs" to make the designs "exploitable ... To the consumer and other adversaries, however, the systems' security remains intact." NSA's budget document (https://embed.documentcloud.org/documents/784285-signit-enabling-project/) includes the following specific goal: "Influence policies, standards and specification for commercial public key technologies". NSA was not just passively recording Internet traffic; it already had a large budget to influence standardization processes so that the resulting standards would be exploitable.

Several years later, the European Court of Human Rights held that GCHQ's activites were illegal. See https://www.theguardian.com/uk-news/2021/may/25/gchqs-mass-data-sharing-violated-right-to-privacy-court-rules. It is unclear what actual effect this court decision had on GCHQ. Certainly the court decision has no power over NSA.

My blog post https://blog.cr.yp.to/20220805-nsa.html covers much more of what is known about NSA's cryptographic sabotage. In particular, when cryptographic standardization began, NSA adopted a secret policy of trying to reduce competition in this space so as to reduce security:

> Narrowing the encryption problem to a single, influential algorithm might drive out competitors, and that would reduce the field that NSA had to be concerned about. Could a public encryption standard be made secure enough to protect against everything but a massive brute force attack, but weak enough to still permit an attack of some nature using very sophisticated (and expensive) techniques?

This is a quote from pages 232–233 of https://archive.org/details/cold_war_iii-nsa, an internal NSA book that was partially declassified in 2013 as a result of journalists forcing declassification-review procedures. There has never been a public statement from NSA revoking the above policy, nor would such a statement be credible given NSA's long history of sabotage.

One cryptographic mechanism that NSA manipulated NIST, ISO, and ANSI into standardizing was Dual EC, a backdoored standard for generating random numbers using elliptic curves. Various other NSA-proposed standards for elliptic-curve cryptography (ECC) turned out to be filled with traps for implementors—traps that continue to cause exploitable problems, as illustrated by CVE-2023-6135 in Firefox. For more about Dual EC, see https://cr.yp.to/papers.html#dual-ec; for many further ECC failures, see https://cr.yp.to/papers.html#safecurves.

In short, despite what one might think from the "National Security Agency" name, NSA has again and again shown its willingness to damage American security for the sake of mass surveillance. This is what NSA did with DES, with export controls, with DSA, with Dual EC, and with any number of unknown targets of NSA's quarter billion dollars a year to make commercial products "exploitable".

NSA's influence on cryptography goes beyond this quarter billion dollars a year. For example, the United States military budget is approaching a trillion dollars per year; NSA sets rules for the cryptographic part of this purchasing (see, e.g., https://web.archive.org/web/20221022163808/https://www.jcs.mil/Portals/36/Documents/Library/Instructions/CJCSI%206510.02F.pdf?ver=qUEnOsWpGPcGGMFTb4yYVA%3D%3D). Some people think that NSA's willingness to damage American security doesn't extend to the American military, but in the end NSA's mission (see https://web.archive.org/web/20250418203700/https://www.archives.gov/federal-register/codification/executive-order/12333.html) is primarily surveillance, not security. NSA has different rules for the data it really cares about: as a public example, https://web.archive.org/web/20220524232250/https://www.nsa.gov/Portals/75/documents/resources/everyone/csfc/threat-prevention.pdf is an NSA program that mandates using two independent encryption layers "to mitigate the ability of an adversary to exploit a single cryptographic implementation".

2 Context, part 2: hybrid ECC+PQ

"Post-quantum cryptography" (I coined the term in 2003) tries to protect against attackers with quantum computers. This isn't easy to get right. For example, 48% of the 69 round-1 submissions in 2017 to the NIST Post-Quantum Cryptography Standardization Project have been broken by now; 25% of the submissions that survived round 1 have been broken by now; and 36% of the submissions selected by NIST for round 2 have been broken by now. See my paper https://cr.yp.to/papers.html#qrcsp for details and references.

One of the broken systems, SIKE, had been applied on a large scale to real user data. To quantify "large scale": https://blog.cloudflare.com/the-tls-post-quantum-experiment/ said that "approximately

one third" of the participating TLS connections used SIKE, and that sampling 5% of the participating TLS connections produced "millions of data samples". In short, tens of millions of user TLS connections were encrypted with SIKE. See https://eprint.iacr.org/2023/376 for an attack taking 11 seconds to break larger SIKE keys. (The first SIKE breaks were slower.)

Fortunately, SIKE was rolled out only as an *extra* layer of defense on top of elliptic-curve cryptography (ECC), rather than as a *replacement* for ECC. ECC+SIKE was failing to protect against quantum computers, but it least it had the strength of ECC against non-quantum attacks, whereas rolling out SIKE by itself would have been an immediate disaster.

More broadly, it's normal for PQ to be rolled out as ECC+PQ, typically called a "hybrid" between ECC and PQ. ECC generally consumes far less Internet traffic than PQ; ECC's computational costs are negligible; ECC software is practically everywhere anyway. So deploying ECC+PQ rather than just PQ is an easy common-sense win, and would remain dominant in a free market. (The situation is different for a market warped by NSA influence: see Section 3.)

The TLS WG adopted an ECC+PQ draft in March 2025 (https://mailarchive.ietf.org/arch/msg/tls/iGLZeIYzsIZPYO1ezA2Eulv5x98/). There have been no objections to the declaration of consensus on adopting that draft. I had pointed out two BCP 79 issues regarding the draft in its current form (stemming from the draft's requirement to use a patented algorithm, Kyber), but I said that the first issue can be fixed after adoption, and I now think that this is also true for the second issue.

3 Context, part 3: NSA's influence on PQ

There have been various messages to the IETF TLS mailing list arguing for IETF approval of a *non-hybrid* PQ option on the basis of this being a condition for the NSA funding mentioned above. For example, https://mailarchive.ietf.org/arch/msg/tls/qFRxBSnEPJcdlt7MO0cIL2kW5qc/ appears to claim that compliance with NSA's document https://web.archive.org/web/20240927153634/https://media.defense.gov/2022/Sep/07/2003071834/-1/-1/0/CSA_CNSA_2.0_ALGORITHMS_.PDF prohibits hybrids by 2033. This specific claim appears to be contradicted by the text of the document—see below—but the claim nevertheless shows that the pursuit of NSA funding is having an influence on IETF decisions.

Sometimes the same argument is stated in more vague terms. For example:

- https://mailarchive.ietf.org/arch/msg/tls/S9Mwv28VEHrG189ZtoubUani7J8/ says the following: "There are people whose cryptographic expertise I cannot doubt who say that pure ML-KEM is the right trade-off for them, and more importantly for my employer, that's what they're willing to buy. Hence, Cisco will implement it; I am essentially just asking for code points."
- The specific non-hybrid draft at issue in this complaint, https://datatracker.ietf.org/doc/draft-connolly-tls-mlkem-key-agreement/, has a total of two sentences in its motivation section, where the first sentence summarizes what ML-KEM is and the second sentence is as follows: "Having a fully post-quantum (not hybrid) key agreement option for TLS 1.3 is necessary for migrating beyond hybrids and for users that need to be fully post-quantum."

Both of these quotes appear to be alluding to the same NSA funding. Certainly "willing to buy" is a statement about funding, evidently from a source large enough to dictate Cisco actions, evidently from a source asking for non-hybrids, evidently from "people whose cryptographic expertise I cannot doubt"; if that isn't NSA, who is it? Similarly, "users that need to be fully post-quantum" avoids mentioning names but certainly isn't excluding the elephant in the room.

On the other hand, the failure to provide clear statements and evidence for those statements makes verification and discussion unnecessarily difficult.

Is NSA on record refusing to purchase hybrids? Remember that NSA controls the cryptographic part of a nearly-trillion-dollar military budget. That's a lot of money. Even a *hint* that this money won't be used for hybrids would be enough to influence many companies to support non-hybrids.

Some individuals from NSA have made public statements that can easily be understood as indicating that NSA will refuse to purchase hybrids. For example, NSA's William Layton wrote in https://mailarchive.ietf.org/arch/msg/tls/ESCdYNwVeF4VkvoORFJLJk_87VU/ that "we do not anticipate supporting hybrid in NSS", and NSA's Mike Jenkins wrote in https://mailarchive.ietf.org/arch/msg/spasm/xUKIoHQwm1BjNZWS2x3xb-BhsLI/ that NSA is "looking for" non-hybrids.

Furthermore, there are official documents from NSA and GCHQ stating anti-hybrid arguments. See https://blog.cr.yp.to/20240102-hybrid.html for quotes and analysis.

On the other hand, recall that an official NSA document https://web.archive.org/web/20220524232250/https://www.nsa.gov/Portals/75/documents/resources/everyone/csfc/threat-prevention.pdf recognized the security value of hybrids. Specifically, it asked for two independent encryption layers "to mitigate the ability of an adversary to exploit a single cryptographic implementation".

Another official NSA document https://web.archive.org/web/20240927153634/https://media.defense.gov/2022/Sep/07/2003071834/-1/-1/0/CSA_CNSA_2.0_ALGORITHMS_.PDF, the document that supposedly prohibited hybrids by 2033, didn't actually say that. On the contrary, it stated that "hybrid solutions may be allowed or required due to protocol standards". This last quote disappeared from a December 2024 version of the document—but a disappearing quote isn't the same as a contrary quote. As far as I know, there are no official statements saying that NSA won't purchase hybrids.

Readers thinking "Don't be obtuse: *obviously* NSA will insist on purchasing non-hybrids even if IETF requires hybrids, and *obviously* this draft is chasing NSA money" are missing the point. The rationale for the draft has to be clearly *documented* to enable discussion—including appeals. Transparency requires this information to be on the public record.

4 Survey of objections to the draft

Turner, also on behalf of Salowey, sent email dated 1 Apr 2025 08:58:01 -0400 that included the following announcement: "This time we are issuing a WG adoption call for the ML-KEM Post-Quantum Key Agreement for TLS 1.3 I-D [1]. If you support adoption and are willing to review and contribute text, please send a message to the list. If you do not support adoption of this draft, please send a message to the list and indicate why. This call will close at 2359 UTC on 15 April 2025."

The cited draft was https://datatracker.ietf.org/doc/draft-connolly-tls-mlkem-key-agreement/. The following paragraphs give quotes showing that various objections were raised during the specified period for the adoption call.

This is just a high-level survey of the objections. These quotes are not intended to convey the full text of objections on the mailing list, and are also not intended to convey how many people were objecting (see Section 5).

The draft creates security risks. See, e.g., my email dated 1 Apr 2025 21:38:16 -0000 (https://mailarchive.ietf.org/arch/msg/tls/2Dfu4x678DEKCzF-fkdvJHJkS-8/): "SIKE was applied to large volumes of user data as part of the CECPQ2 experiment in 2019. SIKE was publicly broken in 2022. [paragraph break] The *only* reason that this didn't immediately give away the user data to attackers is that CECPQ2 was ECC+SIKE, rather than just SIKE. [paragraph break] Should we keep rolling out post-quantum cryptosystems to *try* to stop future quantum attacks? Yes, of course. But, just in case this goes horribly wrong *again*, let's make sure to keep ECC in place. Any draft violating this should be rejected as a security risk not just by WGs but also by the ISE."

The draft violates BCP 188. See, e.g., my email dated 15 Apr 2025 22:33:23 -0000 (https://mailarchive.ietf.org/arch/msg/tls/xJqwB3Ob5wf3GVlAiIP3O4tuBIE/): "To the extent that this is an allusion to NSA purchasing, it violates BCP 188 ('IETF Will Work to Mitigate Pervasive Monitoring')."

The draft violates the WG charter. See, e.g., my email dated 15 Apr 2025 22:33:23 -0000 (https://mailarchive.ietf.org/arch/msg/tls/xJqwB3Ob5wf3GVlAiIP3O4tuBIE/): saying that "the draft's regression from ECC+PQ to just PQ" is "a contravention of the 'improve security' goal in the WG charter".

There are no principles supporting the adoption decision. See, e.g., Stephen Farrell's email dated 1 Apr 2025 15:30:02 +0100 (https://mailarchive.ietf.org/arch/msg/tls/toxVUv_d1pdDspbfo8OxcJeC_QU/): "I don't see what criteria we might use in adopting this that wouldn't leave the WG open to accusations of favouritism if we don't adopt other pure PQ national standards that will certainly arise".

The draft's motivation section is circular. For example, my email dated 3 Apr 2025 16:18:57 -0000 (https://mailarchive.ietf.org/arch/msg/tls/YNSu6ZO5e0JMJ1cRlnxh6oIjyZg/) said that there is "a preliminary step that has been skipped here, namely identifying why the proposal is claimed to be adding something important. The draft's motivation sentence consists of rearranging buzzwords without answering the question: 'Having a fully post-quantum (not hybrid) key agreement option for TLS 1.3 is necessary for migrating beyond hybrids and for users that need to be fully post-quantum.' "

The draft increases software complexity. See, e.g., Andrey Jivsov's email dated 15 Apr 2025 13:49:52 -0700 (https://mailarchive.ietf.org/arch/msg/tls/uOmcMEqlyekrvcOgdsf7GtIlf3w/): "The main stated benefit of using a standalone ML-KEM is complexity reduction, but with the current progress in the deployment of the ML-KEM + ECC hybrid method, a standalone ML-KEM method actually increases overall complexity in software stacks."

As context: Thomas Bellebaum, in email dated 1 Apr 2025 15:18:16 +0000 (https://mailarchive.ietf.org/arch/msg/tls/YyemGJF-4-hRVwOcJ47Rw4Nu8Js/), had quoted "users that need to be fully post-quantum", and had asked for "a specific example of such users and their motivation". The draft author sent a reply dated 1 Apr 2025 11:31:55 -0400 saying "A specific example is moving to a compute / dependency base that is minimalist to only PQ primitives they wish to maintain, such as those that have long update / deployment cycles, as well as those that want a minimalist PQ interop target". Jivsov's objection says that adding the non-hybrid option actually makes software *more* complicated overall. The non-hybrid option doesn't exist in a vacuum: it is on top of the already deployed hybrid option.

5 Lack of consensus

A consensus declaration by a standardization organization does not require unanimity. However, it must be preceded by all of the following:

- general agreement;
- fair consideration of each comment;
- a process of attempting to resolve each objection; and
- documentation—for any objection that was not resolved but that was instead overridden by general agreement—of *why* that objection was overridden.

As this section shows, none of these criteria were met for the consensus declaration at issue in this complaint.

There was not general agreement. During the adoption-call period, there were statements from 20 people unequivocally supporting adoption: David Adrian, Joseph Birr-Pixton, Uri Blumenthal, GCHQ's "Flo D", NIST's Quynh Dang, Viktor Dukhovni, Scott Fluhrer, NSA's Rebecca Guthrie, Russ Housley, Alicja Kario, Kris Kwiatkowski, Andrei Popov, Tirumal Reddy, Yaroslav Rosomakho, Jan Schaumann, Sophie Schmieg, Martin Thomson, Filippo Valsorda, Loganaden Velvindron, and Thom Wiggers.

There were also statements from 2 people *conditionally* supporting adoption: Yaakov Stein ("I support adoption of pure PQC KEMs drafts with Intended status: Informational (meaning that the IETF is not recommending using)") and John Mattsson ("I support adoption as long as reuse of ephemeral keys is normatively forbidden, i.e. MUST NOT reuse").

However, there were statements from 7 people unequivocally opposing adoption: Thomas Bellebaum (https://mailarchive.ietf.org/arch/msg/tls/YyemGJF-4-hRVwOcJ47Rw4Nu8Js/: "I agree with Stephen on this one and would not support adoption of non-hybrids"), Andrey Jivsov (https://mailarchive.ietf.org/arch/msg/tls/uOmcMEqlyekrvcOgdsf7GtIlf3w/: "I am opposed to the adoption of ML-KEM at this time"), Stephen Farrell (https://mailarchive.ietf.org/arch/msg/tls/toxVUv_d1pdDspbfo8OxcJeC_QU/: "I'm opposed to adoption, at this time"), Rich Salz (https://mailarchive.ietf.org/arch/msg/tls/0f6XBGPElMLNoiS1Eh3u7EhzoGc/: "I was all set to say that I am in favor of adoption, but Stephen's post changed my mind. [paragraph break] The conservative and safe thing is to stick to hybrids and that is what the IETF should do for now"), Rob Sayre (https://mailarchive.ietf.org/arch/msg/tls/uhWI53zIjWJT5ZiS_UthyBe2k9Y/: "I oppose adoption"), Sun Shuzhou (https://mailarchive.ietf.org/arch/msg/tls/EzKcwjagajQqcRpH4TDTn70W9hc/: "I'm opposed to adoption"), and me. (There was much more text stating reasons for the objections.)

Even assuming that the 2 statements of conditional support are treated as positive votes, the overall situation here—22 positive votes and 7 negative votes—does not qualify as general agreement. "General" means "shared by or affecting most people, or most of the people in a group" (https://www.ldoceonline.com/dictionary/general); "most" means "nearly all of the people or things in a group, or nearly all of something" (https://www.ldoceonline.com/dictionary/most); the phrase "general agreement" means that nearly everyone agrees. Merely having three quarters agree is not good enough.

Standards-development organizations do not have authority to deviate from the legal definition of consensus. Within that definition, ultimately courts will interpret the words "general agreement" according to the plain meaning of the words, and will not allow "general" to be interpreted as merely "majority"—if regulators had meant "majority" then they would have written "majority". When there are gray areas, courts will allow organizations to issue rules that reduce the ambiguities (for example, ISO requires "the absence of sustained opposition to substantial issues by any important part of the concerned interests"), but will not allow those to override the plain meaning.

In particular, if a standards-development organization issues a rule declaring that "general agreement" exists even when a quarter of the votes are in opposition, courts will reject this as being inconsistent with the plain meaning of "general agreement". Anyway, IETF has not attempted to issue such rules: on the contrary, IETF claims in, e.g., https://web.archive.org/web/20250603130154/https://www.ietf.org/about/introduction/ that WG decisions are *not* taken by voting.

There was not fair consideration of each objection. Within the statements in favor of adoption, most of the statements were very short: e.g., just the words "I support adoption" with no further comments.

Some statements in favor of adoption did say more, such as stating circular arguments for the draft (e.g.: "as time progresses, non-hybrid key exchanges will become more and more commonplace, so why not have it already defined?"), or expressing concerns about key reuse (e.g.: "I also share John's concerns about key reuse, but would prefer to litigate that in the working group, rather than during adoption"), without responding to the content of the objections.

There was a response to one word in the lack-of-principles objection. The response was as follows: "The NIST competition was international, and Kyber was developed by an international team. I struggle to understand

how adopting this document would somehow be 'favoritism'." A brief note by one supporter tangentially related to one objection falls far short of fair consideration of each objection by the group as a whole.

There was not a process of attempting to resolve each objection. I tried to engage that supporter in discussion. I started by quoting the following earlier statement in the commentator's message: "I find it to be cognitive dissonance to simultaneously argue that the quantum threat requires immediate work, and yet we are also somehow uncertain of if the algorithms are totally broken. Both cannot be true at the same time." I responded as follows:

> Rolling out PQ is trying to reduce the damage from an attacker having a quantum computer within the security lifetime of the user data. Doing that as ECC+PQ instead of just PQ is trying to reduce the damage in case the PQ part is broken. These actions are compatible, so how exactly do you believe they're contradictory?
>
> Here's an analogous example of basic risk mitigation: there's endless work that goes into having planes not crash, not hit turbulence, etc., but we still ask airplane passengers to keep their seatbelts on whenever they're in their seats.

My email was dated 1 Apr 2025 21:38:16 -0000 (https://mailarchive.ietf.org/arch/msg/tls/2Dfu4x678DEKCzF-fkdvJHJkS-8/). By the time the adoption call closed two weeks later, there was still no reply.

The broader pattern was that objectors were engaging in discussion while supporters were not. The majority process wasn't "attempting to resolve each objection"; it was simply collecting positive votes.

There was not documentation, for each objection, of why that objection was overridden. When there's an objection, consensus requires not just fairly considering the objection, and not just attempting to resolve the objection, but—if resolution fails—having the group agree on the contents of a *response* to the objection. That's an official statement of *why* the objection was overridden.

For the proposal at hand—the proposal for the WG to adopt a particular draft—there was an objection saying, in a nutshell, that the draft creates security risks. Where's the official TLS WG statement of why this objection was overridden? Answer: The statement doesn't exist. Same for all of the other objections.

Fundamentally, consensus evaluation was replaced by a majority-voting process. A majority-voting process allows the majority to *override* objections from the minority without even *answering* those objections, let alone trying to *resolve* them. That's what happened here.

The chairs asked objectors to state their objections ("If you do not support adoption of this draft, please send a message to the list and indicate why"). Meanwhile the chairs asked supporters merely to state their support ("If you support adoption and are willing to review and contribute text, please send a message to the list"). The chairs didn't ask supporters to respond to objections. Unsurprisingly, there were detailed statements of objections, while the majority simply cast their votes without responding to the objections.

If this *wasn't* a majority-vote process, what's the supposed dividing line between this process and a majority-vote process? How would minority interests ever be protected against being overrun by the majority?

The call for adoption succeeded in achieving majority support, but it failed to achieve consensus. The chairs should have clearly explained from the outset that consensus was required, and should have accurately explained what this means. (Presumably the objections would then have been discussed—perhaps resolved one way or the other.) But the chairs didn't do this.

6 The erroneous consensus declaration, and questioning it

After a message dated 14 Apr 2025 00:02:15 -0400 saying "Just a reminder that this WG adoption call closes tomorrow", Turner sent email dated 15 Apr 2025 13:26:43 -0400 (before the announced closing date of "2359 UTC on 15 April 2025") saying "It looks like we have consensus to adopt this draft as a working group item". There were some notes on followup procedures, but there was no explanation of the rationale for this claim of consensus.

I sent email dated 15 Apr 2025 19:53:51 -0000 quoting "It looks like we have consensus to adopt this draft as a working group item" and continuing as follows:

> Um, what? There were several people (including me) raising objections on list to basic flaws in this draft, such as (1) the failure to provide an ECC backup to limit the damage from further security problems in the PQ layer, (2) the failure to provide an engineering justification for this option, and (3) the lack of any principles that would justify saying no to options selected by other governments if this option is allowed.
>
> Your message doesn't explain how you came to the conclusion that there's consensus. Surely you aren't relying on some tally of positive votes to ram this document through while ignoring objections; voting isn't how IETF is supposed to work. So how _did_ you come to this conclusion?
>
> As a procedural matter, this lack of explanation is in violation of "IETF activities are conducted with extreme transparency, in public forums". Please rectify this violation immediately. Also, please state the procedures for appealing your action. Thanks in advance.

7 AD disruption

One of the ADs, Paul Wouters, sent email dated 16 Apr 2025 09:36:17 -0400 regarding the dispute.

Recall the procedure specified in BCP 9 (RFC 2026), Section 6.5.1: someone who disagrees with a WG recommendation *first* discusses the matter with the WG chairs. *If* "the disagreement cannot be resolved in this way", then the procedure authorizes anyone involved to ask the ADs to "attempt to resolve the dispute".

I had summarized my disagreement and had started a discussion with the WG chairs. Instead of waiting for the discussion with the chairs to finish as per BCP 9, the AD was jumping into the discussion. (The chairs hadn't even posted a reply yet!) Furthermore, the contents of the AD's message (see below) were simply taking sides, rather than attempting to resolve the dispute.

Basic flaws. Regarding my comment on "basic flaws in the draft", the AD wrote the following:

> The term "basic flaws" here is mis-used. There are no known "basic flaws" in pure ML-KEM. If you know of a flaw, please present evidence in the form of a proper reference to the flaw. Your preference for hybrid over pure is not a "basic flaw", and those preferring hybrids can choose to only use hybrids.

This AD comment, like the voting process that preceded it, is non-responsive to the content of the objections.

The first objection that I had highlighted was to the **security risk** incurred by the draft leaving out the common-sense protection of a hybrid. As I said, "the failure to provide an ECC backup to limit the damage from further security problems in the PQ layer" is a basic flaw in the draft. Claiming that this is merely a "preference" is not even acknowledging, let alone responding to, the core point of the objection.

The second and third objections that I had highlighted were to critical gaps in the rationale provided for the draft. Both of these, like the first objection, are foundational issues challenging the notion that the draft is a good idea in the first place, so the word "basic" is proper terminology.

The AD did move on to quoting the specific objections, but did not respond to the content of the objections. See below.

Failure to provide an ECC backup. Regarding the objection to "the failure to provide an ECC backup to limit the damage from further security problems in the PQ layer", the AD wrote the following:

> Not being a hybrid KEM is not a "basic flaw".

The only reason that CECPQ2 didn't expose user data to pre-quantum attackers is that it had the common sense to include an ECC layer.

> The additional security from hybrids comes at a complexity cost that people have different opinions about.

Costs are facts, not opinions. Upgrading from ECC to ECC+PQ is only marginally more expensive than switching to non-hybrid PQ; see https://blog.cr.yp.to/20240102-hybrid.html for quantification. Furthermore, given the fact that the ecosystem includes ECC+PQ anyway, adding non-hybrid PQ as another option makes the ecosystem *more* complicated.

> There will also obviously be differences of opinion on when hybrids will have outlived their security premise in the future,

The proposal was to adopt the draft *now*. The objections were to that proposal. The AD's argument about the future is not responsive to the objections, and in particular is not responsive to "the failure to provide an ECC backup to limit the damage from further security problems in the PQ layer".

As a side note, the AD's argument starts with the claim that hybrids will eventually disappear. This *could* be correct, but maybe not; see the "cheaper to attack" paragraph in https://blog.cr.yp.to/20240102-hybrid.html.

> and so supporting both now and letting implementers make their own choices on which defaults to use now and when to migrate in the future is up to them.

This is not responsive to the objections; it is another circular argument that the document is good for people who think it's good.

> The TLS WG, along with CFRG backed by the larger cryptography community will continue to play an advisory role here over the next years.

This generic comment is not responsive to the objections.

Failure to provide justification. Regarding the objection to "the failure to provide an engineering justification for this option", the AD wrote the following:

> This is your own made up condition.

No, it isn't: "Rather than bringing a fully-formed solution and looking for a use, begin by articulating *what issue or gap needs to be addressed*. ... In other words, *don't put the cart before the horse*: first convince the group that there's an important problem to solve."

These quotes aren't from something binding on the TLS WG—they're quotes from a CFRG process document (see https://web.archive.org/web/20250325135726/https://wiki.ietf.org/en/group/cfrg/CFRG-Process)—but they're still doing a nice job of pinpointing one of the basic flaws in the draft at issue.

IETF claims (https://www.ietf.org/blog/ietf-llc-statement-competition-law-issues/) that "IETF participants use their best engineering judgment to find the best solution for the whole Internet, not just the best solution for any particular network, technology, vendor, or user". The available evidence indicates that this claim is not true in this case: see Section 3, or simply consider the AD's claim that it's okay to not provide an engineering justification. If there *is* an engineering justification for the draft, then this should have been spelled out before the adoption call.

Anyway, claiming that an objection is a "made up condition" doesn't change the fact that the objection was raised. Consensus requires each objection to be addressed: see Section 5.

> People who do not wish to rely on pure PQ can already use a hybrid PQ. There are those who wish to use pure PQs, and your reasons for not letting them are not widely supported within the IETF or the TLS WG, as can be seen by other protocols also implementing pure PQ algorithms.

This is another circular argument that the document is good for people who think it's good. This is again not responding to the objections.

As for "other protocols also implementing pure PQ algorithms", certainly the effects described in Section 3 are creating *some* of this regression, but leaping from such examples to the claim of hybrids being "not widely supported" is wildly inaccurate. The post-quantum connections from Chrome etc. to Cloudflare—a third of Cloudflare's HTTPS connections this year—are hybrid ECC+PQ. ANSSI requires hybrids. BSI requires hybrids. Remember that https://web.archive.org/web/20220524232250/https://www.nsa.gov/Portals/75/documents/resources/everyone/csfc/threat-prevention.pdf asked for two independent encryption layers "to mitigate the ability of an adversary to exploit a single cryptographic implementation".

As for "reasons ... not widely supported within the IETF or the TLS WG": The TLS WG wasn't polled (and certainly IETF as a whole wasn't polled) regarding "reasons", so how does the AD claim to know what the support fraction is for these reasons?

The TLS WG chairs merely asked supporters for votes, not for explanations. Some of the supporters nevertheless tempered their votes with text communicating concerns (for example, https://mailarchive.ietf.org/arch/msg/tls/hz-BtcGhXX2eN_rbVKMypP8XhW8/ said "I might oppose Recommended: Y"; https://mailarchive.ietf.org/arch/msg/tls/Yvjdn-wpF440E-6kjVHsrQMU5Nk/ said "we should be very careful"; https://mailarchive.ietf.org/arch/msg/tls/ppDcmr9twLRiMh-3hGYc0S_t66U/ supported adoption but only if IETF is "not recommending using"). Presumably security was the top source of concerns. Certainly security featured prominently in the stated objections.

Meanwhile far fewer people said that they *weren't* concerned about security. Perhaps the people who didn't say anything one way or the other weren't concerned, but *they didn't say that*. Perhaps the data on this point was biased by the nature of the call (again, the chairs didn't ask supporters to explain their votes), but the AD's claims here *definitely* aren't backed by evidence. It's improper to report guesses as facts.

Failure to provide selection principles. Regarding the objection to "the lack of any principles that would justify saying no to options selected by other governments if this option is allowed", the AD wrote the following:

> This document does not set policy for other documents or governments, so this "reason" is out of scope for the IETF.

Non sequitur. Supporting endless options is a systemic security problem, so the WG shouldn't take every option that's proposed—but then there should be principles for the dividing line. This is entirely about what the WG is endorsing, not about the level of WG power over anyone else.

NSA's selection of non-hybrid Kyber has been repeatedly, sometimes explicitly, cited as justification for the draft in question. NSA is a United States government agency. Meanwhile other governments are making different, often incompatible, selections: for example, hybrid FrodoKEM is one of the selections by BSI, a German government agency.

If the TLS WG is adopting the U.S. government selection, will the TLS WG also adopt the German government selection, the Chinese government selection, etc.? There have to be principles for the answer—engineering principles, not giving special power to the United States government.

The data flow here is from the governments *to* the TLS WG. What the AD is talking about is (1) in the opposite direction and (2) considering only the extreme case of setting policy.

The consensus question. Regarding my question of how the chairs came to the conclusion that there's consensus, the AD wrote the following:

> I have reviewed the responses to this WGLC. There is clearly consensus based on the 67 responses to the adoption call.

There were only 29 people responding to the adoption call (even if some of them, including me, sent multiple messages). Only 22 stated support for adoption, and only 20 of those were unequivocal, while 7 stated unequivocal opposition.

In short, the adoption proposal didn't even reach general agreement, never mind the other criteria for consensus. See Section 5. It's wrong to equate a clear *majority* with clear *consensus*.

By highlighting the number 67, and by claiming that consensus was "clear", the AD's message discourages reviewers from checking the facts. This is inappropriate.

> I support the TLS WG Chairs decision on calling consensus.

As I stated above, the AD's message is taking sides, not attempting to resolve the dispute.

> If you wish to appeal the TLS WG Chairs decision based on RFC 2026, Section 6.5.1 you may do so by contacting me using a working email address.

This is violating the BCP 9 procedure. The procedure clearly states that "A person who disagrees with a Working Group recommendation shall always first discuss the matter with the Working Group's chair(s)". I had briefly summarized the top three objections to the draft, and had asked the chairs how they came to the conclusion that there was consensus. I was still waiting to hear back from the chairs. The AD should have allowed that process to proceed, rather than disrupting it.

> If you present no new information to your appeal to the chairs, I would deny your appeal.

So we have one of the ADs preemptively declaring, before even receiving an appeal, that the appeal will be rejected (unless the appeal somehow digs up "new information" beyond the archives of the responses to the adoption call). This is a violation of "The Area Director(s) shall attempt to resolve the dispute", and a violation of due process.

```
My decision could then be appealed with the IESG.
```

The existence of a subsequent appeal stage does not remove the "shall attempt to resolve the dispute" requirement placed upon the ADs.

```
The vast majority was in favour of adoption,
```

"Vast majority" means "almost all of a group of people or things" (https://www.ldoceonline.com/dictionary/the-vast-majority-of-something). The 22 people in favor (including 2 with conditions) are not "almost all" of the 29 people who spoke up.

```
and this included several vendors who stated they have implementations.
```

I noticed only two such statements. Also, if we're counting vendors, then why were Google's David Adrian and Google's Sophie Schmieg allowed to cast separate votes?

Anyway, IETF claims that "Participants engage in their individual capacity, not as company representatives". More to the point, objection #1 here is not to the draft's implementability, but to its security risks.

```
There were further no raised technical issues.
```

I don't know what this claim means, and I don't know why it's supposed to be relevant. The concept of consensus puts constraints on how *all* comments and objections are handled, not just "technical issues".

```
There were a few dissenting opinions that preferred pure PQ should not be done
at all.
```

Recall the earlier text claiming that there is "clearly consensus based on the 67 responses to the adoption call". Notice the contrast between "67 responses" and "a few dissenting opinions". The reader is being told that the supporters-to-objectors ratio was something like 20 to 1.

But this is simply not true. The actual tallies are as follows: during the specified adoption-call period, 20 people expressed unequivocal support, 2 people expressed conditional support, and 7 people unequivocally objected. See Section 5.

It is astounding that the AD still has not issued an erratum regarding "67 responses ... vast majority was in favour of adoption ... There were a few dissenting opinions". The gap between the AD's text and the facts was pointed out six weeks ago; see Section 10.

```
Note that this document does not set a mandatory to implement or RECOMMENDED Y
option, allowing those who wish to avoid pure PQ to keep avoiding these in the
future. Your arguments on whether hybrid is more secure than pure would be valid
arguments in a discussion about MTI or RECOMMENDED status. However, this is not
that discussion.
```

This is not responsive to the objections. The objections were stated as objections to the actual proposal on the table, the proposal to adopt the draft. They were not merely objections to a potential proposal to recommend or require the draft.

The transparency violation. Regarding my statement that the lack of explanation for the consensus call was in violation of "IETF activities are conducted with extreme transparency, in public forums", the AD wrote the following:

> There is no such violation,

The public was not provided records showing how the chairs concluded that there was consensus. That's not "extreme transparency".

This is also not compliant with the record-keeping requirements in BCP 9, Section 8, which requires a public record of "complete and accurate minutes of meetings" along with "all written contributions from participants that pertain to the organization's standards-related activity". The chairs must have discussed the consensus question, whether by a virtual meeting (telephone, Zoom, etc.) or by a physical meeting or by email; so where are the records of the email, or the minutes of the meeting?

> and you cherry-picking when to call consensus evaluation "voting"

The word "voting" accurately describes what actually happened in this incident. See Section 5. If it walks like a vote and quacks like a vote then it's a vote.

> depending on whether misnaming this is in your advantage (eg recently on the ssh list) is dishonestly manipulative and has no place on this list or anywhere else at the IETF.

This is not responsive to any of the objections at hand, and also doesn't answer the question of how the chairs arrived at their conclusion that there was consensus. The AD is also violating IETF's code of conduct by issuing this ad-hominem attack.

> Your insinuation that this WGLC was not conducted with "extreme transparency" is in itself a violation of our code of conduct

No. It's properly filing a complaint about a transparency violation: "As a procedural matter, this lack of explanation is in violation of 'IETF activities are conducted with extreme transparency, in public forums'. Please rectify this violation immediately."

As a side note, the word "insinuation" means "something that someone says which seems to mean something unpleasant, but does not say this openly" (https://www.ldoceonline.com/dictionary/insinuation).

> through insinuations and a continuation of behaviour you have been warned about recently by the TLS WG chairs, confirmed via me as the TLS AD, and the IESG[1]. I recommend you voluntarily stop this kind of behaviour to avoid triggering measures under the terms of RFC3934 which is part of BCP25.
>
> You are free to voice your dissent. You are not free to make up accusations against process or individuals.

The specific citation "[1]" is to https://datatracker.ietf.org/group/iesg/appeals/artifact/126, which actually said that an appeal to IESG was misdirected, without commenting on the content of the appeal. The same topic is now under appeal to the IAB. Anyway, this text from the AD is again not responding to the topics at hand; it is just another ad-hominem attack.

8 Questions about the AD disruption

I sent email dated 16 Apr 2025 15:10:18 -0000 with clarification questions about portions of the AD's message. I started by quoting "Responding as AD" and continuing as follows:

> Hmmm. I thought that a "person who disagrees with a Working Group recommendation shall always first discuss the matter with the Working Group's chair(s)", whereas AD involvement is only if "the disagreement cannot be resolved in this way". This provides multiple levels of opportunities to resolve disagreements.
>
> So I posed a question to the chairs: specifically, asking how they came to the conclusion that there was consensus here. I also explained why I was asking. (Procedurally, I also shouldn't have to ask.)
>
> Does the new AD interruption mean that the chairs are no longer obliged to engage in discussion of their action? In other words, has the AD single-handedly destroyed a mandated opportunity for resolution? If so, what's the authorization for this under IETF procedures?
>
> The situation was already a bit messy before this (for example, were the chairs deterring input when they issued a consensus declaration before the end of the call period?), but at this point it's very difficult to figure out how the situation relates to how the IETF standardization process is supposed to work.
>
> I'm also not sure how this can be brought back to the proper procedures. Withdrawing the AD message isn't going to magically restore independent evaluation by the chairs.

I then quoted the AD's "There is clearly consensus based on the 67 responses to the adoption call" and "The vast majority was in favour of adoption" vs. "a few dissenting opinions". I asked the following questions:

> I have an easy question and a harder question.
>
> The easy question, just to make completely sure that I'm not missing something: You're saying that the numbers here, such as "67" and "a few", were considered as part of your forming a conclusion that there's consensus here?
>
> (I assume the answer is simply "yes"---why else would the numbers have been brought up?---but I'd just like to make sure.)
>
> The harder question: For transparency, please explain how many different people you're referring to in saying "67 responses" and "vast majority" and "a few", and please provide details so that the rest of us can check your tallies.
>
> My impression from watching the list is that the actual ratio between the numbers of objectors and supporters is vastly larger than the ratio between "a few" and "67", for any reasonable understanding of "a few".

Finally, regarding the AD's claim that there were "no raised technical issues", I wrote the following:

> Can you please clarify what exactly you mean by "technical" here, why this criterion factors into the question of whether there's consensus, and why the issues raised (e.g., the security risks of non-hybrids) don't qualify as "technical"? Thanks in advance.

I also included short responses to specific AD comments on the objections that I had highlighted as flaws 1, 2, and 3. Those responses are included in the more comprehensive text in Section 7, so I won't repeat them here.

9 AD disruption, continued

The AD sent email dated 16 Apr 2025 20:43:21 -0400 as follows.

The violation of BCP 9's resolution procedures. I had written the following: "I thought that a 'person who disagrees with a Working Group recommendation shall always first discuss the matter with the Working Group's chair(s)', whereas AD involvement is only if 'the disagreement cannot be resolved in this way'. This provides multiple levels of opportunities to resolve disagreements." The AD wrote the following:

> If you look at an individual issue, then yes that is the regular procedure. In your case, you seem to object to most WG decisions not in your favour and question motivations of every decision and individual involved in the decision chain. And frankly, it is already a denial of service on the time of many volunteers within IETF, from WG chair to the IESG.
>
> To make it more clear and blunt, you calling into question this consensus call of the WG chair is abusive and follows a repetitive pattern. Nevertheless, for now this is your right, and we will walk through the process.

Here the AD is engaging in further ad-hominem attacks, again violating IETF's code of conduct. The reader is forced to wade through all of this to see that the AD isn't responding to the point at hand, namely that the AD violated BCP 9.

The transparency question, again. Regarding my paragraph "So I posed a question to the chairs: specifically, asking how they came to the conclusion that there was consensus here. I also explained why I was asking. (Procedurally, I also shouldn't have to ask.)", the AD wrote the following:

> Unfortunately, it looks like you are attempting to bait the chairs to say they took inventory of the public emails and then throw in some quotes about "you counted votes but IETF does not vote".

In fact, my first message questioning the claim of consensus had said the following: "Surely you aren't relying on some tally of positive votes to ram this document through while ignoring objections; voting isn't how IETF is supposed to work. So how *did* you come to this conclusion?"

So the AD's "bait" claim makes no sense. I had *already* pointed out that voting isn't how IETF is supposed to work. I was asking for transparency regarding how the chairs had arrived at their conclusion that there was consensus. At this point there still wasn't an answer from the chairs; there was just the AD disruption.

> My previous email explained the obvious way the consensus was validly called. This can be independently verified by anyone reading the email thread.

No. Repeatedly declaring something to be clear and obvious doesn't make it so, nor does it answer the transparency question about how the chairs had arrived at their conclusion.

(The chairs later ended up providing information that can't be reconciled with what the AD claimed was "the obvious way the consensus was validly called". See Section 12.)

> The fact that you are the only one questioning the consensus should be an indication that your reasoning to doubt the consensus call might in fact be erroneous.

When an objection is raised, the content of the objection should simply be addressed. Sometimes people speak up with clarifications to the objection, supplements to the objection, etc., but if the objection is clear and complete in the first place then having other people speak up merely to reiterate the objection is neither required nor desirable. This is supposed to be the Internet Engineering Task Force, not the Internet Politics Task Force.

Replying "you are the only one questioning the consensus", and claiming that this is an indication of error, is both procedurally improper and factually unsupported. In this case it's also factually incorrect, both in the premise and in the conclusion. (In response to the AD, Thomas Bellebaum wrote "He is not the only one"; see Section 10. Regarding the conclusion, see Section 5.)

The violation of BCP 9's resolution procedures, again. Regarding my questions "Does the new AD interruption mean that the chairs are no longer obliged to engage in discussion of their action? In other words, has the AD single-handedly destroyed a mandated opportunity for resolution? If so, what's the authorization for this under IETF procedures?", the AD began by issuing a threat:

> Dan, there comes a point where you will be prevented from further playing these games. There are processes for that, that we really try hard to avoid invoking. But as some point you leave us no choice.

The AD continued by claiming that consensus was not just "obvious" but "*very* obvious":

> This consensus call was _very_ obvious based on the email thread content, again as I explained in my previous message. Whether the TLS Chairs feel obliged to send you another message repeating the obvious is pretty irrelavant other than taking up valuable time an energy of an entire WG in playing a process game with you. Unless you are invoking an appeal as per RFC 2026 Section 6.5.1 against the WG chairs decision that there is consensus to adopt, they are under no obligation to answer you with something they deem obvious. It is completely up to the chairs to make their own decision here. Either way is acceptable in our process.

I had questioned the consensus call, briefly explaning why and asking the chairs for explanation. This was following the BCP 9 procedure ("A person who disagrees with a Working Group recommendation shall always first discuss the matter with the Working Group's chair(s)").

But I hadn't explicitly *pointed* to BCP 9. The AD's paragraph seems to be saying that, because of this, the chairs were not obliged to engage in discussion. I later followed up with an explicit reference to BCP 9; see Section 11. Regarding "game", see above.

> Once you send an RFC 2026 Section 6.5.1 appeal to them, according to process they MUST respond to you. Presumbly once denied - if they are not convinced by your arguments in your appeal - you can then send the same text, with your usual disclaimer that in your opinion I need to recuse myself, to me as TLS AD, and I will reply with "based on the public discussion on the list, with the overwhelming majority being in favour of adoption as long as the MTI/RECOMMENDED values would remain "NO", with a few dissenting views of wanting to block all pure PQ in all IETF protocols in favour of IETF only adopting hybrids, and with no technical flaws pointed out in the specified protocol by anyone, considering there are already a number of interoperable implementations based on early code points, it is unmistakenly clear that the TLS Chairs correctly called consensus on adoption of this document. Your appeal is denied". Upon which you can file another appeal of my decision to the IESG.

Content-wise, this is similar to the previous claims about "67 responses" with just "a few dissenting", and about there being "no technical flaws". Some of the wording is different: "vast majority" has changed to "overwhelming majority", and "clear" has changed to "unmistakenly clear".

The violation of the specified call period. Regarding my comment "The situation was already a bit messy before this (for example, were the chairs deterring input when they issued a consensus declaration before the end of the call period?)", the AD wrote the following:

> According to
> https://datatracker.ietf.org/doc/draft-connolly-tls-mlkem-key-agreement/history/
> a two week adoption call went out on 2025-04-01 and the document status was set
> to adopted on 2025-04-15. (datatracker provides no finer granularity in its
> History tab) This matches the email dates on the respective TLS email messages:
>
> Start: 01 April 2025 12:58 UTC
> https://mailarchive.ietf.org/arch/msg/tls/PpVAwrBTuRb5pR6D0C1ipdQuvYc/
>
> End: 15 April 2025 17:27 UTC
> https://mailarchive.ietf.org/arch/msg/tls/_AWy51BSgX1ipv0hfnAzLrDrTYI/
>
> You are correct that Sean did say in the announcement that the call would close
> at "2359 UTC on 15 April 2025", so indeed technically speaking it was called 6
> hours too early. However, usually adoption and last calls are send out for a
> period of weeks and usually chairs send out a message on which day a call ends
> without further hourly granularity. Regardless, it was obvious that at the time
> no active discussion about fundamental issues was taking place and calling this
> adoption ended on the last day of the adoption call period caused no stiffling
> of discussion. I am further confident that if any real discussion had taken
> place, the chairs would have not called it and would have extended the adoption
> call to give any active discussions more time to settle. I also see no valid
> reason to extend the adoption call by 6 hours.

The claim that no "real discussion had taken place" is correct in the sense that supporters were ignoring the content of the objections. But this was improper—recall from Section 5 that consensus requires addressing each objection—and certainly cannot justify further procedural violations.

The lack of procedural clarity. Regarding "at this point it's very difficult to figure out how the situation relates to how the IETF standardization process is supposed to work", the AD wrote the following:

> As in all cases regarding WG level document disagreements on WG chairs
> decisions, you should follow RFC 2026 Section 6.5.1 as indicated to you a number
> of times over the last few months. If you feel the WG Chairs or AD is giving you
> conflicting information, you should stick to RFC 2026.

I don't find BCP 9 (RFC 2026) so clear—for example, when BCP 9 says "A person who disagrees with a Working Group recommendation shall always first discuss the matter with the Working Group's chair(s)", it doesn't say "the person shall cite this provision", whereas the AD seemed to think this was a critical requirement—but, as noted above, I did end up explicitly invoking it after the AD disruption. See Section 11.

The lack of independence. Regarding "I'm also not sure how this can be brought back to the proper procedures. Withdrawing the AD message isn't going to magically restore independent evaluation by the chairs", the AD wrote the following:

> I disagree we are deviating from existing procedures. You just had a glimpse of the obvious continuation of the process, were you to invoke process from RFC 2026 Section 6.5.1. You have not yet invoked that process as far as I know. You have until June 15 17:27 UTC to appeal.

This is not responding to the point about independence.

Under BCP 9, a disagreement is supposed to be discussed with the chairs first. That's a first stage that could resolve the dispute.

If that doesn't settle things, anyone can contact the AD, who "shall attempt to resolve the dispute". That's a second independent stage trying to resolve the dispute.

What happened here was instead that the AD replaced both steps of this procedure with preemptively taking one side of the dispute, not even giving the chairs a chance to resolve the dispute.

The AD's wrong numbers. Regarding my question "You're saying that the numbers here, such as '67' and 'a few', were considered as part of your forming a conclusion that there's consensus here?", the AD wrote the following:

> The use of the number 67 refered to the number and content of all messages in the adoption call email thread on the TLS list - which is the entire information base upon which the consensus call by the chairs took place, and also constitutes the information on why I believe the chairs reached the right conclusion.

This is dodging the question.

I had asked how the chairs had arrived at their conclusion regarding consensus. The AD had jumped in saying "There is clearly consensus based on the 67 responses to the adoption call. . . . The vast majority was in favour of adoption . . . There were a few dissenting opinions"—along with some other text, but the statements about "67 responses" and "vast majority" and "a few" are prominently placed and look like important parts of the AD's argument that there was consensus.

Having watched the actual responses to the adoption call, I was under the impression that the AD's numbers were divorced from reality. However, I hadn't done the work to tally the actual numbers at that point. So, before doing that work, I wanted to check that, yes, the AD's claim of consensus was based in part on these numbers.

Regarding my parenthetical comment "I assume the answer is simply 'yes'—why else would the numbers have been brought up?—but I'd just like to make sure.", the AD wrote the following:

> I am not answering your question as a boolean. See my previous paragraph.

This is an easy yes-or-no question. Yes means that these numbers are part of the rationale for the AD's conclusion. No means that they aren't.

Leaving out the numbers would have sounded very different—"clearly consensus based on the responses"; "There were dissenting opinions"; "The majority was in favour of adoption"—basically telling readers that the AD was defining "consensus" as "majority". What the AD actually wrote sounded much more lopsided (as one would expect given the requirement of general agreement): "vast majority" of "67 responses" vs. "a few dissenting". So I was expecting a quick response saying, yes, of course these numbers are part of the rationale.

Instead the AD refused to give a straight answer. Being ambiguous about the answer is a transparency violation, and sabotages the appeal process. Imagine doing the work to challenge the numbers, and *then* receiving a response saying "No, those numbers were never part of the rationale".

Regarding "please explain how many different people you're referring to in saying '67 responses' and 'vast majority' and 'a few', and please provide details so that the rest of us can check your tallies", the AD wrote the following:

> The only way to win is not to play. I am not playing your game of forcing me to use numbers only to have you call out "counting is voting".
>
> The content of 67 messages was produced by the WG. Based on the entirety of the content of those messages, consensus was determined.

This is dodging the question, and violating the requirement of transparency.

Regarding my stated impression "that the actual ratio between the numbers of objectors and supporters is vastly larger than the ratio between 'a few' and '67', for any reasonable understanding of 'a few' ", the AD wrote the following:

> I never put "a few" against "67". That is a misleading construct you devised, not me, nor the chairs.

This is baffling. Where is the AD's "vast majority" claim coming from, if not those numbers? Why didn't the AD provide exact numbers in the first place? Did the AD ever actually do the work to go through the messages?

Flaw 1, again. Regarding "The only reason that CECPQ2 didn't expose user data to pre-quantum attackers is that it had the common sense to include an ECC layer", the AD wrote the following:

> This is not new information. The WG heard your statement and people took it into consideration when they expressed their opinion on whether to adopt the document.

Saying that an objection isn't new (1) isn't responding to the objection, and (2) isn't defending the specific AD claim at issue, namely the AD claim that omitting hybridization isn't a "basic flaw".

Flaw 2, again. Regarding "the failure to provide an engineering justification for this option", the AD had falsely claimed that this "is your own made up condition", and in response I had quoted a CFRG document saying "begin by articulating *what issue or gap needs to be addressed*". Instead of admitting error, the AD wrote the following:

> If you had expressed these views at the start of the adoption call, people could have taken this into account. Some of the people that participated on the adoption call were undoubtedly already aware of these quotes.
>
> Regardless, "providing an engineering justification" is not something that one individual (you) can add to the TLS charter in an adhoc matter.

This is not responsive.

Certainly the charter can be a source of objections—and in fact it was in this case: I objected that this draft was in contravention of the "improve security" goal in the charter. But this doesn't mean that the charter is the *only* allowed source of objections. On the contrary, consensus requires general agreement *and* addressing *each* objection. See Section 5.

Furthermore, a WG charter cannot override IETF's broader claim that "IETF participants use their best engineering judgment to find the best solution for the whole Internet, not just the best solution for any particular network, technology, vendor, or user". It is puzzling that the AD persists in claiming that an engineering justification isn't required.

Flaw 3, again. Regarding "the lack of any principles that would justify saying no to options selected by other governments if this option is allowed", the AD had claimed that this was out of scope since the draft "does not set policy for other documents or governments", and I had said "Non sequitur. Supporting endless options is a systemic security problem, so the WG shouldn't take every option that's proposed—but then there should be principles for the dividing line. This is entirely about what the WG is endorsing, not about the level of WG power over anyone else". The AD wrote the following:

> This is not about "endless options". This is about pure ML-KEM. It is clear your view on pure ML-KEM is not universally agreed upon.

The objection here is explicitly looking beyond the particular draft at hand, and asking for principles regarding what to include and what to exclude. Making ad-hoc decisions is a due-process violation. Saying that this particular draft is about just one option is not addressing the objection.

"Technical issues". Finally, regarding the AD's claim that there were "no raised technical issues", I had asked "Can you please clarify what exactly you mean by 'technical' here, why this criterion factors into the question of whether there's consensus, and why the issues raised (e.g., the security risks of non-hybrids) don't qualify as 'technical'?" The AD wrote the following:

> I meant a concrete issue or flaw. Not a hypothetical one. Nor the number of airbags deemed not enough or too much in your hypothetical car with seatbelts.

Two preliminary notes seem warranted here. First, regarding vocabulary: "hypothetical" means "based on a situation that is not real, but that might happen" (https://www.ldoceonline.com/dictionary/hypothetical).

Second, regarding security: A tremendous amount of research and development has gone into systematically considering and proactively eliminating large classes of potential attacks—or at least proactively reducing the damage—rather than merely reacting to demonstrated attacks. This is what happens in every paper on security proofs. This is the motivation for a wide range of standard security tools: for example, key erasure reduces the damage if a device is compromised, and password hashing reduces the damage if a backup is compromised. This is also the motivation for hybrids, reducing the damage if post-quantum systems are compromised—as happened in the case of SIKE. Ignoring hypothetical attacks would be a remarkable regression from the state of the art.

Now back to the discussion. The AD's text here is very far from a clear answer to the three questions I had asked. If the AD is claiming that security failures of non-hybrids are hypothetical, how does the AD explain SIKE?

Furthermore, it is completely unclear how non-hypothetical is supposed to be connected to "technical" and, via that, to the question of whether there was consensus.

10 Terminating the AD disruption

Thomas Bellebaum sent email dated 17 Apr 2025 10:01:30 +0000 that started "I am sorry for interrupting your argument, but as you are discussing this on-list:", quoted the AD's "you are the only one questioning the consensus" paragraph, and then wrote the following (plus line breaks suppressed here):

> He is not the only one. Using the independently verifiable mail thread, I actually did count by a rough look over the messages (sorry if I missed/misinterpreted someone):

> Pro Adoption: - Alicja Kario - Andrei Popov - David Adrian - Filippo Valsorda - Flo D - Jan Schaumann - John Mattson - Joseph Birr-Pixton - Kris Kwiatkowski - Loganaden Velvindron - Martin Thomson - Quynh Dang - Rebecca Guthrie - Russ Housley - Scott Fluhrer - Sophie Schmieg - Thom Wiggers - Tirumal Reddy - Uri Blumenthal - Viktor Dukhovni - Yaakov Stein - Yaroslav Rosomakho
>
> Against Adoption: - Andrey Jivsov - Dan Bernstein - Rich Salz - Rob Sayre - Stephen Farrell - Sun Shuzhou - Thomas Bellebaum
>
> I am counting 22 expressions in favor of adoption and 7 opposing adoption. This amounts to about every fourth person objecting the draft in its current state at this time, which seems more than can be explained by mere blocking of few individuals.

The AD sent a followup dated 17 Apr 2025 09:04:10 -0400 saying "Note that the consensus call was for Working Group Adoption. Not publishing as is". This is non-responsive. The call was for adoption of the draft, and all seven people listed above (including me, obviously) were stating unequivocal objections to adoption of the draft. (See Section 5 for quotes and links.)

Bellebaum continued with procedural objections, which the AD never quoted and never replied to:

> I am not questioning that this is a sound majority, but consensus is a harsh word. Neither am I threatening to appeal, but I do share the view that merely declaring concerns such as "hybrids are way more conservative" as hypothetical/irrelevant to whether or not to publish this is not a reasonable way forward. The feeling (I am not saying "the fact") of this happening is valid. However, openly accusing others of playing games or ignoring procedures does not result in good specifications.
>
> Raised points should be discussed and adequately addressed to reach a consensus (i.e. significantly better than 3 out of 4). We are not making a black-or-white decision on publishing or not, we are influencing many aspects of the document.

Bellebaum finished by stating a wishlist for "the new WG item". The AD did quote that part, and claimed that "This sounds like you are not objecting to adoption, but objecting only to publication as is?". Bellebaum responded to the contrary: "I still believe that not adopting this would have been better, but I am willing to follow along and help improve the document."

Of course, standards-development organizations cannot force participants to withdraw an objection to a document as a condition for participating in further development of the document. More to the point, this complaint is challenging the validity of the consensus declaration in the first place. There was a specified period for the adoption call, and that call failed to produce consensus on adoption; see Section 5. Consequently, this draft was never a valid WG draft. The only way to change this would be to obtain general agreement while properly addressing every remaining objection.

Amazingly, the AD *still* has not withdrawn his text about "67 responses" with the "vast majority" supposedly in favor and just "a few dissenting". But Bellebaum's tallying of the facts did seem to stop the AD from commenting further, while it triggered further messages regarding the question of how to evaluate consensus.

11 Explicitly invoking RFC 2026

I sent email dated 18 Apr 2025 14:02:55 -0000 quoting Bellebaum's "I am counting 22 expressions in favor of adoption and 7 opposing adoption" and continuing as follows:

> Thanks for doing the work to tally this, and for posting the details so that people can check your message and post any necessary adjustments.

> These numbers sound radically different from the AD's portrayal ("67 responses ... vast majority was in favour ... a few dissenting opinions"). My own impression, from having read all messages as they came in, was about a quarter of the people opposing, so I will be very surprised if adjustments end up big enough to rescue the AD's portrayal.
>
> So: Can we please now have an explanation from the chairs of how they arrived at "It looks like we have consensus to adopt this draft as a working group item"?
>
> To prevent any confusion about the procedures: Based on what I've seen (the whole discussion, not just the fragmentary information conveyed by numbers), I disagree with this declaration of consensus. I am therefore invoking the "first discuss the matter with the Working Group's chair(s)" provision of RFC 2026, Section 6.5.1. I ask for this discussion to be on-list for transparency.
>
> Within that, what I'm suggesting---both because I think it's the natural way forward, and because of transparency considerations; I'm not saying this is the only possibility under RFC 2026---is for the chairs to start by explaining to the WG how they evaluated consensus, so that we can all consider the explanation, rather than starting with a bunch of conflicting guesses from the rest of us regarding how consensus might have been evaluated.

By now I have double-checked Bellebaum's list and concur with the tallies, except that the support statements from Mattsson and Stein were conditional (see Section 5 for quotes). To be clear, even if there had been 22 unequivocal supporters vs. 7 unequivocal objectors, that would not have constituted "general agreement", never mind the other requirements for declaring consensus.

12 Conflating consensus with interest

Turner sent email dated 18 Apr 2025 11:27:36 -0400, apparently on behalf of the chairs, as follows:

> Joe and I, as WG chairs and with Deirdre recusing as she is an author, declared consensus to adopt draft-connolly-tls-mlkem-key-agreement. We did this because there is clearly sufficient interest to work on this draft. Different working groups have different styles with respect to how much work is done by the individual author, versus how much work is done by the WG after adopting the work. Now that the draft is a WG draft, we will follow WG process by discussing concerns, already raised and new ones, under IETF change control and progressing after there is consensus.

I sent email dated 18 Apr 2025 16:47:14 -0000 starting by responding to the first two sentences as follows:

> Thanks for your message.
>
> "Sufficient interest to work on this draft" is ambiguous (sufficient for what?), and in any case clearly not the correct criterion for declaring consensus to adopt a draft.
>
> As an extreme example, this criterion would allow a draft to be adopted over amply justified objections of almost all WG participants, simply because the chairs and a few participants say they have enough interest in working on the draft! That's more extreme than what happened here, but it shows that the criterion stated above is procedurally improper.
>
> So I'm guessing that you had some further points in mind in deciding that there was consensus to adopt this draft. For transparency, can you please, without

> omissions, say why you declared consensus to adopt? Or, if the above really is the complete explanation, can you please say so explicitly, to enable an appeal saying that this was improper? Either way, can you please clarify what "sufficient" is referring to? Thanks in advance.

Regarding the "different working groups" sentence, I wrote the following:

> This generic background information about WG work allocation seems off topic (the topic being the disagreement regarding consensus). Certainly this background information doesn't say anything about the draft at hand. If I'm missing some connection, please elaborate.

Finally, regarding the "we will follow WG process" sentence, I wrote the following:

> This also isn't addressing the consensus question, plus it seems to be denying the existence of the active RFC 2026 Section 6.5.1 procedure challenging the chairs' decision to adopt in the first place.

Turner sent email dated 25 Apr 2025 15:04:39 -0400, apparently on behalf of the chairs, as follows:

> "Sufficient" to Joe and I means that there were enough people willing to review the draft. WGs groups have adopted drafts with much less support than this one received.
>
> Now that the document is adopted by the WG, consensus, as judged by the WG chairs (minus Deirdre because she is an author), is needed to progress the draft.
>
> Joe and I have reviewed the WG adoption call messages for ML-KEM Post-Quantum Key Agreement for TLS 1.3 [0] and stand by our consensus call. You can appeal this with the AD: Paul Wouters, but also consider his reply here [1].

Reference "[0]" was to https://datatracker.ietf.org/doc/draft-ietf-tls-mlkem/, and reference "[1]" was to https://mailarchive.ietf.org/arch/msg/tls/nqouPVfPtU7hm-RF0lSDHCfze54/.

Let's now review the procedure that the chairs say they used for evaluating consensus to adopt:

- The procedure asks purely whether there is "sufficient interest to work on this draft", meaning "enough people to review the draft". (Presumably the number of people is compared to the complexity of reviewing the draft.)
- The procedure does not place a minimum requirement on the number of people stating support. "Much less" than 22 is fine.
- This procedure does not consider objections to adoption.

Notice that, structurally, this procedure is an ad-hoc procedure for evaluating consensus to adopt, rather than being a general procedure to evaluate consensus on any decision.

Notice also that this procedure is radically different from what the AD had claimed was the obvious reason for the consensus declaration (67 responses, "vast majority" supporters, just "a few dissenting", etc.). Of course, that claim was before Bellebaum posted the actual numbers (22 supporters, 7 objectors).

Most importantly, this procedure is missing all of the consensus requirements reviewed in Section 5. The procedure does not even say that a majority is required, let alone general agreement, never mind any of the specific requirements for handling comments and objections.

I am not saying that this procedure, applied to the situation at hand, would say that there were not enough reviewers for the draft. I am saying that this procedure is not evaluating consensus. Properly evaluating consensus, as in Section 5, shows a variety of reasons that this adoption call failed to reach consensus.

13 Next steps

As stated at the top, this complaint is that the chairs erred when they declared consensus to adopt this draft. The procedure that they say they used was improper, having essentially nothing to do with the concept of consensus; see Section 12. Their conclusion that there was consensus is simply wrong; see Section 5.

BCP 9 authorizes bringing this "to the attention of the Area Director(s) for the area in which the Working Group is chartered". There are two ADs for this area, and I am bringing it to the attention of both of them.

BCP 9 then says that "The Area Director(s) shall attempt to resolve the dispute". Since there are two ADs, this procedure requires *both* of the ADs to attempt to resolve the dispute.

I have already, before this complaint, laid out a case that both ADs should recuse themselves from various matters. (The behavior by one AD regarding this matter is further evidence on point.) This complaint is within that scope: i.e., the ADs should turn this matter over to two neutral arbiters. However, my understanding is that the ADs are refusing to do so.

Either way, secret discussions among the ADs or other arbiters are not permitted by the record-keeping requirements in BCP 9, and are not permitted by IETF's requirement of "extreme transparency". The relevant IETF mailing list for this matter is the TLS mailing list, and I request that *all* discussion of this complaint be cc'ed to that list.

Beyond what BCP 9 says, all interested parties have a legal right to participate in standards-development activity, and in particular to appeal decisions by standards-development organizations, while receiving due process. The factually incorrect "vast majority" claim from one AD (compare Sections 7 and 9 to Section 10) perfectly illustrates the importance of due process. Neutral arbiters taking reasonable steps to avoid error would never have ended up making such a claim.

Finally, I have recently become aware that IETF Administration LLC believes that it can force parties to trade away other rights in exchange for exercising their rights to appeal. Concretely, IETF Administration LLC appears to believe that it is free to post modified versions of complaints, and that it is free to falsely attribute those modified versions to the original author, without regard to copyright law, moral-rights law (e.g., integrity rights), fraud law, etc. To be clear, those beliefs are incorrect. I have never consented to, and do not consent to, any such trade.