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

Daniel J. Bernstein, 2026-04-08

Contents

1 Overview

This is a complaint to the "Internet Architecture Board" (IAB) within the "Internet Engineering Task Force" (IETF). This complaint is self-contained and includes "a detailed and specific description of the facts of the dispute" as required by BCP 9. Various URLs are provided for background.

Let me be clear at the top about what's going on here. NSA is trying to buy IETF endorsement of weakened cryptography. IETF management is non-consensually ramming a particular NSA-driven document through the IETF process; issuing fake claims of consensus; and making up a series of retroactive, ad-hoc, selectively enforced excuses to avoid addressing objections.

This complaint is specifically regarding the following fake claim of consensus: Joseph Salowey and Sean Turner, in their roles as chairs of an IETF "Working Group" (WG) named "Transport Layer Security" (TLS), 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".)

I have three reasons for complaining about this specific action:

- The chair message claiming consensus did not present the basis for this claim. This lack of information obstructed followup procedures.
- The procedure that the chairs (eventually) said they used for the consensus declaration was improper, having essentially nothing to do with the concept of consensus. See Section 14 below.
- The chair conclusion that there was consensus is simply wrong. See Section 7 below.

Other sections of this complaint provide more information regarding the relevant background, the events during the adoption call issued by the chairs, and what happened when I challenged their claim of consensus.

It is important to set the record straight: to issue an erratum for the misinformation issued by the chairs on this topic. The chairs called for adoption; that call *failed* to reach WG consensus on adoption; the chairs were wrong in claiming consensus.

Uncorrected misinformation tends to snowball. For example, I am told that Paul Wouters, who at the time was an IETF "Area Director" (AD), wrote "it is possible that 30+ TLS participants, 2 Working Group Chairs, 13 IESG members and 13 IAB members are all corrupt and only djb is right" in a public posting in October 2025. Compared to the actual levels of support and opposition that were produced by the adoption call in question during the specified call period (see Section 7 below), the AD is overstating the level of support and understating the level of opposition. The underlying claim of consensus was an official claim by the chairs and needs a similarly official correction.

It is, of course, also important to undo the unauthorized chair action that directly resulted from the erroneous consensus claim, namely the adoption of the document as "draft-ietf-tls-mlkem", along with all followup actions regarding "draft-ietf-tls-mlkem". It is also clear that the chairs need training in how to evaluate consensus; the procedures that they are following threaten to do damage far beyond the specific incident at hand.

2 Procedural background and authorization for this complaint

BCP 9 (RFC 2026), Section 6.5.1, includes the following four 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.
>
> If the disagreement cannot be resolved by the Area Director(s) any of the parties involved may then appeal to the IESG as a whole. The IESG shall then review the situation and attempt to resolve it in a manner of its own choosing.
>
> If the disagreement is not resolved to the satisfaction of the parties at the IESG level, any of the parties involved may appeal the decision to the IAB. The IAB shall then review the situation and attempt to resolve it in a manner of its own choosing.

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

The rest of this section reviews my previous efforts to resolve this matter with the chairs, the ADs, and the "Internet Engineering Steering Group" (IESG), all of which are continuing to dodge (1) the question of what consensus means and (2) the question of whether consensus was achieved in this case. This history also shows that the preconditions for an appeal to IAB have been met.

I explicitly invoked the first of the above 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 claim and that I "can appeal". (The sequence of events is described in more detail below.)

The second BCP 9 paragraph has a prerequisite that "the disagreement cannot be resolved in this way". Procedurally, is this triggered when the chairs say that I "can appeal"? Or would ADs be able to evade addressing the content of a complaint by saying "The chairs were wrong in saying that you can appeal; there's no proof that the disagreement cannot be resolved by the chairs; go back to them"?

Despite the lack of procedural clarity, I sent email to the ADs dated 5 Jun 2025 18:51:36 -0000 to file https://cr.yp.to/2025/20250605-non-hybrid.pdf as a complaint under the second paragraph. I cc'ed the TLS mailing list, and requested that all discussion take place on the TLS mailing list, including, but not limited to, any discussions of this matter among IESG members, IAB members, agents of IETF Administration LLC, etc.

There are two ADs. They didn't recuse themselves, despite my pointing out conflicts of interest. Also, despite BCP 9 clearly stating "The Area Director(s) shall attempt to resolve the dispute", it seems that one of the ADs ignored my complaint, on the grounds of the other AD being "responsible" for the TLS WG. I have received responses only from that other AD, namely Wouters.

That AD sent email dated 12 Jun 2025 15:58:44 -0400 under a different subject line ("AD response to message on WG chair consensus call draft-connolly-tls- mlkem-key-agreement by D. J. Bernstein of 2025-06-05"), also not following IETF's standards for marking the message as a reply. The AD's email refused, for a variety of reasons (covered below), to "get to the content of the complaint (aka appeal)".

Despite the AD's mislabeling of the email, I did see the email. I followed up by email dated 14 Jun 2025 01:15:38 -0000 (switching back to the original subject line while following IETF's standards for marking my message as a reply). I responded point by point to the AD's refusal to address the complaint, and I asked the AD to "please go ahead with answering the contents".

The AD did not respond.

The third BCP 9 paragraph has a prerequisite that "the disagreement cannot be resolved by the Area Director(s)". Is this triggered by a non-responsive AD? Or would IESG be able to evade addressing the content of a complaint by saying "There's no proof that the diagreement cannot be resolved by the ADs; go back to them"?

Further time passed, still with no response from the AD. I filed an appeal https://cr.yp.to/2025/20250812-non-hybrid.pdf with IESG shortly before the two-month deadline for appeals under BCP 9.

IESG did not respond beyond acknowledging receipt. More time passed.

On 1 October 2025, without addressing the actual content of my complaint, IESG wrote that it "directs the appellant to file a valid complaint to the SEC ADs for consideration". IESG rejected one of the AD's excuses for not handling my complaint in the first place but supported another one, saying that a particular paragraph in my complaint made the complaint "invalid". I disagree but, in order to move things forward, decided to remove that paragraph.

On 6 October 2025, I filed a revised complaint https://cr.yp.to/2025/20251006-non-hybrid.pdf.

On 7 October 2025, the AD refused to process my complaint, claiming that "appeals must be in a specific format and this appeal does not conform to that so it will not be processed". The AD cited a document https://datatracker.ietf.org/doc/statement-iesg-statement-on-the-conflict-resolution-and-appeals-processes/. I asked various followup questions on 9 October 2025:

> Let me make sure I understand. You're refusing to obey RFC 2026's "shall attempt to resolve the dispute" requirement, and your reason is that I didn't use a "specific format" described in some IESG statement?
>
> Did IESG announce its statement for IETF consideration? When? Where? I certainly hadn't seen any such announcement when I prepared and filed my complaint.
>
> Maybe I missed an announcement, but if this is an _ex post facto_ rule ("now that you've done the work to file a revised complaint, we'll tell you about new

> rules that we just posted, and retroactively throw your complaint away on this basis, ha ha ha") then it's glaringly unethical.
>
> Beyond timeline questions, why exactly do you believe that you and/or the full IESG have authority to make exceptions to RFC 2026's "shall attempt to resolve the dispute" requirement?
>
> Also, you objected to my email normatively citing what you called "a remotely hosted PDF", but then you didn't answer my followup question: "For the record, is draft-ietf-tls-mlkem going to be banned because it normatively cites FIPS 203, which is a remotely hosted PDF?"
>
> I filed the actual content of my complaint more than four months ago. With all due respect: It looks terrible for an AD and IESG to have been making up retroactive, ad-hoc, selectively enforced excuses to not address the content of the complaint.

There was still no reply from the AD by 13 October 2025. I took the time to review https://datatracker.ietf.org/doc/statement-iesg-statement-on-the-conflict-resolution-and-appeals-processes/. That statement is dated 1 October 2025. As far as I can tell from various searches, the statement was not announced anywhere until after I had filed my 6 October 2025 complaint. The statement has 1314 words and imposes a variety of *content* requirements, not merely formatting requirements.

To move things forward, I revised my complaint to comply with all of IESG's latest demands, and filed the revised complaint by email on 13 October 2025. My understanding was and is that the AD and IESG would have refused to consider any further complaints regarding this matter if I did not file that complaint by 15 October 2025. (The AD rejected an appeal from someone else regarding the same matter on the grounds of the appeal being filed too late.)

The ietf.org mail system accepted my message (receipt 7168872B5B8C for tls@ietf.org, receipt 7458772B5B8E for sec-ads@ietf.org). However, the message did not promptly appear on tls@ietf.org.

On 14 October 2025, I posted my complaint at https://web.archive.org/web/20251014135826/https://cr.yp.to/2025/20251014-non-hybrid.md and sent another message linking to that. Formally, the normative link meant that this did not comply with IESG's demands, but IESG cannot demand that appeals be filed through a non-functional mail system.

The original message was delivered after an 18-hour delay inside mail2.ietf.org (according to the "Received" lines in the message). Presumably the delivery happened only because someone had taken manual action to allow it.

The same AD refused to handle my complaint. I sent simple yes/no clarification questions by email dated 15 Oct 2025 00:05:25 -0000; the AD did not reply.

I filed a complaint with IESG by email dated 15 Oct 2025 19:21:15 -0000 regarding the ADs refusing to handle my complaint regarding the chair consensus claim.

IESG sent email dated 01 Nov 2025 03:05:04 -0700 linking to a "response" to my 15 October 2025 complaint. In that "response", IESG refused to address the procedural violations that I was complaining about.

However, in a step away from its previous dodging, IESG wrote that it had "determined that a specific, unresolved question remains implicit" in my complaints. IESG's formulation of the question was as follows: "Was rough consensus to adopt draft-connolly-tls-mlkem-key-agreement in the TLS Working Group appropriately called by the WG chairs?" IESG directed the same AD to "directly investigate and provide an answer to the above question".

The AD sent email mere hours later (email dated 1 Nov 2025 15:40:57 -0400) with the following conclusion: "I agree with the TLS WG Chairs that the Adoption Call result was that there was rough consensus to adopt the document". The AD's message had three major gaps:

- The AD did not address my complaint about the chair action on this topic. Structurally, the AD didn't quote and answer the points in my complaint; the AD instead placed the entire burden on the reader to try to figure out, for each point in my complaint, what in the AD's text might have supposedly been answering that point, and to realize that many points remain unanswered. Of course, IESG as a whole is at fault for instructing the AD to answer the question formulated by IESG rather than to try to resolve my complaint.
- The AD did not address the chairs claiming that "we have consensus to adopt this draft", or the previous claim from the AD that there was "clearly consensus". Instead IESG and the AD quietly shifted from "consensus" to "rough consensus". RFC 2418 says that "rough consensus" suffices; however, BCP 9 says that "consensus" is required. IETF even claims that "decision-making requires achieving broad consensus"; "broad consensus" is even stronger than "consensus", since it's saying that there's consensus in a broad group. Each IESG-approved RFC claims to have "consensus of the IETF community". You can't have it both ways, allowing an admittedly non-consensual action on the basis of "rough consensus" but then claiming that the result was "consensus".
- The AD's argument was not structured as saying (1) here's the definition of "consensus" (or "rough consensus" if we're switching to that) and (2) let's now apply the definition to the situation at hand. Nobody reading the AD's message can figure out what the AD believes "consensus" (or "rough consensus") means. Not being clear about the procedural rules being applied is unacceptable: it means that IETF management has the power to make arbitrary decisions. Such power invites corruption.

Section 17 below comments point by point on what the AD did write.

A few days later (email dated 05 Nov 2025 10:51:13 -0800), the TLS WG chairs announced a "last call" setting a deadline of 26 November 2025 for objections to "draft-ietf-tls-mlkem-05", as if unresolved objections had not already been raised. Since there was no consensus to adopt this "ietf-tls" document in the first place, the chairs were not permitted to issue "last call" for the document.

People who don't understand the importance of following proper procedures might claim that there is no need to resolve a complaint about improper adoption if a subsequent "last call" shows consensus for the document. But this is improperly shifting burdens to dissenters to repeat their objections.

Furthermore, even more people objected to this document during "last call" than during adoption; fewer people supported the document during "last call" than during adoption; and the central objections to the document remain unresolved. The chairs did not report these facts, but the chairs did eventually admit (email dated 7 Dec 2025 20:39:00 -0800) that "we do not have consensus to publish the document as is", so it seems unnecessary to review the details here. Unfortunately, it was also clear that the chairs were planning to issue another "last call", once again improperly shifting burdens to dissenters regarding this document.

On 23 December 2025, I filed a complaint https://web.archive.org/web/20251223163731/https://cr.yp.to/2025/20251223-non-hybrid.md with IESG. In the email to IESG providing this link, I wrote "The IETF mail system does not allow large messages through by default, so I am using a link rather than placing the same information inline. Please acknowledge receipt, and please confirm that you will 'attempt to resolve' the situation as required by BCP 9." I cc'ed the TLS mailing list for transparency, but this message was censored by the TLS chairs and never appeared on the TLS mailing list.

By email dated 15 Jan 2026 10:41:28 -0800, IESG finally acknowledged that it had received my December 2025 complaint. Unfortunately, in violation of BCP 9, IESG refused to process the complaint. Concretely, IESG quoted its prohibition of normative "URLs to non-IETF websites"; I fully agree that my link to https://web.archive.org/web/20251223163731/https://cr.yp.to/2025/20251223-non-hybrid.md was a URL to a non-IETF website, but I don't agree that IESG has any authority to disregard the complaint on that basis.

I sent another message, this time including the complaint as inline text. IESG received that too. I again cc'ed the TLS mailing list for transparency, but this was another message that was censored by the TLS chairs and that never appeared on the TLS mailing list.

IESG stalled for two more months. During that period, the chairs issued a second "last call" for objections to

publishing the document as an RFC. By my count, the second "last call" produced 22 statements opposing publication and 21 statements supporting publication. See https://blog.cr.yp.to/20260405-votes.html for names, quotes, and links for verification.

There's a giant gap between (1) WG documents claiming to have "consensus" and (2) the latest "last call" not even managing to produce *majority* support for this document. So it might seem unnecessary to resolve the dispute regarding the original claim of consensus on adoption. But the chairs haven't removed this document from the WG pile. The chairs have never admitted the level of opposition to this document. The chairs are pretending that there are merely some requests for tweaks to the document. The chairs have already announced plans for a third "last call"—which again improperly shifts burdens to dissenters. These "last calls" are predicated on the fiction that the document was properly adopted.

IESG sent email dated 13 Mar 2026 08:30:58 -0700 linking to a "response" to my December 2025 complaint. This "response" has the same central flaws as what the AD had already written; the details are covered in Section 19 below. Clearly this triggers the BCP 9 condition that "the disagreement is not resolved to the satisfaction of the parties at the IESG level". Also, my complaint to IAB easily meets the two-month deadline.

3 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.

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 sabotaging security.

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". See https://blog.cr.yp.to/20250930-stealth.html.

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".

4 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. Almost the entire cost of ECC+PQ is the PQ communication, not the ECC communication or computation. 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 5.)

In November 2025, the TLS WG asked IESG to publish an ECC+PQ document as Proposed Standard. The same AD sat on that document until January 2026 for no apparent reason; IESG approved publication in February 2026. (I am concerned about this specific ECC+PQ document's use of a patented PQ algorithm, namely Kyber/ML-KEM, but this concern is orthogonal to the topic at hand. The document for which I am complaining about a fake claim of consensus uses the same patented PQ algorithm.)

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

draft-connolly-tls-mlkem-key-agreement was first posted in March 2024. Of course someone asked "what the motivation is for being 'fully post-quantum' rather than hybrid". The draft author responded: "FIPS / CNSA 2.0 compliance guidelines ... currently are a big 'maybe' at best for 'hybrid solutions', and the timetables for compliant browsers, servers, and services are to exclusively use FIPS 203 at level V (ML-KEM-1024) by 2033. I figure there will be demand for pure ML-KEM key agreement, not hybrid (with no questions that come along with it of whether it's actually allowed or not)." https://mailarchive.ietf.org/arch/msg/tls/qFRxBSnEPJcdlt7MO0cIL2kW5qc/

In June 2024, NSA's William Layton wrote that "we do not anticipate supporting hybrid in NSS". https://mailarchive.ietf.org/arch/msg/tls/ESCdYNwVeF4VkvoORFJLJk_87VU/

In December 2024, a Cisco employee wrote 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." https://mailarchive.ietf.org/arch/msg/tls/S9Mwv28VEHrG189ZtoubUani7J8/

In June 2025, NSA's Mike Jenkins posted the following: "As the CNSA 2.0 profiles should make clear, we are looking for products that support /standalone/ ML-DSA-87 and /standalone/ ML-KEM-1024. If there is one vendor that produces one product that complies, then that is the product that goes on the compliance list and is approved for use. Our interactions with vendors suggests that this won't be a problem in most cases." https://mailarchive.ietf.org/arch/msg/spasm/xUKIoHQwm1BjNZWS2x3xb-BhsLI/

These quotes show NSA employees requesting non-hybrids, and show some companies complying, in particular leading to the draft at issue in this complaint, a specification of non-hybrid ML-KEM in TLS. Obviously NSA does not need IETF endorsement for NSA's purchases; what is important is the influence of IETF endorsement on what the rest of the world does.

6 Survey of objections to the draft during the adoption-call period

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 7).

(The quotes here focus on objections raised during the adoption-call period. See https://blog.cr.yp.to/20260221-structure.html for a newer chart of arguments and counterarguments.)

6.1 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."

6.2 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')."

6.3 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".

6.4 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".

6.5 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.' "

6.6 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.

7 Lack of consensus

https://www.ldoceonline.com/dictionary/consensus says "consensus" is "an opinion that everyone in a group agrees with or accepts". There are ample resources available such as https://www.seedsforchange.org.uk/shortconsensus explaining the process and value of building consensus, while emphasizing that consensus means unanimous acceptance.

With this concept of consensus, obviously the chairs erred in claiming consensus. But this is not the end of the analysis. The word "consensus" can be used in less stringent ways: for example, https://www.merriam-webster.com/dictionary/consensus says "consensus" can be "general agreement : unanimity" or "the judgment arrived at by most of those concerned".

Legitimate standards-development organizations need clear, well-documented procedures to protect against errors and abuse. These organizations converged many years ago on a concept of "consensus" that *can* allow non-unanimous standards but that still imposes important procedural constraints to protect the interests of minorities. The central requirements are as follows:

- 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.

For example, the ISO/IEC Directives say "Committees are required to respond to all comments received"; define "consensus" as "General agreement, characterized by the absence of sustained opposition to substantial issues by any important part of the concerned interests and by a process that involves seeking to take into account the views of all parties concerned and to reconcile any conflicting arguments"; and say "Consensus, which requires the resolution of substantial objections, is an essential procedural principle and a necessary condition for the preparation of International Standards that will be accepted and widely used". Similar comments apply to ANSI, ASME, etc.

The point of this section is that *none* of these requirements were met when the TLS WG chairs claimed "consensus" to adopt draft-connolly-tls-mlkem-key-agreement. There was not general agreement. There was not fair consideration of each comment. There was not a process of attempting to resolve each objection, There was not documentation, for each objection, of why that objection was overridden. Fundamentally, consensus evaluation was replaced by a majority-voting process. What happened here also violates a variety of promises specifically from IETF; see Section 7.6.

7.1 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 Florence Driscoll, 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/a

rch/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.

Even for readers who understand "consensus" as meaning merely "general agreement" or "the judgment arrived at by most of those concerned" without any further constraints, the chairs were communicating false information when they claimed consensus.

7.2 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.

7.3 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.

7.4 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.

7.5 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 support of a majority of the voters, 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.

7.6 What IETF says about consensus

Some people tell me that, beyond saying (1) that the chairs communicated false information when they claimed "consensus" and (2) that non-consensually ramming this document through standardization processes is improper, I should be saying (3) that what happened here isn't consistent with IETF's promises regarding how IETF works.

A prominent IETF statement says the following: "IETF activities are conducted with extreme transparency, in public forums. Decision-making requires achieving broad consensus via these public processes." Another prominent IETF statement says that WG decisions are *not* taken by voting. BCP 9 lists "consensus" as one of its "requirements" and recognizes "the importance of establishing widespread community consensus". For many years, every RFC produced by IETF WGs has been marked with a notice that the document "represents the consensus of the IETF community". Furthermore, there is ample evidence of IETF participants using the word "consensus" in the strong sense of unanimity.

On the other hand, RFC 2418, part of BCP 25, says that WG decisions require merely "rough consensus" within the WG rather than having "all participants agree". RFC 2418 also says "It is up to the Chair to

determine if rough consensus has been reached". Does this mean that the prominent claims of consensus are false advertising by IETF? Chairs can do whatever they want, simply ignoring objections?

Well, no, RFC 2418 has some perfectly clear rules that the chairs violated here. For example:

- After saying that unanimity is not required, RFC 2418 says that "51% of the working group does not qualify" as "rough consensus". This does *not* say merely 51% of the voters; it says 51% *of the working group*. IETF says that "Anyone can participate by signing up to a working group mailing list", and says that all "official work" of a WG is carried out on the WG list; the adoption call certainly did not produce positive votes from 51% of the TLS WG participants; this rule from RFC 2418 says that even 51% is not enough for "rough consensus". The chairs violated this rule when they claimed consensus.
- As for handling of objections, RFC 2418 says the following: "As much as possible the process is designed so that compromises can be made, and genuine consensus achieved; however, there are times when even the most reasonable and knowledgeable people are unable to agree. To achieve the goals of openness and fairness, such conflicts must be resolved by a process of open review and discussion." People objecting to draft-connolly-tls-mlkem-key-agreement were trying to engage in this resolution process, but people supporting the draft were skipping the process. The chairs violated this "must" from RFC 2418 when they moved ahead without the conflicts having been resolved.

Beyond these clear rules, RFC 2418's examples of "rough consensus" have *vastly* higher levels of agreement than the 22-vs.-7 case at hand. For example, if "100 people in a meeting" reach "consensus" but then "a few people" on the mailing list object, then RFC 2418 says that "the consensus should be seen as being verified"; evidently this means that there is still "rough consensus" so the group goes ahead. As another example, RFC 2418 says that "99% is better than rough".

When RFC 2418 says that it is "up to the Chair to determine if rough consensus has been reached", it's assigning to chairs the clerical responsibility of tracking whether "rough consensus" exists. This does *not* say that "rough consensus" means whatever the chairs want it to mean. Here are two examples of IETF statements promising that chairs have merely clerical power:

- IETF claims in https://web.archive.org/web/20260217212708/https://www.ietf.org/support-us/endowment/ that IETF is "the primary *neutral* standards body because participants cannot exert influence as they could in a pay-to-play organization where members, companies, or governments pay fees to set the direction. IETF standards are reached by rough consensus, allowing the ideas with the strongest technical merit to rise to the surface". Companies can and do purchase "NomCom" appointments, IESG appointments, and chair appointments, by the mechanisms explained in Section 4.1 of RFC 9389; the only way for this to not turn IETF into a pay-for-play organization is for chairs to have purely clerical roles.
- Back in 2014, the IETF subsidiary IRTF refused to remove an NSA employee as co-chair of CFRG. This refusal was by the IRTF chair, who quoted IRTF rules stating that co-chairs "perform the administrative functions of the group", and who concluded that "co-chairs are little more than group secretaries. Their ability to influence the technical work of the group is little different from that of any other group participant". (See https://mailarchive.ietf.org/arch/msg/cfrg/Aqe9HaZQ4JStGeXeWujt6hLS6uU/.) IETF rules, just like IRTF rules, state that co-chairs merely "perform the administrative functions of the group"; see RFC 2418.

If the TLS WG chairs can declare "consensus" on a controversial document, then it's *not* true that IETF is a "neutral standards body", it's *not* true that the chairs are merely performing "administrative functions", it's *not* true that they are "little more than group secretaries", and it's *not* true that their "ability to influence the technical work of the group is little different from that of any other group participant". Several other group participants spoke up during the adoption call to indicate that this draft should be discarded.

7.7 Standards and other specifications

Many of the rules for legitimate standards development and for IETF are motivated by the importance of standards. Does this mean that it's okay to violate those rules if the results aren't labeled as a "standard"?

IETF recognizes and proudly advertises its influence when influence sounds good: for example, IETF claims to be "the premier standards development organization (SDO) for the Internet". But IETF also engages in systematic doublespeak allowing IETF to disclaim influence when influence sounds bad. IETF places the label "standard" on approximately 0% of its standards; IETF instead uses a variety of gentler labels such as "*proposed* standard" (emphasis added) and "request for comments" and "draft" and "informational".

When IETF issues an RFC prominently claiming IETF "consensus", a typical purchasing manager understands this as an IETF-endorsed standard. People responding "the RFC doesn't say it's a standard" or "the RFC says it's informational" or "there's a warning in the RFC saying don't use this" are missing the point. What matters isn't the specific word "standard"; what matters is the *endorsement*, the claim of *consensus*.

Even without an RFC, the mere *adoption* of this document is already being advertised as endorsement. For example, in March 2026, as an argument for supporting "pure" ML-KEM in IEEE's wireless standards, Dan Harkins wrote that the TLS WG "is proposing to define pure PQC ML-KEM ciphersuites for TLS 1.3": https://archive.cr.yp.to/2026-04-06/21:20:16/JKLN7ND3b-wJ0LH9oKZfFcTtNe17yAWArv8z-FcogFo/https/grouper.ieee.org/groups/802/11/email/stds-802-11-tgbt/msg00083.html

In fact, the TLS WG did not make any such proposal. The *chairs* called for adoption, and that call failed to achieve consensus.

8 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". I continued 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.

9 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.

IESG has quoted BCP 9's text "The AD has the authority and the responsibility to assist in making those decisions at the request of the Chair or when circumstances warrant such an intervention" as justifying the AD's disruption here. But that text is only for "matters of working group process and staffing" (which IESG failed to quote); obviously the AD's comments went far beyond those limits, never mind the point that the circumstances warranted the AD *not* interrupting.

9.1 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.

9.2 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 upgrading from ECC 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.

9.3 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 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 5, 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 7.

> 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 5 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 by the end of 2025—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-wpF44OE-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.

9.4 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.

9.5 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 7. 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".

> 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 says that "Participants engage in their individual capacity, not as company representatives". More to the point, the top objection 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. To repeat the actual tallies (see Section 7 for details): during the specified adoption-call period, 20 people expressed unequivocal support, 2 people expressed conditional support, and 7 people unequivocally objected.

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 in April 2025; see Section 12.

> 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.

9.6 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?

[I filed a complaint with IAB regarding the transparency failure here. IAB did not reply to any of the specific points I had made, and did not quote or analyze any of the applicable rules; IAB simply declared that there was "no process violation". I have filed a complaint with the Internet Society Board of Trustees regarding this matter.]

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

The word "voting" accurately describes what actually happened in this incident. See Section 7. 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. That document said that an appeal to IESG was misdirected; it didn't comment on the content of the appeal, let alone taking the position that the AD attributes to the IESG here. Anyway, this text from the AD is again not responding to the topics at hand; it is another ad-hominem attack.

10 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 9, so I won't repeat them here.

11 AD disruption, continued

The AD sent email dated 16 Apr 2025 20:43:21 -0400. This section looks at what the AD wrote.

11.1 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.

11.2 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 14.)

> 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 12. Regarding the conclusion, see Section 7.)

11.3 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 claim, 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 13. 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".

11.4 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/PpVAwrBTuRb5pR6DOC1ipdQuvYc/
>
> 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 7 that consensus requires addressing each objection—and certainly cannot justify further procedural violations.

11.5 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 13.

11.6 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. (Theoretically independent, at least.)

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.

11.7 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?

11.8 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".

11.9 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 7.

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.

11.10 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.

11.11 "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.

12 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 7 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 7. 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.

13 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 7 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.

14 Chairs 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 7. 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 7, shows a variety of reasons that this adoption call failed to reach consensus.

15 Complaining to the ADs

I filed a complaint with the ADs on 5 June 2025. I quoted the authority under BCP 9 to bring this matter "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 explicitly brought it to the attention of both of them.

I also quoted BCP 9 saying that "The Area Director(s) shall attempt to resolve the dispute". I pointed out that, since there are two ADs, this procedure requires *both* of the ADs to attempt to resolve the dispute. (BCP 9 does not authorize ADs, or IESG as a whole, to make any exceptions to this requirement.)

I also noted that I had already, before the 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. Observers with no prior knowledge regarding the AD, simply looking at how the AD has behaved here, will correctly infer that the AD has a conflict of interest. Conflict-of-interest policies protect the organization by requiring recusal when there is even the *appearance* of a conflict.) My complaint was within that scope: i.e., the two ADs should have turned this matter over to two neutral arbiters. I said, however, that my understanding was that the ADs were refusing to do so.

16 AD evasion

The AD who had disrupted the earlier discussion (see Section 9) sent email on 12 June 2025 refusing to handle the contents of my complaint. I sent email on 14 June 2025 addressing the AD's excuses for refusing to handle the complaint. Neither AD replied.

This section goes through the AD's excuses for not handling the complaint. This section is similar to my email on 14 June 2025.

The AD's email started as follows:

> Before we can get to the content of the complaint (aka appeal), I need to clarify some process issues with your message.

I replied as follows: "Happy to discuss. However, the 'need' statement is incorrect (see below), so please go ahead with answering the contents."

> First, the Security Area Directors have divided their work based on Working Groups, with me being the responsible AD for the TLS WG so as per the Security Area workflow decided by the Security Area Directors,

RFC 2026 says "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".

There is a dispute about the mid-April declaration that there was TLS WG consensus to adopt this non-hybrid Kyber draft. I explicitly brought this dispute to the attention of both ADs. Ergo, both ADs "shall attempt to resolve the dispute".

There is nothing in RFC 2026 allowing one of the ADs to shirk this responsibility by declaring that the responsibility falls solely upon the other AD.

This RFC 2026 requirement is not overridden by generic language in RFC 2026 saying that appeals bodies can define their appeals procedures. The requirement is also not overridden by whatever might be said in RFC 3710, which describes IESG: RFC 3710 is an Informational RFC, not an IETF policy.

RFC 2026 has a precondition involving discussion with the WG chairs. In the case at hand, that discussion led to the WG chairs (1) sticking to their claim of consensus and (2) explicitly authorizing an appeal.

I continued my reply as follows: "Structurally, the message that I'm replying to doesn't appear to be arguing that the situation at hand is somehow outside RFC 2026's 'shall attempt to resolve the dispute' requirement. It's raising various side issues—which, again, I'm happy to discuss, but this discussion doesn't remove this RFC 2026 requirement."

> I will be the only Area Director handling your message at this point,

I asked the following clarification questions: "To clarify, you're saying that the other AD won't attempt to resolve this dispute? Surely you wouldn't be able to make such a statement without discussion with the other AD; can you please point to the records of that discussion? (Dates, minutes, email records, etc.)"

There was never any answer to these questions. It's clear that the ADs intentionally violated the record-keeping requirements in BCP 9. (As noted above, this is now on appeal to IAB.)

> which is presumably aimed to be a message under BCP 9 (RFC 2026) Section 6.5.1.

The document explicitly says it's invoking Section 6.5.1 of RFC 2026, and pinpoints the specific paragraph it's invoking. Consequently, the word "presumably" is inaccurate.

> Second, I am unable to respond privately

In context, this appears to be saying that Wouters *wants* to reply privately rather than on-list, but is unable to.

I began my reply to this as follows: "You can send private email to djb@math.uic.edu, but you shouldn't: that would be a transparency violation. This is a complaint about an action by TLS WG chairs regarding TLS WG activity, so I requested that all discussion of this complant be cc'ed to the TLS mailing list."

Anyway, "The Area Director(s) shall attempt to resolve the dispute" does not make an exception for ADs who ask for non-transparency.

> The TLS WG list should not be used as a backup for not being able to receive direct emails.

This statement has nothing to do with my complaint. My complaint cited various transparency requirements, and on that basis asked for the discussion of this TLS WG consensus claim to take place on the relevant mailing list, the TLS list.

> Note that as per Section 6.5.1, it is the Working Group Chairs who may decide whether or not to involve the WG as a whole:

No. A quote saying that a WG chair "may" do something doesn't say that other people can't.

> That is to say, WG Chairs may decide a complaint is or isn't suitable for further discussion on the WG list.

RFC 2026 doesn't say that. Furthermore, transparency requires IETF to be public in how it handles complaints about IETF standardization activity.

> By continuing your participation using Contributions, you are agreeing to operate under this Note Well.

See above.

> This raises concerns since there is no guarantee of the permanence of the material behind that link, and the content will not be part of the IETF public mail archive.

https://web.archive.org/web/20250613195523/https://cr.yp.to/2025/20250605-non-hybrid.pdf has a copy, so the archiving concern seems misplaced.

More to the point, "The Area Director(s) shall attempt to resolve the dispute" is not limited to disputes that were brought to AD attention via permanently archived documents.

> The PDF format also discourages participation

Such a broad general statement is certainly not true. For example, other formats are more likely than PDFs to be displayed in mangled form; when that happens, those formats are discouraging participation.

I'm not saying PDF is always better than other formats. The evaluation depends on many factors, such as the type of material being presented. I often post PDFs, but I often post information in other formats. PDF is also one of the formats that IETF habitually uses, although certainly not the only one.

Anyway, "The Area Director(s) shall attempt to resolve the dispute" does not make an exception for disputes documented as PDFs.

```
as the content cannot be easily replied to preserving context via email threads
on the TLS WG mailing list.
```

The AD's 12 June 2025 email is replying to my 5 June 2025 email—but for some reason *doesn't* use the standard threading header fields to show this. So it's puzzling to see the message also recognizing that threading is a positive feature.

A reply can and should be posted under the original subject line, in the same thread, for easy tracking, even if the reply is simply a link to another document.

It seems that the AD's mail software doesn't make it as easy for the AD to quote a PDF as to quote something else, but typing time is a very small part of normal procedures for dispute resolution.

```
Thus, an email to the TLS WG mailing list consisting of a (link to) PDF does not
"involve the Working Group as a whole in the discussion".
```

It provides the requisite transparency. It provides other people an opportunity to comment if they would like to.

```
This is not a valid use of the TLS WG mailing list.
```

My complaint is on topic for the TLS WG mailing list. See above.

The AD's message was unnecessarily "meta": in particular, it explicitly avoided the actual content of the complaint. But that doesn't make my messages off-topic, nor would it have justified hiding the AD's message from interested TLS WG participants. Furthermore, since the AD's message was portrayed as a step towards addressing the content, my reply to that message—trying to move things forward—was also on-topic. Interested parties have a right to see what's going on, and the TLS mailing list is the obvious place for that.

```
Furthermore, previous efforts of converting your remotely hosted PDFs to a local
location within the IETF archives for the community by the IESG
```

What PDFs is this referring to?

My guess here was and is that the AD was talking specifically about one particular earlier PDF, namely https://cr.yp.to/2025/20250501-bcp-79.pdf. IESG did *not* simply repost a copy of that PDF. In particular, that document has a central diagram with two main items and various subitems; IESG posted a modified version of the document that destroyed this structure, instead showing a dozen main items in the diagram.

Eventually I noticed this and complained, and it was fixed—but what other changes did IESG make? The bigger problem, before and after the fix, is that IESG placed a burden on the reader (and on me!) to (1) realize that something changed and (2) figure out what changed.

The issue here isn't permanence. The issue is IESG turning a simple situation of a single document into an unnecessarily complicated situation of (1) an original document and (2) an unauthorized IESG-modified version of the document. There has never been a statement of why IESG modified the document in the first place.

```
have resulted in claims on your end of "gross misrepresentation"
```

It's true that the words "gross misrepresentation" appear in the email that's linked at the end of this paragraph in the AD's message. However, that email wasn't from me. Amazingly, the AD still hasn't issued an erratum here.

One basic part of proper dispute resolution is simply tracking who said what. One of my reasons to ask for multiple people to handle this complaint is that this reduces the risk of error.

> merely for containing formatting changes:

No, the changes were not just "formatting changes": IESG changed a central aspect of what the document communicates to readers.

> https://mailarchive.ietf.org/arch/msg/admin-discuss/y6eNBaogfeCZ2oEVyVmdrEPrjJg/

Having this link is what made the AD's comment about the linked message particularly easy to debunk. (This is a link to a message from Jay Daley claiming that something I had written was a "gross misrepresentation".) My reply was as follows: "Thank you for the link."

> Your remotely hosted PDF furthermore contains text disallowing the content to be reformatted and thus quoted.

See above.

> This prevents me and others from discussing the content.

No, it does not. For example, copyright has a "fair use" exception, which I'm comfortably relying upon here in giving many quotes and commenting individually on each quote, allowing readers to easily track which comment is in reply to which quote.

> Additionally, as you did not mail the PDF to any IETF mailing list, there is now a question as to whether this link to a remotely hosted PDF is considered a Contribution under the IETF Note Well.

Again: "The Area Director(s) shall attempt to resolve the dispute" does not make an exception for disputes documented as PDFs.

> If it is not a Contribution, it cannot be a complaint (appeal) under RFC 2026 Section 6.5.1.

That section says no such thing. Regarding the specific complaints about PDF etc., see above.

> If it is a Contribution, it cannot come with its own stipulated restrictions.

See above.

> The text, which you did not consent me to quote, does not comply with your agreed participation under the Note Well.

"The Area Director(s) shall attempt to resolve the dispute" does not make an exception for ADs issuing complaints about alleged violations of other procedures.

> BCP 9 (RFC 2026) Section 10.2 states: 10.2 Confidentiality Obligations No contribution that is subject to any requirement of confidentiality or any restriction on its dissemination may be considered in any part of the Internet Standards Process

My complaint is not confidential. Creating derived works, as IESG did with a previous PDF, is modification, not dissemination.

> This text states that due to your dissemination modifier, your PDF file cannot be considered a Contribution in any part of the Standards Process.

No, the text certainly doesn't state this, nor does it imply it, nor do any of these references to confidentiality provisions create exceptions to "The Area Director(s) shall attempt to resolve the dispute". I'll stop commenting on the "Contribution"-related claims after this.

> This is, however, further updated by RFC3978 Section 3.2 which states: 3.2. Confidentiality Obligations

Again, this is not a confidential document.

I asked the AD two questions at this point: "Are your discussions with other IETF LLC agents regarding this matter confidential? Where are the records of those discussions?" There were never any answers to these questions.

> However, it is also possible to argue that remotely linked content has not in fact been submitted under the IETF Note Well and are thus not Contributions as per RFC3978. In that case, you have not actually submitted a complaint under RFC 2026 Section 6.5.1 either.

Again, RFC 2026 says "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 doesn't allow ADs to put limitations on the mechanism of bringing the dispute to their attention.

> Either way, this prevents me from processing your complaint under the Internet Standards Process.

No, it doesn't. See above.

> Fourth, your email to the TLS WG list (thus per definition a Contribution) contains additional instructions that you are attempting to force

I replied by quoting "IETF participants have impersonal discussions."

> onto the Internet Standards Process that are not codified in any RFCs:

At this point the AD continued by quoting my request for transparency: "For transparency, please carry out all discussion of this matter on the relevant public mailing list (tls@ietf.org), including, but not limited to, any discussions of this matter among IESG members, IAB members, agents of IETF Administration LLC, etc."

The AD is claiming that my transparency request lacks authority from RFCs. Most readers will understand this as claiming that my transparency request lacks authority from IETF. Both parts of this are incorrect.

IETF LLC says that IETF operates with "extreme transparency". RFC 2026 says that "Each of the organizations involved in the development and approval of Internet Standards shall publicly announce, and shall maintain a publicly accessible record of, every activity in which it engages, to the extent that the activity represents the prosecution of any part of the Internet Standards Process".

More specific language in RFC 2026 requires records of "complete and accurate minutes of meetings" (my complaint here is that IESG members are violating the word "complete") and records of "all written contributions from participants that pertain to the organization's standards-related activity" (my complaint here is that IESG members are hiding their email on this topic).

These RFC 2026 requirements, like another requirement that I commented upon above, are not overridden by generic language in RFC 2026 saying that appeals bodies can define their appeals procedures. They are also not overridden by whatever might be said in RFC 3710, which describes IESG: RFC 3710 is an Informational RFC, not an IETF policy.

For discussion of TLS WG activity in particular, the obvious place for these records is the TLS mailing list. Burying the records in some place that's theoretically public but hard to find is not a reasonable interpretation of "publicly accessible" and is certainly not "extreme transparency".

> You have been notified that this is inappropriate before, by the IETF Executive Director:
> https://mailarchive.ietf.org/arch/msg/admin-discuss/y6eNBaogfeCZ2oEVyVmdrEPrjJg/

I hadn't seen that message before the AD linked to it—looks like it was buried on some obscure mailing list, which, again, is contrary to the point of transparency rules.

Anyway, the contents of the message don't say what the AD claims they're saying, and in any event they're not relevant to "The Area Director(s) shall attempt to resolve the dispute".

> By insisting on including this language, you are misleading participants about their responsibilities and obligations under the Internet Standards Process as set forth by the IETF (again via the Note Well that you voluntarily comply to by sending messages to IETF mailing lists).

See above.

> You are purposefully amplifying this misleading text with a "Thanks in advance" modifier, that further exacerbates the misleading message by giving it a false aura of authority.

I have no idea how putting "Thanks in advance" after "Please ..." is conveying a "false aura of authority".

> This is not appropriate use of the TLS WG mailing list.

I've complained about an erroneous declaration of TLS WG consensus. The discussion of the complaint belongs on the TLS WG mailing list.

> Summary Your message to the TLS WG list on June 5 2025 does not constitute a valid submitted Contribution or complaint under RFC 2026 Section 6.5.1, but you can rectify this by sending a compliant email to the Area Director and/or TLS WG mailing list.

Again: RFC 2026 says "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".

> Your unique personal process

This is another point where I replied by quoting "IETF participants have impersonal discussions."

> negatively affects the IETF Standards Process by confusing participants in what they can and cannot do with the content of your emails, including hyperlinks to off-site material that you sent to the list.

Hyperlinks appear all the time on IETF mailing lists.

> They may further feel prohibited from discussing your email content - in email threads or otherwise - by your disclaimers and stipulations on how to behave.

Vague comments about "stipulations on how to behave" aren't addressing what actually happened here. As noted above, I posted a complaint with a central diagram having two main items and various subitems; IESG posted a modified version of the document that destroyed this structure, instead showing a dozen main items. I complained about this. Someone then fixed that particular problem, but there wasn't even an apology, let alone assurance that such incidents won't recur.

> They might also be discouraged from engagement for safety reasons

"Safety reasons"?

I asked what the AD was talking about here. The AD never answered.

My best guess is that the AD has somehow acquired the notion that the dividing line between links that are safe to click on and links that aren't is the dividing line between links ending ".html" and links ending ".pdf".

> and/or because following a discussion via attached PDF files is too cumbersome.

Regarding the oversimplified view of PDF as a bad thing, see above. Regarding "too cumbersome", there are many factors influencing the decisions of individual participants regarding which discussions to follow. Of course, IETF LLC agents in particular are under more stringent obligations, such as "The Area Director(s) shall attempt to resolve the dispute".

> This in turn, might lead to a false sense of consensus in the WG.

I responded as follows: "This unsubstantiated speculation about conceivable future events is not relevant to the complaint at hand, which is about an erroneous consensus call from mid-April 2025. Can you (and the other AD) please attempt to resolve the dispute, as required by RFC 2026?" This was never answered; the ADs continued to violate RFC 2026.

> You are requested by me as Area Director to stop engaging in this behaviour.

I responded with the following questions: "I understand that you're asking me to stop linking to PDFs. To clarify, am I correctly understanding that 'request ... as Area Director' means 'demand', with a threat that you will use your position as AD to impose punishment for non-compliance? What exactly do you believe

gives ADs authority to ban links to PDFs? [Then after a paragraph break:] Note that draft-connolly-tls-mlkem-key-agreement normatively cites NIST's ML-KEM standard—which is a PDF. Is that also banned now?"

There was no answer. PDFs are so widely used for professional communication that the AD's anti-PDF narrative comes across as bizarre.

17 AD promotion of draft-connolly-tls-mlkem-key-agreement

After I objected to the AD evasion described in Section 16, there was further dodging from the AD and from IESG as a whole (see Section 2 for details), but eventually IESG asked the AD to address the consensus question.

The AD's email on 1 November 2025 claimed that "there was rough consensus to adopt the document". The AD did briefly consider the tallies of supporters and opponents, but then left the reader entirely in the dark as to what the AD thinks consensus (or "rough consensus") means. The AD's message failed to compare the facts to a definition of consensus (or "rough consensus"), failed to address the gap between "rough consensus" and the original claim of "consensus", and failed to respond to my complaint. Instead the AD spent most of the email trying to argue that the specification is desirable. This section goes point by point through what the AD wrote.

17.1 Still getting the numbers slightly wrong

> The IESG has requested that I evaluate the WG Adoption call results for ML-KEM
> Post-Quantum Key Agreement for TLS 1.3 (draft-connolly-tls-mlkem-key-agreement).
> Please see below.

As noted above, IESG had instructed the AD to answer the following question: "Was rough consensus to adopt draft-connolly-tls-mlkem-key-agreement in the TLS Working Group appropriately called by the WG chairs?"

> ExecSum
> -------
> I agree with the TLS WG Chairs that the Adoption Call result was that there was
> rough consensus to adopt the document.

As noted above, the TLS WG chairs had claimed "consensus", and the AD had claimed that there was "clearly consensus". The AD and IESG as a whole are quietly shifting to a weaker claim of "rough consensus".

> Timeline
> --------
> April 1: Sean and Joe announce WG Adoption Call [about 40 messages sent in the
> thread]

"About 40"? What happened to the AD previously writing "There is clearly consensus based on the 67 responses to the adoption call"? And why is the number of messages supposed to matter in the first place?

> April 15: Sean announces the Adoption Call passed. [another 50 messages are
> sent in the thread]

Messages after the specified adoption-call deadline can't justify the claim that "the Adoption Call result was that there was rough consensus to adopt the document". The adoption call *failed* to reach consensus. As noted above, subsequent discussion has been even more opposed to the document, but this complaint is about the fake claim of consensus.

> April 18 to today: A chain of (attempted) Appeals by D. J. Bernstein to the AD(s), IESG and IAB, parts of which are still in process.

The fact that the security ADs and the rest of IESG stonewalled in response to complaints doesn't mean that they were "attempted" complaints.

> Outcome
>
> -------
>
> 30 people participated in the consensus call, 23 were in favour of adoption, 6 against and 1 ambivalent (names included at the bottom of this email).

These numbers are *much* closer to reality than the AD previously writing "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".

Also, given that the AD is continually making claims that aren't true (see examples below) and seems generally allergic to providing evidence (the text I'm quoting below has, amazingly, *zero* URLs), it's a relief to see the AD providing names to back up the claimed numbers here.

However, the AD still doesn't get the numbers exactly right. The AD slightly overstates the number in favor while slightly understating the number opposed.

The actual numbers were 20 people unequivocally supporting adoption, 2 people conditionally supporting adoption, and 7 people unequivocally opposing adoption. Clearly 7 is close to 6, and 20+2 is close to 23, but, hmmm, not exactly. Let's check the details:

- How does the AD end up with 6 negative votes rather than 7? By falsely listing Thomas Bellebaum as "ambivalent" and falsely attributing a "prefer not, but okay if we do" position to Bellebaum. In fact, Bellebaum had written "I agree with Stephen on this one and would not support adoption of non-hybrids." (This was in reply to Stephen Farrell, who had written "I'm opposed to adoption, at this time.")
- How does the AD end up with 23 positive votes rather than 22? By falsely listing the document author (Deirdre Connolly) as having stated a pro-adoption position during the call. The AD does not seem familiar with conflict-of-interest concepts and probably doesn't find it obvious that an author *shouldn't* vote, but the simple fact is that the author *didn't* vote. She sent three messages during the call period, all of which were merely commenting on specific points, not casting a vote on the adoption question: https://mailarchive.ietf.org/arch/msg/tls/i7OLt0FkBULQy-ghuPaOp_oUuYM/ https://mailarchive.ietf.org/arch/msg/tls/6ATuzoPpN1_BXMgfbs01OMzQQkE/ https://mailarchive.ietf.org/arch/msg/tls/10-4k8ZmcodjWsyqBCaNZSjcl3s/

The document author didn't object to the AD fudging the numbers. Bellebaum did politely object (https://mailarchive.ietf.org/arch/msg/tls/FmhDavgLrGR2y920WWbM80yMge8/). The AD's response (https://mailarchive.ietf.org/arch/msg/tls/zz5twhe4oIivPxVciHr04omUwMc/) didn't argue, beyond writing "Thanks for the clarification", which mischaracterizes the AD's misrepresentation as being Bellebaum's fault.

More to the point, the AD has never explained whether or how the tallies of positive and negative votes are supposed to be relevant to the "rough consensus" claim.

17.2 Incorrectly arguing for the draft

As context for the following, recall that IETF claims 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".

> In favour argument summary
>
> While there is a lack of substantiating why adoption is desired - which is typical -

This text seems to be admitting that this document flunks the "engineering judgment" criterion. The AD tries to defend this by saying that other documents flunk too.

> the big use case seems to be to support those parties relying on NIST and FIPS for their security requirements.

Wrong. Anything+PQ, and in particular ECC+PQ, complies with NIST's standards when the PQ part does.

Concretely, NIST SP 800-227 says "This publication approves the use of the key combiner (14) for any $t > 1$ if at least one shared secret (i.e., S_j for some j) is generated from the key-establishment methods in SP 800-56A [1] or SP 800-56B [2] or an approved KEM." For example, if the PQ part is ML-KEM as per FIPS 203, then NIST allows ECC+PQ too.

What's next: claiming that using PQ in an Internet protocol would violate NIST standards unless NIST has standardized that particular Internet protocol?

> This encompasses much more than just the US government as other certification bodies and other national governments have come to rely on the outcome of the NIST competition, which was the only public multi-year post-quantum cryptography effort to evaluate the security of proposed new post-quantum algorithms.

Nonsense. The premier multi-year effort by cryptographers to "evaluate the security of proposed new post-quantum algorithms" is the cryptographic literature.

I'm addressing this specific bit of nonsense since it's repeated later. I won't bother addressing the other errors in the AD's argument here about what governments are doing, since the AD's bottom-line claim here is orthogonal to the issue at hand. The TLS WG already has an ECC+PQ document using NIST-approved PQ; the question is whether to also have a document allowing the ECC seatbelt to be removed.

> It was also argued pure PQ has less complexity.

What would be even less complicated is encrypting with the null cipher! Simplicity per se is not the goal; the goal is security.

There was a claim that PQ is less complex than ECC+PQ. There was no response to Andrey Jivsov objecting that having a PQ option makes the ecosystem *more* complicated. The basic error in the PQ-less-complex claim is that it ignores ECC+PQ already being there.

17.3 Incorrectly describing the objections

> Opposed argument summary
>
> Most of the arguments against adoption are focused on the fact that a failsafe is better than no failsafe, irrespective of which post-quantum algorithm is used,

This is the closest that the AD comes to acknowledging the central security argument for ECC+PQ. What is remarkable here is how little emphasis this AD places upon security, despite this being a "security area director" commenting on a "Transport Layer Security" working group.

For comparison, my own objection to adoption started with SIKE as a concrete example of the dangers and continued with "SIKE is not an isolated example: https://cr.yp.to/papers.html#qrcsp shows that 48% of the 69 round-1 submissions to the NIST competition have been broken by now".

> and that the practical costs for hybrids are negligible.

By listing this as part of an "opposed argument summary", is the AD suggesting that this was disputed? When and where was the dispute?

I've seen unquantified NSA/GCHQ fearmongering about costs, but that was outside IETF. If NSA and GCHQ tried the same arguments on a public mailing list then they would end up being faced with questions that they can't answer.

> It was also argued that having an RFC gives too much promotion or sense of approval to a not recommended algorithm.

When I wrote my own summary of the objections, I provided a quote and link for each point. The AD doesn't do this. If the AD is accurately presenting an argument that was raised, why not provide a quote and a link? Is the AD misrepresenting the argument? Making up a strawman? The reader can't tell.

> I have expanded some of the arguments and my interpretation of the weight of these below.

This comment about "weight" is revealing. What we'll see again and again is that the AD is expressing the weight that *he* places on each argument (within the arguments selected and phrased by the AD), i.e., the extent to which *he* is convinced or not convinced by those arguments.

Given that IESG has power under IETF rules to unilaterally block publications approved by WGs, it's unsurprising that the ADs, in their roles as IESG members, will end up evaluating the merits of WG-approved documents. But *that isn't what this AD was instructed to do here*. There isn't a WG-approved document at this point. Instead the AD was instructed to evaluate whether the chairs "appropriately" called "rough consensus" to "adopt" the document. The AD is supposed to be evaluating procedurally what the WG decision-makers did. Instead the AD is putting his thumb on the scale in favor of the document.

17.4 Failing to manage risks

> Non-hybrid as "basic flaw"
>
> The argument by some opponents that non-hybrids are a "basic flaw" seems to miscategorize what a "basic flaw" is. There is currently no known "basic flaw" against MLKEM.

I think that this text is trying to make some sort of claim about ML-KEM not having been attacked, but the wording is so unclear as to be unevaluatable.

Why doesn't KyberSlash count? That was widespread news; best-paper award at CHES 2025; patches to the majority of ML-KEM libraries. And why doesn't Clangover count? And what about the continuing advances in lattice attacks that have already reduced ML-KEM below its claimed security targets?

More importantly, claiming that ML-KEM isn't "known" to have problems is utterly failing to address the point of the ECC seatbelt. Imagine an SDO saying the following: "We're standardizing cars with seatbelts.

Also, recognizing generous funding from the National Morgue Association, we're standardizing cars without seatbelts as another option, ignoring the safety objections. This specific car hasn't crashed, so the absence of seatbelts isn't a basic flaw."

> As was raised, it is rather odd to be arguing we must immediately move to use post-quantum algorithms while at the same time argue these might contain fundamental basic flaws.

Here the AD is reasonably capturing a statement from one document proponent (original wording: "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").

But, back when that statement was issued, I promptly followed up explaining the error (see Section 7.3): "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?"

There was no reply at the time. It's inappropriate for the AD to repeat the erroneous argument without acknowledging and addressing the counterargument.

> As TLS (or IETF) is not phasing out all non-hybrid classics,

"Non-hybrid classics" is weird terminology. Sometimes pre-quantum algorithms (ECC, RSA, etc.) are called "classical", so I guess the claim here is that using just ECC in TLS isn't being phased out. That's a bizarre claim. There are intensive efforts to roll out ECC+PQ in TLS to try to protect against quantum computers. Cloudflare reports the usage of post-quantum cryptography having risen to about 50% of all browsers that it sees (compared to 20% a year earlier); within those connections, 95% use ECC+MLKEM768 and 5% use ECC+Kyber768.

The AD also gives no explanation of why the "not phasing out" claim is supposed to be relevant to the consensus question at hand.

> I find this argument not strong enough

See how the AD is saying the weight that the AD places on each argument (within the arguments selected and phrased by the AD), rather than evaluating whether there was consensus to adopt the document?

> to override the consensus of allowing non-hybrid standards from being defined

Circular argument. There wasn't consensus to adopt the document in the first place.

> especially in light of the strong consensus for marking these as "not recommended".

I think many readers will be baffled by this comment. If something is "not recommended", wouldn't that be an argument *against* standardizing it, rather than an argument *for* standardizing it?

The answer is that "not recommended" doesn't mean what normal people think it means: the AD is resorting to confusing jargon. I don't think there's any point getting into the weeds on this: the AD's claim about a "not recommended" consensus does nothing to support the claim that there was consensus to adopt the document.

17.5 Failing to plan properly for the future

> Non-hybrids are a future end goal
>
> Additionally, since if/when we do end up in an era with a CRQC, we are ultimately designing for a world where the classic components offer less to no value.

If someone is trying to argue for removing ECC, there's a big difference between the plausible scenario of ECC having "less" value and the extreme scenario of ECC having "no" value. It's wrong for the AD to be conflating these possibilities.

As I put it almost two years ago: "Concretely, think about a demo showing that spending a billion dollars on quantum computation can break a thousand X25519 keys. Yikes! We should be aiming for much higher security than that! We don't even want a billion-dollar attack to be able to break *one* key! Users who care about the security of their data will be happy that we deployed post-quantum cryptography. But are the users going to say 'Let's turn off X25519 and make each session a million dollars cheaper to attack'? I'm skeptical. I think users will need to see much cheaper attacks before agreeing that X25519 has negligible security value."

Furthermore, let's think for a moment about the idea that one will eventually want to transition to the specific proposal that the AD is portraying as the future, namely non-hybrid ML-KEM. Here are three ways that this can easily be wrong *even if* one eventually wants non-hybrids:

- Maybe ML-KEM's implementation issues end up convincing the community to shift to a more robust option, analogously to what happened with ECC.
- Maybe the advances in public attacks continue to the point of breaking ML-KEM outright.
- Maybe the cliff stops crumbling and ML-KEM survives, but more efficient options also survive. At this point there are quite a few options more efficient than ML-KEM. (Random example: SMAUG. The current SMAUG software isn't as fast as the ML-KEM software, but this is outweighed by SMAUG using less network traffic than ML-KEM.) Probably some options will be broken, but ML-KEM would have to be remarkably lucky to end up as the most efficient remaining option.

Does this AD think that all of the more efficient options are going to be broken, while ML-KEM won't? Sounds absurdly overconfident. Does the AD even *realize* that there are more efficient options? For anyone thinking "presumably those newer options have received less scrutiny than ML-KEM": we're talking about what to do long-term, remember?

Taking ML-KEM as the PQ component of ECC+PQ is working for getting something rolled out now. Hopefully ML-KEM will turn out to not be a security disaster (or a patent disaster). But, for guessing what will be best to do in 5 or 10 or 15 years, picking ML-KEM is premature.

> When and where to exactly draw the line of still using a classic component safeguard is speculation at best.

Here the AD is *clearly* attacking a strawman.

> Already supporting pure post quantum algorithms now to gain experience

How is rolling out PQ supposed to be gaining experience that isn't gained from the current rollout of ECC+PQ?

Also, I think it's important to call out the word "pure" here (not specifically from the AD) as incoherent, indefensible marketing. A hybrid of ECC and ML-KEM isn't making *any* changes to ML-KEM; it's simply taking the *output* of ML-KEM, the ML-KEM session key, and hashing that together with other inputs. Is ML-KEM no longer "pure" when it's plugged into TLS, which also hashes session keys?

> while not recommending it at this time seems a valid strategy for the future, allowing people and organizations their own timeline of deciding when/if to go from hybrid to pure PQ.

Here we again see *the AD making a decision to support the document*, rather than evaluating whether there was consensus in the WG to adopt the document.

17.6 Misevaluating complexity

> Added complexity of hybrids
>
> There was some discussion on whether or not hybrids add more complexity, and thus add risk, compared to non-hybrids. While arguments were made that proper classic algorithms add only a trivial amount of extra resources, it was also pointed out that there is a cost of implementation, deployment and maintenance.

Here the AD is again making the same mistake explained earlier: ignoring the fact that ECC+PQ is already there, and thus getting the complexity evaluation backwards.

The "thus add risk" logic is also wrong. Again, all of these options are more complex than the null cipher.

> Additionally, the existence of draft-ietf-tls-hybrid-design and the extensive discussions around "chempat" vs "xwing" vs "kitchensink" shows that there is at least some complexity that is added by the hybrid solutions.

No, the details of how to combine ECC with PQ in TLS are already settled and deployed.

Looking beyond TLS: Chempat hashes the transcript (similarly to TLS), making it robust for a wide range of protocols. The other options add fragility by hashing less for the sake of minor cost savings. Each of these options is under 10 lines of code. The AD exaggerates the complexity by mentioning "extensive discussions", and spends much more effort hyping this complexity as a risk than acknowledging the risks of further PQ attacks.

Anyway, it's not as if the presence of this document has eliminated the discussions of ECC+PQ details, nor is there any credible mechanism by which it could do so. Again, the actual choice at hand is whether to have PQ as an option *alongside* ECC+PQ. Adding that option *adds* complexity. The AD is getting the complexity comparison backwards by instead comparing (1) PQ in isolation to (2) ECC+PQ in isolation.

17.7 Misevaluating human factors

> RFCs being interpreted as IETF recommendation
>
> It seems there is disagreement about whether the existence of an RFC itself qualifies as the IETF defacto "recommending" this in the view of IETF outsiders/ implemeners whom do not take into account any IANA registry RECOMMENDED setting or the Mandatory-To-Implement (MTI) reommendations.

I would expect a purchasing manager to have instructions along the lines of "Buy only products complying with these RFCs", and to never see IETF's confusing jumble of further designations. See Section 7.7.

> This is an area where we recently found out there is little consensus on an IETF wide crypto policy statement via an RFC. The decision on whether an RFC adds value to a Code Point should therefor be taken independently of any such notion of how outsiders might interpret the existence of an RFC.

From a security perspective, it's a big mistake to ignore the human factor, such as the impact of a purchasing manager saying "This is the most efficient standard so I'll pick that", or the impact of purchasing managers not understanding the distinctions between RFCs and standards.

> In this case, while Section 3 could be considered informative, I believe Section 4 and Section 5 are useful (normative) content that assists implementers.

Is this supposed to have something to do with the consensus question?

> And people have proposed extending the Security Considerations to more clearly state that this algorithm is not recommended at this point in time. Without an RFC, these recommendations cannot be published by the IETF in a way that implementers would be known to consume.

This "known to consume" claim does not pass the laugh test.

In the real world, even if an implementor does see a "This document is a bad idea" warning, this simply doesn't matter when the implementors are chasing contracts issued by purchasing managers who simply care what's in an RFC and haven't seen the warning.

It's much smarter for the document to (1) eliminate making the proposal that it's warning about and (2) focus, starting in the title, on saying why such proposals are bad. This makes people *more* likely to see the warning, and at the same time it removes the core problem of the bad proposal receiving an endorsement.

17.8 Incorrectly describing country actions

> Say no to Nation State algorithms
>
> The history and birth of MLKEM from Kyber through a competition of the international Cryptographic Community, organized through US NIST can hardly be called or compared to unilateral dictated nation state algorithm selection.

NIST repeatedly refused to designate the "NIST Post-Quantum Cryptography Standardization Process" as a "competition". It even wrote that the process "should not be treated as a competition".

Certainly there were competition-like aspects to the process. I tend to refer to it as a competition. But in the end the selection of algorithms to standardize was made by NIST, with input behind the scenes from NSA.

> There has been no other comparable public effort to gather cryptographers and publicly discuss post-quantum crypto candidates in a multi-years effort.

Once again: Nonsense. The premier multi-year effort by cryptographers to "publicly discuss post-quantum crypto candidates" is the cryptographic literature. See Section 17.2.

> In fact, other nation states are heavily relying on the results produced by this competition.

Here's the objection from Stephen Farrell that the AD isn't quoting or linking to: "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".

After reading this objection, you can see how the AD is sort of responding to it by suggesting that everybody is following NIST (i.e., that the "certainly arise" part is wrong).

But it's not true that everybody is following NIST. NIST's selections are controversial. For example, ISO is considering not just ML-KEM but also

- Classic McEliece, where NIST has said it's waiting for ISO ("After the ISO standardization process has been completed, NIST may consider developing a standard for Classic McEliece based on the ISO standard"), and
- FrodoKEM, which NIST said "will not be considered further for standardization".

ISO is also now considering NTRU, where the advertisement includes "All patents related to NTRU have expired" (very different from the ML-KEM situation).

BSI, which sets cryptographic standards for Germany, recommends not just ML-KEM but also FrodoKEM (which it describes as "more conservative" than ML-KEM) and Classic McEliece ("conservative and very thoroughly analysed"). Meanwhile China has called for submissions of new post-quantum proposals for standardization.

I could keep going, but this is enough evidence to show that Farrell's prediction was correct and that the AD is wrong.

> The use of MLKEM in the IETF will not set a precedent for having to accept other nation state cryptography.

Notice how the AD is dodging Farrell's point. If NSA can pressure the TLS WG into standardizing non-hybrid ML-KEM, why can't China pressure the TLS WG into standardizing something China wants? What criteria will IETF use to answer this question without leaving the WG "open to accusations of favouritism"? The only way to get people to believe that this isn't about the money is to provide a *really* convincing alternative story.

17.9 Denouement

> Not recommending pure PQ right now
>
> There was a strong consensus that pure PQ should not be recommended at this time, which is reflected in the document. There was some discussion on RECOMMENDED N vs D, which is something that can be discussed in the WG during the document's lifecycle before WGLC. It was further argued that adopting and publishing this document gives the WG control over the accompanying warning text, such as Security Considerations, that can reflect the current consensus of not recommending pure MLKEM over hybrid at publication time.

This is just rehashing earlier text, even if the detailed wording is a bit different.

> Conclusion
> ----------
> The pure MLKEM code points exist.

Irrelevant. The question is whether IETF is endorsing them.

> An international market segment that wants to use pure MLKEM exists

"International"? Like Swedish company Ericsson setting up its "Ericsson Federal Technologies Group" in 2024 to receive U.S. military contracts?

> as can be seen by the consensus call outcome

How?

> along with existing implementations of the draft on mainstream devices and software.

Yes, NSA waving around money has convinced some corporations to provide software. How is this supposed to justify the claim that "there was rough consensus to adopt the document"?

> There is a rough consensus to adopt the document

Repeating a claim doesn't make it true.

> with a strong consensus for RECOMMENDED N and not MTI, which is reflected in the draft.

Again, no point in getting into the weeds on this, since the AD's claim here does nothing to support the claim that there was consensus to adopt the document.

> The reasons to not publish MLKEM as an RFC seem more based on personal opinions of risk and trust not shared amongst all participants as facts.

This sort of dismissal might be more convincing if it were coming from someone providing more URLs and fewer easily debunked claims. But it's in any case not addressing the consensus question.

> Based on the above, I believe the WG Chairs made the correct call that there was rough consensus for adopting draft-connolly-tls-mlkem-key-agreement

The chairs claimed that "we have consensus to adopt this draft" (based on claiming that "there were enough people willing to review the draft", never mind the number of objections). That claim is wrong. The call for adoption failed to reach consensus.

The AD claimed that "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". These statements wildly misrepresent what happened, but still haven't been retracted. Again, the actual tallies were were 20 people unequivocally supporting adoption, 2 people conditionally supporting adoption, and 7 people unequivocally opposing adoption.

Without admitting error, the AD has retreated to a claim of "rough consensus". The mishmash of ad-hoc comments from the AD certainly doesn't demonstrate any coherent meaning of "rough consensus".

18 Complaining to IESG

As noted in Section 2, I filed a complaint with IESG on 23 December 2025. That complaint's review of the previous events is similar to the previous sections of this complaint, although I have now added some clarifications and further points.

19 IESG evasion

IESG issued a "response" dated 2026-03-12 to my December 2025 complaint (mislabeled by IESG as January 2026). This "response", like the AD's earlier text, failed to compare the facts to a definition of consensus (or "rough consensus"). This section goes through what IESG wrote.

> The IESG received an appeal from Dan Bernstein on 16 Jan 2026 related to a dispute processing action taken by the responsible AD on the decision of the TLS Working Group (WG) Chairs declaring rough consensus to adopt draft-connolly-tls-mlkem-key-agreement.

The chairs declared "consensus to adopt". My complaint was and is that this was a fake claim of consensus. The conclusion of the declaration is wrong: the call to adopt did not in fact reach consensus. The mechanism that the chairs used to arrive at the declaration is wrong: they looked only at the positive votes, the "interest" in the document, while ignoring objections.

My complaint also pointed out procedural problems with sliding between "consensus" and "rough consensus". For example, WG documents end up claiming "consensus", which would be fraudulent if the underlying procedures allowed something weaker than consensus. IESG's "response" says nothing about this gap.

> This is the third time that the appellant has filed an appeal related to this WG decision. The first appeal response can be found here, and the second appeal response can be found here. The responsible AD's consensus review at IESG request can be found here.

IESG's "third time" text makes the reader think that this was a redundant appeal regarding something IESG had already addressed (even addressed twice). Not true.

IESG's "first appeal response" (1 October 2025) invented a procedural excuse for not addressing the substance of my complaint that the chairs had improperly declared consensus. IESG asked me to "file a valid complaint to the SEC ADs for consideration".

After further AD stonewalling, IESG's "second appeal response" (31 October 2025) *also* refused to address the substance of my complaint that the chairs had improperly declared consensus. IESG instead passed the following question to the AD: "Was rough consensus to adopt draft-connolly-tls-mlkem-key-agreement in the TLS Working Group appropriately called by the WG chairs?"

The AD's "consensus review" was labeled as a response to IESG's question—but in fact addressed neither the concept of consensus nor the content of my complaint. So I complained to IESG. This was the December 2025 complaint, which IESG didn't "respond" to until March 2026.

Yes, this was the third time I was asking IESG to address the substance of my complaint—but that's because IESG had refused to address the substance of my complaint the first two times.

As a separate matter, it's wrong for IESG to say that this complaint is related to a "WG decision". The complaint is about the chairs falsely claiming WG consensus. Non-consensual decisions are not WG decisions.

> The responsible AD, Paul Wouters, did not participate in the processing of this appeal.

Really? It seems that IESG's "response" copies easily debunked misinformation from the AD without making any effort to independently check accuracy. See below.

> The IESG notes that this appeal was submitted with text that is non-conformant to RFC5387. The IESG again provides notice to the appellant that future appeals should not have this non-conformant text.

On the contrary: my text is exercising the provisions in RFC 5387 to opt out of modifications. IESG has no authority to place limits on those provisions.

> This appeal repeats and raises procedural matters on submitting and handling disputes.

Reviewing the history in detail is mandated by BCP 9's requirement of "a detailed and specific description of the facts of the dispute". BCP 9 places some procedural requirements upon complaints, and IESG has been inventing further procedural excuses to avoid handling complaints; omitting a description of the procedural aspects of the history would invite further evasion.

> These include:
>
> On formatting and content of appeals - see IESG Statement on Conflict Resolution and Appeals
>
> On contributions and derivative rights - see the IESG Statement on Clarifying Derivative Works Rights
>
> On transparency of appeal processing - see the IESG response to the first appeal and IAB appeal response on this topic.
>
> On processes that all IETF participants need to know to participate - see IETF Note Well
>
> The IESG declines to address these procedural matters further.

This part of IESG's text is explicitly non-responsive to my complaint. In particular, IESG obviously isn't addressing the consensus question.

> The appellant has raised a question regarding the adoption of "draft-connolly-tls-mlkem-key-agreement" by the TLS WG.

No. This adoption was action by the chairs *without* authorization from the WG. I'm complaining that the chairs issued a fake claim of WG consensus; this is raising a question about a chair action, not a question about a WG action.

> Background:
>
> 1 April 2025 Call for adoption by TLS WG Chair Sean Turner:
> https://mailarchive.ietf.org/arch/msg/tls/PpVAwrBTuRb5pR6DOC1ipdQuvYc/

Yes, that was the adoption call.

> 2 April 2025 TLS WG Chair Deirdre Connolly posted that the IANA setting is Recommended = N
> https://mailarchive.ietf.org/arch/msg/tls/6ATuzoPpN1_BXMgfbs01OMzQQkE/

How is this supposed to be relevant to the question of whether there was consensus to adopt?

Also, side note regarding "TLS WG Chair Deirdre Connolly": Is there evidence that this was an official message in her role as chair? The message doesn't say so (although it also doesn't say the opposite). https://mailarchive.ietf.org/arch/msg/tls/k7_PEqKZulXMtm-D_9pEvfqcxAU/ indicates that Connolly as document author had recused herself from at least some aspects of the chair handling of this spec, but in general it's unnecessarily difficult for readers to figure out which messages are official chair statements and which messages can be appealed. The chairs often fail to label their messages as official or as non-official.

> Many messages were shared on the TLS mailing list; most of them professed willingness to review and contribute to the draft.

IESG provides no evidence to support this claim. Anyone who checks the facts finds that IESG's claim is ludicrously inaccurate. There were exactly 3 people during the adoption-call period stating willingness or potential willingness to review: Uri Blumenthal wrote "I support adoption. Might be able to review"; "Flo D" wrote "I support adoption of this draft and am happy to review"; Scott Fluhrer wrote "I support adoption" and later "I am willing to review". None of the other supporters expressed willingness to review.

More importantly, the question at hand is *not* whether some people were willing to review the document. The question is whether there was consensus to adopt the document.

> 15 April 2025 TLS WG Chair Sean Turner declared rough consensus to adopt: https://mailarchive.ietf.org/arch/msg/tls/_AWy51BSgX1ipv0hfnAzLrDrTYI/

That's when the chairs issued their fake claim of consensus. See above regarding "rough".

> 25 April 2025 TLS WG Chair Sean Turner confirmed the rough consensus to adopt: https://mailarchive.ietf.org/arch/msg/tls/25uD_umU7uOL9LQ1sirFxyA9xmA/

Why does IESG skip the AD's intermediate "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" claim? Why does IESG not admit the actual tallies: 20 supporters, 2 conditional supporters, and 7 opponents, with a variety of unanswered objections? And what does IESG mean by "confirmed the rough consensus"?

The crux of this 25 April 2025 message from Turner was asking whether there were "enough people willing to review the draft". That question simply ignores how many objections there were. This isn't confirming consensus; it's shifting away from the consensus question to something that is *obviously* the wrong question. See Section 14.

> When looking for guidance on working group adoption of work, RFC2418 is silent on the adoption of drafts.

Irrelevant. RFC 2418 has rules covering *all* WG decisions (for example: "Decisions reached during a face-to-face meeting about topics or issues which have not been discussed on the mailing list, or are significantly different from previously arrived mailing list consensus MUST be reviewed on the mailing list"; "Working groups make decisions through a 'rough consensus' process"; "the Working Group Chair and Document Editor positions are filled by different individuals to help ensure that the resulting documents accurately reflect the consensus of the working group"). Adoption decisions are covered by these rules in the same way as other WG decisions. Furthermore, IETF labels all WG documents as "the consensus of the IETF community".

> There is non-normative guidance in Section 2 of RFC7221.

Irrelevant. My complaint is about a fake consensus claim by the chairs, not about violations of RFC 7221.

It seems that, instead of trying to defend the chair claim that there was consensus on adoption, IESG is trying to shift to claiming that the chairs were free to simply dictate the WG's adoption decisions, ignoring IETF's normal consensus requirements. But IESG doesn't even state this claim clearly, let alone provide any justification for it. The fake consensus claim by the chairs should be withdrawn in any case.

> In this specific call for adoption the chairs stated the following: '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.'

Indeed, the chairs were posing (at least) two different questions, one about support for adoption, one about reviewing text. Blumenthal's response "I support adoption. Might be able to review" illustrates that these aren't the same question.

Looking at the full message from the chairs shows much more emphasis on the adoption question—starting with "WG Adoption Call" in the subject line—so it's not surprising that the responses consistently addressed the adoption question.

Anyway, none of this is responsive to my complaint about the fake consensus claim by the chairs.

> The IESG concurs with the chairs and the responsible AD that the draft is within the TLS charter.

No idea where this is coming from. IESG doesn't address any of the arguments that the draft violates the charter, nor do I see where the chairs or AD have addressed those arguments.

More importantly, how is this supposed to be relevant to the consensus question? Consensus requires general agreement, trying to resolve objections, etc. What happened here flunks all of these criteria; see Section 7. Retroactively declaring that one of the objections was wrong does nothing to support the claim that there was consensus.

> The appellant highlights the objections to the document content which were raised during the adoption call.

Well, yes, for reasons that I explained. Declaring consensus in non-unanimous cases requires—among other things—a process of attempting to resolve each objection; see Section 7. Various objections to adopting this document were simply ignored, as my complaint reviewed in detail.

> However, as Section 2.2 of RFC 7221 highlights, "adopting a document does not automatically mean that the working group has agreed to all of its content."

Indeed, the people who objected to adoption were saying not merely that they disagreed with something in the document, but that they didn't want this document to be adopted by the WG.

How is this basic observation supposed to be responding to my complaint? I'm unable to figure out what IESG thinks it's contradicting by writing "However".

> The question being asked at adoption is whether the document represents a topic on which the working group wishes to work.

This is again non-responsive to my complaint.

> While many technical concerns were raised about the approach taken in that version of the document, even some participants who disagreed with the content expressed a preference that the WG have a hand in publishing the document so as to ensure appropriate context would be added.

IESG makes this statement without explaining why the statement is supposed to be relevant. Let me spell out (1) the relevant point that this statement in context insinuates to readers and (2) why this insinuation is a falsehood originating from the AD.

The full history of the adoption call and the two failed publication calls shows more and more people objecting to this document, most importantly because using non-hybrid PQ is a step backwards from using hybrid

ECC+PQ. Hybrid ECC+PQ reduces security risks (including mathematical risks and implementation risks) while adding negligible extra cost.

In response, some people who were already on record as document proponents have given the following argument that the objections are a reason for the WG to *publish* the document rather than to *reject* the document: *publishing* the document lets the WG insert a warning along the lines of "Don't use this; use hybrids instead".

Is IESG *endorsing* this argument? The argument obviously has a fatal flaw. The only way for the WG to insert a "Don't use this" warning would be to have consensus on the warning (along with the rest of the document)—but if the WG has consensus on "Don't use this" then it can issue that statement as a separate RFC, just like previous RFCs deprecating various other problematic cryptographic choices. Publishing a new problematic RFC along with a "Don't use this" warning is strictly worse than publishing the warning and rejecting the RFC.

Readers who understand this won't expect IESG to be endorsing this argument, nor is an endorsement clear from the IESG text, nor would endorsement do anything to justify the fake consensus claim by the chairs.

Is IESG merely pointing out that document proponents have issued a flawed argument? This would make more sense, but the IESG text doesn't clearly say this, and this still would have no evident relevance to the fake consensus claim by the chairs.

Is IESG instead saying that some of the seven people I listed as *objecting to adoption* were actually *supporting* adoption ("expressed a preference that the WG have a hand in publishing the document") while merely objecting to specific document contents ("disagreed with the content")? Aha: *this* statement is obviously relevant to the question of whether there was consensus to adopt. This is saying that there weren't so many people objecting to adoption.

This is just an insinuation by IESG, not something that IESG is stating clearly, but it's what typical readers will gather from what IESG wrote.

Now let's compare this insinuation to the facts. Here are quotes again from the other six people (see Section 7 for names and links): "I oppose adoption"; "I'm opposed to adoption"; "I'm opposed to adoption, at this time"; "I am opposed to the adoption of ML-KEM at this time"; "I agree with Stephen on this one and would not support adoption of non-hybrids"; "I was all set to say that I am in favor of adoption, but Stephen's post changed my mind".

All of these are completely clear in objecting to adoption, not merely objecting to specific document contents. There wasn't general agreement to adopt the document, never mind the other requirements included in the concept of consensus. So what IESG is insinuating is simply false.

Now let me explain why I'm saying that this false insinuation originates from the AD.

After the adoption call ended, the AD posted his claim that there were "67 responses to the adoption call. ... The vast majority was in favour of adoption ... There were a few dissenting opinions". The AD's portrayal of the responses collapsed when Thomas Bellebaum posted the actual tallies. After listing names, Bellebaum wrote "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 quoted "This amounts to about every fourth person objecting the draft in its current state at this time" out of context—omitting the list of names, the tallies, and, most importantly here, the words "opposing adoption". The AD replied as follows, as noted in Section 12: "Note that the consensus call was for Working Group Adoption. Not publishing as is."

Bellebaum's original message is clear and explicit in counting people who had stated opposition to *adopting* the draft at that time. The AD ripped this context away and grossly mischaracterized the statement as being merely about people opposing *publishing* as is.

Could someone opposed to publishing as is be in favor of adoption? Yes, of course. But is that the position of these 7 people? No: these 7 people stated opposition *to adoption*.

Anti-corruption organization Transparency International has a "Complaint Mechanisms" guide https://knowledgehub.transparencycdn.org/kproducts/ti_document_-_guide_complaint_mechanisms_final.pdf requiring that "two staff conduct independent reviews of the complaints" and requiring an "appeal process whereby a review by an independent panel can verify the result of a complaint". An independent reviewer would have checked the facts and would have caught this falsehood from the AD.

I had already pointed out the problem with this AD argument in email dated 13 Oct 2025 20:33:37 -0000 to the ADs and in the December 2025 complaint that I filed with IESG. Specifically, I quoted the AD writing "Note that the consensus call was for Working Group Adoption. Not publishing as is"; I replied as follows: "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.)"

> Adoption is not consensus to publish the document as it is or that the document will be progressed for publication as an RFC. Adoption places the work into the working group's control to develop (i.e., it is no longer 'owned' by the authors) and provides the venue for the work. It is then up to the working group to decide how and whether to progress the document.

This IESG text, like the earlier quote from RFC 7221, is just repeating basic facts about what adoption means—which, in context, feeds into the same falsehood covered above.

> Conclusion
>
> It is the opinion of the IESG that there was sufficient interest in reviewing and contributing to the work to justify adoption by the WG.

This is again dodging the consensus question, a question that forces consideration of objections; it's shifting to a sufficient-interest question, a question that ignores objections. See Section 14.

> Therefore, both the TLS WG chairs and TLS responsible AD decisions are correct. The appeal is thus denied.

The chairs claimed that "we have consensus to adopt this draft as a working group item". Where has IESG defended this claim? IESG never compares a definition of consensus to the facts of the case at hand. The closest that IESG's "response" comes to being on topic is through insinuating, falsely, that there weren't actually 7 people expressing opposition to adoption.

20 Notices

This document may not be modified, and derivative works of it may not be created, and it may not be published except as an Internet-Draft. (That sentence is the official language from IETF's "Legend Instructions" for the situation that "the Contributor does not wish to allow modifications nor to allow publication as an RFC". I'm fine with redistribution of copies of this document; the issue is with modification. Legend language also appears in, e.g., RFC 5831.)