

3GPP · PCG · ITU-T · ISO/IEC · IETF · W3C

Singapore → PCG#56bis → Madrid · post-plenary edition

AI Training-Data Governance in 3GPP: the ULBC Codec, PCG#56bis and the Madrid Decision

The first 3GPP work item to standardise a trained neural model, and the dataset question it forces - the SA#112 referral, the SA WG4 survey and reflector debate, the PCG#56 bis, Ad Hoc and its minutes, the TSG SA#113 plenary week in Madrid and the liaison it did not send, how the ITU-T, ISO/IEC, JET, MPEG, JPEG, the IETF, W3C, IEEE and ETSI govern training data, and the legal, regulatory, product and IPR conditions bearing on a speech corpus - read through to the Release 20 and Release 21 calendars and the stakeholders a normative dataset reaches.

C20004178 · 19 September 2026

Marquee Brief

3GPP is about to standardise a trained neural network for the first time: the ULBC speech codec, whose weights would be part of the specification. That makes the recordings it learns from an upstream input to a 3GPP standard, and puts to the partnership a question no working group can settle. Two bodies were asked in turn and neither answered. **PCG#56bis chose none of three scenarios and created a legal ad hoc group [P06]; at TSG SA#113 in Madrid a draft liaison giving SA4 guidance was rewritten four times in three days and noted, so nothing was sent [M07] [M08].**

New here? A standard is about to contain a trained model, so the data behind it becomes a standards question - passed back twice in three weeks, while the release calendar does not pause (Table 12).

The question in six numbers

ULBC documents at Madrid item 3.4	7
Of those, approved or endorsed	0 of 6
Drafts of the guidance, then noted	4
Companies objecting on the list	5 of 9
Selection Pdocs without a version	4 of 11
Days from Madrid to PCG#57	36

Dispositions are the TSG SA#113 register's own; the list count is of organisations whose messages object to sending the liaison, not of companies holding a view. Bars run to the stated whole.

Source Audits to Anchor Critical Interpretation

Every claim-bearing paragraph and every mixed exhibit opens with one of four marks; a paragraph takes the weakest class it depends on.

SRC Source Record. Traceable to an identified record - a contribution, a meeting report, a PCG input, a reflector message. Wrong only if the record is wrong.

CNT Counted. Computed from the packaged data by a counting rule stated in Part 13. Wrong only if the rule is wrong or misapplied.

INT Interpretation. This brief's reading of the record and the counts, stated outside the quotation, with the alternatives that fit less well beside it.

CTX Context. Law, regulation, markets or other standards bodies' practice that the 3GPP record does not itself establish. Wrong if it does not transfer.

Wording from a record is a "verbatim quotation" - that term throughout - always in green italic with its source, so the record's own words are unmistakable. Interpretation is kept outside it and names the document, reference and live link it rests on.

Confidence is stated where a finding is first made and never rises later: established, well supported, or conditional, each conditional with the observable signal that would test it.

What the record establishes

Where it came from - the referral and SA4's answer

1. One edited word framed SA4's question. Singapore changed 'common training dataset' to 'technical training dataset' before asking SA4 whether requirements are needed. SA4 was never asked whether 3GPP should adopt a common corpus; its answer, no consensus need, answers the edited question (3.2).

2. SA4 split 13 to 9, and one safeguard drew broad support. Of 24 companies, 13 favoured requirements, 9 did not and 2 took no position. Separately, 20 supported keeping test material out of training - survey support, not an agreed rule, and it reaches the plenary only through SP-260919 (section 3.5).

Where it stands - PCG#56bis and the Madrid plenary

3. The three OP inputs agree on the requirement and part on the mechanism. TTA, the legal paper attributed to ETSI and ATIS all require datasets to meet legal requirements naming licensing and privacy; they part on a common corpus, on who verifies and on scope beyond ULBC (Table 23, section 4.5).

4. PCG chose no scenario, and its own minutes now say so. The PCG#56bis draft report records no clear conclusion on the legal analysis, and creates an ad hoc group on the legal aspects led by ETSI reporting to the next PCG meeting [P06]. Six partners spoke; no two hold the same combination of positions (Table 11, section 2.6).

5. Madrid sent no guidance: four drafts in three days, then noted. Seven ULBC documents were taken at item 3.4 and six asked for approval or endorsement; none received either, while the same meeting approved 156 documents elsewhere [G41]. Five companies objected on the public list, on three grounds - consensus, utility and competence (Table 7, Figure 4).

Where it is heading - the layer without precedent

6. Three calendars, and no interval between them. Release 20 targets stage 3 for March 2027 and freezes coding in June; ULBC opens selection testing on 19 March and files specifications in June; Release 21 freezes stage 3 in December 2028 [G47] [G48] [M01]. Four of eleven Pdocs - the selection rules and both test plans - have no agreed version. EVS reached approved specifications in 4.5 years and IVAS in 6.0, against ULBC's planned one: programmes of different scope, printed apart and never as a ratio (Figure 14, Table 62).

7. No other body decides as 3GPP does. None of the 8 bodies read here governs a corpus legally: 3GPP is consensus-first with a 71% vote fallback; JET Chair-assessed, then two parent approvals; MPEG and JPEG expert consensus, then ISO/IEC ballots; the IETF rough consensus (Figure 8, Table 60).

8. Reproducibility is two questions, and model hosting is a third. Implementing a fixed model is conformance: two implementations of the same weights can differ. Regenerating them needs the recipe, not the data. MPEG leaves open which object is normative: structure, weights, or a hosted artefact (6.1, Table 31).

9. The impact is structural. A normative artefact opens a second licensing gate: a corpus is held by copyright, database right and contract, which no FRAND undertaking reaches. The cost falls on fourteen verticals, and those who carry it are not those who decide (Figure 16, Table 72, Table 73).

10. The outside calendar runs ahead of 3GPP's. The revised Product Liability Directive must be transposed by 9 December 2026 and treats software as a product; the German Federal Court of Justice rules on text and data mining on 17 December. Both fall before ULBC selection starts (8.5).

Who reads what

Five readerships, five entry points.

Inside SA and SA4 delegates, PCG and OP staff Sections 1.3, 4.5; Table 2, Table 23: the decision set, and where the inputs meet.

RAN, CT and the other groups delegates and standards researchers Sections 10.4, 10.5; Table 55, Figure 12: where each group already meets a model.

Legal and policy OP counsel, regulators, policy staff Sections 4.3, 8.1; Figure 9, Table 54: the burden scale and the questions for counsel.

Product and operations operators, device, chipset and network vendors Sections 6.4, 9.3; Table 33, Table 49: maintenance, and the evidence a product needs.

Other bodies and research JET, MPEG, JPEG and IETF participants, curators Sections 7.6, 4.7; Figure 8, Table 21: the cross-body pattern and the corpora's terms.

Why this concerns RAN, CT and SA, not only SA4

A model is what it learned from, so every performance figure 3GPP writes about a model is a figure about a dataset. All three TSGs already specify models - RAN measurements, CT data between functions, SA requirements and security - and in each the material stays with operators and vendors, reached by no 3GPP rule.

ULBC is the first 3GPP item in which the **model weights** themselves are normative text, while the **training** corpus that produced them is standardised nowhere - so what is settled here, or left unsettled, is the **precedent** every later group inherits. It reaches a second, closely watched track: **ULBC-MED** is specified for IMS voice over NB-IoT non-terrestrial access through geostationary **satellite NTN** [G01], where the languages, accents and channel conditions the training data covers stop being a governance abstraction and become whether the call connects.

Nothing in the record says a precedent was intended; it is the consequence of being first (Table 55, section 13.6).

What it means for the industry

Four conditional outlooks, read both ways, each with the signal that would move it. None is a prediction.

A requirement before a corpus. Turns on whether the new PCG group drafts requirements before choosing a scenario. **Holds:** a declaration any corpus must satisfy. **Fails:** the common-against-proprietary argument returns at PCG#57. **Signals:** the group's terms.

The question returns to SA4 unanswered. Turns on whether SA4 reopens it now that neither PCG nor the plenary has settled it. **Holds:** SA4 spends selection-rule time on it again. **Fails:** the rules close without it. **Signals:** the SA4#139 agenda.

The margin is in the rules, not the tests. Turns on agreed versions of the qualification and selection Pdocs by February. **Holds:** testing opens on 19 March. **Fails:** the Release 20 coding freeze in June is reached first. **Signals:** the Pdoc versions at SA4#139.

Implementers will ask about the weights. Turns on whether PCG treats weights and data apart. **Holds:** rights questions answered once for every option. **Fails:** each implementer seeks its own evidence. **Signals:** weights in PCG's terms.

Where the consensus stands after Madrid

Five rooms hold this question and none has closed it. **SA4** found no consensus need. **PCG**, where the Organisational Partners carry 3GPP's legal personality, chose none of three scenarios and created a legal ad hoc group. **TSG SA** rewrote a liaison four times and noted it, so the question went back to SA4 as asked. **RAN and CT** stand apart: a rule reaches them as metadata, not custody (Table 3, Figure 2, Table 60).

Nothing was rejected at Madrid and no position was defeated; what did not happen is an instruction. Each step from here is dated: PCG in October, TSG SA#114 in December, SA4#139 in February (Table 68).

Contents

Fourteen parts, numbered part.section. Every heading below is a live link, as is every exhibit in the two indexes beside it.

1 The position at Madrid, in one page 4

1.1 The dataset question at Madrid, in one page	4
1.2 From Fukuoka to Madrid: the chain of records	4
1.3 The decision set in front of TSG SA#113	5
1.4 Two clocks: SA4's selection calendar and PCG's open date	5
1.5 Who governs what: the bodies this brief reads, and why	5
1.6 How the question arrived: three eras, and five years of 3GPP records	6

2 What Madrid decided, and what it left open 9

2.1 Madrid in plain terms: what was asked, and what the plenary did . .	9
2.2 The nine ULBC documents, and what the register records for each	10
2.3 Four drafts in three days: the guidance that was not sent	10
2.4 The argument on the email reflector, message by message	10
2.5 Where the companies stand, across two meetings	10
2.6 PCG#56bis, now on the record: six partners and no scenario	11
2.7 Three calendars: Release 20, Release 21 and the ULBC plan	11

3 Where the question came from: Singapore to SA4#137-e . 15

3.1 What Fukuoka and Singapore settled, and what they passed on . .	15
3.2 One edited word: 'common' becomes 'technical'	15
3.3 The SA WG4 survey: 24 companies, two questions, no consensus	15
3.4 Arguing over the wording: the SA WG4 email reflector, 27-28 August	16
3.5 What the survey found that the report to TSG SA does not carry . .	16

4 PCG#56bis: three Organisational Partner inputs, read in full. 18

4.1 A dedicated agenda item, and who tabled what	18
4.2 TTA: a common public corpus for due diligence and maintenance	19
4.3 The legal paper: three levels of involvement, and the burden of each	19
4.4 ATIS: common legal requirements, verified by the Organisational Partners	19
4.5 Sixteen elements, three inputs: where they agree and where they part	20
4.6 What PCG#56bis concluded, as reported	20
4.7 Cross-checking the inputs: corpora, precedents and claims	20
4.8 Three partners, three institutional views - and four that said nothing	20

5 What SA4 is building, and what Madrid is asked 24

5.1 Eleven Pdocs, twenty declared candidates and a thin margin . . .	24
5.2 The open design issues, and which of them a corpus must serve . .	24
5.3 The sequencing proposal: hold the selection rules for PCG	24
5.4 The transparency proposal: requirement dimensions and maintenance.	24

6 The technical question beneath the legal one 27

6.1 What is standardised when the weights are normative	27
6.2 Fair comparison: four instruments and what each controls	27
6.3 Robustness across languages and conditions: data or requirements	27
6.4 Maintenance against interoperability: who can refrain, and what changes	27
6.5 Regeneration needs more than the data	28

7 How other standards bodies govern training data 29

7.1 JPEG AI: a mandated common set and normative parameters . . .	29
7.2 JVET: common sets in exploration, disclosure in the 2026 call . .	30
7.3 MPEG-AI and neural network coding: common conditions and licence terms	30
7.4 The IETF: separation in mlcodec, preferences in AIPREF, receipts in SCITT	30
7.5 Usage signals and provenance: the SEI message, C2PA, CAWG and TDMRep	31
7.6 The pattern across bodies, and the layer none has taken on	31

8 Law and regulation that bear on a speech corpus 33

8.1 Eight layers of rights in one recording	33
8.2 Copyright and text-and-data mining across jurisdictions	33
8.3 Voice as personal data, as biometric data and as an attribute of the person	34
8.4 The AI Act after the Omnibus: what reaches a codec and what does not	34
8.5 Liability, product safety and cyber resilience	34
8.6 The ETSI IPR Policy, and what it does not reach	35

9 Product, supply chain and assurance 38

9.1 Who carries which exposure along the chain	38
9.2 Documenting a dataset: the instruments that already exist	38
9.3 A ULBC dataset declaration, field by field	38
9.4 Emergency calls, lawful interception and what a decoder must support	38

10 The option space for 3GPP 40

10.1 Six arrangements, read against seven criteria	40
10.2 Who could verify, and what each would need	40
10.3 What an external counsel could be asked	40
10.4 Precedent: a rule for ULBC or for 3GPP's AI work	41
10.5 What 3GPP has already approved about training data	41

11 How the consensus is forming, and who is standing where 45

11.1 How each body decides, and why the same word can mislead . . .	45
11.2 What crosses between bodies: the liaison record	45
11.3 The company lens: seven roles, and what each asks for	46
11.4 The government lens: five questions, eight answers	46
11.5 Why positions sit where they do: six standing conditions	46

12 What the record establishes, and what could follow 50

12.1 What the record already establishes	50
12.2 Conditional readings, each with what would weaken it	51
12.3 Three scenarios for Madrid and PCG#57	51
12.4 Observable signals to watch after Madrid	51
12.5 Eleven seats, eleven readings of the same record	51

13 If the dataset becomes normative: what it would move . . 53

13.1 What would change the moment the weights are normative . . .	53
13.2 Standard-essential patents, and an artefact patents do not reach .	53
13.3 Innovation, research budgets and product strategy	53
13.4 Regulation, policy and the geopolitics of one global corpus . . .	53
13.5 Fourteen seats: interests, hard lines and tradeable ground	54
13.6 Industrial verticals: the first-order effect, and the one after it . . .	54
13.7 Four lenses on the same effect: liability, IP, dependency and position.	54

14 Method, counting rules and limits 58

14.1 Scope and documentary base	58
14.2 How a number is counted, and what the count excludes	58
14.3 The record and the reading: how this brief cites	59
14.4 Boundary conditions and what this brief does not claim	59

Appendix A. Source ledger and verbatim ledger 60

Methodology and interpretation criteria 66

Update triggers: what would prompt a new edition 66

References 66

Glossary, in two layers 69

Index of tables

Table 1. The ULBC documents in the TSG SA#113 export	6
Table 2. The Madrid decision set: three questions, each outcome read both ways	6
Table 3. Dates the record gives, with their precision	7
Table 4. The bodies this brief reads: what each is, and what each cannot decide	8
Table 5. How the question arrived, in date order	9
Table 6. The nine ULBC documents at TSG SA#113, as the meeting's register records them	12
Table 7. The guidance to SA4, draft by draft, and what was said at each . .	12
Table 8. The TSG SA email reflector thread of 16-18 September, message by message: the TSG SA#113 Madrid plenary week	12
Table 9. The same companies across two meetings: SA WG4 in August, TSG SA in September	13
Table 10. What the meeting report records each speaker saying, and on which draft	14
Table 11. PCG#56bis as its own draft report records it: six partners, and no scenario chosen	14
Table 12. Three calendars: Release 20, Release 21 and the ULBC work plan	15
Table 13. The survey tally on both bases	16
Table 14. The codec in the TSG SA exports across three plenaries	17
Table 15. What the two closed plenaries recorded about the codec	17
Table 16. The SA WG4 email reflector thread of 27-28 August, message by message: the SA4#137-e online meeting	18
Table 17. Reasons given on each side, and where each is recorded	18
Table 18. PCG#56 bis_Ad Hoc: the agenda and the documents	21
Table 19. The PCG#56 bis_Ad Hoc register: what the meeting recorded about its own documents	21
Table 20. TTA's input, point by point	21
Table 21. The four corpora TTA names, checked against their 2026 terms	22
Table 22. The legal paper: risks, levels of involvement and the framework for a 3GPP dataset	22
Table 23. Sixteen elements across the three OP inputs	23
Table 24. The three inputs read as institutional positions, and the four partners that tabled nothing	23
Table 25. Statements checked against the record, and how this brief carries each	23
Table 26. ATIS's four statements, and how each can be read	24
Table 27. The eleven Pdocs as the status report prints them	25

Table 28. SA4's open technical issues, and which of them bear on training data	26
Table 29. The two Madrid proposals side by side, and four readings of 'transparent'	26
Table 30. Where each organisation appears in the dataset record	26
Table 31. What a classical and a neural codec specification each contain .	29
Table 32. Four instruments for a fair comparison, and what each controls	29
Table 33. Retraining occasions, and what each needs	29
Table 34. What regenerating a trained model needs, and which input names it.	31
Table 35. Usage signals and provenance instruments	31
Table 36. JPEG AI: the training-data rules, and the standard they produced	32
Table 37. JVET's 2026 joint call: the training-data provisions	32
Table 38. MPEG-AI and related ISO/IEC work	32
Table 39. The training sets in JVET's NNVC common test conditions . . .	33
Table 40. IETF documents bearing on training data	33
Table 41. Copyright and text-and-data mining, by jurisdiction	36
Table 42. What governments are deciding, by mechanism rather than by country	36
Table 43. Voice as personal data, biometric data and an attribute of the person.	37
Table 44. Court decisions on training data and voice, 2024-2026	37
Table 45. The AI Act after the Omnibus, and comparable instruments elsewhere	37
Table 46. Liability, product safety and cyber resilience	37
Table 47. The ETSI IPR Policy against the ULBC artefacts	39
Table 48. Dataset documentation and data-quality instruments	39
Table 49. A ULBC dataset declaration, field by field	39
Table 50. Who carries which exposure along the chain	40
Table 51. Assurance questions a learned codec raises	40
Table 52. Six arrangements read against seven criteria	42
Table 53. Who could verify, and what each could check	43
Table 54. Questions an external counsel could be asked	43
Table 55. Every TSG already specifies models: what each depends on, and whether a ULBC rule would transfer	43
Table 56. Six levels of dependence, and what attaches at each	43
Table 57. The approved 3GPP record that names AI model training data, by group family	44
Table 58. One fixed vocabulary against 1,000 approved records	44
Table 59. What a signal of progress establishes, and what it does not	47
Table 60. How each body decides, whether it votes, and what that means for a rule about training data	48
Table 61. The codec programmes behind ULBC, in two lineages	48
Table 62. Three clocks, kept apart: what each span measures and why it does not transfer	49
Table 63. The liaison record: what crossed between bodies, and when . . .	49
Table 64. EVS and IVAS milestone by milestone, as the record dates them	50
Table 65. Six standing conditions behind each actor's position	51
Table 66. What the record establishes: the first layer	52
Table 67. Conditional readings, with what would weaken each	52
Table 68. Three scenarios, as bundles of the conditional readings	52
Table 69. Observable signals to watch after Madrid	53
Table 70. Standard-essential patents against a normative dataset: what each instrument reaches	55
Table 71. Fourteen seats: what each wants, what each fears, and what each will not trade	56
Table 72. Twenty verticals: the first-order effect, the second-order effect, and whom the record names	56
Table 73. The same verticals under four lenses: liability, intellectual property, dependency and position	58
Table 74. Eleven seats, eleven readings of the same record	59
Table 75. The documentary base, by band and by how each record was read	60
Table 76. Every counting rule, stated once	60
Table 77. What this brief could not read, and what was done instead	60
Table A1. Source ledger: every record this brief reads	61
Table A2. Verbatim ledger: every quotation as recorded	63

Index of figures

Figure 1. One calendar, six tracks: 3GPP, the ITU-T and ISO/IEC coding bodies, MPEG and JPEG, the IETF and W3C, from Fukuoka to TSG SA#116	8
Figure 2. Which body writes rules about what, and where nothing is written	8
Figure 3. Decision proximity: nothing about training data has reached a rule	9
Figure 4. Four drafts in three days: what each revision changed, what was said about it, and where the chain ended	15
Figure 5. The SA WG4 survey: every answer drawn against all 24 respondents	17
Figure 6. Five levels of intervention in training data, and where each input stands	22
Figure 7. From recorded speech to a deployed decoder: the artefacts and what attaches to each	28
Figure 8. Eight bodies, eight provisions: the training-data and model-artefact practices visible in the public record	32
Figure 9. Eight layers of rights in one recording, and which input names each.	36
Figure 10. The regulatory and judicial calendar around the ULBC selection	36
Figure 11. Six arrangements placed by burden on 3GPP and by maintenance independence	42
Figure 12. When training data becomes a 3GPP question: six levels of dependence	42
Figure 13. Five decision architectures: where each family of bodies puts its decisive gate	47
Figure 14. Thirty years of codec standardisation in two lineages, how long each took, and how long 3GPP took after it	50

Figure 15. Who can move the outcome, and who carries it: fourteen seats by influence and exposure55

Figure 16. The impact hierarchy: one decision, five industry groups, twenty verticals, and the effect that follows the first one57

How to use this brief

One question, read in order: the position at Madrid in part 1; what the plenary decided, and what it left open, in part 2; how the question was framed and answered in SA4 in part 3; the three Organisational Partner inputs in part 4; what SA4 is building and what Madrid was asked in part 5; the technical question beneath the legal one in part 6; how other standards bodies govern training data in part 7; the law that bears on a speech corpus in part 8; product, supply chain and assurance in part 9; the option space in part 10; how the consensus is forming and who is standing where in part 11; what the record establishes against what could follow in part 12; and what a normative dataset would move, and for whom, in part 13. The method closes the brief in part 14.

Evidence status: what could be read, and its limits

What the outcomes are read from.

Both meetings are now read from their own records: TSG SA#113 from its register, its draft report and its list thread, and PCG#56bis from its draft minutes, which became available during the Madrid week. Where an earlier edition read the PCG outcome through two company papers, it is a recorded outcome from this edition. The SA4 survey is read from the Chair's report and, for its annex figures, from SP-260919.

What it buys. One rule for every source: a statement is attributed to the document that makes it, so a reader can see which document carries each fact, and weigh a company's report of a meeting as a company's report.

What it does not capture. Anything said at PCG#56bis that neither paper reports; SP-260913, whose text was not read; the new PCG group's terms; and any company's survey answer, since replies were anonymous. Of 68 source records, 29 were read in full (Table 75).

The direction of the error. Both reports of PCG#56bis come from companies arguing for a course of action, so they may state agreement where it helps the argument and omit what does not bear on it. Legal context resting on secondary summaries may lag the primary texts by weeks.

The meetings behind and ahead. Madrid has been held and its record is read here. The dispositions are the register's own; the report is a draft and may be revised before TSG SA#114 approves it. Nothing about SA4#139, PCG#57 or TSG SA#114 is an outcome.

The full statement is in the [methodology](#) and [interpretation criteria](#) at the back of the brief.

Reading guide

Four routes through the same evidence. Each names the sections, exhibits and anchors that carry the answer for that reader.

Delegate route SA, SA4 and PCG delegates

Section 1.3 and Table 2 for the decision set; section 3.5 for the survey figure the report omits; sections 4.5 and 4.6 with Table 23 for PCG; sections 5.3 and 5.4 for the two proposals.

Legal route OP counsel, legal and policy staff

Section 4.3 and Table 22 for the burden scale; Part 7 from section 8.1, with Figure 9; sections 10.2 and 10.3 with Table 54 for who verifies and what to ask.

Product route operators and implementers

Sections 6.1 to 6.5 for the technical question; section 9.1 with Table 50 for exposure along the chain; section 9.3 with Table 49 for the declaration; section 9.4 for emergency and interception.

Cross-body route JVET, MPEG, JPEG and IETF participants, researchers

Part 6 from section 7.1, with Figure 8 for the pattern; section 4.7 and Table 21 for the corpora; section 10.4 for 3GPP's other AI/ML work.

Three things agreed, three not, three not yet asked

Agreed. Test material stays out of training - 20 of 24 in the survey. Datasets must meet legal requirements naming licensing and privacy - all three OP inputs. The weights will be normative - every document.

Not agreed. Whether one common corpus is the way to meet the requirements; who verifies compliance; whether SA4's selection rules should wait for PCG.

Not yet asked. Whether the published weights are personal data; whether rights in the recordings flow into them; what a codec needs, beyond its data, to be regenerated.

Why a codec's training data reaches PCG

A working group decides what a codec does. It cannot decide what the partnership warrants or answers for: that sits with the Organisational Partners. Once trained weights are normative, the data behind them brings copyright, performer, database, contract and personal-data questions into a specification. SA4 said as much: its survey put legal and licensing aspects out of scope before asking anything, and Singapore had already split the question in two.

Reading the two tracks together is not a choice of method: the technical answer is incomplete without the legal one, and the legal one has no object without the technical.

What happens next

Four dates the record points at, all inside 3GPP; the outside dates are in Table 69.

15-18 Sep 2026, Madrid (TSG#113)

TSG SA#113: the SA4 report, the sequencing paper, the transparency paper and a draft liaison. Table 2 carries the decision set.

21 Oct 2026, online

PCG#57: the first record that can show what the new ad hoc group was asked to do.

Nov 2026, Calgary

SA4#138: Pdocs to be finalised, the funding estimate, the qualification decision - the point at which sequencing would bite.

Jun 2027, TSG SA#116

Approval of the ULBC specifications on the current plan, three months after the selection tests launch on 19 March 2027.

Every date is the record's own; none is this brief's forecast.

Cross-references in this brief

Every pointer is a live link, and its colour says what it leads to.

Section 4.5 a numbered section of the argument, part first and section second.

Figure 6 a drawn exhibit, numbered in order of appearance and indexed on this page.

Table 23 a set exhibit, indexed the same way; an appendix table keeps its letter: **Table A2**.

[P03], SP-260912, [L13] a source this brief read, leading to the ledger row or reference that carries it.

Wording taken from a record is set in *green italic* wherever it appears, with the record it came from in the same short form throughout.

The brief in fourteen parts

1. The position at Madrid, in one page 1.1-1.6

For: everyone; read this first

2. What Madrid decided, and what it left open 2.1-2.7

For: everyone; start here if you have not followed the meeting

3. Where the question came from: Singapore to SA4#137-e 3.1-3.5

For: everyone; the baseline for the rest

4. PCG#56bis: three Organisational Partner inputs, read in full 4.1-4.8

For: PCG and OP delegates, legal staff

5. What SA4 is building, and what Madrid is asked 5.1-5.4

For: SA and SA4 delegates, proponents

6. The technical question beneath the legal one 6.1-6.5

For: codec engineers, operators, test laboratories

7. How other standards bodies govern training data 7.1-7.6

For: other standards bodies, researchers, curators

8. Law and regulation that bear on a speech corpus 8.1-8.6

For: OP counsel, regulators, policy staff

9. Product, supply chain and assurance 9.1-9.4

For: implementers, operators, regulators

10. The option space for 3GPP 10.1-10.5

For: PCG, OP legal staff, strategists

11. How the consensus is forming, and who is standing where 11.1-11.5

For: everyone interpreting the record

12. What the record establishes, and what could follow 12.1-12.5

For: everyone acting on the brief

13. If the dataset becomes normative: what it would move 13.1-13.7

For: IPR and licensing readers, research and policy staff

14. Method, counting rules and limits 14.1-14.4

For: anyone checking a number

Where to start, by seat

The guide above routes by what a reader needs to find; this one by what a reader is responsible for. Table 74 carries eleven seats in full, each with what to watch.

Operator Maintenance over a service lifetime first, then the Madrid transparency proposal. 6.4, 5.4.

Implementer The evidence a product needs, and who could supply it under each option. 9.1, 9.3.

OP legal counsel The burden scale, the rights layers and the questions for counsel. 4.3, 8.1, 10.3.

Codec proponent The unwritten rules, the calendar and the fairness instruments. 5.1, 5.3, 6.2.

Regulator Fidelity under the AI Act, and assurance for emergency calls. 8.4, 9.4.

What each body hands the others

SA4 to TSG SA. An answer to two questions - no consensus need for technical requirements, no additional work - and, through a company paper, the annex it kept back.

PCG to TSG SA. No conclusion, a new ad hoc group, and three analyses that between them cover what data, what burden and who verifies.

TSG SA to both. At Madrid, a choice between noting, sequencing and sending requirements - each of which changes what the other body is waiting for.

From a document to a reading, in seven steps

Every interpretation in this brief can be walked back to the document it rests on, and each step names where in the brief it happens.

1. The document A tabled contribution, a PCG input, a reflector message or a report, each a row of Table A1.

2. The words Its wording, quoted only where the reading leans on it, verified and listed in Table A2.

3. The count Every number by a rule stated once - 14.2 and Table 76.

4. The comparison Inputs and bodies read against one list of elements - 4.5, 7.6.

5. The context Law, regulation and other bodies' practice, marked as context wherever it appears - Part 7.

6. The reading Stated with its tier, and never inside a quotation - 12.1.

7. The conditional With its assumption, both outcomes and the signal that would test it - 12.2.

Seven words this brief uses narrowly

Each has a wider meaning in ordinary use; the narrow one is the one this brief relies on, and the two-layer glossary gives the full set.

Common dataset One corpus all candidates train on for the results that are compared; not a corpus 3GPP owns or hosts, unless the brief says so.

Technical requirement A rule SA4 can specify and test - separation, coverage, disclosure format. Never a legal requirement.

Legal requirement A condition on rights and personal data that only lawyers can verify.

Transparent Left undefined by SP-260919; this brief names four readings and chooses none.

Maintenance Changing the standard over its life - correction, extension, continuity.

Reproduction Regenerating the weights from data and recipe; distinct from reproducing a decoder's output, which is conformance.

Conditional A reading that depends on something outside this evidence. It carries its assumption and its signal, and never a probability.

How to read a conditional

Every conditional reading is written in the same five parts, so a reader can disagree with one part without discarding the whole. The form, on one of the nine:

The baseline what the record establishes and nothing more: four selection Pdocs have no agreed version, and SP-260912 asks that they wait for PCG.

The assumption what has to hold outside the record: that TSG SA does not endorse the sequencing proposal.

If it holds SA4 finalises its rules at SA4#138 and the January contracts proceed, with the legal risk carried towards selection.

If it fails finalisation waits for a PCG conclusion, and the margin SP-260907 calls thin becomes PCG's to spend.

The observable signals the disposition of SP-260912 in the TSG SA#113 report, and whether any liaison to SA4 mentions PCG.

No conditional carries a probability, a ranking or a committed date - section 12.2 carries all of them in this form.

What this edition adds

This brief expands two base briefings of 10 and 11 September 2026 and follows the companion plenary brief C20004166. What it adds:

The Madrid documents SP-260907, SP-260909, SP-260912 and SP-260919 read in full and set against the PCG inputs - Part 4.

The public SA4 debate The reflector thread of 27-28 August read message by message - [section 3.4](#).

Verified wording 136 quotations verified against their sources and set apart from the reading - [Table A2](#).

Other bodies' records JVET's joint call, NNVC documents and SEI text read in full; IETF, MPEG and JPEG status checked - Part 6.

Reconciliation Statements in the base briefings checked against the record - [Table 25](#).

Four layers Findings, conditionals, scenarios and dated signals kept apart - Part 11.

How to read this brief: four layers, and the five moves inside each section

How every quotation and figure in this brief is sourced

A quotation is set in green italic with its source in one short form - the document number and the place in it - and every quotation is repeated in [Table A2](#). A bracketed key such as [P03] or [L13] is a live link to the ledger row or the reference that carries the source, and every reference prints its address in full. Every record read for this brief is a public one: 3GPP and PCG documents from the public 3GPP file server, the SA WG4 discussion from the public list, other bodies' documents from their own public pages, and law, regulation and judgments from official sources - the addresses in the references are where a reader can open each of them. No participant in the standards process is named: people are cited by the role the record gives them.

The brief keeps four kinds of statement apart and never lets one pass for another. Each layer has its own place and its own exhibits, and a reader can stop at any layer with everything above it intact.

Layer 1 - what the record establishes

Direct statements of a record and counts from it: the survey tally, the Pdoc register, what each PCG input states. Parts 1 to 4, gathered in [Table 66](#).

Layer 2 - what it means

This brief's reading, marked as interpretation and tiered: the edited word, the combinable inputs, the weights as the common object. Parts 2 to 9.

Layer 3 - what could happen

Conditional readings, each with its assumption, both outcomes and a signal, bundled into three scenarios. [Tables 67](#) and [68](#).

Layer 4 - what to watch

Dated observables, each something that will or will not appear in a record. [Table 69](#).

Inside each section the same five moves recur: the record, its own wording where the reading leans on it, the reading, its tier, and - where the reading looks forward - the signal that would test it.

Before section 1: who decides what, from a codec to a dataset

3GPP, its seven Organisational Partners, and PCG

3GPP is a partnership of seven standards organisations - ARIB, ATIS, CCSA, ETSI, TSDSI, TTA and TTC - known as the Organisational Partners (OPs). Companies take part through one of them, and the OPs, not 3GPP, carry legal personality and publish the specifications. The technical work happens in three Technical Specification Groups (TSGs) and their working groups; the Project Coordination Group (PCG) sits above them and handles the partnership's own affairs. A question about what the partnership may take responsibility for - a dataset, a warranty, a liability - therefore goes to PCG rather than to a working group [G11] [G12].

How a 3GPP speech codec is chosen

Codecs are the business of SA WG4 (SA4) under TSG SA. A codec is selected by competition: SA4 agrees design constraints, performance requirements and test plans, candidate companies submit codecs, independent laboratories paid for by the candidates run qualification and selection listening tests, and the winning candidate's description and reference software become specifications that TSG SA approves. For this codec the rules of the competition are being written as permanent documents, Pdocs, which [section 5.1](#) reads one by one.

What makes this codec different

Earlier 3GPP codecs were defined by their algorithm: a bitstream, a decoding procedure and reference code that anyone could reimplement. A neural codec is defined in part by trained numbers - model weights - that are the result of training on recorded speech. If those weights are standardised, the recordings that shaped them become an upstream input to a 3GPP specification for the first time, with their copyright, performer, database, contract and privacy conditions attached. That is the question this brief follows.

Where the question sits at the cutoff of 11 September 2026

In June 2026 TSG SA split it in two: SA4 was asked whether technical training-dataset requirements are needed, and PCG was to be told about 3GPP's involvement and the legal requirements. SA4 has answered: no consensus. PCG met on 3 September and reached none either, and its own minutes now record that it created an ad hoc group on the legal aspects instead. TSG SA met in Madrid on 15-18 September, rewrote a liaison statement to SA4 four times in three days, and noted the last of them, so no guidance was sent. Part 1 sets out the position; part 2 is the meeting itself.

How to read the marks in this brief

Every paragraph that carries a claim opens with a chip saying what kind of claim it is: SRC for a fact in an identified document, CNT for a number counted by a stated rule, INT for this brief's reading, CTX for law, markets or other standards bodies that the 3GPP record does not itself establish. Words quoted from a record are set in green with their source. The key is on page one; the reading routes are on page two.

1 The position at Madrid, in one page

For: everyone; read this first.

1.1 The dataset question at Madrid, in one page

SRC The ULBC-MED work item - an ultra-low bitrate speech codec for IMS voice calls over NB-IoT non-terrestrial access through geostationary satellites, in Release 20 and owned by SA WG4 - was approved at TSG SA#112 in Singapore on SP-260696 [S07] [G01]. SA4 has since fixed the operating points at two bitrates, "*Bitrates: 1.0 kbps and 1.4 kbps, and only these two*" SP-260907, slide 5. What turns a codec project into a governance question is stated in TTA's input to PCG: the standardised trained weights "*embed the outcome of the training process into the codec definition*" PCG#56bis(26)02, section 2, problem statement. Once the weights are part of what 3GPP specifies, the data that produced them sits upstream of the specification.

SRC Three records now stand on that question. SA4 surveyed its members and reported that it "*does not identify a consensus need for technical training dataset requirements*" SP-260909, slide 6, conclusion, on a split of 13 companies in favour of requirements or guidance, 9 not

in favour and 2 without a position [M02]. The PCG ad hoc of 3 September read three Organisational Partner inputs - from TTA, from ATIS and a legal paper that SP-260912 and SP-260919 attribute to ETSI - and, as SP-260912 reports it, "*no consensus was reached on whether a common or a proprietary training dataset approach*" SP-260912, section 1 [M03]. Madrid is the TSG#113 round - RAN#113, SA#113 and CT#113 in one week, 14-18 September 2026 [B04] - and it is SA#113, sitting 15-18 September, that has the SA4 answer and three company papers in front of it: SP-260907, SP-260909, SP-260912 and SP-260919, with a draft liaison, SP-260913. Where this brief says Madrid it means that round; where a decision belongs to one plenary it names it.

CNT ESTABLISHED The documentary base is small and complete for what it covers: 68 source records in 7 bands, 29 of them read in full, 136 verbatim quotations each verified as an exact substring of its source, and 213 external references with their addresses printed in full. Every OP input to PCG#56bis, every ULBC document tabled for Madrid that could be opened, the meeting's own register, report and list thread, the PCG#56bis minutes, and the public reflector thread in which SA4 argued over its own report were read end to end.

INT WELL SUPPORTED The dispute is narrower than it looks. All three OP inputs state that any training dataset must meet legal requirements, and all three name licensing and privacy documentation among them ([Table 23](#)). What divides them is whether one common corpus is the way to meet those requirements, who verifies compliance, and - since Madrid - whether SA4's selection timetable should wait for PCG. Beneath the legal question sits a technical one the legal inputs cannot settle: the maintainability of a trained model that only its proponent can retrain, set against the interoperability of a decoder that others cannot retrain without changing it ([section 6.4](#)).

CTX Outside 3GPP the pattern is consistent. JPEG AI mandated a common training set; JVET's 2026 call requires disclosure of training material and access for verification or retraining; MPEG requires public data under commercial-use licences; the IETF's neural speech-codec work keeps test material out of training. None of these bodies has taken on legal or liability governance of a training corpus ([Figure 8](#)). That layer is what PCG has been asked about, and it has no precedent to copy.

The position is narrow; how the record reached it is the next thing to establish.

1.2 From Fukuoka to Madrid: the chain of records

SRC The record begins at Fukuoka. At TSG SA#111 in March 2026 the plenary approved SP-260324, a liaison on support of the codec in Release 20 drafted by Verizon [S01], after a discussion recorded against SP-260258 in which the SA WG4 Chair said some work was intended in Release 20 but "*whether this will be the full ULBC Codec is still to be determined*" [S02]. The question at that point was scope and schedule, not data.

SRC Singapore turned it into a work item. TSG SA#112 approved the work item on SP-260696, a revision of SP-260640 and SP-260417 [S07], approved the study report TR 26.940 v1.0.0 on SP-260421 [S03], and noted two opposing principle papers: SP-260609 from Samsung, Verizon and KT [S04] and SP-260623 from Dolby, Nokia, Ericsson and Fraunhofer IIS [S05]. It edited the guidance Samsung tabled on SP-260641 before recording it ([section 3.2](#)), and asked the SA Chair to inform PCG.

SRC Between Singapore and Madrid the chain ran through five steps: a call for interest to which twenty

companies answered (15-19 June), 6 offline Audio SWG calls from 24 June to 12 August, SA4#137-e on 24-28 August with the survey result, a two-day reflector argument over how that result was worded (27-28 August), and the PCG#56bis ad hoc on 3 September. The documents for Madrid were uploaded on 7 and 8 September [G04].

CNT The companion brief's TSG SA exports carry 22 rows naming the codec across the three plenaries: 3 at #111, 13 at #112 and 6 at #113 (Table 14). Of the #112 rows, 8 were noted and 2 approved. The close of Singapore and the opening of Madrid are 95 days apart; PCG#56bis and Madrid, 12; Madrid and PCG#57, 36. Figure 1 draws the whole chain on one calendar.

INT ESTABLISHED · The referral created two tracks with different clocks. The technical track has reported; the legal track has not, and at the cutoff the 3GPP file server carries no PCG#57 folder [G03]. Madrid is therefore the first plenary at which the technical answer is available and the legal answer is not, and every document tabled for it is, in one way or another, a proposal about what to do in that interval.

A chain of records ends in a set of decisions, and Madrid has three in front of it.

1.3 The decision set in front of TSG SA#113

SRC SP-260909 is the SA WG4 Chair's answer to the two questions TSG SA#112 put, tabled for information: SA4 *"does not identify a consensus need for technical training dataset requirements"* and proposes no work beyond the approved item [M02]. It asks the plenary for nothing. SP-260907, the Chair's status report, is tabled for presentation and carries the work plan through to TSG SA#116 in June 2027 [M01].

SRC Two papers ask for decisions. Samsung's SP-260912 asks the plenary to endorse that work governing qualification and selection *"should proceed only after PCG has reached a conclusion on the legal issues"* SP-260912, section 3, proposal, and that *"SA4 is encouraged to re-evaluate the need for a common dataset"* SP-260912, section 3, proposal [M03]. Verizon's SP-260919 asks it to endorse requirements on *"the number of languages / language families, speaker demographics, acoustic environments"* SP-260919, section 3, proposal, and that *"The codec training dataset should be transparent such that 3GPP can maintain its specifications"* SP-260919, section 3, proposal, and to send those two points to SA4 in a liaison [M05]. Samsung has also tabled a draft liaison, SP-260913 Guidance on ULBC, whose text this brief did not read; the companion brief's scan records 273 words naming SA WG4 7 times [M04].

INT CONDITIONAL · The decision set is therefore three questions, not one. Whether to note the SA4 answer is procedural. Whether to gate the Pdocs on PCG is a decision about SA4's calendar taken above SA4. Whether to send requirement dimensions and a transparency principle to SA4 is a decision about content that SA4 did not reach by consensus. Table 2 sets out what follows from each outcome, both ways; none of them is a forecast of what Madrid will do.

INT One distinction matters for reading the outcome. A plenary that endorses a direction is not a working group that agrees a requirement, and a liaison is a request, not a specification. If Madrid endorses either paper, the effect is to change what SA4 is asked to consider at Shanghai and Calgary; it does not write anything into a Pdoc. Figure 3 places each element of the dataset question on the rung its record supports.

Each of the three decisions runs on a different clock.

1.4 Two clocks: SA4's selection calendar and PCG's open date

SRC The work plan in SP-260907 runs from the Audio SWG ad hoc in Shanghai (28-30 September) to SA4#138 in Calgary in November 2026, where the Pdocs are to be finalised, the funding estimate completed and the qualification phase kept or dropped; binding letters of intent in December; test-laboratory contracts in January 2027; qualified candidates at SA4#139 in February; selection tests from 19 March; the winning candidate at SA4#140 in May; approval at TSG SA#116 in June 2027 [M01]. The rapporteur's own reading is *"Overall completion rate for the WI was determined by the Rapporteur to be at 10%"* SP-260907, slide 7, and the Chair adds that *"Completing ULBC WI for Rel-20 is still possible, but the safety margin is thin."* SP-260907, slide 4

CNT ESTABLISHED · SP-260907 prints 11 Pdocs. Of these, 7 have an agreed version and 4 do not; the 4 without a version are the qualification and selection rules and the two test plans - the documents SP-260912 proposes to hold. None of the eleven concerns training data (Table 27). A training-data Pdoc was proposed at SA4#137-e in S4-261509 [W09], as the base briefing records; it is not among the agreed set.

INT CONDITIONAL · The limiting clock is PCG's, and PCG has no date for an answer. PCG#57 falls 36 days after Madrid opens. If the new ad hoc group reports there and PCG records a conclusion, SA4#138 can still finalise the Pdocs in November and the plan holds even if Madrid endorses sequencing. If PCG needs longer, sequencing moves the qualification rules past November and the January testing with them. The signal that would weaken this reading is a PCG#57 record that closes the legal requirements; the signal that would strengthen it is a PCG#57 record that extends the ad hoc group's mandate.

INT WELL SUPPORTED · The calendar carries a process dependency as well as a document dependency. A Pdoc is not completed by counting votes: SA4 has to reach agreement on the rules that candidates will be bound by. An unresolved objection can therefore consume time while the group works towards consensus, even where many participants are ready to proceed. The Working Procedures provide a fallback - if consensus cannot be achieved the Chair may call a vote, carried at 71% of the votes cast - so an objection is not an automatic veto [G26]. The thin margin in the Chair's report is nevertheless a property of the decision process as much as of the technical work left, and the record has a worked example: in 2004 an unresolved SA WG4 disagreement over two H.264 adoption issues sent the change requests to the plenary by another route rather than ending the work [G36] (section 11.1, Table 61).

CTX Two outside calendars run alongside. JVET evaluates responses to its beyond-VVC call in January 2027 [K02] [K03], the month SP-260907 pencils in for the first ULBC tests, so the two largest competitive selections of neural coding tools in standards will run in the same quarter under different training-data rules. And the EU Product Liability Directive must be transposed by 9 December 2026 [L13], before any ULBC candidate is selected. Table 3 lists the dates.

1.5 Who governs what: the bodies this brief reads, and why

CTX Nine kinds of body appear in this brief, and a reader outside standards work is owed the reason for each. 3GPP is not an organisation but a partnership: seven Organisational Partners - ARIB, ATIS, CCSA, ETSI, TSDSI, TTA and TTC - own the specifications jointly, publish them as their own deliverables and carry the legal personality the partnership itself does not have

[G11] [G12]. That is why a question about what may lawfully stand behind a specification cannot be settled by the working group that writes it, and why it arrived at PCG, the body where the partners meet as partners rather than as engineers.

CTX The other bodies are read because each has addressed an adjacent part of the learned-codec problem - training conditions, model artefacts, reproducibility, disclosure or test separation. Their technical practice is comparable; their decision machinery is not necessarily so (Table 60, Figure 13). JPEG - a joint activity of ISO/IEC JTC 1/SC 29 and the ITU-T - published the first learned image coding standard with normative trained parameters, published by both bodies at once - as ITU-T T.840.1 and as ISO/IEC 6048-1 [K07] [K08]. JVET, the joint team of ITU-T Study Group 21, Question 6, and the same ISO/IEC subcommittee [K20], runs the neural video exploration and the 2026 call that will compete learned tools; the ITU, ISO and IEC issued that call together on 28 July 2026 [K21]. MPEG, also SC 29, sets the common conditions those competitions use, and in July 2026 opened its own call on media authenticity and provenance, for evaluation in January 2027 [K25]. Beyond the coding bodies the ITU-T has a second seat at this question: Study Group 13's Y.3170-series standardises machine learning in future networks - a pipeline with data-collector and preprocessor nodes, a marketplace for models, a sandbox - and, exactly as 3GPP does, standardises the handling and not the data [K22] [K23] [K24]. The IETF is writing the only other neural speech codec in a standards body, and separately the signals by which a rightholder can express a training preference. W3C carries the machine-readable reservation those signals build on. ISO/IEC JTC 1/SC 42, IEEE and ETSI's own security committee write the data-quality and AI-management standards a declaration could be built from, and the provenance formats - SPDx, CycloneDX, Croissant, C2PA - are where a dataset or model manifest already has a vocabulary (Table 4).

INT ESTABLISHED · Read as a map rather than a list, the pattern is regular (Figure 2). Rules about training and evaluation method are written everywhere: in 3GPP by SA4, SA5, SA2 and the radio groups, in JVET and MPEG as common conditions, in the IETF as a normative prohibition on training with the test clips. Rules about the model artefact are written wherever a body publishes one. Rules about the recordings themselves are written almost nowhere inside the standards system, and where they exist they come from outside it - from copyright and data-protection law. The empty cells in the first column are the subject of this brief.

CNT ESTABLISHED · The same emptiness shows in 3GPP's own decided record. Of the 1,000 approved, agreed or endorsed 3GPP documents that name AI, a model and training data, not one carries the words licence, copyright or personal data in the text the export reproduces (section 10.5). The technical vocabulary is dense and the rights vocabulary is absent, across work spanning August 2021 to August 2026.

CTX One step in the argument is worth making explicitly, because it decides who this brief is for. A trained model has no behaviour of its own. What it can do, how well it does it, which languages and conditions it holds up in, and where it fails are properties of the material it learned from; the architecture fixes the shape of the function, the data fixes the function. A performance figure written into a standard about a model is therefore a figure about a dataset, and a requirement written about a model is a requirement someone must meet with data.

INT WELL SUPPORTED · Which is why the question is not SA4's alone. 3GPP already specifies models in all three Technical Specification Groups: RAN

standardises the measurements collected for training and inference [G13], CT the AI/ML data and the jointly trained models that move between network functions [G14], and SA their service requirements, core-network analytics, management and security [G15] [G16] [G17] (Table 55). In every one of them the material has stayed with operators and vendors, and no 3GPP rule has reached it - which is workable exactly as long as nobody has to answer for the material. ULBC changes one thing and changes it for the partnership rather than for a group: it is the first item whose trained parameters would be normative, so it is the first time 3GPP has to state what it requires of the data behind a model. What it states, or declines to state, becomes the reference the next group inherits (Figure 12, section 10.4).

INT CONDITIONAL · What would weaken this reading: a standards instrument outside the 8 bodies and efforts examined here that does allocate corpus clearance, verification and downstream liability for a normative trained model. Every claim of absence in this brief is bounded by that sample and never says that no such instrument exists anywhere. *Signal:* a new AI-governance instrument from a body whose deliverables 3GPP references.

1.6 How the question arrived: three eras, and five years of 3GPP records

CTX The question is younger than it looks and older than this year. A conventional codec is standardised as a bitstream syntax, a decoder description, conformance vectors and reference software; whatever recordings its designers tuned on stay in the laboratory, because no implementer needs them. That held for every 3GPP speech codec from the first to EVS, and it is the arrangement ULBC ends.

CNT ESTABLISHED · The second era is already five years old inside 3GPP. The earliest record in this brief's export of approved work naming AI, a model and training data is from August 2021 [B06]; within ten days of it four different groups had records of the same kind - a security study, the radio framework study, a management data-collection change and a service

requirement for model transfer (Table 5). None of them standardises a model or a corpus: they standardise how data is collected, how a model is managed, transferred and retrained, and how the result is secured. By August 2026 that record had reached 1,000 decided documents across 43 specifications.

CTX The learned-coding bodies moved in the same window and in the same order: technique first, conditions second, rights last. MPEG issued its call for video coding for machines in 2022 and JVET asked in January 2023 whether the datasets behind it could be used at all [J11] [J14]; a signalling proposal for AI use of video followed in April 2024 and a usage-restriction message in January 2025 [J18] [J17]; JPEG AI was approved as ITU-T T.840.1 in March 2025 with its trained parameters normative [K07] [K08]; and by July 2026 the joint call beyond VVC required disclosure of training material and access for re-training [J01].

CTX The law arrived on its own schedule, and mostly from outside telecommunications. Text and data mining was given its European exception in 2019, with the reservation mechanism that goes with it [L11]; the AI Act entered into force in August 2024 and its general-purpose obligations applied from August 2025 [L01]; the revised Product Liability Directive must be transposed by 9 December 2026 [L13]; and the Cyber Resilience Act's reporting duties began on 11 September 2026, the day this brief's documentary cutoff falls [L16]. The courts began deciding in the same period and are deciding still (Table 44).

INT ESTABLISHED · What changed in 2026 was not the technology but the object. A codec whose trained parameters are part of the specification makes the recordings an upstream input to a 3GPP deliverable, for the first time in the record read here. That is why the question travelled from a working group to a plenary to the partnership in six months, and why the bodies that settled the technical half of it a year earlier are no help with the other half.

What this part established

The dataset question reaches Madrid as three questions: whether to note SA4's answer, whether to hold SA4's selection rules for PCG, and whether to send requirement dimensions and a transparency principle to SA4 (Table 2). None is a forecast of what Madrid will do.

Behind it stands a chain of 22 codec rows in three plenary exports, two recorded dispositions at Singapore that framed the question, a survey, a public thread, three OP inputs and four Madrid documents (Figure 1).

Two clocks run: SA4's plan to June 2027 and PCG's, which has no date. PCG#57 falls 36 days after Madrid opens.

Nine kinds of body write about training data and only the law writes about the recordings; the map of who writes what is Figure 2, and the bodies are in Table 4.

The technical era is five years old inside 3GPP and the object changed only in 2026, when trained parameters became the thing being standardised (Table 5).

Hands forward: what the Madrid plenary did with the question - part 2.

Table 1. The ULBC documents in the TSG SA#113 export

TDoc	Source as recorded	Type / for	Agenda item	Export status	What it asks	Read
SP-260790	SA WG4	CR pack / Approval	20.4	available	Change requests on the closed study, FS_ULBC	title only
SP-260907	SA WG4 Chair	report / Presentation	4.4	reserved	Presents status and work plan	read in full
SP-260909	SA WG4 Chair	report / Information	3.4	available	Reports no consensus need; no additional work	read in full
SP-260912	Samsung	discussion / Endorsement	3.4	available	Two endorsements: sequencing behind PCG; re-evaluate common data	read in full
SP-260913	Samsung	LS out / Approval	3.4	available	A liaison: three bullets of guidance to SA4	read in full
SP-260919	Verizon Switzerland AG	discussion / Endorsement	3.4	reserved	Two endorsements and a liaison to SA4	read in full

Every document naming the codec in the Madrid export at the cutoff, with its recorded status. 'Reserved' means a number allocated before its file was uploaded. The last column says how far this brief could read each.

Table 2. The Madrid decision set: three questions, each outcome read both ways

Question	Document	What is asked	If endorsed or sent	If noted	Signal	Class
Note SA4's answer	SP-260909 - SA WG4 Chair, for information	Nothing: reports no consensus need for technical requirements and no work beyond the approved item	Not for endorsement	SA4 continues on the approved work item	Disposition 'noted' in the SA#113 report	SRC
Hold the selection rules for PCG	SP-260912 - Samsung, for endorsement	Finalise the Pdocs governing qualification and selection only after PCG concludes; SA4 to re-evaluate a common dataset	SA4#138 finalisation waits for PCG; PCG#57 becomes the gate	SA4 proceeds; the legal risk is carried towards selection	Disposition of SP-260912; any liaison citing PCG	INT
Send requirement dimensions and transparency to SA4	SP-260919 - Verizon, for endorsement	SA4 to define requirements on languages, demographics and environments; training data transparent for maintenance; a liaison to SA4	SA4 receives testable dimensions and a principle it cannot implement without PCG	The dimensions can still reach ULBC-3 as contributions	An outgoing liaison and its wording on 'transparent'	INT
A draft liaison on guidance	SP-260913 - Samsung, liaison for approval	Three bullets of guidance to SA4 on dataset requirements, maintainability and the work plan	SA4 receives a plenary instruction on what to settle and in what order	The question returns to SA4 in the form it was asked	Outcome: rewritten four times and noted; no liaison sent (Table 7)	SRC
The status report	SP-260907 - SA WG4 Chair, for presentation	Nothing: the work plan to TSG SA#116 in June 2027	-	-	Presented; any comment on the Rel-20 margin	SRC

For delegates and anyone briefing one. Each row is a document tabled for TSG SA#113 and what follows if the plenary endorses or sends it and if it notes it. Neither column is a forecast; the signal column names what the meeting report will show.

Figure 1. One calendar, six tracks: 3GPP, the ITU-T and ISO/IEC coding bodies, MPEG and JPEG, the IETF and W3C, from Fukuoka to TSG SA#116

Solid bars were held before the cutoff; outlined bars are scheduled; dashed outlines are marked to be confirmed or given only to the month, and a month-only date is drawn across its month. Dates from SP-260907, the PCG record, the ITU-T Study Group 21 and MPEG meeting records and the IETF meeting page; Table 3 lists them with their precision.

Every track this brief reads, on one linear calendar, so a delegate in any of them can see where their own body stands against the others. The 3GPP plenary round is RAN, SA and CT together; SA WG4 runs the codec competition; PCG holds the legal question with no date on it; JVET - ITU-T Study Group 21 with ISO/IEC SC 29 - runs the beyond-VVC call whose conditions of entry already reach training material; MPEG's own authenticity call and the IETF's test-clip rule move on their own clocks. The technical track has reported; the legal track has not, and PCG#57 is the first date at which it could. Nothing to the right of the cutoff is an outcome.

Table 3. Dates the record gives, with their precision

Date	Precision	Lane	Event	Status	Anchor
------	-----------	------	-------	--------	--------

For planners. A date given only to the month is printed as a month and drawn across it in Figure 1; no day is supplied where the record gives none. Status is as the record states it at the cutoff.

10-13 Mar 2026	day	3GPP plenary	TSG SA#111 Fukuoka: LS on support of ULBC in Rel-20 approved	held	[S01]
12 May 2026	day	3GPP PCG	PCG#56, Sophia-Antipolis: no dataset referral yet	held	[B02]
9-12 Jun 2026	day	3GPP plenary	TSG SA#112 Singapore: WID approved, guidance edited, PCG to be informed	held	[S07]
15-19 Jun 2026	day	3GPP SA4	Call for interest: twenty companies declare a candidate	held	[M01]
24 Jun - 12 Aug 2026	day	3GPP SA4	Six offline Audio SWG calls	held	[M01]
Jul 2026	month	IETF, W3C	IETF neural speech codec: no training on the test clips, in the draft's own text	held	[I15]
7-15 Jul 2026	day	JVET, ITU-T	JVET#43 Geneva: joint call beyond VVC dated 17 July	held	[J01]
13-17 Jul 2026	day	MPEG, JPEG	MPEG 155 Geneva: call for proposals on media authenticity and provenance	held	[K25]
28 Jul 2026	day	JVET, ITU-T	ITU, ISO and IEC publish the joint call	held	[K21]
1 Aug - 1 Sep 2026	day	JVET, ITU-T	Joint call: registration window	held	[K21]
24-28 Aug 2026	day	3GPP SA4	SA4#137-e: survey reported, Pdocs and work plan agreed	held	[M02]
27-28 Aug 2026	day	3GPP SA4	Reflector debate on the report to SA	held	[W04]
3 Sep 2026	day	3GPP PCG	PCG#56 bis_Ad Hoc: no consensus; new ad hoc group reported	held	[P01]
14-18 Sep 2026	day	3GPP plenary	RAN#113 and CT#113: the same Madrid round; the ULBC documents are at SA#113	scheduled	[B04]
15-18 Sep 2026	day	3GPP plenary	TSG SA#113 Madrid: the four ULBC documents and SP-260913	scheduled	[M02]
28-30 Sep 2026	day	3GPP SA4	Audio SWG ad hoc, Shanghai: progress the Pdocs	scheduled	[M01]
14 Oct 2026	day	3GPP SA4	Audio SWG telco post 137_2	scheduled	[M01]
17-23 Oct 2026	day	JVET, ITU-T	JVET#44, Hangzhou	scheduled	[K20]
21 Oct 2026	day	3GPP PCG	PCG#57, online	scheduled	[B02]
26 Oct 2026	day	JVET, ITU-T	Joint call: main submission package due	scheduled	[J01]
28 Oct 2026	day	3GPP SA4	Audio SWG telco post 137_3	scheduled	[M01]
Nov 2026	month	3GPP SA4	SA4#138 Calgary: finalise Pdocs, funding estimate, qualification decision	scheduled	[M01]
14-20 Nov 2026	day	IETF, W3C	IETF 127, San Francisco	scheduled	[I29]
Dec 2026	month	3GPP SA4	Binding letters of intent; funding commitments	scheduled	[M01]
Jan 2027	month	3GPP SA4	Test-laboratory contracts; qualification testing to be confirmed	to be confirmed	[M01]
Jan 2027	month	JVET, ITU-T	Joint call evaluated at the January 2027 JVET meeting, Brisbane	scheduled	[K02]
13 Jan 2027	day	MPEG, JPEG	MPEG media-authenticity proposals due; evaluated at MPEG 157	scheduled	[K25]

Date	Precision	Lane	Event	Status	Anchor
------	-----------	------	-------	--------	--------

For planners. A date given only to the month is printed as a month and drawn across it in Figure 1; no day is supplied where the record gives none. Status is as the record states it at the cutoff.

Feb 2027	month	3GPP SA4	SA4#139: qualified candidates, to be confirmed	to be confirmed	[M01]
19 Mar 2027	day	3GPP SA4	Selection tests launched	scheduled	[M01]
May 2027	month	3GPP SA4	SA4#140: winning candidate; TS and TR drafts	scheduled	[M01]
Jun 2027	month	3GPP plenary	TSG SA#116: TS and TR for approval	scheduled	[M01]

Figure 2. Which body writes rules about what, and where nothing is written

	The recordings and the corpus	Training and evaluation method	The trained model artefact	The product, and the law
The partnership <i>PCG and the seven Organisational Partners</i>	PCG: asked, open		OPs publish the specification	Each OP's IPR policy
3GPP technical groups <i>TSG SA, RAN and CT, and their working groups</i>	SA4: no rule	SA4 test plans SA5, SA2, RAN and CT interfaces	ULBC weights, normative	
Coding and network bodies <i>JPEG, JVET, MPEG and ITU-T Study Group 13</i>	JVET: disclose and give access	JPEG AI common set MPEG common conditions ITU-T Y.3172 pipeline	JPEG AI parameters	
Internet and web <i>IETF working groups and W3C community work</i>	AIPREF, TDMRep signals	mlcodec: not on the test clips	SCITT receipts	
Assurance and quality <i>ISO/IEC SC 42, IEEE, ETSI security, and the manifest...</i>	ISO/IEC 5259, SPDx, Croissant	ISO/IEC 42001, IEEE 2801	ML-BOM, C2PA credentials	ETSI EN 304 223
Law and the courts <i>European and national instruments, and litigation</i>	Copyright and mining Data protection		AI Act duties	Liability, safety, cyber

 binds someone
 a condition of entry, or a voluntary format
 describes only

A chip is a body that writes about that object at that layer; the tint says how far the writing goes. Rows are layers and columns are objects: neither carries distance or rank. Read from the bodies in Table 4.

The institutional map, for a reader meeting this question from outside standards work or from another 3GPP group. Method rules are written everywhere - by SA4, SA5, SA2, CT3 and the radio groups inside 3GPP, by JVET and MPEG in ISO/IEC and the ITU-T, and by ITU-T Study Group 13 for networks - artefact rules wherever a body publishes an artefact, and rules about the recordings almost nowhere inside the standards system. Every body in the top four rows depends on material governed only in the bottom one.

Table 4. The bodies this brief reads: what each is, and what each cannot decide

Body	What it is	Its writ in this question	What it cannot decide	Where it is read
3GPP PCG	the partnership's coordination group; the seven Organisational Partners meet there as partners	the question was referred to it: 3GPP involvement, and the legal requirements for AI/ML training datasets	technical requirements for a codec	Part 3, section 4.6
The seven OPs	ARIB, ATIS, CCSA, ETSI, TSDSI, TTA, TTC: they own and publish the specifications and carry the legal personality	each OP's own IPR policy binds its members; three of them tabled the inputs read here	one OP cannot bind the partnership alone	Part 3, section 8.6
TSG SA and SA WG4	the plenary that owns the work item, and the codec group that writes the competition rules	selection rules, test plans, requirements and the deliverables that will carry the weights	what the partnership warrants or answers for	Parts 2 and 4
Other 3GPP groups	SA5, SA2, SA3, SA6 and the radio groups	AI/ML management, analytics and federated learning, security, and data collection for training	nothing about a corpus 3GPP never holds	section 10.5
JPEG	a joint ISO/IEC JTC 1/SC 29 and ITU-T activity; its learned codec is published in both bodies, as ITU-T T.840.1 and ISO/IEC 6048-1 [K07] [K08]	the first learned coding standard with normative trained parameters, and a mandated common training set	rights in the images it trained on	section 7.1
JVET	the joint team of ITU-T Study Group 21, Question 6 (VCEG), and ISO/IEC JTC 1/SC 29 [K20]	common training sets in exploration; disclosure and access for re-training as conditions of its 2026 call, issued jointly by the ITU, ISO and IEC [K21]	liability for the material a proponent used	section 7.2
MPEG	ISO/IEC JTC 1/SC 29 working groups	common test and training conditions, and licence conditions on the data a call may use	the rights position of a corpus	section 7.3
IETF	internet standards working groups	a normative rule against training on the test clips in its neural speech work, and the signals a rightholder can set	enforcement of those signals	section 7.4
W3C	web standards community	the machine-readable reservation the signals build on	whether a reservation is legally effective	section 7.5
ITU-T Study Group 13	the network study group, outside the coding bodies	the Y.3170-series on machine learning in future networks: a pipeline with data-collector and preprocessor nodes, a marketplace for models, a sandbox [K22] [K23] [K24]	the material in the pipeline, which stays with the operator	sections 7.6 and 10.4
ISO/IEC SC 42, IEEE	the AI and data-quality committees	data quality, data life cycle, AI management systems, and dataset quality management in adjacent fields	anything specific to a speech corpus	section 9.2
Provenance formats	SPDX, CycloneDX, Croissant, C2PA and the transparency work at IETF	the vocabulary a dataset or model manifest would be written in	whether the manifest is true	sections 9.2 and 7.5
Law and the courts	European and national legislatures, regulators and courts	copyright and its exceptions, personal data, product liability, cyber duties	how a standards body allocates work	Part 7


Who can settle which half of the question. The partnership's own bodies are at the top because the referral put the question there.

Table 5. How the question arrived, in date order

Era	When	What happened	What it changed	Source
Classical codecs	to 2020	3GPP speech codecs are standardised as syntax, decoder description, conformance vectors and reference software	training material, where used, stays with the designer and no implementer needs it	[G01]
AI in the system	Aug 2021	the earliest approved 3GPP records naming AI, a model and training data: a security study, the radio AI framework, a management data-collection change and a model-transfer requirement	3GPP begins standardising how data is collected, models managed and results secured - never the data	[B06]
AI in the system	2022-2026	the same record grows to a thousand decided documents across the management, core, radio, security, requirements and media groups	an operational vocabulary forms - training data, collection, retraining - with no rights vocabulary in it	[B06]
Learned coding	2022-2023	MPEG calls for video coding for machines; JVET asks whether the datasets behind it may be used at all	dataset licensing becomes a question inside a coding body for the first time	[J11] [J14]
Learned coding	2024-2025	JVET proposes signalling of AI use, then a usage-restriction message; JPEG AI is approved as ITU-T T.840.1 with normative trained parameters	a standards body publishes trained parameters, and another writes the first signals about AI use of content	[J18] [J17] [K07]
Learned coding	Jul 2026	the JVET and MPEG joint call beyond VVC requires disclosure of training material and access for re-training	disclosure and access, not a common corpus, become the condition of entry	[J01]
The law	2018-2024	GDPR; the mining exception and its reservation; the AI Act in force August 2024	the conditions on collecting and using recordings are set outside telecommunications	[L10] [L11] [L01]
The law	2026	AI Act general-purpose duties applied from August 2025; Cyber Resilience Act reporting from 11 September 2026; Product Liability Directive transposition due 9 December 2026	the evidence a product must carry arrives before ULBC selection begins	[L01] [L16] [L13]
The threshold	Mar-Jun 2026	SA#111 asks SA WG4 about ULBC in Rel-20; SA#112 approves the work item and edits its guidance from 'common' to 'technical' training dataset	the codec whose weights would be normative is approved, and the question is narrowed on the way to the working group	[S01] [S07] [S06]
The threshold	Aug-Sep 2026	SA WG4 surveys its members and reports no consensus need; the reflector argues over the wording; PCG#56 bis_Ad Hoc reads three Organisational Partner inputs and agrees no scenario	the question reaches the partnership, and Madrid receives it open	[M02] [W04] [P02] [P03] [P04]

Three eras and one threshold. The technical work is five years old, the legal work older still, and the object of standardisation changed only in 2026.

Figure 3. Decision proximity: nothing about training data has reached a rule



Each element rests on the highest rung its record supports; the rungs are ordinal and carry no distance. Placements follow Tables 2, 13 and 23.

Where each element of the dataset question stands at the cutoff, for a reader who wants the state of play in one view. Two technical agreements and one reported priority are the only elements above the proposal rung; every governance element is at or below it.

2 What Madrid decided, and what it left open

For: everyone; start here if you have not followed the meeting.

2.1 Madrid in plain terms: what was asked, and what the plenary did

SRC Between 15 and 18 September 2026 the 3GPP System Aspects plenary, TSG SA, met in Madrid. One item on its agenda concerned a speech codec called ULBC and the data used to train it. To a reader outside the room the question is simple to state: 3GPP is about to standardise a trained neural network for the first time, and nobody has yet decided what rules, if any, apply to the recordings it learns from. The plenary was asked to give its working group direction. It did not, and how it did not is the subject of this part.

CNT ESTABLISHED · 7 documents on ULBC were taken at agenda item 3.4 [G41] [M07] [M08]. Six of them asked the plenary for approval or endorsement. None received either. Four were noted and three were revised into the next draft of the same text, which was itself noted at the end. Two further ULBC documents were taken elsewhere on the agenda: the work item's progress report, noted, and a pack of change requests on the study, approved in a block (Table 6).

CNT ESTABLISHED · The contrast is inside one meeting rather than between meetings. Across all 299 documents tabled at TSG SA#113, the register records 156 approved and 6 endorsed [G41]. The meeting was able to approve and to endorse; it did neither here. The one ULBC item it did approve was the change-request pack, which carries the study's technical text and asks nothing about data.

INT ESTABLISHED · For a reader deciding whether this matters beyond 3GPP, the useful way to read 'noted' is as neither yes nor no. It is the disposition a plenary records when a document has been presented and

discussed and no agreed outcome follows. Nothing was rejected at Madrid: no company's position was defeated, no rule was struck down, and the work item continues on its existing terms. What did not happen is that the working group received an instruction. The question it asked in June returns to it in the same form it was asked.

CTX This is the first occasion on which 3GPP has had to decide anything about training data as a matter of its own procedure rather than as a matter of engineering, and the bodies that write the neighbouring rules - the video coding groups in ITU-T and ISO/IEC, the IETF, the ISO committees that publish dataset documentation standards - have not settled it either (part 7). A reader following AI standardisation more broadly can treat Madrid as an early, well-documented instance of a question every standards body with a trained model in its specifications will meet: who is responsible for the provenance of the data a normative model learned from, and at what level of the organisation that responsibility sits.

The plenary met, discussed the question for four days, and sent nothing.

2.2 The nine ULBC documents, and what the register records for each

SRC The plenary's own register prints, for every document, the agenda item, the type, the purpose it was tabled for, the source, a comment and a result [G41]. The nine ULBC rows are reproduced in [Table 6](#) exactly as the register records them, including the comment column, which is where a revision chain is recorded. The meeting report gives the same nine documents in narrative form, with the discussion under each [G40] [M07].

SRC The three papers that opened the item set out the position as their authors saw it. The SA WG4 Chair's report on training-data requirements was taken first and noted; on it, Qualcomm observed that *"the use of 'consensus' in this is not in line with the use of consensus in 3GPP"* TSG SA#113 report [M07], Qualcomm on SP-260909, Samsung that there is no consensus on either question but that the result indicates general support for further work, and the SA WG4 Chair that *"there is no basis for obtaining a consensus or way forward at this point"* TSG SA#113 report [M07], SA WG4 Chair on SP-260909. Verizon's paper on deployability and maintenance was noted without a recorded comment. On Samsung's paper on next steps the report records *"the work in SA WG4 should not be stopped while discussing this issue and should aim to meet the schedule for the work"* TSG SA#113 report [M07], plenary on SP-260912.

CNT ESTABLISHED · The abstract of that third paper carries the PCG outcome into the plenary record in one line: at PCG#56bis *"no consensus was reached on whether a common or a proprietary training dataset approach is appropriate"* TSG SA#113 report [M07], abstract of SP-260912. That sentence is the hinge of the week. The plenary was being asked to give guidance on a question that the partnership's own coordination body had, two weeks earlier, declined to answer ([section 2.6](#)).

INT ESTABLISHED · A reader should be careful with the word 'noted' in the other direction too. Noting the SA WG4 Chair's report is not a finding that the report is wrong, and noting Verizon's and Samsung's papers is not a finding against their arguments. The disposition attaches to the document, not to the case it makes; the cases themselves are read in parts 4 and 5 and are not disposed of by what Madrid did with the paper that carried them.

Every disposition in that item is in the meeting's own register.

2.3 Four drafts in three days: the guidance that was not sent

SRC The substance of the week is in one document that was rewritten four times. Samsung tabled SP-260913, a draft liaison statement to SA WG4 carrying three bullets of guidance: that *"SA4 should strive to develop a set of technical requirements for datasets that any ULBC candidate codecs shall be trained/tested against"* SP-260947, section 1, first bullet; that *"SA4 should consider/evaluate the maintainability of the specs when using AI training datasets"* SP-260947, section 1, second bullet; and that *"ULBC work plan (e.g., governing candidate qualification and selection phase) should take into account any potential guidance on the training dataset and legal requirements expected from PCG"* SP-260947, section 1, third bullet [M12]. It was tabled for approval on the Monday. By the Thursday it had become SP-260947, then SP-260988, then SP-261002, and the last of these was noted ([Table 7](#), [Figure 4](#)).

SRC The objections are recorded draft by draft. On the first: Nokia that there had been no input to TSG SA

from PCG, so the plenary could not state what had happened there; MediaTek that the first two bullets belong to SA WG4 and the third in the meeting report; Ericsson that *"an SA WG4 colleague had considered this draft and did not think this would be useful to SA WG4 at this time"* TSG SA#113 report [M07], Ericsson on SP-260913. On the second: Nokia that its concerns had not been taken up, Ericsson that the text would slow SA WG4 rather than help it, Verizon that it would not, because SA WG4 must define the requirements for the codec evaluations. On the third: Nokia that *"each time this LS is presented there are the same comments received"* TSG SA#113 report [M07], Nokia on SP-261002; Qualcomm that *"the text of the third bullet is misleading as to the scope of the PCG discussions and the responsibilities of SA WG4 regarding AI work"* TSG SA#113 report [M07], Qualcomm on SP-261002; Apple that an official statement to SA WG4 that PCG is working on the legal implications would be useful [M07].

SRC What replaced the liaison statement is recorded in the same paragraph of the report. The SA WG4 Chair was asked to inform SA WG4 of the PCG discussions directly, and *"PCG has decided to further review these aspects in October 2026 and updated information is expected to be reported at the December TSG SA#114 meeting"* TSG SA#113 report [M07], close of agenda item 3.4 [M07].

CNT ESTABLISHED · The chain is 4 drafts in 3 working days, tabled alternately for approval and for endorsement, ending in a disposition that sends nothing. The purpose column of the register records the shift: the first draft was for approval, the second for endorsement, the third and fourth for approval again ([Table 6](#)).

INT WELL SUPPORTED · Two readings of that sequence are open and the record does not choose between them. On the first, the text was simply not ready: each revision drew new comments, the drafting ran out of meeting, and a fifth draft in a longer week might have carried. On the second, the obstacle was not the drafting but the object: the plenary was being asked to state a position on a question on which it had no agreed position, and no amount of rewording supplies one. The clearest support for the second reading is in the record itself - Nokia's observation that the same comments were received at each presentation, and its later point that the plenary should have consensus within itself before telling a working group it hopes for consensus there. What would weaken it is a liaison statement on the same subject carried at TSG SA#114 in December without the underlying disagreement having changed.

One document carries the whole week, and it was rewritten four times.

2.4 The argument on the email reflector, message by message

SRC The drafting was done in public. Between 16 and 18 September the TSG SA list carried 27 messages on the draft, from 9 organisations [G45] [M11]. [Table 8](#) prints every message with its timestamp, the organisation the record gives for the sender, what the message does to the draft and a summary. The thread is the closest thing in the record to a transcript of the disagreement, and it is public in the same way the August SA WG4 thread is ([section 3.4](#), [Table 16](#)).

CNT ESTABLISHED · Of the 27 messages, 5 are the editor's drafts and revisions, 12 object to sending a liaison statement at all, 7 argue for sending one, one proposes text without opposing it, and two concern where the draft could be read. The messages arguing for the statement come from two organisations; the objections come from five - Dolby, Ericsson, Nokia, Novamint and Orange.

SRC The objections are not variations on one argument, and the differences matter. Nokia's is procedural and about consensus: the text *"the text goes beyond what can reasonably represent a consensus view"* TSG SA reflector [M11], message 4, Nokia, and *"Before indicating that SA is hopeful of a consensus-based decision in SA4, we should first have consensus within SA on that position"* TSG SA reflector [M11], message 21, Nokia. Ericsson's is about utility: *"SA4 do not need any guidance from SA at this point in time"* TSG SA reflector [M11], message 5, Ericsson, and a softened text would still not help because *"Even if you soften the wording a lot it is not not helping SA4 to progress and will add confusion what SA plenary consensus is"* TSG SA reflector [M11], message 25, Ericsson. Dolby's is about the premise: *"there is no consensus on what problem technical training dataset requirements will solve"* TSG SA reflector [M11], message 13, Dolby, so SA4 time spent on it is time taken from the Release 20 schedule; and guidance from the plenary is not read neutrally, because *"Guidance received by SA may also be understood as corrective directions to SA4"* TSG SA reflector [M11], message 17, Dolby. Orange's is about competence: *"the LS simply means that SA overrules some technical decisions and even decisions that have been or to be made by SA4"* TSG SA reflector [M11], message 10, Orange.

SRC The case for the statement is equally specific. Samsung, which drafted every revision, holds that dataset requirements and maintainability are critical to implementers and operators, that the same concerns will simply be raised again in SA4 if the plenary does not address them, and that *"SA4 is not making sufficient progressing and it is the responsibility of SA to guide"* TSG SA reflector [M11], message 26, Samsung. Verizon asked whether any company would say it does not want maintainable specifications, and supplied the wording the editor applied to the first bullet. vivo, the only message that moved the text towards agreement, proposed narrowing the first bullet to transparency of training datasets and replacing the plenary's expression of hope with an encouragement; the editor accepted the edits the following morning.

INT WELL SUPPORTED · The thread shows the two sides answering different questions, which is why the rewording did not converge. One side is arguing about whether the content is correct - whether dataset requirements and maintenance are real problems for a deployed codec. The other is arguing about whether the plenary is entitled to say so - whether a body without an agreed position can address a working group at all. A text can be amended until the first objection is met; the second is not met by amendment, and every revision in [Table 7](#) was an answer to the first kind.

The drafting happened in public, and the thread survives it.

2.5 Where the companies stand, across two meetings

SRC Twelve organisations are named in the Madrid record: 9 on the list and seven in the meeting report, with five appearing in both ([Table 10](#), [Table 8](#)). Seven were also named in the SA WG4 thread of 27 and 28 August ([section 3.4](#)). [Table 9](#) puts the two meetings side by side, company by company, with what each record shows and nothing added.

CNT ESTABLISHED · Five companies appear in both records: Dolby, Ericsson, Nokia, Qualcomm and Samsung. Two appear only at SA WG4 and seven only at Madrid. The position each of the five takes does not change between August and September: the two records agree, which is the more notable result given that the question moved from a working group to a plenary and from the content of a survey report to the wording of a liaison statement.

INT WELL SUPPORTED · The division is not the one a reader might expect. It does not separate operators from vendors: an operator and a handset vendor are on the same side of it, and two vendors with audio-codec businesses of long standing are on opposite sides. Nor does it separate companies with declared ULBC candidates from companies without one: of the twenty declared candidate proponents, 12 are absent from both records entirely (Table 30). What the two records do separate is companies that want the dataset question settled above the working group from companies that want it settled inside it, and that is a question about where authority sits rather than about what the answer should be.

INT WELL SUPPORTED · The absence of most candidate proponents from both records is worth stating plainly, because it bears on how the outcome should be read. Twelve companies have declared an intention to submit a codec and said nothing on the record about the rules the submission would be judged against. This brief does not read silence as a view, in either direction (section 14.3). It does mean that the disagreement documented here is carried by a small number of companies, and that a reader estimating how representative it is has only the record of those who spoke.

The same companies had been here three weeks earlier.

2.6 PCG#56bis, now on the record: six partners and no scenario

SRC The coordination body's own minutes became available during the Madrid week, and they change the evidential standing of everything this brief says about PCG#56bis [G44] [P06]. Until this edition the outcome of that meeting was known only through two company papers written by participants arguing for a course of action. It is now a recorded outcome, from the meeting's own draft report, and the TSG SA Chair told the plenary that the report was on the public file server for all Working Group members to consult [M07].

SRC The report records the conclusion in the Chair's words: *"No clear conclusion on the legal analysis provided"* PCG#56bis report [P06], Chair's closing summary, additional legal aspects need clarification, and *"An ad-hoc group in PCG will be created for the clarification of the legal aspects of the training of AI data models"* PCG#56bis report [P06], Chair's closing summary. That group is decision D-PCG56-Bis/01; ETSI leads it and legal representatives of the other Organisational Partners are invited to join. Its action, A-PCG56-Bis/01, is to provide legal guidance on the training of models for datasets, as an input to the next PCG meeting.

SRC The minutes also record what the six partners argued, which the company papers did not. Table 11 sets out each partner's input, its position on a common dataset, its view of who should decide, and what it said about the clock. TTA proposed a common dataset of publicly available voice material under permissive licensing and held that *"no consensus in SA4 technically speaking, we cannot leave the decision to SA4 since they cannot have consensus"* PCG#56bis report [P06], TTA in discussion. ATIS held that all candidates must be trained on a dataset meeting agreed legal requirements, that a common dataset is one means to that end rather than the objective, and that the priority is not the timeline. ETSI set out three scenarios and declined to choose between them, holding that *"Legal and technical decisions should be separate"* PCG#56bis report [P06], ETSI in discussion and that the technical choice belongs to SA and SA4; one ETSI participant recorded that *"the legal risks associated with non common datasets are higher"* PCG#56bis report [P06], ETSI in discussion, another that *"there is no consensus on having a common dataset in SA4 and ETSI"* PCG#56bis report [P06], ETSI in discussion. CCSA asked that AI/ML work not bring legal risk to 3GPP

companies, that the guidance cover all Working Groups rather than SA4 alone, and that *"Rel-20 should not be delayed"* PCG#56bis report [P06], CCSA in discussion. TSDSI asked PCG to pick one scenario. TCCA asked for at least a common set for validating correct operation.

CNT ESTABLISHED · Six Organisational Partners are recorded as speaking, and no two of them hold the same combination of positions. The SA Chair's closing summary is the plainest statement of where that left the question: there are legal requirements for all scenarios, and there is *"no consensus on the way forward for a common training dataset and left to SA4 to decide"* PCG#56bis report [P06], SA Chair's summary.

INT ESTABLISHED · The sequence matters for anyone reading the Madrid outcome as a failure of the plenary. PCG declined to choose a scenario and returned the choice to SA4; the plenary was then asked to give SA4 guidance on the same choice, and declined in turn. Each body gave the same reason in its own vocabulary: the decision needs an expertise the body holding it does not have, or a consensus the body holding it has not reached. The question has now been returned twice to the group that reported, in writing, that it has no consensus on it either.

The body that was asked first had already declined to answer.

2.7 Three calendars: Release 20, Release 21 and the ULBC plan

CTX Three calendars run at once, and they are published separately. Release 20 froze its service requirements in June 2025 and its architecture in September 2026; its protocol detail is targeted for March 2027 and its protocol coding for June 2027 [G47]. Release 21, whose timeline the June 2026 plenaries approved, freezes its own stage 1 in March 2027, stage 2 in June 2028, stage 3 in December 2028 and its coding in March 2029 [G48] [G49]. The ULBC work plan places the start of selection testing on 19 March 2027, a winning candidate at SA4#140 in May 2027 and the specifications at TSG SA#116 in June 2027 [M01] (Table 12).

CNT ESTABLISHED · Laid against each other, the three calendars leave no interval. The month in which ULBC selection testing opens is the month Release 20 stage 3 is targeted, and the month the ULBC specifications are due is the month Release 20 protocol coding freezes. The ULBC plan does not run ahead of the release it belongs to; it lands on it.

CNT ESTABLISHED · The comparison the record supports is with the two codecs 3GPP has developed by the same route. EVS took 4.5 years from work-item start to approved specifications and IVAS 6.0; measured from the first qualification milestone the figures are 3.0 and 6.7 years (section 11.5, Figure 14, Table 64). The ULBC work item was approved at TSG SA#112 in June 2026 and plans approved specifications in June 2027: one year (Table 62). This brief does not draw a conclusion from that comparison, because the three programmes differ in ways Table 62 sets out - a selection between submitted candidates is not the same exercise as a development from a call - and the comparison is printed with those differences attached rather than as a ratio.

INT WELL SUPPORTED · What the record does support is a statement about what is not yet fixed. Four of eleven Pdocs have no agreed version, and they include the selection rules and both test plans - the documents that define the competition opening on 19 March (Table 27, section 5.1). The dataset question those documents would have to settle has now been returned unanswered by PCG and by TSG SA within three weeks. A reader can put those two facts together without this brief doing it for them; what would change the reading is an agreed

version of the qualification and selection Pdocs at SA4#139 in February.

INT WELL SUPPORTED · Release 21 is the part of the picture least discussed in the record and most consequential for a reader outside SA4. If ULBC does not close in Release 20, the next structural landing place is Release 21, whose stage 3 freezes in December 2028 - two years and nine months after the date ULBC currently targets for its specifications. No document read for this brief proposes moving ULBC to Release 21, and nothing here should be read as predicting it. The interval is stated because it is the size of the decision: in a body whose releases freeze on published dates, the practical cost of an unresolved question is not measured in meetings but in the distance to the next freeze that can carry the work.

CTX For a reader whose interest is in AI standardisation rather than in speech coding, this is the transferable part. A trained model entering a standard inherits the standard's release calendar, and the governance questions the model raises - provenance of the training material, reproducibility of the result, maintenance of the weights - do not run on that calendar at all. They run on the schedule of whichever legal or policy body is asked to answer them, which at the time of writing has no date attached to it (Figure 1). Madrid is the first well-documented case of the two clocks meeting, and it will not be the last.

What remains is a calendar, and it is published.

What this part established

TSG SA#113 took seven ULBC documents at one agenda item. Six asked for approval or endorsement and none received either; the one ULBC document the meeting approved was a pack of change requests taken elsewhere on the agenda (Table 6).

One liaison statement to SA4 was rewritten four times in three days and noted. The objections did not move with the text: five companies opposed sending it, on three distinct grounds - that the plenary has no consensus to express, that guidance would not help SA4, and that the subject belongs to SA4 (Table 7, Table 8, Figure 4).

Five companies appear in both the August SA WG4 record and the Madrid record, and none changes position between them (Table 9). Of the twenty declared candidate proponents, 12 appear in neither.

PCG#56bis is a recorded outcome from this edition: its own draft minutes report no clear conclusion on the legal analysis and create an ad hoc group on the legal aspects, led by ETSI (Table 11).

The calendars are published and they do not leave an interval: Release 20 targets stage 3 in the month ULBC selection testing opens and freezes its coding in the month the ULBC specifications are due; Release 21 freezes stage 3 in December 2028 (Table 12).

Hands forward: where the question came from, and what the edited word did - part 3.

Table 6. The nine ULBC documents at TSG SA#113, as the meeting's register records them

TDoc	Item	Type and purpose	Title as the register prints it	Source	Result	Comment as recorded
SP-260909	3.4	REPORT for information	SA WG4 Chair report on Training Data Requirements for ULBC Candidate Codecs	SA WG4 Chair	Noted	Noted
SP-260919	3.4	DISCUSSION for endorsement	Ensuring a deployable and maintainable ULBC codec	Verizon	Noted	Noted
SP-260912	3.4	DISCUSSION for endorsement	Next steps on ULBC	Samsung	Noted	Noted
SP-260913	3.4	LS OUT for approval	[DRAFT] Guidance on ULBC	Samsung	Revised	Revised to SP-260947
SP-260947	3.4	LS OUT for endorsement	[DRAFT] Guidance on ULBC	TSG SA	Revised	Revision of SP-260913. Revised to SP-260988.
SP-260988	3.4	LS OUT for approval	[DRAFT] Guidance on ULBC	TSG SA	Revised	Revision of SP-260947. Revised to SP-261002.
SP-261002	3.4	LS OUT for approval	[DRAFT] Guidance on ULBC	TSG SA	Noted	Revision of SP-260988. Noted
SP-260907	4.4	REPORT for presentation	ULBC-MED Progress Report to TSG SA#113	SA WG4 Chair	Noted	Noted
SP-260790	20.4	CR PACK for approval	CR Pack on WI_FS_ULBC	SA WG4	Approved	Block approved.

For everyone. Every column is the register's own: the agenda item the document was taken at, the type, the purpose it was tabled for, the title as printed, the source, the recorded result and the comment. Six of the seven documents at item 3.4 were tabled for approval or endorsement and none received either. The one ULBC document approved at this meeting is the change-request pack at item 20.4, which carries technical text and asks nothing about data.

Table 7. The guidance to SA4, draft by draft, and what was said at each

Draft	What changed in it	What the meeting report records was said	Result
SP-260913 draft 1 of 4 for approval	Three bullets: technical requirements for datasets; maintainability of the specifications; work plan to take account of PCG guidance	Nokia: no input to TSG SA from PCG, so what happened there cannot be stated. MediaTek: the first two bullets are debatable and belong to SA WG4; the third belongs in the TSG SA report. Verizon: asked whether any company would not want maintainable specifications. Ericsson: an SA WG4 colleague did not think it would be useful to SA WG4 at this time	Revised to SP-260947
SP-260947 draft 2 of 4 for endorsement	First two bullets reworded after the comments received; third bullet unchanged	Nokia: the same concerns, its earlier comments not taken into account. Ericsson: concern that it would slow SA WG4 work rather than help. Verizon: it would not slow the work, because SA WG4 must define the requirements for codec evaluation. Samsung: if no guidance can be provided, the discussions continue in SA WG4. SA WG4 Chair: if an LS is sent, record that the Chair report was received and that PCG is working on the technical impacts	Revised to SP-260988
SP-260988 draft 3 of 4 for approval	First bullet amended at Verizon's request; a further revision reworded the text as information to SA WG4 rather than guidance	Clarifications to the first bullet requested off-line. Nokia: each time this LS is presented the same comments are received. Qualcomm: the third bullet is misleading as to the scope of the PCG discussions and the responsibilities of SA WG4 on AI work. Apple: an official statement to SA WG4 that PCG is working on the legal implications would be useful. TSG SA Chair: the PCG draft report is on the public file server for all WG members	Revised to SP-261002
SP-261002 draft 4 of 4 for approval	Final text on the table; edits from vivo accepted by the editor on 18 September	No further text was agreed. The SA WG4 Chair was asked instead to inform SA WG4 of the PCG discussions; PCG will review in October 2026 and report to TSG SA#114 in December	Noted

For delegates and anyone reconstructing the week. One document rewritten four times between Monday and Thursday. The third column is the meeting report's own account of the discussion, with the companies named as the report names them; nothing is added to it.

Table 8. The TSG SA email reflector thread of 16-18 September, message by message: the TSG SA#113 Madrid plenary week

Msg	Timestamp	Organisation	What the message does	This brief's reading	Verbatim quotation
-----	-----------	--------------	-----------------------	----------------------	--------------------

For everyone. The public email reflector on which every revision of the guidance was drafted and argued, set in the same representation as [Table 16](#), the SA WG4 email reflector for the SA4#137-e online meeting, so that the two threads can be read row against row. The fourth column classifies what each message does to the draft and never a view imputed to its sender; the fifth is this brief's reading; the sixth is the sender's own words, set in green italics and verified as an exact substring of the thread before it entered the store. Organisations are given as the record gives them.

1	2026-09-16 14:30	Samsung	Circulates or amends the draft	Circulates the draft LS in the meeting drafts folder; states that the text reflects Monday's discussion, allows SA4 to progress the technical work, and carries no reference to a common dataset	<i>"I tried to reflect the discussion on Monday with some meaningful guidance that will allow SA4 to progress the technical work"</i>
2	2026-09-16 17:45	Dolby	Concerns where the draft can be read	Asks that the draft be placed on the public meeting server or sent by email, because the local address is not reachable by remote participants	<i>"on 10.10.10.10 it is not accessible for remote participants"</i>
3	2026-09-16 19:04	BT	Concerns where the draft can be read	Gives the public server path at which the draft can be read	<i>"Can I suggest you look here?"</i>
4	2026-09-16 19:40	Nokia	Objects to the draft as it stands	The text goes beyond what can represent a consensus view; the first bullet reflects neither the lack of consensus in SA4 nor the continued lack of consensus at SA level, and conflates training data with testing and evaluation; the last bullet implies SA has assessed a timeline impact	<i>"the text goes beyond what can reasonably represent a consensus view"</i>
5	2026-09-16 19:51	Ericsson	Objects to the draft as it stands	Shares Nokia's view; SA4 needs no guidance from SA at this point and the matters were not discussed in the detail the draft assumes; an LS from the plenary would not add value, though PCG recommendations, once made, would have to be taken into account	<i>"SA4 do not need any guidance from SA at this point in time"</i>
6	2026-09-17 02:10	Novamint	Objects to the draft as it stands	Shares the two preceding views; the draft goes beyond what was discussed and no LS is needed at this stage; SA4 should continue within its area of expertise and take PCG input into account when available	<i>"the draft LS goes beyond what was discussed and that, at this stage, there is no need for an LS from the SA plenary"</i>
7	2026-09-17 04:13	Nokia	Objects to the draft as it stands	Reads the revision as a rewording of the same concepts; there is no consensus in SA to spend SA4 agenda time on matters SA4 has already discussed without agreement; retraining of a selected codec is speculative at this stage	<i>"There is no consensus in SA on providing guidance to spend precious agenda time allocated to this work item"</i>
8	2026-09-17 05:47	Ericsson	Objects to the draft as it stands	Cannot see how the plenary can give SA4 guidance that helps it progress; the LS may instead consume SA4 meeting time in interpreting what the plenary meant	<i>"It may in reality rather consume SA4 meeting time to try to interpret what SA plenary wanted to give guidance on"</i>
9	2026-09-17 07:27	Samsung	Argues for the draft	States that technical dataset requirements and maintainability are critical to ULBC for implementers and operators; asks the objecting companies whether they agree with the issues and on what rationale; notes that the same concerns will be raised again in SA4	<i>"Samsung views technical dataset requirements and maintainability as critical issues to the development of ULBC"</i>
10	2026-09-17 07:34	Orange	Objects to the draft as it stands	Reads the LS as SA overruling technical decisions made or to be made by SA4, given the survey SA4 has already carried out	<i>"the LS simply means that SA overrules some technical decisions and even decisions that have been or to be made by SA4"</i>
11	2026-09-17 09:25	Samsung	Circulates or amends the draft	Circulates a revision modifying the first two bullets on the feedback received, and records the third bullet as generally acceptable	<i>"I modified the first two bullets based on the feedback I received"</i>

Msg	Timestamp	Organisation	What the message does	This brief's reading	Verbatim quotation
<i>For everyone. The public email reflector on which every revision of the guidance was drafted and argued, set in the same representation as Table 16, the SA WG4 email reflector for the SA4#137-e online meeting, so that the two threads can be read row against row. The fourth column classifies what each message does to the draft and never a view imputed to its sender; the fifth is this brief's reading; the sixth is the sender's own words, set in green italics and verified as an exact substring of the thread before it entered the store. Organisations are given as the record gives them.</i>					
12	2026-09-17 10:37	Verizon	Argues for the draft	Asks whether the objecting company agrees that maintainable specifications should be developed	<i>"I assume you agree we should develop maintainable specifications?"</i>
13	2026-09-17 12:01	Dolby	Objects to the draft as it stands	There is no consensus on what problem technical training-dataset requirements would solve; allocating scarce SA4 time to an unclear topic would be counterproductive for the Rel-20 timeline; maintainability is a technical topic for SA4 once the technology is selected	<i>"there is no consensus on what problem technical training dataset requirements will solve"</i>
14	2026-09-17 12:15	Samsung	Argues for the draft	Replies that if SA4 is already covering the points, stating them in the LS does no harm and keeps everyone on the same page	<i>"by communicating the same points will at least keep everyone on the same page"</i>
15	2026-09-17 12:28	Nokia	Objects to the draft as it stands	Answers the maintainability question: changes to a selected codec do not necessarily preserve backward compatibility, in which case a new codec is added rather than an existing one maintained; guidance cannot be given on a question not yet concluded	<i>"updating a codec does not necessarily preserve full backward compatibility"</i>
16	2026-09-17 12:34	Samsung	Argues for the draft	Clarifies that the proposal asks SA4 only to consider and evaluate maintainability, not how to handle it	<i>"we are only advising SA4 to"</i>
17	2026-09-17 12:53	Dolby	Objects to the draft as it stands	Guidance from SA is read in SA4 as a corrective direction and takes priority; if it is not based on broad consensus it could disturb SA4's work and affect the Rel-20 timeline; the topics can be brought to SA4 directly	<i>"Guidance received by SA may also be understood as corrective directions to SA4"</i>
18	2026-09-17 12:53	Verizon	Argues for the draft	Proposes wording for the first bullet: that SA4 allocate time to discuss technical requirements for datasets against which candidates are trained and tested	<i>"SA4 should allocate time to discuss a set of"</i>
19	2026-09-17 13:23	Samsung	Circulates or amends the draft	Applies the wording proposed in the preceding message and circulates the next draft	<i>"I have applied the change requested by VZ to the first bullet"</i>
20	2026-09-17 15:21	Samsung	Circulates or amends the draft	Circulates a revision reworded as information to SA4 rather than guidance, following the Thursday afternoon discussion	<i>"It is significantly reworded to be of"</i>
21	2026-09-17 16:06	Nokia	Objects to the draft as it stands	Objects that the reworded text still indicates a non-consensus position to SA4, which is not usually appropriate in a liaison statement; consensus within SA should come before expressing hope of consensus in SA4	<i>"this represents a non-consensus-based indication to SA4, which is not usually appropriate to report in a liaison statement"</i>
22	2026-09-17 20:04	Samsung	Argues for the draft	States that the objections offer no text to work from and that the concerns of companies on requirements and maintainability remain unaddressed; remains open to text proposals	<i>"we will remain open to any text proposals that can find a way forward"</i>
23	2026-09-18 06:00	vivo	Proposes a change without objecting	Proposes specific edits to the draft, narrowing the first bullet to transparency of training datasets and replacing the expression of hope with an encouragement to SA4	<i>"I further update the LS (highlighted in yellow) considering the comments by Alessio"</i>
24	2026-09-18 06:54	Samsung	Circulates or amends the draft	Accepts the proposed edits	<i>"Many thanks for your constructive suggestions, these are good and we can take them on-board"</i>
25	2026-09-18 07:25	Ericsson	Objects to the draft as it stands	Still sees no need for a further LS; the urgency for Rel-20 has already been communicated and softened wording would add confusion about what plenary consensus is	<i>"Even if you soften the wording a lot it is not is not helping SA4 to progress and will add confusion what SA plenary consensus is"</i>
26	2026-09-18 07:38	Samsung	Argues for the draft	Replies that the aim is to capture the SA discussion and help SA4 steer the work, and that it is the plenary's responsibility to guide	<i>"SA4 is not making sufficient progressing and it is the responsibility of SA to guide"</i>
27	2026-09-18 07:51	Ericsson	Objects to the draft as it stands	Does not see the SA discussion as conclusive or the information as helpful; is concerned the LS would be counterproductive for SA4	<i>"I do not see that the SA discussion have been conclusive and that the information provided help SA4"</i>

Table 9. The same companies across two meetings: SA WG4 in August, TSG SA in September

Company	SA WG4, 27-28 August: this brief's reading	TSG SA#113, 15-18 September: this brief's reading	What the two records together show	Verbatim quotation
Samsung	Asked for the reasons in favour to be listed beside those against, and for timing to be attributed to the Rel-20 schedule rather than to the companies; accepted the Chair's v2 and asked that the balanced slide stand	Source of both company papers and editor of all four drafts of the guidance; states that technical dataset requirements and maintainability are critical for implementers and operators, that the same concerns will return to SA4 if the plenary cannot settle them, and that it is the plenary's responsibility to guide	Consistent across both meetings: wants the dataset and maintenance questions settled above SA4	<i>"SA4 is not making sufficient progressing and it is the responsibility of SA to guide"</i>
Nokia	Recorded that common datasets and private training have both been used in earlier audio work and that ML codecs raise no new technical dataset question	The most sustained objection at Madrid, on the floor and on the list: the draft goes beyond what can represent a consensus view, conflates training data with testing, rests on speculation about retraining, and would put a non-consensus position into a liaison statement	Consistent across both meetings: opposes a plenary instruction, and grounds the opposition in the absence of consensus rather than in the substance	<i>"Before indicating that SA is hopeful of a consensus-based decision in SA4, we should first have consensus within SA on that position"</i>
Ericsson	Disagreed that the timeline was the main hurdle and listed technical arguments against common training data	Objected at every revision: SA4 needs no guidance at this point, an LS would consume SA4 time in interpretation, and softened wording would add confusion about what plenary consensus is	Consistent across both meetings: the question belongs to SA4, and PCG recommendations are the input it would accept	<i>"Even if you soften the wording a lot it is not is not helping SA4 to progress and will add confusion what SA plenary consensus is"</i>
Dolby	Asked that no technical need be listed first, and that the report answer only the question SA put	Objected on substance: no consensus on what problem technical dataset requirements would solve, and guidance not based on broad consensus would take SA4 cycles and bear on the Rel-20 timeline; also asked that drafts be placed where remote participants can read them	Consistent across both meetings: contests the framing of the question rather than the schedule	<i>"there is no consensus on what problem technical training dataset requirements will solve"</i>
Qualcomm	Seconded the objection to the argument from earlier audio practice and thanked the Chair	Intervened twice on the record, both times on precision: that 'consensus' in the Chair report is not used as 3GPP uses it, and that the third bullet is misleading as to the scope of the PCG discussions and the responsibilities of SA WG4	Consistent across both meetings: intervenes on how the record is worded rather than on which dataset policy to adopt	<i>"the use of 'consensus' in this is not in line with the use of consensus in 3GPP"</i>
Verizon	-	Source of the operator paper; asked on the floor whether any company would not want maintainable specifications, held that the LS would not slow the work because SA4 must define the evaluation requirements, and supplied the wording the editor applied to the first bullet	Present only at the plenary: an operator's deployment and maintenance case	<i>"whether there were any companies who would not want to have maintainable specifications"</i>
Tencent	Circulated the draft survey report and issued v2 with the reasons restructured for balance; declined to go further into who wants what	-	Present only at SA WG4, and in the Chair's role: the record shows the office, not a company position	<i>"3GPP (including SA4 and SA) is contribution driven"</i>
InterDigital	Supported v2, rejected the argument from earlier audio practice, and asked SA4 to state that the decision sets no precedent	-	Present only at SA WG4	<i>"There are means to consider private datasets in conjunction with a common public dataset"</i>

Company	SA WG4, 27-28 August: this brief's reading	TSG SA#113, 15-18 September: this brief's reading	What the two records together show	Verbatim quotation
---------	--	---	------------------------------------	--------------------

For anyone assessing how settled the disagreement is. The reading and the record are kept apart across the row, as they are in **Tables 16 and 17**: the middle columns are this brief's reading of each meeting, the last is the company's own words in green italics. Five companies appear in both records and none changes position between them. A company appears in a column only where that meeting's own record names it; a dash means the record does not, which is not a view and is never read as one.

MediaTek	-	Held that the first two bullets are debatable because they belong to SA WG4, and that the third should be handled by a note in the TSG SA report	Present only at the plenary: places each bullet with the body it belongs to	<i>"the first 2 bullets are debatable as they are related to SA WG4 discussions"</i>
Apple	-	Asked for an official statement to SA WG4 that PCG is working on the legal implications of AI training sets	Present only at the plenary: asks for information rather than guidance	<i>"it would be good to make an official statement to SA WG4 that the PCG is working on the AI training sets legal implications"</i>
Orange	-	Read the LS as SA overruling technical decisions made or to be made by SA4, given the survey SA4 had already carried out	Present only on the Madrid list	<i>"the LS simply means that SA overrules some technical decisions and even decisions that have been or to be made by SA4"</i>
vivo	-	Proposed the edits the editor accepted: narrowing the first bullet to transparency of training datasets and replacing the expression of hope with an encouragement to SA4	Present only on the Madrid list, and the only message that moved the text toward agreement	<i>"I further update the LS (highlighted in yellow) considering the comments by Alessio"</i>
Novamint	-	Shared the objections and held that no LS is needed at this stage	Present only on the Madrid list	<i>"the draft LS goes beyond what was discussed and that, at this stage, there is no need for an LS from the SA plenary"</i>
BT	-	Gave the public server path at which the draft could be read	Present only on the Madrid list, on procedure	<i>"Can I suggest you look here?"</i>

Table 10. What the meeting report records each speaker saying, and on which draft


Speaker, as the report names them	On	What the report records
Qualcomm	SP-260909	That the use of 'consensus' in the report is not in line with the use of consensus in 3GPP
Samsung	SP-260909	That there is no consensus for either question, but that the result can be taken as an indicator of general support for doing the additional work
SA WG4 Chair (role)	SP-260909	That the situation is that there is no basis for obtaining a consensus or way forward at this point
Plenary (role)	SP-260912	A number of comments, and a view that the work in SA WG4 should not be stopped while this issue is discussed and should aim to meet the schedule for the work
Nokia	SP-260913	That there has been no input to TSG SA from PCG, so the plenary cannot state what has happened there or what conclusions it has reached; and, at each later revision, that the same comments are received and its concerns have not been taken into account
MediaTek	SP-260913	That the first two bullets are debatable because they relate to SA WG4 discussions, and that the third should be handled by a note in the TSG SA report or similar
Verizon	SP-260913	Whether there were any companies who would not want to have maintainable specifications; and, at the next revision, that the LS would not slow the work because SA WG4 needs to define the requirements for the codec evaluations
Ericsson	SP-260913	That an SA WG4 colleague had considered the draft and did not think it would be useful to SA WG4 at this time; and, at the next revision, that it would slow the SA WG4 work rather than help
Samsung	SP-260947	That if no guidance can be provided, these discussions will continue in SA WG4; and, at the close, that it is very concerned about the divergence of views on technical requirements for training datasets and about codec maintenance for AI-based solutions
SA WG4 Chair (role)	SP-260947	That if an LS is to be sent it should record that the SA WG4 Chair report was received and that PCG is working on the technical impacts of AI training datasets
Qualcomm	SP-261002	That the text of the third bullet is misleading as to the scope of the PCG discussions and the responsibilities of SA WG4 regarding AI work
Apple	SP-261002	That it would be good to make an official statement to SA WG4 that PCG is working on the legal implications of AI training sets
TSG SA Chair (role)	SP-261002	That the PCG draft report is available on the 3GPP file server to be consulted by all Working Group members

For everyone. Companies are named only where the meeting report names them as speaking; two rows are roles rather than organisations and are marked as such. No view is attributed to any organisation that the report does not record.

Table 11. PCG#56bis as its own draft report records it: six partners, and no scenario chosen

Organisational Partner	Input	On a common dataset	On who decides	On the clock
TTA	PCG56Bis_02	For. Recommends defining and using a common training dataset for ULBC: clearly defined, publicly available voice datasets usable for standardisation under permissive licensing terms, obtained with transparent and publicly documented privacy details	PCG to give guidance before TSG#113; SA to define the candidate list where possible and SA4 to make the final selection; states that SA4 cannot reach consensus, so the decision should not be left to it	Urgency for completion in Rel-20 is a stated reason for guidance before the plenary
ATIS	PCG56Bis_04	Open. All candidates must be trained on a dataset meeting the legal requirements; a common dataset could be a means to that end, but there is no clear motivation for one among its members	Legal requirements to be verified and further developed by a group of legal representatives from the Organisational Partners that wish to take part, which would also verify that the datasets meet them	States that the priority is not the timeline but doing the right thing from the legal point of view
ETSI	PCG56Bis_03	Open, as three scenarios: proprietary or no common dataset; a common dataset chosen from those already available; a common dataset defined by 3GPP. One participant put the legal risk of non-common datasets higher; another recorded that there is no consensus on a common dataset in SA4 or ETSI	The technical choice belongs to SA and SA4; legal and technical decisions should be separate; declines to select a scenario, on the basis that technical experts should make that decision	No assessment of the legal resources had been made, and the scale of the work was not clear enough to estimate a time frame
CCSA	-	No common view among its members on proprietary against common; asks that legal risks be mitigated for either option so that SA4 can hold its technical discussion	PCG should not enter technical discussions; general 3GPP principles and guidelines for AI/ML work should cover all Working Groups, not SA4 alone	States that Rel-20 should not be delayed
TSDSI	-	Prefers PCG to pick one scenario so that the guidance is clear; asked about the use of synthetic training data	Asked for more governance content in the legal analysis	Notes that a regular PCG meeting falls in October and a response could be made by then
TCCA	-	For at least a common set for validation: states a concern about trusting the end-to-end product to work correctly, and accepts that implementers may use additional training material	-	-

For OP delegates and legal staff. From this edition the outcome of PCG#56bis is a recorded outcome, taken from the meeting's own draft minutes rather than from the two company papers that described it. No two partners hold the same combination of positions, and the meeting chose none of the three scenarios.



The shape of the week in one line. A reader who takes nothing else from this part should take this: the text moved four times, the objections did not, and the chain ends in a disposition that sends nothing. What replaced the liaison statement is in the last box - a request to the SA WG4 Chair to inform the working group directly, and a PCG review in October reporting to the December plenary.

Table 12. Three calendars: Release 20, Release 21 and the ULBC work plan

Release	Milestone as 3GPP publishes it	Date	Source	Where ULBC sits against it
Rel-20	Stage 1 (service requirements) frozen	June 2025	[G47]	Behind ULBC: the requirements layer was fixed before the work item started
Rel-20	Stage 2 (architecture) final freeze	September 2026	[G47]	The month Madrid was held
Rel-20	Stage 3 (protocol detail) target	March 2027	[G47]	The month ULBC selection testing is scheduled to begin, on 19 March
Rel-20	ASN.1 and OpenAPI freeze	June 2027	[G47]	The month the ULBC specifications are scheduled for TSG SA#116
ULBC	Selection testing begins	19 March 2027	[M01]	Launch, not completion: the competition opens on this date
ULBC	Winning candidate scheduled	May 2027, SA4#140	[M01]	One plenary cycle after the launch
ULBC	Specifications scheduled	June 2027, TSG SA#116	[M01]	Coincides with the Rel-20 protocol-coding freeze
Rel-21	Timeline approved	June 2026, Singapore plenaries	[G48]	Approved at the plenary that also referred the dataset question
Rel-21	Stage 1 freeze	March 2027	[G48]	The next release's requirements close as ULBC testing opens
Rel-21	Stage 2 freeze	June 2028	[G48]	An 80% checkpoint falls in March 2028
Rel-21	Stage 3 freeze	December 2028	[G48]	The next structural landing place for work that does not close in Rel-20
Rel-21	ASN.1 and OpenAPI freeze	March 2029	[G48]	The outer edge of the next release

For everyone, and the one table in this part that needs no 3GPP vocabulary. Release milestones are as 3GPP publishes them; ULBC dates are as its work plan states them. The two coincidences in the middle rows are not this brief's observation but arithmetic: selection testing opens in the month Release 20 stage 3 is targeted, and the specifications are due in the month its protocol coding freezes.

3 Where the question came from: Singapore to SA4#137-e

For: everyone; the baseline for the rest.

3.1 What Fukuoka and Singapore settled, and what they passed on

SRC Fukuoka left the scope open. The SA#111 discussion on SP-260258 records Verizon proposing a liaison to SA4 for timely standardisation of geostationary voice in Release 20, the SA WG4 Chair confirming that some work was intended, and Nokia commenting that the study was 43% complete; the liaison was approved as SP-260324 after two revisions [S01] [S02]. The record at that point says nothing about training data.

CNT Singapore's export carries 13 rows on the codec. By type: five discussion papers, four work-item texts of which one was approved, one draft study report approved, one Chair's status report and two versions of the guidance paper. By status: 2 approved, 8 noted and three revised (Table 14). The two principle papers - SP-260609 and SP-260623 - were both noted, as was an alternative work-item text from Samsung and CMCC, SP-260611 [S08].

INT ESTABLISHED · The plenary decided the procedural questions and left the principle open. It approved a work item and a study report, recorded two opposing principle papers without choosing between them, and passed the principle to two other bodies. That is the ordinary shape of a plenary record; what is unusual is that one of the two bodies was PCG, which does not

normally receive technical questions from a TSG. The referral is itself the first sign that the question was seen to exceed a working group's competence.

Singapore's most consequential act was not a vote but an edit.

3.2 One edited word: 'common' becomes 'technical'

SRC The draft report of TSG SA#112 records how Samsung's guidance on SP-260641 was edited before it was recorded: "It was decided to change 'common training dataset' to 'technical training dataset.'" TSG SA#112 draft report, disposition of SP-260641 The same entry records that the word 'potential' was removed, that Qualcomm asked for SA4's time on the subject to be limited, and that "It was decided to leave this to the SA WG4 Chair" [S06]. The guidance that reached SA4, as SP-260909 quotes it, asks whether there is any need for technical training dataset requirements, and says "SA WG4 should provide feedback to TSG#113 if any additional work is needed" SP-260909, slide 2, the SA#112 guidance as quoted.

INT WELL SUPPORTED · One edited word changed the question SA4 was asked. 'Common training dataset' names a solution; 'technical training dataset requirements' names a class of rule that may or may not include one. SA4 was therefore never asked whether 3GPP should adopt a common corpus, and its answer - no consensus need for technical requirements - is an answer to the edited question. The common-dataset question reached Madrid again only through company papers, and it reached PCG through TTA.

INT This matters for how Madrid reads the SA4 result. A reader who takes SP-260909 as SA4 rejecting a

common dataset has read an answer to a question that was not put. A reader who takes it as SA4 endorsing proponent-selected data has made the same error in the other direction. The survey went wider than the edited question, and section 3.5 shows where the wider answers went.

The edited question went to SA4, which answered it in writing.

3.3 The SA WG4 survey: 24 companies, two questions, no consensus

SRC The SA WG4 Chair did not put the question to a plenary session. SP-260909 gives three reasons - a summer break, a single e-meeting with few online slots, and a work item with urgent technical content - and describes a written survey, S4-261382, whose "Anonymized responses were requested and received in private" SP-260909, slide 3 by the Chair [M02] [W01]. "The survey was designed to be comprehensive and deliberately went beyond the SA request" SP-260909, slide 4; it covered the general need for requirements or guidance, common training data, separation of training data from test material, visibility and disclosure of training data, and possible SA4 outputs. Its scope note reads "legal and licensing aspects were out of scope" SP-260909, slide 4.

CNT ESTABLISHED · 24 companies responded: 16 that had declared an intention to submit a candidate and 8 that had not. Against the 20 declarations recorded in SP-260907, that is a response from 80% of declared proponents. On the need for technical requirements or guidance, 13 were in favour, 9 not in favour and 2 took no position: 54.2% of all respondents and 59.1% of those who took a position. On work beyond the

approved item, 10 favoured some output and 9 none, a margin of 1, with 5 (20.8%) taking no position (Figure 5, Table 13).

CNT Three things the tally cannot distinguish. It counts one response per company whatever that company's role, so a candidate and a company with no candidate weigh the same. It is anonymous, so no answer can be joined to a company or to the declared or undeclared group. And the report does not quantify how the 10 favouring some output split between informal guidance and a mandatory requirement; it says only that *"they split between informal guidance and a mandatory requirement, mainly for fair comparison and robustness"* SP-260909, slide 6.

INT ESTABLISHED · The Chair's conclusion is the conservative one the tally supports. A question that divides 13 to 9 among companies of a working group that decides by consensus has no consensus, and the report says exactly that and no more. It does not say that SA4 opposes requirements, and a majority in a survey is not a decision in 3GPP.

The answer was argued over in public before it was sent.

3.4 Arguing over the wording: the SA WG4 email reflector, 27-28 August

SRC The report's wording was argued in public. On 27 and 28 August the SA4 list carried 11 messages from 7 senders on the Chair's draft [W04] [G09]. Samsung asked for the reasons in favour to be listed beside those against and argued that 'timing and schedule impacts' *"is not a technical reason but a general concern about the tight Rel-20 timeline"* SA4 reflector [W04], message 2 (Samsung). Ericsson disagreed that the timeline was the main hurdle and listed technical objections: *"A common training data does not guarantee a certain performance, e.g. across different languages."* SA4 reflector [W04], message 3 (Ericsson); *"Access to common training data is not sufficient to replicate training."* SA4 reflector [W04], message 3 (Ericsson); *"if re-training is foreseen, it raises questions on maintaining the interoperability of the standard"* SA4 reflector [W04], message 3 (Ericsson); and *"Why should non-public data be forbidden to be used, if it can improve the codec performance?"* SA4 reflector [W04], message 3 (Ericsson)

SRC Dolby asked that no technical need be listed first and wrote, of its own answer, that *"the timeline aspect was definitely not the reason"* SA4 reflector [W04], message 7 (Dolby). The Chair issued a restructured v2 and declined to go further, noting that *"3GPP (including SA4 and SA) is contribution driven"* SA4 reflector [W04], message 6 (SA WG4 Chair). Samsung accepted v2 and asked that it stand, because *"the survey responses were anonymous and only the Chair has visibility of all 24 replies"* SA4 reflector [W04], message 8 (Samsung). A participant whose organisation the thread does not state supported v2, observed that *"There are means to consider private datasets in*

conjunction with a common public dataset" SA4 reflector [W04], message 9, and asked SA4 to make clear that *"this decision SHALL NOT set a precedent for other codecs"* SA4 reflector [W04], message 9; a second participant seconded that. Nokia closed the thread: *"ML codecs do not present a specific new technical dataset problem"* SA4 reflector [W04], message 11 (Nokia).

CNT Of the 11 messages, 2 are the Chair's; the others come from 6 organisations - 9 messages carry the organisation in the thread itself, and for 2 the thread names only the sender, with the organisation supplied for this brief (Table 16). The reasons each side gave, on the reflector, in the survey report and in the company papers, are collected in Table 17 with where each is recorded.

INT WELL SUPPORTED · The argument was about whether 'timing' is a technical reason, and it was not settled. SP-260909 kept the formulation *"the two main reasons include timing/schedule impacts and no technical need"* SP-260909, slide 6. For Madrid this matters because SP-260919 reads the same result as showing that *"the principal reservation among those not in favor was not a technical objection"* SP-260919, section 2, a reading two of the companies not in favour had contested on the reflector a week earlier. Both readings are available from the same slide; only the reflector shows that the companies concerned said timing was not their reason.

INT The thread also shows where the two sides touch. No message argues that test material may be used for training. The participant who objected to the precedent argument proposed combining private data with a common public set, and Dolby's survey reply, reproduced as an image in the thread, answers robustness with codec requirements enforced in selection tests - which is where Verizon's requirement dimensions also point (section 5.4).

The thread argued about what the report carried; what it left out comes next.

3.5 What the survey found that the report to TSG SA does not carry

SRC The survey collected more than SP-260909 reports. The report's annex says the additional views were synthesised in S4-261674, intended only to improve members' common understanding, and states: *"It does not form part of the feedback provided to TSG SA in this report."* SP-260909, slide 8, annex [M02] [W02]. Verizon's SP-260919 reports two of those additional figures to the plenary: *"near-unanimous support (20 of 24 companies)"* SP-260919, section 2, quoting S4-261674 for an explicit requirement that qualification, selection and characterisation material must not be used for training, and on common training data that *"approximately 60% were in favor"* SP-260919, section 2, quoting S4-261674, counting

12 in favour and 1 supporting an optional common set against 9, with 2 expressing no position [M05].

CNT ESTABLISHED · The two annex figures check against their own bases. 13 of the 22 companies that expressed a position on common data is 59.1%, which is what 'approximately 60%' describes. The test-material figure is 83.3% of all respondents. The split on common data, 13 to 9 with 2 abstaining, is identical to the split on the need for requirements; the anonymity of the survey means that coincidence is consistent with, but does not show, the same companies on each side.

INT WELL SUPPORTED · The most widely supported technical rule in the survey is the one figure the report to TSG SA does not carry. Madrid will receive a Chair's report that excludes the annex by design and a company paper that quotes it, and the two are different evidence: one is the working group's feedback, the other a member's reading of a working-group document. Neither is wrong. But the one point on which the survey approaches unanimity - keeping test material out of training - reaches the plenary only through a paper tabled for endorsement by one company.

What this part established

Singapore edited the question before SA4 answered it, and SA4's answer - no consensus need - answers the edited question, not whether 3GPP should adopt a common corpus.

The survey split 13, 9 and 2 on the need for requirements, and 10 to 9 with 5 abstaining on additional work. Its most widely supported rule, 20 of 24 for keeping test material out of training, sits in an annex the report excludes by design (Table 13).

The thread shows the two sides disagreeing about whether timing is a technical reason and agreeing, without saying so, that test material stays out (Table 17).

Hands forward: the three OP inputs PCG read on 3 September - part 4.

Table 13. The survey tally on both bases

Question	Where reported	In favour	Not in favour	No position	Share of all 24	Share with a position
Need for technical requirements or guidance	SP-260909	13	9	2	54.2%	59.1%
Additional work beyond the approved item	SP-260909	10	9	5	41.7%	52.6%
Common training data (favour includes 1 optional)	SP-260919, annex figure	13	9	2	54.2%	59.1%
Test material excluded from training	SP-260919, annex figure	20	-	-	83.3%	-

Two bases, both printed: all 24 respondents, and those who took a position. The annex rows are attributed to SP-260919, which reports them, not to the survey synthesis it quotes. The favour count on common data includes one company supporting an optional common set.

Table 14. The codec in the TSG SA exports across three plenaries


Meeting	TDoc	Title as recorded	Source as recorded	Type	Recorded status
#111 Fukuoka	SP-260257	[DRAFT] LS on Support of ULBC in Release 20	Verizon Spain	LS out	revised
#111 Fukuoka	SP-260283	[DRAFT] LS on Support of ULBC in Release 20	Verizon Spain	LS out	revised
#111 Fukuoka	SP-260324	LS on Support of ULBC in Release 20	TSG SA	LS out	approved
#112 Singapore	SP-260417	New WID on Ultra Low Bitrate Speech Codec for Voice over GEO	SA WG4	WID new	revised
#112 Singapore	SP-260421	TR 26.940 v1.0.0	SA WG4	draft TR	approved
#112 Singapore	SP-260602	Considerations and proposals for ULBC	vivo Mobile Communication Co.,	discussion	noted
#112 Singapore	SP-260609	Principles for standardizing an AI based codec for ULBC	Samsung, Verizon, KT Corp.	discussion	noted
#112 Singapore	SP-260611	New WID on Ultra Low Bitrate Speech Codec for IMS Voice Call over NB-IoT NTN via GEO Satellite	Samsung, CMCC	WID new	noted
#112 Singapore	SP-260614	SA WG4 chair status report on ULBC	SA WG4 Chair	report	noted
#112 Singapore	SP-260619	Clarification on ULBC codec bitrates support in the Rel-20 ULBC-MED Work Item	Dolby Laboratories, Inc., Nokia, Novamint, Terrastar, Viasat, Eutelsat, Sateliot, Omnispace, ESA	discussion	noted
#112 Singapore	SP-260623	Thoughts on proposed principles for standardizing an AI based codec for ULBC	Dolby Laboratories, Inc., Nokia, Ericsson LM, Fraunhofer IIS	discussion	noted
#112 Singapore	SP-260627	Views and Suggestions on Rel-20 ULBC (Ultra Low Bitrate Speech Codec)	China Mobile (Rapporteur of FS_ULBC)	discussion	noted
#112 Singapore	SP-260640	New WID on Ultra Low Bitrate Speech Codec for Voice over GEO	SA WG4	WID new	revised
#112 Singapore	SP-260641	Guidance for SA WG4 ULBC Work	Samsung	other	revised
#112 Singapore	SP-260671	Guidance for SA WG4 ULBC Work	Samsung	other	noted
#112 Singapore	SP-260696	New WID on Ultra Low Bitrate Speech Codec for Voice over GEO	SA WG4	WID new	approved
#113 Madrid	SP-260790	CR Pack on WI FS_ULBC	SA WG4	CR pack	available
#113 Madrid	SP-260907	ULBC-MED Progress Report to TSG SA#113	SA WG4 Chair	report	reserved
#113 Madrid	SP-260909	SA WG4 Chair report on Training Data Requirements for ULBC Candidate Codecs	SA WG4 Chair	report	available
#113 Madrid	SP-260912	Next steps on ULBC	Samsung	discussion	available
#113 Madrid	SP-260913	Guidance on ULBC	Samsung	LS out	available
#113 Madrid	SP-260919	Ensuring a deployable and maintainable ULBC codec	Verizon Switzerland AG	discussion	reserved

The rows of the companion brief's exports that name the codec, as recorded. Status is the export's own word; 'revised' means replaced by a later number in the same meeting.

Table 15. What the two closed plenaries recorded about the codec

Meeting	TDoc	Title as recorded	Class	What the record states
#111	SP-260257	(LS OUT) [DRAFT] LS on Support of ULBC in Release 20	decision	"This LS was approved."
#111	SP-260258	(DISCUSSION) Views and Suggestions on progressing NB IoT GEO voice	endorsement	"Verizon clarified that they propose to send a liaison to SA WG4 requesting timely standardization to ensure commercial readiness of GEO-based voice services in Rel-20. The SA WG4 Chair confirmed that some work is [...]"
#111	SP-260264	(REPORT) 3GPP Support Team Report	decision	Support team table, SP-260324: "Revision of SP-260283. This LS OUT was approved"
#112	SP-260404	(REPORT) 3GPP Support Team Report	decision	Support team table, SP-260696: "Revision of SP-260640. This WID NEW was approved"
#112	SP-260417	(WID NEW) New WID on Ultra Low Bitrate Speech Codec for Voice [...]	decision	"This new WID was approved."
#112	SP-260421	(DRAFT TR) TR 26.940 v1.0.0	decision	"This draft TR was approved."
#112	SP-260641	(OTHER) Guidance for SA WG4 ULBC Work	endorsement	"It was decided to change 'common training dataset' to 'technical training dataset'. It was decided to also remove 'potential'. Qualcomm asked to ask SA WG4 to allocated limited time for this discussion. Including 'limited' [...]"

Every disposition the two draft reports record against a codec document, in green where the record's own wording is quoted; long entries are shortened at a marked point. The edited guidance on SP-260641 is the entry [section 3.2](#) reads.

Figure 5. The SA WG4 survey: every answer drawn against all 24 respondents

One answer per company, anonymous. Hatched segments are answers whose split the record does not report. Table 13 gives the shares on both bases.

What the survey found, drawn so that no-position answers stay visible as part of the base. For a reader weighing a majority claim: the need for requirements splits 13 to 9; keeping test material out of training does not split at all, and it is the row the report to TSG SA leaves in the annex.

Table 16. The SA WG4 email reflector thread of 27-28 August, message by message: the SA4#137-e online meeting

Msg	Timestamp	Organisation	What the message does	This brief's reading	Verbatim quotation
1	27 Aug 00:48 2026	Tencent (Chair)	Circulates or amends the draft	Circulates the draft revision of S4-261675 after the dedicated session on the survey	-
2	27 Aug 04:00 2026	Samsung	Proposes a change without objecting	Asks for the reasons in favour to be listed, for timing to be attributed to the Rel-20 schedule rather than presented as a technical reason, for the survey pie chart in place of large numbers, and for an annex on common training data	<i>"is not a technical reason but a general concern about the tight Rel-20 timeline"</i>
3	27 Aug 07:24 2026	Ericsson	Objects to the draft as it stands	Disagrees that the timeline is the main hurdle; lists technical arguments against common training data	<i>"A common training data does not guarantee a certain performance, e.g. across different languages."</i>
4	27 Aug 10:17 2026	Samsung	Argues for the draft	Clarifies that the proposed wording was taken from the Chair's own summary in S4-261674	-
5	27 Aug 12:49 2026	Dolby	Proposes a change without objecting	Keeps slide 6; reasons on both sides are non-representative examples; the report should answer only SA's question; refers delegates to the Audio SWG report	-
6	28 Aug 00:17 2026	Tencent (Chair)	Circulates or amends the draft	Issues v2 with slide 6 restructured for balance; declines to go further into who wants what	<i>"3GPP (including SA4 and SA) is contribution driven"</i>
7	28 Aug 08:14 2026	Dolby	Objects to the draft as it stands	Objects that v2 presents timeline as the main concern; proposes that no technical need is listed first	<i>"the timeline aspect was definitely not the reason"</i>
8	28 Aug 14:01 2026	Samsung	Accepts the draft as it stands	Accepts v2 and asks that slide 6 stay as it is; the bullet aggregates anonymous replies only the Chair has seen	<i>"the survey responses were anonymous and only the Chair has visibility of all 24 replies"</i>
9	28 Aug 15:01 2026	InterDigital	Accepts the draft as it stands	Supports v2; rejects the argument from earlier audio practice; asks SA4 to state that the decision sets no precedent	<i>"this decision SHALL NOT set a precedent for other codecs"</i>
10	28 Aug 17:11 2026	Qualcomm	Seconds another message	Seconds message 9 and thanks the Chair	-
11	28 Aug 14:51 2026	Nokia	Objects to the draft as it stands	Common datasets and private training have been used in earlier audio work; ML codecs raise no new technical dataset problem	<i>"ML codecs do not present a specific new technical dataset problem"</i>

The public argument over how SA4's report was worded, message by message, in the order the thread records. Six columns, and the same six as [Table 17](#), the TSG SA email reflector for the Madrid plenary week, so that the working group's thread and the plenary's can be read row against row: what the message does to the draft, this brief's reading of it, and the sender's own words in green italics. Times are from the messages' own headers, in UTC where the header states a zone. The Chair's two messages are cited by role; four messages carry no quotation because none was verified for them.

Table 17. Reasons given on each side, and where each is recorded

Reason	Where recorded	Class	This brief's reading
In favour of requirements or guidance			
Fair comparison	SP-260909 (the 10 in favour of output); SP-260919	SRC	Addressed by a common set for primary results, or by disclosure
Robustness	SP-260909; SP-260919	SRC	Addressed by requirement dimensions verified in testing
Bias avoidance	SP-260919	SRC	A coverage requirement on the corpus or outcome
Reproducibility	SP-260919	SRC	Two meanings: decoder conformance, or regeneration of the weights
Maintenance independent of the proponent	PCG#56bis(26)02; SP-260919	SRC	The only reason at PCG that does not depend on law
Against, or not in favour			
Timing and schedule impacts	SP-260909	SRC	Contested on the reflector as not a technical reason
No technical need	SP-260909; message 11	SRC	Earlier audio work used common and private data without a rule
A common set does not guarantee performance across languages	message 3	SRC	Answered by requirement dimensions rather than data
Access to data is not enough to replicate training	message 3	SRC	Consistent with Table 34 : a recipe is needed too
Retraining raises interoperability questions	message 3	SRC	A retrained model is a new decoder version
Better non-public data may improve the codec	message 3	SRC	Met by the hybrid: primary results on common data, extra data disclosed

The reasons the survey report, the reflector and the Madrid papers record, grouped by side, with this brief's reading beside each. The reading column is interpretation; the first two columns are the record.

4 PCG#56bis: three Organisational Partner inputs, read in full

For: PCG and OP delegates, legal staff.

4.1 A dedicated agenda item, and who tabled what

SRC PCG#56 bis_Ad Hoc - the name the meeting's own register gives it - was an ad hoc electronic meeting of the Project Coordination Group on 3 September 2026 with one substantive item. The agenda lists the opening, approval of the agenda with an antitrust and competition law compliance statement, any other business and the closing around agenda item 3: *"3GPP involvement and legal requirements for AI/ML training datasets"* PCG#56bis(26)01, agenda item 3 [P01]. Three contributions were tabled under that item: TTA's, marked for decision; a paper on legal considerations; and ATIS's recommendations [P02] [P03] [P04]. The register and the documents number them differently - PCG56Bis_01 to PCG56Bis_04 in the register, PCG#56bis(26)01 to

(26)04 on the documents themselves - and this brief prints both rather than correcting either ([Table 19](#)). After the first use of the full name it writes PCG#56bis, as the Madrid papers do.

CNT ESTABLISHED · The register is the meeting's own record of what was tabled, and it fills in three things the documents do not. It gives every document a source: the agenda and the legal paper from ETSI, the problem statement from TTA, the recommendations from ATIS [P05] [G02]. It shows how quickly the meeting was assembled: every one of the 4 numbers was reserved between 12:44 and 14:03 on 2 September, the day before the meeting, and the last file reached the folder at 11:02 on the morning of the meeting itself. And it records a status against each document: all 4 stand as available.

INT ESTABLISHED · That third field matters twice over. The legal paper carries no Source line on its face, and two Madrid papers had to be used to attribute it: SP-260912 calls it the ETSI legal analysis and SP-260919 describes the ad hoc as held on the basis of contributions from TTA, ETSI and ATIS [M03] [M05]. The register now agrees with them, which is as far as

attribution can be taken from records: the paper is ETSI's as a tabled document, and its positions are still read here as the paper's rather than as any Organisational Partner's formal policy. The status field is the second point. In the 3GPP vocabulary a document that has been disposed of carries the disposition - approved, agreed, endorsed, noted, not pursued - and none of these four does ([Table 18](#)).

INT ESTABLISHED · The agenda wording repeats the Singapore instruction. TTA's introduction records that the TSG SA Chair was to inform PCG of the issue of 3GPP involvement and legal requirements related to datasets used to train AI models, using ULBC as the example [P02]. PCG was therefore asked two things - how far 3GPP itself should be involved, and what the legal requirements are - and neither of them is which corpus to use. The three inputs answer the two questions in different proportions, and that difference, rather than a disagreement about any one corpus, is what the rest of this part reads.

CTX The division of labour the referral assumes is the partnership's own. Under the 3GPP Working Procedures the Organisational Partners carry legal personality and

publish the specifications, PCG handles matters of the partnership, and the Technical Specification Groups and their working groups own technical content [G11] [G12]. A question about what the partnership may warrant, host or be liable for cannot be settled by SA4 even if SA4 wished to settle it - which is why the SA WG4 Chair's survey put legal and licensing aspects out of scope before it asked anything (section 3.3).

The first input, and the fullest, is TTA's.

4.2 TTA: a common public corpus for due diligence and maintenance

SRC TTA's paper is the fullest of the three and the only one marked for decision. Its problem statement starts from the codec: the standardised trained weights *"embed the outcome of the training process into the codec definition"* PCG#56bis(26)02, section 2, problem statement, which creates implementation risks not met with traditional speech codecs. It lists six risk dimensions - source provenance, licensing and intellectual property, privacy compliance, diversity and bias, test-set leakage, and rights to use training material in standardisation - and observes that *"3GPP currently has no established governance framework specific to AI training data in its specifications"* PCG#56bis(26)02, section 2. A self-declaration by proponents gives some information but does not allocate responsibility if a problem is found after implementation; a 3GPP-defined or approved dataset would be *"a best-effort mechanism to reduce such risks"* PCG#56bis(26)02, section 2 [P02].

SRC The second leg of the argument is maintenance, not law. For a conventional codec, TTA writes, everything needed to maintain it is in the specification and its deliverables; for a codec with standardised trained parameters, maintenance may require investigation, modification or retraining. If the selected model was trained on material only the winning proponent holds, retraining depends on that company, *"creating a long-term dependency for the maintenance of a 3GPP specification"* PCG#56bis(26)02, section 3.3. TTA adds that availability of the material does not by itself guarantee exact reproduction, but calls it the necessary basis. It is careful about roles: *"Defining a dataset does not mean that 3GPP must own or host the dataset"* PCG#56bis(26)02, section 3.2, and *"selection of the dataset based on technical suitability remains within the responsibility of SA4"* PCG#56bis(26)02, section 3.3 [P02].

SRC TTA supports the proposal with practice elsewhere - JPEG AI's mandated training set, its CC0 licence and its rights clause; JVET's NNVC exploration, which it expects to follow JPEG AI; and MPEG's Video Coding for Machines, whose conditions it describes as requiring publicly available training data under licence terms allowing commercial use - and with four public speech corpora, of which it recommends Mozilla Common Voice for its CC0 licence and its privacy framework, noting that *"None of the other public voice dataset has an explicit contributor consent framework and documentation"* PCG#56bis(26)02, section 4.1 [P02]. Section 4.7 checks the four corpora against their 2026 terms and the three precedents against the records of the bodies concerned.

SRC The proposals are three. PCG should give clear guidance before TSG#113 in September, given the Rel-20 urgency and that ULBC is *"the first case of specifying an AI/ML model in 3GPP"* PCG#56bis(26)02, section 5, proposal. SA4 should define and use a common training dataset of clearly defined public voice datasets with permissive licences, such as CC0 or CC BY 4.0, obtained with transparent and publicly documented privacy details, and select it against those conditions and the technical requirements. And for other AI/ML functionality PCG should continue to study whether guidance is needed [P02] (Table 20).

INT WELL SUPPORTED · The two legs carry different weight and should be read apart. The legal leg is shared by all three inputs and can be met without a common corpus, by requirements and verification. The maintenance leg is TTA's alone at PCG - SP-260919 takes it up at Madrid - and it does not depend on the legal one: a proprietary dataset that is lawful in every respect still leaves retraining with its owner. A reader who finds the legal case for a common corpus weak has not yet answered the maintenance case, which section 6.4 examines on its own terms.

TTA says what data; the legal paper asks what each answer would cost the partnership.

4.3 The legal paper: three levels of involvement, and the burden of each

SRC The legal paper opens by setting its own limits: *"This contribution does not address which of those options 3GPP should choose."* PCG#56bis(26)03, section 1 It distinguishes two scenarios - technical dataset requirements with each proponent training on its own data, and a common reference dataset, either an existing third-party one or one developed within 3GPP - and finds third-party rights relevant to both [P03]. It then restates how 3GPP handles rights today: members contribute text in good faith and hold copyright in it, essential IPR is disclosed, and third-party material is incorporated only with authorisation, a diligence that is easy because such material is usually limited.

SRC A training dataset changes the scale. The paper names six classes of risk - copyright in underlying recordings, performer rights, database rights, confidentiality obligations, contractual licensing restrictions and personal-data protection for voice - and calls the provenance work *"similar to that of the establishment of appropriate Bills of Material for software"* PCG#56bis(26)03, section 2. Because the dataset shapes the model that implementers ship, its legal quality matters to them under the EU AI Act and the General Product Safety Regulation, the latter cited as recognising digital components *"while reversing the burden of proof on manufacturers"* PCG#56bis(26)03, section 2. Hence *"legal due diligence cannot be treated as a secondary issue"* PCG#56bis(26)03, section 2, and *"Robust dataset due diligence should therefore be considered an industry requirement."* PCG#56bis(26)03, section 2 [P03].

SRC On governance the paper separates an existing dataset from one 3GPP builds. For an existing one, 3GPP would need to assess whether it was lawfully developed, whether its terms permit industry use for implementation, and whether maintenance and long-term availability are arranged; the same assessment applies where proponents use their own data, and proprietary non-public data raises a question the paper leaves open. For a 3GPP dataset it names three objectives - provenance and legal compliance, transparency and auditability, maintenance and remediation - and three cautions: *"3GPP cannot formally provide guarantees"* PCG#56bis(26)03, section 3; *"contributor self-declarations alone would likely be insufficient"* PCG#56bis(26)03, section 3; and *"3GPP legal resources are limited and such activities would constitute a new area of responsibility"* PCG#56bis(26)03, section 3. For defects found later, *"a mechanism comparable to Coordinated Vulnerability Disclosure (CVD) could be considered"* PCG#56bis(26)03, section 3 [P03].

SRC The liability section is the one the other two inputs do not have. Disclaimers work today, the paper says, because OPs and implementers deal business to business and a specification is not a product component. A common dataset developed by 3GPP would instead become *"a de facto determinant element on the final product"* PCG#56bis(26)03, section 4, so OPs could face

greater exposure to civil claims than with ordinary specifications, and *"seeking advice from external counsel may be appropriate"* PCG#56bis(26)03, section 4. Its conclusion is a three-level table: requirements with proponent-chosen data, a more limited burden; an existing external dataset referenced for all, medium; a dataset created and maintained by 3GPP, high - with seven framework elements for the last, from admission criteria to terms of use (Table 22) [P03].

INT ESTABLISHED · The paper's contribution to the decision is to make the burden a function of 3GPP's role rather than of the data. The same corpus carries a more limited burden if proponents choose it under 3GPP requirements and a medium one if 3GPP names it for all. That turns the question PCG was handed from common against proprietary into how much responsibility the partnership takes on, which is a question PCG can answer and SA4 cannot. The paper does not recommend a level, and nothing in it should be read as ETSI preferring one.

CTX One citation in the paper is better carried by a different instrument. The General Product Safety Regulation, applicable since 13 December 2024, requires the safety assessment to consider evolving, learning and predictive functionalities; the provisions that ease a claimant's burden of proof, through rebuttable presumptions of defectiveness and causation, are in Article 10 of the revised Product Liability Directive, which member states must transpose by 9 December 2026 [L14] [L13]. The paper's point - that implementers carry exposure a dataset can create - stands on the second instrument; section 8.5 carries both.

ATIS answers a third question: what must be true of any answer, and who checks it.

4.4 ATIS: common legal requirements, verified by the Organisational Partners

SRC ATIS's input is two slides and four statements. The datasets used to train models in the ULBC work should meet at least two legal requirements: *"Available to be used for standardisation, implementation, and deployment purposes with appropriate licensing terms"* PCG#56bis(26)04, slide 2, first requirement, and *"Dataset obtained with transparent and publicly documented privacy details"* PCG#56bis(26)04, slide 2, second requirement. The requirements are to be verified and further developed by *"a group of legal representatives from all OPs"* PCG#56bis(26)04, slide 2, verification that wish to take part, for example ETSI legal and ATIS legal, and that group would also verify that datasets meet them. *"All ULBC candidate codecs must be trained with a dataset that meets the legal requirements."* PCG#56bis(26)04, slide 2 And *"A common dataset could be a means to ensure"* PCG#56bis(26)04, slide 2 that all candidates are trained with a compliant dataset [P04] (Table 26).

INT ESTABLISHED · ATIS makes compliance mandatory and the mechanism optional. The fourth statement offers a common dataset as a means, not the means, so the input cannot be summarised as a mandate for one corpus; the base briefing makes the same caution [B02]. What ATIS does require of every candidate is the outcome - a dataset that meets the requirements - and what it adds that neither other input has is a verifier: a group of OP legal representatives, outside 3GPP's technical structure and above any single OP.

INT WELL SUPPORTED · The first requirement is the widest in scope of any at PCG. It asks for licensing that covers standardisation, implementation and deployment - the whole chain from the meeting room to an operator's network - where TTA speaks of standardisation and commercial implementation and the

legal paper of industry use for implementation. Deployment is the term an operator's counsel reads first, and it is the term that brings the service lifetime, and therefore maintenance, back into a legal requirement.

INT The input is silent on four things the other two address: provenance as a separate requirement, maintenance and long-term availability, a transparency template, and how the work would be resourced. Silence in a two-slide paper is not a position, and this brief does not read it as one; the four points fall to the group ATIS proposes, which is also where ATIS leaves the requirements themselves to be developed.

Three inputs to three questions can be read against one list.

4.5 Sixteen elements, three inputs: where they agree and where they part

CNT **ESTABLISHED** · Read against one list of 16 elements, the three inputs state 10, 10 and 4 of them respectively - TTA, the legal paper and ATIS - and state a further 1, 4 and 2 partly or conditionally (Table 23). Three are stated by all three: legal requirements that any training dataset must meet, licensing that permits standardisation and commercial implementation, and publicly documented privacy information. Eleven are stated by one input only, and five of those by the legal paper alone - incident reporting and remediation after release, Organisational Partner liability and external counsel, an assessment of legal and operational resources, terms and conditions for a dataset 3GPP distributes and application beyond ULBC to other 3GPP AI work.

INT **ESTABLISHED** · The common core is the legal requirement itself, and its two named components are licensing and privacy. That is exactly what SP-260912 calls undisputed, and on this reading the report is accurate: the three inputs differ on how the requirement is met, not on whether it applies. Where they part is on mechanism - a common corpus stated by TTA, offered as one means by ATIS and declined by the legal paper - on who verifies, which only ATIS states, and on scope beyond ULBC, which the legal paper states and TTA leaves to further study.

INT **WELL SUPPORTED** · The inputs are better read as answers to three different questions than as three positions on one line. TTA answers what data should be used; the legal paper answers what each answer would cost 3GPP; ATIS answers what must be true of any answer and who checks it. Figure 6 places them on one five-level scale of intervention in training data, with the legal paper's own three levels set beside it: the status quo at the cutoff sits on the lowest level, TTA's proposal on the fourth, and nothing tabled asks for the fifth.

INT **CONDITIONAL** · Because the answers are to different questions, they can be combined without any input withdrawing: ATIS's requirements and verifier, the legal paper's scale to decide which level 3GPP accepts, and TTA's corpus as one candidate at the level of an external reference. *Signal that would weaken this reading:* a PCG#57 record that adopts one input's mechanism to the exclusion of the others, or that asks the ad hoc group to choose between the inputs rather than to draft requirements.

What the inputs share is clear; what PCG made of them its own minutes now record.

4.6 What PCG#56bis concluded, as reported

SRC No minutes of PCG#56bis were available at the cutoff, and no PCG document records its outcome. The meeting's own register is a second silence rather than a source: it carries a status field, and against all 4 documents that field still reads available, where a document the meeting had disposed of would carry the

disposition [P05] (Table 19). The outcome reaches the reader through two company papers tabled for Madrid. SP-260912 reports that different scenarios were discussed, that *"no consensus was reached on whether a common or a proprietary training dataset approach"* SP-260912, section 1, and that more discussion is needed, for which *"a new PCG ad hoc group has been established"* SP-260912, section 1; it adds that the discussion concluded with a common understanding that establishing the legal requirements is the first and highest priority, so that *"the open question in PCG concerns only how the legal requirements should be satisfied"* SP-260912, section 2 [M03]. SP-260919 reports that no conclusion was reached and that the discussion will continue in a newly created PCG ad hoc group [M05].

INT **WELL SUPPORTED** · The two reports agree on the facts they share - no consensus on the scenario, a new ad hoc group - and differ in what they add. SP-260912 reports a common understanding on priority; SP-260919 reports no conclusion. A common understanding is not a PCG decision, and until a PCG record states it this brief treats it as Samsung's account of the discussion, consistent with the matrix in Table 23 but not established by a record of the meeting itself.

INT **ESTABLISHED** · What the outcome leaves is as important as what it says. There is no requirement text, no verifier, no date, and no statement of the new group's membership or mandate in any document this brief read (Table 77). The TTA request that PCG give guidance before TSG#113 has therefore not been met, and Madrid receives the legal question in the state it left Singapore, with three written analyses attached.

INT **CONDITIONAL** · Where PCG goes next depends on how its new group is framed. *If* the group is asked to draft requirements, the combination section 4.5 describes is available to it and PCG#57 on 21 October could record a first text. *If* it is asked to choose a scenario first, the division between the inputs returns unchanged. *Signals:* the PCG#57 agenda and any report of the ad hoc group - the first public record of its terms of reference.

A report of a meeting is only as good as the claims it carries, which the next section checks.

4.7 Cross-checking the inputs: corpora, precedents and claims

CNT **ESTABLISHED** · TTA's corpus table was accurate for copyright and incomplete for access (Table 21). Common Voice's recordings are CC0-1.0, but since October 2025 the corpus is distributed only through the Mozilla Data Collective, whose terms prohibit re-identifying speakers and re-hosting or re-sharing the dataset; release 26.0 of 17 June 2026 carries 42,388 recorded hours, 28,893 validated, in 294 languages, against the 19,000 or more TTA cites [V01] [V02] [V03]. LibriSpeech remains CC BY 4.0 [V04]. Of the four corpora, 3 carry a single licence in 2026 and 1 does not: the DNS Challenge set is documentation under CC BY 4.0, code under MIT and data under the terms of each component, one of them share-alike [V07]. VCTK carries two licence statements for two versions [V05] [V06].

INT **WELL SUPPORTED** · For TTA's recommended corpus the access terms decide the design. A waiver of copyright settles the recordings; it does not let 3GPP host or pass on a subset, and it waives nothing of the speakers' personal-data rights. Each participant would download the corpus from the platform under its terms, and the specification would name a release, a dataset identifier and checksums rather than carry data. TTA's line that defining a dataset does not mean owning or hosting it is therefore, for this corpus, not a choice but a condition.

CTX The IETF's neural speech-codec work adds a caution from the other side. Its speech-enhancement draft for Opus bars training on the clips used in its tests and discourages training on their source datasets, Common Voice and VCTK [I17]. If ULBC's qualification or selection material were drawn from a corpus also named for training, the separation rule that 20 of 24 survey respondents support would bind the choice of corpus before any legal question did - a point section 6.2 develops.

SRC The precedents check out with qualifications. JPEG AI's call does mandate its training set, describe it as CC0 and require rights for JPEG [K05]; the dataset page, as read, does not print a licence [K06]. JVET's current joint call requires disclosure of training material and access for re-training, not a common set [J01], so TTA's expectation that JVET would follow JPEG AI describes the NNVC exploration rather than the call. The MPEG licence condition is carried here as TTA and the Apex export quote it, because the call document could not be opened [J11] [B03].

SRC Table 25 checks ten statements - from the base briefings, the PCG inputs and the Madrid papers - against the record, and says how this brief carries each. The one with most practical weight concerns the calendar: the base briefing lists an SA4 meeting in Prague on 12 October, whereas SP-260907 schedules online Audio SWG telcos on 14 and 28 October and SA4#138 in Calgary in November [B02] [M01]; this brief follows SP-260907. None of the ten changes a count in this brief.

4.8 Three partners, three institutional views - and four that said nothing

SRC Read as institutions rather than as arguments, the three inputs line up with the three regions that tabled them (Table 24). TTA, the Korean partner, writes the paper that asks 3GPP to act: a common public corpus for due diligence and maintenance, and PCG guidance before Madrid [P02]. The legal paper, sourced by ETISI, writes from the partner that holds the publication machinery and the IPR database, and its subject is the partnership's own exposure: three levels of involvement, and what each would cost in governance, maintenance and remediation [P03]. ATIS, the North American partner, writes the shortest text and the one that moves least: two legal requirements every candidate dataset must meet, and a group of OP legal representatives to develop and verify them [P04].

INT **WELL SUPPORTED** · The three read the same problem through the instrument each body actually has. A partner whose members compete in codec selection is drawn to the fairness of the comparison; a partner that publishes the specification and hosts the declarations is drawn to what the publisher warrants; a partner whose legal culture runs on documented compliance is drawn to a verifiable requirement. None of the three positions is a national legal position, and none of the papers argues from its own jurisdiction's law: the divergence is institutional, which is why the three can be combined without any being withdrawn (Table 23).

CNT **ESTABLISHED** · Four of the seven partners tabled nothing. ARIB and TTC in Japan, CCSA in China and TSDSI in India are named in no document this brief read on this question, and the register records no contribution from them [P05]. Nothing follows from silence: the ad hoc was called at a day's notice, and a partner may hold a view it has not written down. But the three written positions are the only ones on the record, and a rule written from them would be written without the partners whose members supply much of the world's device volume.

CTX Their national settings do differ, even where their papers do not invoke them. The European partners work under a mining exception with a reservation mechanism and the strictest personal-data regime; the North American partner under fair use decided case by case and a patchwork of state biometric statutes; the Korean partner under a fair-use guide published in February 2026; and the partners that said nothing under Japan's non-enjoyment exception, China's rules for public generative services and India's unsettled consultation (Table 41, Table 42). A common corpus assembled once has to be lawful in all of them at once, which is the practical reason the legal paper puts governance before scenario.

INT **CONDITIONAL** · *Signal*: whether the new PCG group's membership is drawn from all seven partners or from the three that wrote. *If* it is drawn from all seven, the requirement text has to survive four more legal environments and will be written more abstractly; *if* from the three, it will be more concrete and less portable. *What would weaken this reading*: a contribution from a silent partner at PCG#57, which would show the silence was procedural rather than positional.

What this part established

PCG#56bis read three inputs that answer different questions: TTA what data, the legal paper what burden for 3GPP, ATIS what any answer must meet and who checks it.

They share three elements - the legal requirement, licensing and privacy - and part on mechanism, verification and scope; the legal paper alone carries liability, resources, remediation and terms of use (Table 23).

PCG reached no consensus and set up a new group, as two company papers report; there are no minutes. TTA's corpus table holds for copyright and not for access: Common Voice is CC0 but may not be re-shared (Table 21).

Read as institutions, the three inputs track the instrument each partner holds rather than the law of its region; four of the seven partners tabled nothing at all (Table 24).

What this part established — part 4

Hands forward: what SA4 is building, and what Madrid was asked - part 5.

Table 18. PCG#56 bis_Ad Hoc: the agenda and the documents

Document	Source	Where the source sits	For	What it contains	Read
PCG#56bis(26)01	PCG Chair	Europe (ETSI as secretariat)	decision	The agenda: opening; approval with an antitrust and competition law statement; item 3, 3GPP involvement and legal requirements for AI/ML training datasets; AOB; closing	read in full
PCG#56bis(26)02	TTA	Korea	decision	Problem statement; practice in other bodies; risks and legal requirements; a common dataset for due diligence and maintenance; four candidate corpora; three proposals	read in full
PCG#56bis(26)03	ETSI, as the meeting register records it; no Source line on the document itself	Europe	not stated	Two scenarios; third-party rights and regulation; governance of an existing or a 3GPP dataset; OP liability; a three-level conclusion	read in full
PCG#56bis(26)04	ATIS	United States	not stated	Two legal requirements; a group of OP legal representatives to develop and verify them; compliance for every candidate; a common dataset as one means	read in full
TDoc list	meeting folder, 7 September 2026	-	-	The meeting's document register, posted as a spreadsheet (Table 19)	read in full

What PCG read on 3 September, as the documents themselves number and title them, with the region each source speaks from. Public source: the meeting folder at https://www.3gpp.org/ftp/PCG/PCG_56_Bis/Docs [G02]. The meeting's own register is the next table.

Table 19. The PCG#56 bis_Ad Hoc register: what the meeting recorded about its own documents

Document	Title as the register records it	Source	Type	Agenda item	Status	Number reserved	File uploaded
PCG56Bis_01	Agenda	ETSI	agenda	2 Approval of Agenda	available	2026-09-02 12:44	2026-09-02 12:52
PCG56Bis_02	Issues with AI model training datasets within 3GPP (ref: ULBC codec) & Way Forward	TTA	discussion	3 3GPP involvement and legal requirements for AI/ML training datasets	available	2026-09-02 12:44	2026-09-02 12:52
PCG56Bis_03	Legal Considerations Relating to Training Datasets for AI models	ETSI	discussion	3 3GPP involvement and legal requirements for AI/ML training datasets	available	2026-09-02 12:50	2026-09-02 12:52
PCG56Bis_04	ATIS recommendations on legal requirements for AI/ML training datasets	ATIS	discussion	3 3GPP involvement and legal requirements for AI/ML training datasets	available	2026-09-02 14:03	2026-09-03 11:02

The meeting's own document register, read from the TDoc list posted with the documents. Public source: https://www.3gpp.org/ftp/PCG/PCG_56_Bis/Docs [G02]. It sources every document, times the meeting's assembly to the minute, and records no disposition against any of the four.

Table 20. TTA's input, point by point

Point	What TTA states	Asks of	This brief's reading
Problem	Standardised weights embed the outcome of training; six risk dimensions	-	Holds for a codec with normative weights (section 6.1)
Governance gap	No 3GPP framework for AI training data; a self-declaration does not allocate responsibility	PCG	Shared by the legal paper
Maintenance	Undisclosed proprietary data makes retraining depend on one company	3GPP and the OPs	Holds for two of five maintenance occasions (Table 33)
Precedent	JPEG AI's mandate; JVET NNVC expected to follow; MPEG licence conditions	-	JPEG AI confirmed; JVET's 2026 call chose disclosure (Part 6)
Corpora	Four public corpora; Common Voice recommended for CC0 and privacy framework	SA4	Copyright holds; access terms differ (Table 21)
Legal requirements	Licensing scope; description and provenance; privacy documentation; rights for 3GPP; a transparency declaration like the EU template	PCG	Five of the fourteen declaration fields (Table 49)
Proposal 1	PCG guidance before TSG#113	PCG	Not met: no conclusion at PCG#56bis
Proposal 2	A common training dataset under CC0 or CC BY 4.0 with documented privacy; SA4 selects it	SA4	The fourth level of Figure 6
Proposal 3	PCG to keep studying other AI/ML functionality	PCG	The precedent question (section 10.4)

For a reader who wants TTA's argument without its layout. The last column is this brief's reading, with the section or exhibit that checks it.

Table 21. The four corpora TTA names, checked against their 2026 terms

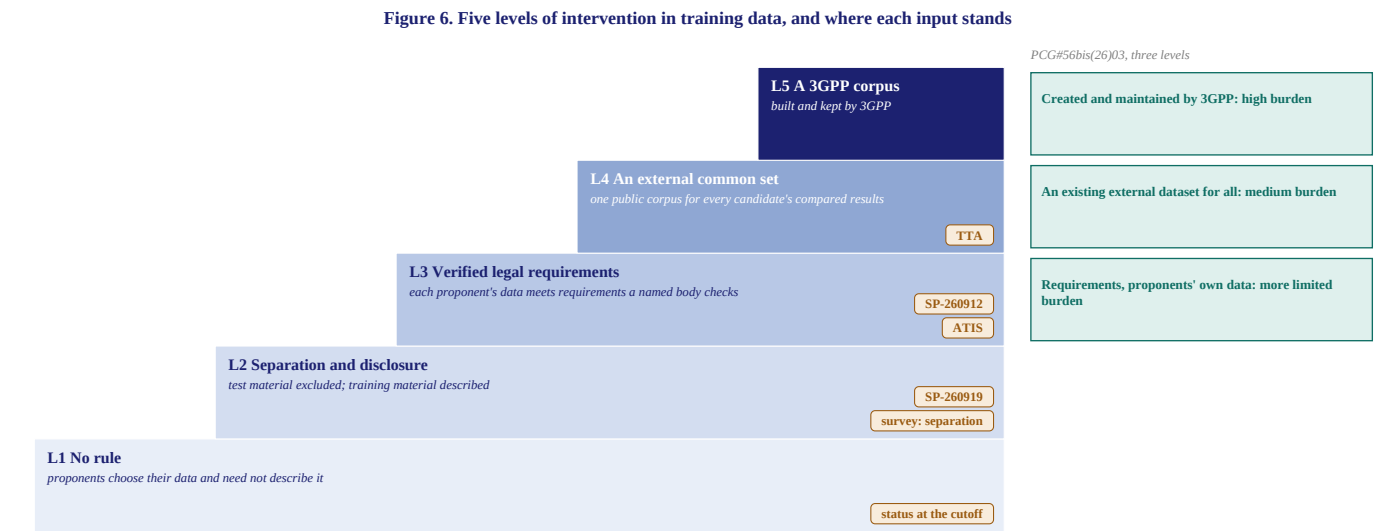
Corpus	As TTA lists it	Licence in 2026	Terms and release in 2026	What follows	Sources
Mozilla Common Voice	19,000+ h; 8-48 kHz; CC0	CC0-1.0 on the recordings	Distributed only through the Mozilla Data Collective since October 2025; its terms prohibit re-identifying speakers and re-hosting or re-sharing the dataset. Corpus 26.0 (17 June 2026): 42,388 h recorded, 28,893 h validated, 294 languages	Copyright is waived; access is contractual, so a participant downloads it rather than receiving a copy from 3GPP	[V01][V02][V03]
LibriSpeech (OpenSLR SLR12)	about 1,000 h; 16 kHz; CC BY 4.0	CC BY 4.0	Attribution; derived from LibriVox public-domain audiobooks. unchanged since 2015	Read English only; attribution travels with any redistributed subset	[V04]
CSTR VCTK	about 44 h; 48 kHz; CC BY 4.0	CC BY 4.0 on the CSTR card for v0.92; ODC-By 1.0 in the DNS Challenge listing	Attribution. v0.92 (13 November 2019), 110 speakers	Two licence statements for two versions; the version selected has to be named	[V05][V06]
ICASSP 2022 DNS Challenge 4	about 1,800 h; 48 kHz; CC BY 4.0	Not one licence: documentation CC BY 4.0, code MIT, data under the terms of each component	Components include CC BY-SA 3.0 (DEMAND), a custom licence (Edinburgh 56-speaker), ODbL and DbCL, Apache 2.0 and YouTube-derived material (VoxCeleb2). as listed in the repository README	A single licence label understates a compound set of terms, one of them share-alike	[V07]

For curators, counsel and anyone choosing a corpus. TTA's figures are as tabled; the 2026 columns are from the distributors' own pages. Copyright terms hold; access and component terms change what 3GPP could do with each.

Table 22. The legal paper: risks, levels of involvement and the framework for a 3GPP dataset

Item	What the paper says	Burden or scope
Six classes of third-party risk		
Copyright in underlying recordings	named	all scenarios
Performer rights	named	all scenarios
Database rights	named	all scenarios
Confidentiality obligations	named, e.g. dataset know-how	all scenarios
Contractual licensing restrictions	named	all scenarios
Personal-data protection (voice)	named	all scenarios
Three levels of 3GPP involvement		
Requirements, each proponent's own dataset	qualification and documentation; proponents fully responsible	more limited
An existing external dataset referenced for all	qualification and documentation; 3GPP takes some responsibility for the choice	medium
A dataset created and maintained by 3GPP	full governance, maintenance and remediation	high
Seven framework elements, for a 3GPP dataset		
Admission criteria	a data contribution package: provenance, rights, licences, personal data	highest level
Standardised documentation	datasheets or transparency records	highest level
Review process	technical and legal assessment before acceptance	highest level
Versioning and change management	traceable updates and corrections	highest level
Incident reporting and remediation	comparable to coordinated vulnerability disclosure	highest level
Maintenance responsibilities	a named entity to maintain and address reports	highest level
Terms and conditions	for use, distribution and maintenance	highest level

The paper's structure in one place, for counsel and PCG delegates. The burden column is the paper's own scale; the paper recommends no level.



The five levels are this brief's ordinal scale; step height carries order only. The right-hand column sets the legal paper's own three levels against it. Placements follow Tables 20, 23 and 26; nothing tabled asks for the fifth level.

Where the tabled positions sit on one scale of intervention, for a reader comparing them. The status quo is the lowest level; TTA's proposal is the fourth; the legal paper's three levels map onto the third, fourth and fifth.

Table 23. Sixteen elements across the three OP inputs

Element	TTA	Legal paper	ATIS	Basis
A common training dataset for ULBC	stated	not stated	partly	TTA proposes it; the legal paper declines to choose; ATIS calls it one possible means
Legal requirements that any training dataset must meet	stated	stated	stated	TTA lists five; the legal paper calls due diligence an industry requirement; ATIS lists two
Licensing that permits standardisation and commercial implementation	stated	stated	stated	CC0 or CC BY 4.0 (TTA); terms permitting industry use (legal); standardisation, implementation and deployment (ATIS)
Publicly documented privacy information	stated	stated	stated	All three name personal data or privacy documentation
Source and provenance documentation	stated	stated	not stated	TTA and the legal paper; ATIS does not name provenance separately
Rights for 3GPP to use the material in standardisation	stated	partly	partly	Explicit in TTA; implied by the authorisation principle (legal) and by standardisation use (ATIS)
A transparency declaration modelled on an existing template	stated	partly	not stated	TTA points to the EU AI Act template; the legal paper to regulators' templates as reference points
A named body to verify compliance	not stated	partly	stated	ATIS proposes a group of OP legal representatives; the legal paper a review process for a 3GPP dataset
Maintenance and long-term availability	stated	stated	not stated	TTA frames it as dependency on one company; the legal paper as lifecycle and maintenance arrangements
Incident reporting and remediation after release	not stated	stated	not stated	The legal paper's mechanism comparable to coordinated vulnerability disclosure
Organisational Partner liability and external counsel	not stated	stated	not stated	Only the legal paper
An assessment of legal and operational resources	not stated	stated	not stated	Only the legal paper
Terms and conditions for a dataset 3GPP distributes	not stated	stated	not stated	Only the legal paper
Technical selection kept with SA WG4	stated	partly	not stated	TTA explicitly; the legal paper leaves the options to 3GPP
Application beyond ULBC to other 3GPP AI work	partly	stated	not stated	TTA asks PCG to continue studies; the legal paper names SA4 the first use case
Guidance before TSG SA#113	stated	not stated	not stated	Only TTA sets a date
Stated, of 16	10	10	4	Partial readings are printed and never added

The matrix behind [section 4.5](#). A cell is stated, partly or not stated, read from the input's text, and the basis of every cell is printed. Only stated cells are counted in the final row; a partial reading is never added to them.

Table 24. The three inputs read as institutional positions, and the four partners that tabled nothing

Partner	Setting, as this brief reads it	What its paper optimises for	Mechanism proposed	What it does not ask for	Where the position is stated
TTA	Korea; a fair-use guide published February 2026 and a fair-use analysis decided case by case [A10]	comparability of the competition, and the ability to maintain the winner	one common public corpus, chosen from four named candidates, plus PCG guidance before Madrid	a verifier, or any allocation of liability	PCG#56bis(26)02, sections 2 to 5 [P02]
ETSI	Europe; a mining exception with a machine-readable reservation, and the strictest personal-data regime [L11] [L10]	what the partnership itself would warrant, host and answer for	three levels of 3GPP involvement, each with its governance burden priced	a corpus, or a decision between the levels	PCG#56bis(26)03, sections 1 to 4 [P03]
ATIS	North America; fair use decided case by case and state biometric statutes [U02] [U05]	a requirement that can be verified, whatever dataset is used	two legal requirements for every candidate dataset, verified by a group of OP legal representatives	a common corpus, or scope beyond ULBC	PCG#56bis(26)04, slides 1 and 2 [P04]
ARIB, TTC	Japan; a non-enjoyment exception with limits on exploiting datasets sold for training [A05]	-	nothing tabled	-	no document read; no entry in the register [P05]
CCSA	China; lawful sources and IP respect for public generative services, and a training-data security standard [A11] [A12]	-	nothing tabled	-	no document read; no entry in the register [P05]
TSDSI	India; no dedicated exception and a working paper proposing a blanket licence [A17]	-	nothing tabled	-	no document read; no entry in the register [P05]

Who wants what, and why each partner's instrument shapes its answer. The settings column is this brief's reading of each partner's legal environment, not a position any paper argues from.

Table 25. Statements checked against the record, and how this brief carries each

Statement	Where it is made	What the record shows	How this brief carries it	Sources
An SA4 meeting in Prague on 12 October 2026	base briefing	SP-260907 schedules Audio SWG telcos on 14 and 28 October and SA4#138 in Calgary in November	follows SP-260907	[B02] [M01]
The Korean fair-use guide for AI training, May 2026	base briefing	Korean original 26 February 2026; English edition 18 May 2026	gives both dates	[B02] [A10]
PCG#56bis(26)03 from ETSI and CCSA	base briefing, as meeting context	no Source line on the document; SP-260912 and SP-260919 attribute it to ETSI	attributed to ETSI as the Madrid papers do, and to no other OP	[B02] [M03] [M05]
The GPSR reverses the burden of proof on manufacturers	PCG#56bis(26)03	the GPSR is a safety regulation; rebuttable presumptions are in the revised Product Liability Directive, Article 10	carried under the Directive	[P03] [L14] [L13]
S4-261712 reported to this meeting	SP-260919	the SA WG4 Chair's report reached the plenary as SP-260909	cited as SP-260909	[M05] [M02] [W03]
Common Voice, 19,000 hours or more, CC0	PCG#56bis(26)02	release 26.0 has 42,388 recorded hours; CC0-1.0; distributed only through the Mozilla Data Collective, whose terms bar re-sharing	both carried; the access terms added	[P02] [V01] [V02] [V03]
CSTR VCTK under CC BY 4.0	PCG#56bis(26)02	CC BY 4.0 on the CSTR card for v0.92; ODC-By 1.0 in the DNS Challenge listing	the version to be named with the licence	[P02] [V05] [V06]
JVET NNVC is expected to follow the JPEG AI approach	PCG#56bis(26)02	NNVC compares tools on common sets; the 2026 joint call requires disclosure and access, not a common set	carried as TTA's characterisation of NNVC	[P02] [J05] [J01]
MPEG VCM requires public training data under commercial-use licences	PCG#56bis(26)02; Apex export	the call document could not be opened for this brief	carried as the two sources quote it	[P02] [J11] [B03]
The status report numbered SP-2606907	SP-260907, title slide	the file name and the meeting folder give SP-260907	cited as SP-260907	[M01] [G04]

A reconciliation register, for anyone comparing this brief with the base briefings or the inputs. None of the entries changes a count in this brief.

Table 26. ATIS's four statements, and how each can be read

Statement	As tabled	What it requires	What it leaves open
Two legal requirements	<i>"Available to be used for standardisation, implementation, and deployment purposes with appropriate licensing terms"; "Dataset obtained with transparent and publicly documented privacy details"</i>	Licensing through deployment; published privacy details	Provenance, maintenance and transparency format
A verifier	<i>"a group of legal representatives from all OPs"</i>	Requirements developed and checked by OP lawyers who wish to take part	Membership, resources, timetable
Compliance for all candidates	<i>"All ULBC candidate codecs must be trained with a dataset that meets the legal requirements."</i>	Every candidate trained on a compliant dataset	How compliance is shown before selection
The mechanism	<i>"A common dataset could be a means to ensure"</i> compliance	A common dataset offered as one means, not required	Which other means would qualify

Two slides read line by line. The wording is ATIS's; the last two columns are this brief's reading of what each statement requires and leaves open.

5 What SA4 is building, and what Madrid is asked

For: SA and SA4 delegates, proponents.

5.1 Eleven Pdocs, twenty declared candidates and a thin margin

SRC The competition opened with twenty declared candidates. A call for interest on the SA4 reflector from 15 to 19 June drew declarations from twenty companies: *"A total of 20 companies declared their interest in submitting a ULBC candidate"* SP-260907, slide 2 [M01]. SP-260907's speaker notes name them - Apple, Bytedance, China Mobile, China Telecom, Dolby, Ericsson, Fraunhofer, Huawei, Lenovo/Motorola, Nokia, NTT, Orange, Panasonic, Philips, Qualcomm, Samsung, Spreadtrum, vivo, VoiceAge and Xiaomi. With no agenda item available before SA4#137-e, the Audio SWG ran six offline calls between 24 June and 12 August on the organisation of the work, the SA request on training data and the preparation of the e-meeting.

CNT ESTABLISHED · The offline calls drew between 22 and 49 participants, 230 attendances over the six calls, a mean of 38 per call. SA4#137-e then agreed the structure of the work as eleven Pdocs with assigned editors and target dates: 2 due at SA4#137-e, 6 by SA4#138 in November, with the September ad hoc as the first target for the three marked top priority, and 3 at SA4#139 in February (Table 27). The rapporteur's maturity readings for the agreed Pdocs run from 3% to 95% (Table 27).

INT ESTABLISHED · What the plan places in March is worth stating precisely, because it is often compressed in summary. SP-260907 launches selection testing on 19 March 2027; a winning candidate is scheduled for SA4#140 in May and the specifications for TSG SA#116 in June [M01]. A reading that treats March as the date by which a codec must be defined or selected is therefore reading the plan differently from the document that sets it out. The distinction changes what an unresolved rule delays: the qualification and selection Pdocs bind the competition that starts in March, so an objection still open at SA4#139 reaches the March launch, while one opened after it reaches the May and June dates instead (Table 3, section 11.1).

INT WELL SUPPORTED · A maturity percentage does not measure the consensus work that remains. A Pdoc can be technically mature and still lack agreement on a disputed rule, and an apparently narrow process question can hold a document open. The remaining calendar therefore cannot be read off the completion figures alone (section 11.1).

SRC The meeting also recorded agreements that bound the candidates: *"Bitrates: 1.0 kbps and 1.4 kbps, and only these two"* SP-260907, slide 5; different input signal

levels for qualification and selection and background noise types and signal-to-noise ratios for qualification; a total cost estimate with cost sharing for test funding; and a reply to SA WG2 on bundling assumptions, S4-261695 [M01]. SP-260907 notes that a direct selection test of all candidates is being studied, which would drop the qualification phase and save testing resources, because the number of proposed qualification conditions keeps rising.

INT ESTABLISHED · The work plan has no slot for training data, and none of the Pdocs is about it. A Pdoc on training-data requirements for candidate submission was proposed at SA4#137-e in S4-261509 by Samsung and Viasat, as the base briefing records, and is not in the agreed set [W09] [B02]. The editors of the Pdocs are assignments made by the working group; this brief treats them as roles, counts no editorship as an organisation's activity, and prints the editors only because SP-260907 does.

INT CONDITIONAL · If a training-data rule is adopted, the structure already says where it would land. A condition on how candidates are trained belongs in the qualification and selection rules, ULBC-5a and ULBC-5b; requirement dimensions of the kind SP-260919 proposes belong in the performance requirements, ULBC-3; and what a winner must hand over - weights, reference software, perhaps a description of training - belongs in the deliverables, ULBC-6, which stands at 5%. *Signal*: a contribution to any of the four at the Shanghai ad hoc that names training material.

The Pdocs fix the procedure; the open design issues fix what any corpus would have to serve.

5.2 The open design issues, and which of them a corpus must serve

SRC SP-260907 lists nine technical issues still open after SA4#137-e: the sampling rate and audio bandwidth, 16 or 32 kHz with the 24 kHz option removed; the frame length, 20, 40 or 80 ms; bundling and packetisation for 80 and 160 ms bundles; how frame length couples to bit rate and packetisation; algorithmic delay and lookahead, with proposals from 0-5 ms to 80 ms; complexity and memory, where *"the metrics themselves still need to be defined"* SP-260907, slide 6; whether there is a qualification phase; the RTP payload format; and jitter buffer management [M01] (Table 28).

INT WELL SUPPORTED · The sampling-rate decision alone determines which corpora could serve. A super-wideband mode at 32 kHz cannot be trained on material recorded at 16 kHz, which is LibriSpeech's rate; Common Voice varies from 8 to 48 kHz by contributor, and VCTK and the DNS set are 48 kHz [P02]. A common-corpus rule written before the rate is fixed would either constrain the rate decision or name a corpus that part of the design could not use.

INT WELL SUPPORTED · The testing agreements reach into training as well. Qualification is to use background noise types and signal-to-noise ratios as test conditions; a rule keeping test material out of training therefore covers noise recordings, and any room or channel material, as well as speech. The DNS set that TTA lists is itself a noise and speech-enhancement corpus, so a candidate trained on it and tested with noise drawn from the same components would breach the separation that the survey supports almost without exception.

INT The open complexity and memory metrics leave the size of the normative weights open too - and with it the size of what 3GPP would publish, host by reference and maintain. Frame length and delay shape the architecture more than the data. Read together, the open issues show that a technical training-data requirement is not a free-standing decision: its content depends on design constraints and test conditions that ULBC-3, ULBC-4 and ULBC-8a have not yet fixed.

INT CONDITIONAL · That makes SA4's answer a statement about timing as much as substance. *If* the design constraints settle at the Shanghai ad hoc or at Calgary, the question of technical training-data requirements can be asked again against fixed parameters, as SP-260912 invites. *If* they do not, any requirement agreed earlier would have to be written as conditional on them. *Signal*: the sampling-rate decision in ULBC-4.

Against that unsettled design, two companies ask the plenary to decide.

5.3 The sequencing proposal: hold the selection rules for PCG

SRC Samsung's SP-260912, for discussion and endorsement, makes three observations. The legal issues are agreed at PCG to be of the highest priority and must be resolved. Implementers of an AI-based ULBC must comply with regulatory frameworks that govern training data, so a codec whose data cannot be shown to meet the eventual requirements risks not being implementable in practice. And proceeding to qualification and selection before PCG concludes risks invalidating the results and would cost more delay than sequencing the work now [M03].

SRC The paper quotes SA4's own survey against SA4's schedule: legal and licensing aspects *"are expected to be handled at the appropriate level outside SA4"* SP-260912, section 2, quoting S4-261382, so SA4 cannot resolve the dependency alone [W01]. It notes that ULBC is the first 3GPP case of a specified AI/ML model, so *"how this dependency is handled will set the precedent for future 3GPP work involving AI"* SP-260912, section 2. Its proposals: *"should proceed only after PCG has reached a conclusion on the legal issues"* SP-260912, section 3, proposal, and *"SA4 is encouraged to re-evaluate the need for a common dataset"* SP-260912, section 3, proposal and the associated technical requirements, taking the PCG discussion into account [M03].

INT ESTABLISHED · The documents the first proposal would hold are identifiable. The Pdocs governing qualification and selection are the rules and test plans, of which none has an agreed version, and the performance requirements and design constraints, both at 15% and marked top priority. Under the plan in SP-260907 their finalisation falls at SA4#138 in November; PCG#57 falls 36 days after Madrid opens, on 21 October. If PCG concludes at PCG#57, the public timetable leaves room for the affected Pdocs to move without changing the March selection-test date; if it does not, the dependency moves directly into the January contracting and March selection sequence. This is a timing analysis and not a feasibility one: it identifies the points at which an unresolved decision would begin to move the recorded plan, and establishes neither that the schedule is achievable nor that it is not.

INT WELL SUPPORTED · The paper's central risk is conditional in a way its wording compresses. Selection results are invalidated only if a candidate's training data later fails requirements not yet written. That risk could be addressed by sequencing, as proposed, or by disclosure and interim criteria applied from the start, as the base briefing's draft PCG wording suggests when it lets SA4 continue while applying any interim criteria the OPs agree [B02]. Both readings are consistent with the three observations; the paper argues for the first and does not discuss the second.

INT CONDITIONAL · If Madrid endorses the proposal, SA4 receives a dependency on a body whose next meeting is online and whose ad hoc group has no published terms, and the Rel-20 margin SP-260907 already calls thin becomes PCG's to spend. If it notes the paper, SA4 proceeds and the legal risk is carried to selection. *Signals:* the disposition of SP-260912 in the TSG SA#113 report, and whether a liaison from Madrid mentions PCG.

Samsung asks Madrid to decide the order of work; Verizon asks it to decide content.

5.4 The transparency proposal: requirement dimensions and maintenance

SRC Verizon's SP-260919 argues from the operator's side of the service lifetime. Operators deploy voice services for decades and rely on 3GPP, not on any single company, to maintain what they deploy; for every codec so far everything needed to maintain it has been

in the specification. Citing TTA's PCG input, it concludes that *"non-disclosed proprietary training data carries a structural maintenance risk for the specification itself"* SP-260919, section 2, whatever a candidate's technical merits [M05].

SRC The paper reads the survey as a mandate for technical work. It cites the 13 to 9 split with 2 without a position, the reasons in favour - fair comparison, robustness, bias avoidance and reproducibility - and the reasons against, timing and schedule and no technical need, and concludes that *"the principal reservation among those not in favor was not a technical objection"* SP-260919, section 2. It adds the annex figures on test material and common data (section 3.5). Its proposals: SA4 should define requirements, for example on *"the number of languages / language families, speaker demographics, acoustic environments"* SP-260919, section 3, proposal, so that candidates can be verified to work across them; *"The codec training dataset should be transparent such that 3GPP can maintain its specifications"* SP-260919, section 3, proposal; and a liaison to SA4 should carry both [M05].

INT WELL SUPPORTED · The two proposals are rules of different kinds. The first is a requirement on the codec - how it must perform across languages, speakers and environments - which can be verified in selection tests without reference to training data at all; it is close to what Dolby's survey reply argued on the reflector, that robustness is a codec requirement to be enforced in testing [W05]. The second is a requirement on the data, justified by maintenance. The first sits inside SA4's remit as SP-260909 drew it; the second is the governance question PCG holds.

INT ESTABLISHED · 'Transparent' is not defined in the paper, and the choice of definition is most of the decision. It could mean disclosure of what was used, publication of the data, access for verification or retraining, or deposit with 3GPP; Table 29 sets the readings beside the two Madrid papers. JVET's joint call already operates one of them - mandatory reporting of training material, and retraining with alternative material available to JVET where rights cannot be granted - which section 7.2 reads.

INT CONDITIONAL · If Madrid sends the liaison, SA4 receives requirement dimensions it can test and a transparency principle it cannot implement without PCG, and the two tracks of Singapore meet in one document for the first time. If it does not, the

requirement dimensions can still reach ULBC-3 as contributions. *Signals:* an outgoing liaison in the TSG SA#113 record, and its wording on 'transparent'.

What this part established

SA4 works through eleven Pdocs with twenty declared candidates; four selection Pdocs are unwritten, none concerns training data, and the Chair calls the margin thin (Table 27).

The open design issues, sampling rate first, decide what any corpus would have to serve (Table 28).

SP-260912 asks Madrid to hold the selection rules for PCG; SP-260919 asks it to send requirement dimensions and a transparency principle to SA4. The first changes SA4's calendar, the second its content (Table 29).

Neither Madrid paper asks the plenary to choose a dataset: one asks for an order of work and the other for dimensions and a principle, and both leave the corpus question with PCG (Table 29).

The margin is thin either way. Qualification testing is timetabled for the quarter after Madrid, and the four Pdocs with no agreed version are the ones a training-data rule would have to enter (Table 27).

Hands forward: the technical question beneath the legal one - part 6.

Table 27. The eleven Pdocs as the status report prints them

Pdoc	Title	Editor, as assigned	Target	Version	TDoc	Maturity	Note
ULBC-1	Overview	Ericsson	SA4#137-e	0.1.0	S4-261687	95%	
ULBC-2	Project plan and timeline	vivo	SA4#137-e	0.1.0	S4-261745	80%	
ULBC-3	Performance requirements	Dolby	Audio SWG ad hoc (Sep 2026), latest SA4#138	0.1.0	S4-261771	15%	top priority
ULBC-4	Design constraints	Fraunhofer	Audio SWG ad hoc (Sep 2026), latest SA4#138	0.2.1	S4-261847	15%	top priority
ULBC-5a	Qualification rules	Nokia	Audio SWG ad hoc (Sep 2026), latest SA4#138	-	-	-	no agreed version; top priority if there is a qualification phase
ULBC-5b	Selection rules	Orange	SA4#139 (Feb 2027)	-	-	-	no agreed version
ULBC-6	Deliverables	Huawei	SA4#138 (Nov 2026)	0.2.0	S4-261729	5%	
ULBC-7	Processing plan	Ericsson	SA4#138 (Nov 2026)	0.1.0	S4-261728	20%	
ULBC-8a	Qualification test plan	VoiceAge	SA4#138 (Nov 2026)	-	-	-	no agreed version
ULBC-8b	Selection test plan	VoiceAge	SA4#139 (Feb 2027)	-	-	-	no agreed version
ULBC-9	RTP header compression and system integration	CMCC	SA4#139 (Feb 2027)	0.0.1	S4-261722	3%	

The rules of the competition and their state. Editors are assignments by the working group and are not counted as any organisation's activity. None of the eleven concerns training data.

Table 28. SA4's open technical issues, and which of them bear on training data

Issue	As the status report records it	Bears on training data	This brief's reading
Sampling rate and bandwidth	16 or 32 kHz; 24 kHz removed; single rate open	yes	Decides which corpora can serve: a 16 kHz corpus cannot train a 32 kHz mode
Frame length	20, 40 or 80 ms, as candidate options	little	Shapes the architecture more than the data
Bundling and packetisation	80 and 160 ms bundles; the interface still open	no	A transport question
Frame length coupling	to bit rate and packetisation, unresolved	no	A design question
Algorithmic delay and lookahead	from 0-5 ms to 80 ms; maximum or fixed disputed	little	Bounds the model's context, not its corpus
Complexity and memory	no convergence; metrics undefined	indirect	Sets the size of the normative weights 3GPP would publish and maintain
Qualification phase	keep or drop; direct 20-candidate selection studied	indirect	Where a training-data condition would be checked
RTP payload format	payload header, DTMF over non-IP, header reduction	no	A transport question
Jitter buffer management	proposal to remove it from design constraints	no	A system question
Test conditions (agreed)	input levels; noise types and SNRs for qualification	yes	Separation must cover noise and room material as well as speech

The issues SP-260907 lists as open, plus the agreed test conditions. The third and fourth columns are this brief's reading of how each would constrain a corpus or a data rule.

Table 29. The two Madrid proposals side by side, and four readings of 'transparent'

	Sequencing paper / reading	Transparency paper / what follows
The two papers		
Source and status	Samsung; for discussion and endorsement	Verizon; for discussion and endorsement
Premise	Legal issues are the highest priority and must be resolved before selection	Undisclosed proprietary data is a structural maintenance risk
Reads the survey as	SA4 itself put legal aspects outside its scope	A majority for technical work; timing as the main reservation
Asks Madrid to	Hold qualification and selection Pdocs for PCG; SA4 to re-evaluate a common dataset	Send requirement dimensions and transparency to SA4 by liaison
Changes	SA4's calendar	SA4's content
Depends on PCG for	the timing of SA4's rules	the meaning of 'transparent'
Four readings of 'transparent'		
Disclosure of what was used	a description in the proposal, as JVET's template asks	lowest burden; no maintenance gain
Access for verification or retraining	as JVET's joint call expects	maintenance gain where rights allow; the fallback is retraining
Deposit with 3GPP or a trustee	escrow of data and recipe for the selected candidate	maintenance gain; contract terms needed
Publication of the data	the corpus made public	highest burden; excludes most proprietary data

The two papers that ask Madrid for decisions - the sequencing paper, SP-260912, and the transparency paper, SP-260919 - set against each other, then the four meanings 'transparent' could take. The readings are this brief's; neither paper defines the word.

Table 30. Where each organisation appears in the dataset record

Organisation	Documents in which the record names it	What the record shows
Samsung	SP-260609 (co-source); SP-260611 (co-source); SP-260641, SP-260671; S4-261500, S4-261501, S4-261509, S4-261510, S4-261511, S4-261703; SP-260912; SP-260913; reflector messages 2, 4 and 8	Guidance at SA#112; common-data process proposals at SA4#137-e; sequencing proposal at SA#113
Verizon	SP-260257, SP-260283, SP-260258 (SA#111); SP-260609 (co-source); SP-260919	Rel-20 readiness at SA#111; principles at SA#112; transparency for maintenance at SA#113
KT Corp.	SP-260609 (co-source)	Principles at SA#112
Qualcomm	S4-261500, S4-261501, S4-261511 (co-source); asked at SA#112 that SA4 time be limited	Common-data process proposals
Viasat	SP-260619 (co-source); S4-261509, S4-261510 (co-source)	Bitrate scope; training-data requirements and common data
Dolby	SP-260619, SP-260623 (co-source); S4-261644; reflector messages 5 and 7	Counter-principles at SA#112; survey reply against SA4 requirements or common data
Nokia	SP-260619, SP-260623 (co-source); S4-261572; reflector message 11; comment at SA#111	Counter-principles; survey reply; earlier audio practice
Ericsson	SP-260623 (co-source); reflector message 3	Counter-principles; technical arguments against common data
Fraunhofer IIS	SP-260623 (co-source)	Counter-principles
Orange	S4-261643	Legal arrangements for selection testing
China Mobile	SP-260611 (co-source, as CMCC)	Alternative work-item text at SA#112; SP-260627 was filed by the study's rapporteur and is not credited to the organisation
vivo	SP-260602	Considerations and proposals at SA#112
TTA	PCG#56bis(26)02	Organisational Partner input: common public dataset
ATIS	PCG#56bis(26)04	Organisational Partner input: legal requirements and OP legal group
ETSI (attributed)	PCG#56bis(26)03	Organisational Partner input: legal and governance analysis

Organisations as the record names them - as source, co-source, editor or sender. Role documents, such as a study rapporteur's, are not credited to the organisation. A presence here is not a count of influence or a position score.

6 The technical question beneath the legal one

For: codec engineers, operators, test laboratories.

6.1 What is standardised when the weights are normative

SRC The work item approved at Singapore is for a codec whose trained parameters are part of the standard [S07] [G01]. TTA's input describes the standardised trained weights as embedding the outcome of training into the codec definition, and SP-260919 speaks of an AI-based codec with standardised trained model parameters [P02] [M05]. What exactly the winner will deliver is not yet written: the deliverables Pdoc, ULBC-6, has an agreed first version and stands at 5% [M01].

CTX Two other bodies have faced the same step. ITU-T T.840.1, the JPEG AI core coding system published by ISO/IEC as 6048-1, treats its trainable model parameters as an integral part of the Recommendation and International Standard, held outside the text because of their size [K07] [K08]. The IETF's Deep Audio Redundancy draft for Opus makes the decoder and its weights normative, distributes the weights outside the document, and leaves the encoder and the training non-normative [I15] [I16]. Neither body standardised its training data; both standardised the result of training.

INT ESTABLISHED · Normative weights raise a conformance question before they raise a data question. A classical decoder is verified against its bitstream and reference code; a neural decoder must also say whether implementations are to match the reference output bit for bit, which fixes the arithmetic, or within a tolerance, which lets implementations differ. JVET's neural-network ad hoc group carries a mandate to *"Investigate bit-exact reproducibility of NN-based methods on various platforms"* JVET-AQ0011, AHG11 mandates [J03] – the same question, asked for video, and not yet closed there.

INT ESTABLISHED · Fixed-model reproducibility is an implementation question as well as a specification one. The same model definition and the same weight values do not by themselves guarantee bit-identical behaviour across inference libraries, numerical precision, hardware and execution environment. A standard therefore has to decide what is normative: the mathematical model, a reference implementation, the weight artefact, the output tolerance, or some combination of them. That decision is separate from any question about the data, and it has to be made whether or not a training-data rule is ever written.

INT ESTABLISHED · 'Reproducibility', a reason given in favour of requirements in the survey, therefore means two things that should not be merged. Reproducing the decoder's output from the published weights is a conformance matter and needs no training data at all. Reproducing the weights from the data is a maintenance matter and needs the data and much else besides (section 6.5). Table 31 sets out what a classical and a neural codec specification each contain, and where training data sits relative to both.

INT CONDITIONAL · The data question is upstream of three artefacts: the weights, the reference software, and possibly a description of how the weights were trained. Which of the three become deliverables is ULBC-6's decision, and it determines what 3GPP would have to keep. *Signal*: the deliverables list in ULBC-6 at SA4#138 – a training description or script named as a deliverable would move part of the maintenance question into SA4's own text.

Once the weights are normative, the next question is whether the comparison that selects them is fair.

6.2 Fair comparison: four instruments and what each controls

SRC The survey's reasons in favour of requirements were, as SP-260919 records them, fair comparison, robustness, bias avoidance and reproducibility [M05]. The reflector recorded the objection in technical form: *"A common training data does not guarantee a certain performance, e.g. across different languages."* SA4 reflector [W04], message 3 (Ericsson) *"Access to common training data is not sufficient to replicate training."* SA4 reflector [W04], message 3 (Ericsson) [W04]. The dispute is not whether comparison should be fair but which instrument makes it so.

INT ESTABLISHED · Four instruments are on the table in one document or another, and each controls something different (Table 32). Keeping test material out of training controls leakage. Disclosure of training material controls opacity and lets others judge what a result rests on. A common training set for primary results controls the data advantage itself, by fixing data as a variable. And requirement dimensions verified in testing control the outcome, whatever data produced it. Only the third removes data from the comparison; the first, second and fourth are all compatible with proprietary data.

CTX The first instrument is close to universal practice. JVET's joint call states that *"No part of test sequences shall be used as a training set"* JVET-AQ2021, coding conditions for submissions [J01]; the IETF's speech-enhancement draft bars training on its test clips [I17]; JPEG AI withholds its test set from participants [K06]. JVET applies it to its own common set as well: a teleconference record notes that one training set proposed for removal *"has a chance to overlap with test sequences"* JVET-AQ0042, clarification on NNVC CTC, question 2 [J04].

INT WELL SUPPORTED · Ericsson's question on the reflector – *"Why should non-public data be forbidden to be used, if it can improve the codec performance?"* SA4 reflector [W04], message 3 (Ericsson) – is the argument for letting data vary, and two bodies have answered it in related ways. JPEG AI requires its common set for the proposals it compares and allows further data only if it is described, licensed and handed over [K05]; JVET's NNVC exploration compares tools trained on its common sets, and its summary report states that *"Comparison is done between tests which use the same sub-set of training data."* JVET-AQ0023, introduction [J05]. TTA cites the JPEG AI model at PCG, and the reflector participant who rejected the precedent argument described a version of it: *"There are means to consider private datasets in conjunction with a common public dataset"* SA4 reflector [W04], message 9 [W04]. The hybrid keeps the comparison clean without forbidding better data.

INT CONDITIONAL · For ULBC the hybrid would need one further rule: which result decides selection. If selection rests on the common-set results, the codec standardised is the one trained on data 3GPP can name, and the maintenance case is met in the same step. If it rests on the supplementary results, the data advantage returns at the point that matters. *Signal*: any proposal for ULBC-5b that distinguishes primary from supplementary results.

A fair comparison selects a codec; robustness is whether it works where it is deployed.

6.3 Robustness across languages and conditions: data or requirements

SRC SP-260919 asks SA4 to define requirements covering *"the number of languages / language families, speaker demographics, acoustic environments"* and similar dimensions, so that candidates can be verified across them [M05]. SA4 has already agreed part of the acoustic dimension as test conditions – input levels for qualification and selection, and background noise types and signal-to-noise ratios for qualification [M01]. The reflector recorded both sides: Ericsson that a common set does not guarantee performance across languages, and Dolby's reply that robustness is a codec requirement to be enforced in the selection tests [W04] [W05].

INT WELL SUPPORTED · A requirement on the outcome can be verified only for the conditions actually tested, and listening tests are where the selection spends its money; the total cost estimate is still open, and a direct test of all candidates is being studied to save resources [M01]. A requirement on the number of languages or language families is therefore a decision about test cost as much as about data, and it is one SA4 can take without reference to PCG.

CTX Corpus coverage is a fact that can be checked, and it differs widely among the four corpora TTA lists. Common Voice's release 26.0 spans 294 languages; LibriSpeech is read English from public-domain audiobooks; VCTK is 110 speakers of English; the DNS set combines speech, noise and room-response components from many sources [V03] [V04] [V05] [V07]. A common-corpus rule fixes coverage at whatever the chosen corpus contains, so a coverage requirement and a corpus choice have to be read against each other.

CTX The AI Act gives a benchmark for what coverage documentation looks like, without applying to a codec. Article 10, which binds only high-risk systems, asks that training data be relevant and sufficiently representative, and that it reflect the geographical, contextual, behavioural and functional setting of use [L01]. A ULBC requirement on languages, demographics and environments would be the same kind of statement, made for a technical rather than a legal reason.

INT CONDITIONAL · Robustness therefore splits cleanly between the two tracks. If SA4 adopts requirement dimensions and tests them, robustness is secured on the outcome and the data question narrows to maintenance and law. If it relies on a common corpus instead, robustness is secured only as far as the corpus reaches. *Signal*: requirement dimensions in ULBC-3, and the language count in the funding estimate. The assurance questions for emergency calls are in section 9.4.

Robustness is judged at selection; maintenance begins after it.

6.4 Maintenance against interoperability: who can retrain, and what changes

SRC TTA names five occasions on which a trained model may need work after approval: to correct a defect or undesirable behaviour, to address systematic performance issues found after deployment, to modify or extend it in a later release, to let members other than the proponent reproduce and investigate it, and to continue maintenance if the proponent can no longer support it [P02]. On the reflector Ericsson answered that *"if re-training is foreseen, it raises questions on maintaining the interoperability of the standard"* SA4 reflector [W04], message 3 (Ericsson), and Nokia closed the thread with the view that *"ML codecs do not present a specific new technical dataset problem"* SA4 reflector [W04], message 11 (Nokia) [W04].

INT ESTABLISHED · Retraining produces new weights, and new weights are a new decoder. For a deployed codec, interoperability then requires the old and the new versions to be distinguishable and the old weights to remain the reference for devices in service. A defect fixed by retraining is therefore closer to a new codec version than to a correction, and the maintenance TTA and SP-260919 describe is mostly evolution. Ericsson's point and TTA's are two halves of the same fact.

INT WELL SUPPORTED · The five occasions do not all need the original data (Table 33). A systematic weakness in one language or condition needs data that covers it, which need not be the original; an extension in a later release needs new data in any case. Reproduction and investigation by others, and continuity after the proponent leaves, need the original data and the training recipe. The original corpus is necessary for two of the five, and so is the recipe - which is where the maintenance argument has its force.

INT WELL SUPPORTED · Under undisclosed proponent data the specification stays implementable and maintainable as fixed weights: text and reference code can be corrected by change request as for any codec. What 3GPP loses is the ability to regenerate the weights on the original basis. SP-260919's concern is therefore specific, and it can be met in more than one way - a common set that the winner actually trained on, an escrow of data and recipe, or a maintenance commitment from the proponent - which section 10.1 compares.

CTX 3GPP's speech codecs have been maintained for many years through change requests to their specifications and reference code across successive releases, and none of that maintenance required anything outside the specification. The new element here is not maintenance but regeneration: for the first time, part of the normative content could not be produced again from what 3GPP holds unless 3GPP decides to hold more.

Maintenance by retraining assumes the weights can be regenerated, which needs more than data.

6.5 Regeneration needs more than the data

INT ESTABLISHED · This section concerns *training reproducibility* - whether the weights can be regenerated from the same data and the same process. It is a different problem from the one in section 6.1, which is whether two implementations of the same fixed model produce the same output. The word is used for both in these communities, and the two carry different

obligations: the first is a maintenance question, the second a conformance question.

SRC TTA states that availability of the training material does not by itself guarantee exact reproduction of the model [P02]. JVET's records show what reproduction takes in practice. In the NNVC exploration, *"Candidates for adoption to NNVC are required to undergo training cross-check"* JVET-AQ0023, EE1 test conditions, and *"training scripts (including model quantizing algorithm) must be made available"* JVET-AQ0023, EE1 test conditions; in the round reported at Geneva, one independent re-training came back within a stated margin - *"After re-training of EE1-1.1.1 the BD-rate difference (Y component) of 0.1%"* JVET-AQ0023, EE1-1.1 results [J05]. The ad hoc group is also mandated to *"Investigate bit-exact reproducibility of NN-based methods on various platforms"* JVET-AQ0011, AHG11 mandates [J03], and a contribution to it that cut a reference model's training from about 4.6 to 3.2 days warns that *"the reported training time and memory usage are representative only on Ampere-class or newer GPUs"* JVET-AQ0200, training configuration note [J09] - the hardware is part of the recipe.

SRC The joint call carries the same discipline into a competition. Its source-code clause says the material a proponent provides *"may include training scripts or equations used to derive parameters of a tool or algorithm"* JVET-AQ2021, source code and IPR considerations, and the proposal template asks proponents to *"describe training materials and how they were used to obtain the learned parameters"* JVET-AQ2022, learning-based elements [J01] [J02]. Reproduction by someone other than the proponent is thus already an operating procedure in one standards body, with scripts required, a cross-checker named and a tolerance reported rather than exact identity assumed.

CNT ESTABLISHED · Reproducing a trained model needs seven kinds of artefact: the exact data version with a manifest and hashes; the preprocessing and augmentation steps, including noise, reverberation and channel simulation; the training code and hyperparameters; the random seeds and data order; the framework, library and hardware versions; the quantisation and export steps; and an agreed tolerance for what counts as reproduced (Table 34). Of the seven, the three OP inputs name the first - the data, its description and its versions - and none of the other six.

INT WELL SUPPORTED · A rule on data alone therefore meets a necessary condition for maintenance and not a sufficient one. A common corpus without the winner's recipe would let 3GPP retrain a model, but not the standardised one. Conversely, an escrow of data and

recipe could deliver regeneration without any common corpus. The maintenance case is about a package, and the package is wider than the dataset the PCG discussion is about.

CTX Most of the package has existing instruments. ETSI EN 304 223, the baseline security standard for AI models and systems, asks for documented sources of training data and cryptographic hashes for shared model components [N01]; SPDX 3.0.1's trainedOn relationship links a model to the datasets it was trained on [N10]; ISO/IEC 15938-17 compresses trained networks for distribution [K14]. What has no standard form is the recipe itself, which calls for proposals handle by requiring scripts or descriptions in the submission.

What this part established

Normative weights make 'reproducibility' two things: conformance of a decoder's output, which needs no data, and regeneration of the weights, which needs data and much more.

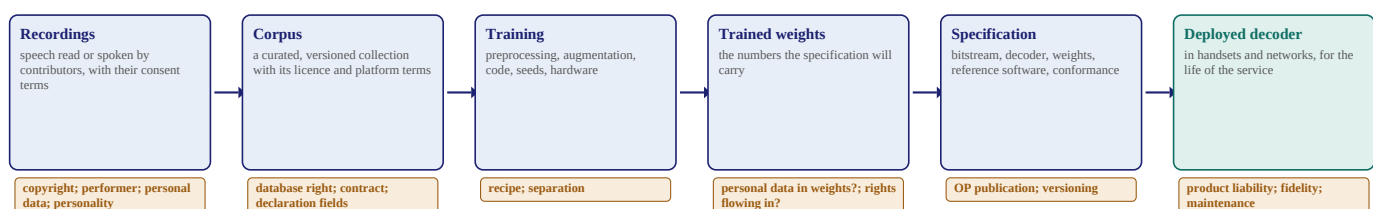
Four instruments make a comparison fair; only a common set removes the data advantage, and JPEG AI's hybrid keeps better data allowed (Table 32).

Retraining creates a new decoder, so maintenance by retraining is evolution, and the original data and recipe matter for two of TTA's five occasions (Table 33). The inputs name only the data (Table 34).

What this part established — part 6

Hands forward: how other standards bodies govern training data - part 7.

Figure 7. From recorded speech to a deployed decoder: the artefacts and what attaches to each



Boxes are artefacts, arrows the processes that produce the next one; the tags beneath name the rights, records and duties that attach at each step. Box width carries no quantity.

The chain the dataset question runs along, for a reader outside codec work. The legal layers attach at the left; the weights carry them into the specification; the product obligations attach at the right.

Table 31. What a classical and a neural codec specification each contain

Element	Classical codec	Neural codec with normative weights	Needs training data
Bitstream and decoding procedure	normative	normative	no
Decoder parameters	tables in the text	trained weights, held outside the text	the weights are the product of training
Reference software	normative or informative code	code plus the weight files	no, once the weights exist
Conformance	bit-exact test vectors	output equivalence has to be defined: bit-exact, equivalent to a reference implementation, or within a stated tolerance - open. Implementation and platform behaviour is part of the same question	no
Model artefact boundary	the algorithm and its tables are specified directly	architecture, trained weights, a separately held artefact and a reference implementation have to be distinguished, and the specification has to say which of them is normative	no, once the artefact is fixed
Encoder	informative	informative; may be trained too	if the encoder is trained
Training data	not applicable	not part of any specification so far	yes
Training recipe	not applicable	a submission item in JVET, not a deliverable	yes

Where training data sits relative to the deliverables. Two rows carry open questions [section 6.1](#) names, and neither is answered by having the data: what counts as conformant output, and which model object the specification makes normative.

Table 32. Four instruments for a fair comparison, and what each controls

Instrument	What it controls	Who proposes it here	Precedent	Compatible with proprietary data
Separation of test material	leakage of test conditions into training	near-unanimous in the survey (20 of 24)	JVET joint call; IETF mlcodec; JPEG AI	compatible
Disclosure of training material	opacity; lets others judge a result	SP-260919 ('transparent')	JVET joint call; JPEG AI extra data	compatible
Common set for primary results	the data advantage itself	TTA; ATIS as one means	JPEG AI; JVET NNVC; MPEG FCM	not for primary results
Requirement dimensions verified in testing	the outcome, whatever the data	SP-260919; Dolby's survey reply	listening tests in every codec selection	compatible

For proponents and test laboratories. Only the common set removes the data advantage; the other three are compatible with proprietary data. The precedent column is context.

Table 33. Retraining occasions, and what each needs

Occasion (TTA's five)	Needs the original data	Needs the recipe	Effect on the deployed decoder
Correct a defect or undesirable behaviour	only for an exact regeneration	only for an exact regeneration	a new version, unless fixed in reference code
Address a systematic weakness found after deployment	no: data covering the weakness	helpful	a new version
Modify or extend in a later release	no: new data in any case	helpful	a new mode or version
Reproduce and investigate by others	yes	yes	none, if investigation only
Continue if the proponent cannot support	yes	yes	a new version, if retrained

TTA's five occasions read against what each needs. The original data and recipe are needed for two of the five, and for a third only if an exact regeneration is wanted; every retraining changes the deployed decoder unless the fix is in reference code.

Table 34. What regenerating a trained model needs, and which input names it

Artefact	Named by an OP input	Existing instrument or practice
The exact data version, manifest and hashes	TTA; legal paper (versioning)	SPDX Dataset; Croissant [N09] [N12]
Preprocessing and augmentation: noise, reverberation, channel	none	JVET: training scripts required [J05]
Training code and hyperparameters	none	JVET joint call: scripts may be required [J01]
Random seeds and data order	none	none dedicated
Framework, library and hardware versions	none	JVET-AQ0200: results representative only on one GPU class [J09]
Quantisation and export steps	none	JVET EE1: quantising algorithm with the scripts [J05]
A tolerance for what counts as reproduced	none	JVET EE1: re-training differences reported [J05]

Seven artefacts, of which the OP inputs name one. The last column shows where JVET already requires the rest in practice. These are inputs to regenerating the weights and they do not answer the separate question of whether two implementations of a fixed model agree, which is governed by the normative model artefact, the arithmetic and the conformance tolerance in [section 6.1](#).

7 How other standards bodies govern training data

For: other standards bodies, researchers, curators.

7.1 JPEG AI: a mandated common set and normative parameters

CTX The boundary of this comparison. This part compares technical instruments, not institutional equivalence. JPEG, JVET, MPEG and the IETF work through calls, experiments, working documents and later approval paths that differ from a 3GPP working-group agreement ([Table 60](#), [Figure 13](#)). A mechanism used by another coding body is therefore evidence that the mechanism can be made to work

technically. It is not evidence that 3GPP can adopt it through the same procedure, or on the same timetable.

CTX The distinction has a long history in this partnership. The video codecs 3GPP uses were developed by successive joint teams of ITU-T and ISO/IEC - JVT for H.264, JCT-VC for H.265, an informal exploration team and then the formal JVET for H.266 - and in each case 3GPP's own act came afterwards, as an evaluation and a service profile [K33] [K34] [K35] [K36]. The speech and audio codecs, by contrast, 3GPP selected for itself, through the qualification and selection sequence ULBC is running now. JVET is therefore the same institution that supplies the training-data instruments read below and a body whose codec 3GPP has only ever adopted at the far end of someone else's process ([Figure 14](#), [Table 61](#)).

CTX JPEG AI is the one learned coding standard with both a mandated common training set and normative

trained parameters. Its final call for proposals required proponents to use the JPEG AI training dataset, and to describe any subset they used; a proponent using further data had to describe it in detail with its licensing conditions, include it in the submission and give JPEG enough rights to use it in standardisation [K05]. The dataset - 5,264 training and 350 validation images, with a withheld test set - is released on request, and common training and test conditions govern its use [K06] [K09] ([Table 36](#)).

CTX The standard that resulted carries its parameters as normative content. ITU-T T.840.1, approved in March 2025 and published by ISO/IEC as 6048-1:2025, treats the trainable model parameters as an integral part of the text even though they are held outside it [K07] [K08]. JPEG's provenance standard, JPEG Trust, reached a second edition in 2026 with attribution assertions and rights declarations [K10], and the JPEG AI training sets

have been made available to JVET for its own neural-network work (Table 39).

INT ESTABLISHED · For the *technical arrangement*, JPEG AI is the nearest published precedent in this brief: the combination TTA proposes - a common set for the comparison and normative weights for the result - has been run to publication once, in a body with ISO/IEC and ITU behind it. It is not a precedent for 3GPP's decision process, which reaches agreement differently and publishes through different owners (Table 60). The call's treatment of extra data, described, licensed and handed over, is also a worked answer to the objection that a common set holds back better data.

INT WELL SUPPORTED · For legal sufficiency it is not a precedent, for three reasons. Photographs carry copyright and sometimes personal data; recorded speech carries a performance and a voice as well, and a copyright licence covers neither (section 8.1). JPEG assembled and distributed its own set, where TTA proposes an external corpus distributed by its curator under the curator's terms. And no JPEG document read for this brief allocates liability for the set; the call's rights clause runs from proponents to JPEG, not from JPEG to implementers.

SRC One detail shows what a rule has to reach. The call describes the JPEG AI set as usable for standardisation because of its permissive CC0 licence; the dataset page, as read for this brief, does not print a licence [K05] [K06]. A licence stated in a call and a licence attached to the data are two different records, and a 3GPP requirement would need the second.

JPEG AI shows one body's full answer; JVET shows two answers in one body.

7.2 JVET: common sets in exploration, disclosure in the 2026 call

SRC JVET's neural-network video coding exploration works from common training sets. Its ad hoc group lists twelve in the NNVC common test conditions, from BVI-DVC to the two JPEG AI sets, and is mandated to update the list; another of its mandates reads *"Promote the call for training materials, distribute it, and actively communicate with content owners."* JVET-AQ0011, AHG11 mandates [J03] (Table 39). A June teleconference settled two entries: DIV2K stays with a comment, because *"This training set is for academic research purposes only."* JVET-AQ0042, clarification on NNVC CTC, question 1, and a set that *"has a chance to overlap with test sequences"* is to be removed [J04]. The exploration experiments compare tools trained on the same sets - *"For all proposals, training sets specified in [1] were used."* JVET-AQ0023, introduction [J05].

CNT ESTABLISHED · Of the 12 listed sets, 4 carry a note: 1 limited to academic research, 1 at risk of overlapping the test sequences, and 2 provided by JPEG. The test material itself is held apart: *"the JVET database of test sequences is only available to accredited members of JVET"* JVET-AQ0004, test material [J06]. Even a common set assembled by standards experts for their own exploration therefore carries licence limits and leakage checks as standing items of work.

SRC The 2026 joint call for proposals on video compression beyond VVC, issued at the July Geneva meeting and approved by the ITU-T SG21 plenary on 15 July, takes a different route [J01] [K02]. It states that *"No part of test sequences shall be used as a training set"* JVET-AQ2021, coding conditions for submissions, and that *"it is mandatory to report which material was used to train corresponding parts of the algorithm"* JVET-AQ2021, coding conditions for submissions. Its source-code clause expects that *"training materials used in the preparation of a proposal would be made available for further*

investigation and verification in JVET" JVET-AQ2021, source code and IPR considerations, or that *"re-training would need to be conducted using alternative material that is available to JVET"* JVET-AQ2021, source code and IPR considerations [J01]. Main submissions are due at the end of October and evaluation is planned for January 2027 [K03].

INT WELL SUPPORTED · JVET's step from the exploration to the call is informative for 3GPP. Among its own experts it uses common sets; for a call open to conventional and neural tools from any proponent, it chose disclosure and verifiability over a mandated corpus, and it wrote a clause for the case ULBC faces - a proponent that cannot make its data available - by making re-training on material JVET holds the fallback. That is a transparency rule in working form, and it is the nearest existing text to what SP-260919 proposes.

CTX These rules sit inside a different process from a 3GPP working-group decision. JVET can place training conditions in an exploration or in a call before the resulting technology reaches either parent's approval path [K27]. The comparison that transfers to ULBC is therefore the content of the rule - separation, disclosure, scripts, availability and re-training - and not the time it took JVET to write it.

INT ESTABLISHED · JVET's record also makes a second distinction this brief carries throughout. A model can be fixed while implementation results remain platform-sensitive, and a training process can be repeated while the regenerated weights and the coding performance they give are not identical. Those are two reproducibility questions, not one, and a rule that answers either leaves the other open (sections 6.1 and 6.5).

INT ESTABLISHED · The two JVET routes also show where each kind of rule lives. The common sets and their notes sit in an ad hoc group with a standing mandate; the call's rules sit in the call. ULBC has the equivalent of the second - its qualification and selection rules - and no equivalent of the first: no group or Pdoc whose standing task is the training material (section 5.1).

MPEG, JVET's partner in the video work, relies on documents of common conditions.

7.3 MPEG-AI and neural network coding: common conditions and licence terms

CTX MPEG's artificial-intelligence family, ISO/IEC 23888, is now a multi-part standard. Part 1, a technical report on vision and scenarios, was published in July 2026; Part 2, Video Coding for Machines, closed its draft-international-standard ballot the same month; Part 4, Feature Coding for Machines, is at committee draft; Part 5 covers AI-based point cloud compression; Part 6, media authenticity and provenance, has an open call with proposals due in January 2027 [K04] [K03] (Table 38).

CTX The limit of the public record here. The material read for this brief establishes the MPEG-AI project stages, the common test and training conditions and the work on neural-network coding. It does not expose the working-level discussion behind them. In particular, the public record read here does not establish a settled MPEG position on whether a learned coding standard should normatively specify model structure alone, trained weights as well, or a separately held model artefact. A blank in Table 38 records what this brief could read, not a decision against a provision.

CTX Common conditions are MPEG's standing instrument. The Feature Coding for Machines work maintains a document titled common test and training conditions, reissued at every meeting from MPEG 144 to MPEG 155 [K04]. For Video Coding for Machines,

TTA and the Apex export both describe the call's conditions as recommending training sets and requiring training data to be publicly available under licence terms that allow commercial use; the call itself is distributed as an archive that could not be opened for this brief, and the condition is carried as those two sources quote it [P02] [J11] [J14].

CTX Two further MPEG instruments bear on a trained codec. MPEG and JVET issued a joint call for training materials for neural-network video coding in January 2024, an attempt to widen the pool of data the bodies can use [K13]. And ISO/IEC 15938-17, neural network coding, specifies how trained networks are compressed for transmission and storage - the step between a weight file and a deployed decoder [K14].

INT WELL SUPPORTED · The MPEG material visible in the public record is strongest on test and training conditions and on representation, and it is not a completed governance answer for the normative model artefact. A versioned conditions document reissued every meeting is a mechanism 3GPP's Pdoc structure could mirror directly - a training-conditions Pdoc with editors and a target - and it would give SA4 a place for the separation rule, the disclosure format and any common set, without any of those deciding the legal question. The licence condition, where MPEG states one, is a condition on proponents; nothing read here shows MPEG taking responsibility for the data.

CTX The joint work of IEC, ISO and ITU on multimedia authenticity adds a map rather than a rule. Its July 2026 technical report groups the relevant standards into provenance, authenticity, asset identifiers, rights declarations and watermarking, and lists the opt-out tools; the companion policy paper surveys the diverging legal treatment of generative-AI training across regions [K15] [K16]. Neither sets training-data rules for a codec.

The IETF works on the same questions from the network's side.

7.4 The IETF: separation in mlcodec, preferences in AIPREF, receipts in SCITT

CTX The IETF's one neural speech-codec effort governs training by separation alone. The mlcodec working group's Deep Audio Redundancy draft for Opus, in working-group last call since August 2026, makes the decoder and its weights normative, distributes the weights outside the document and leaves encoder and training non-normative [I15] [I16]. Its speech-enhancement draft bars training on the test clips and discourages training on their source datasets, Common Voice and VCTK [I17]. Neither draft sets a provenance, licence or disclosure rule (Table 40).

CTX The AI Preferences working group standardises signals, not governance. Its charter covers a vocabulary for usage preferences and ways to attach it to content, and puts enforcement, registries and the auditing or transparency of training out of scope [I01]. The vocabulary draft of August 2026 has two categories, training AI models and search, the first now limited to models that generate synthetic content; the attachment draft defines a Content-Usage header and a robots.txt rule and would update RFC 9309 [I02] [I03] [I05]. Both drafts have passed their milestones, and a further interim met on 24-26 August [I04].

CTX Two other strands supply mechanism. RFC 9943, published in June 2026, defines an architecture for signed statements registered in append-only transparency services, and its reference API is approved and awaiting publication; no profile for datasets or model weights has been adopted [I06] [I07]. The vCon working group defines a container for recorded conversations, with individual drafts adding lawful-basis, consent-revocation and provenance

records [I10] [I11] [I12] [I21]. Individual drafts propose attestation from training data to inference output and dataset integrity records, none with IETF consensus [I23] [I25].

INT WELL SUPPORTED · Read together, the IETF offers 3GPP three usable pieces and no rule. Separation, as mlcodec applies it, is the same rule the SA4 survey supports. A transparency log built on RFC 9943 could record signed statements about a dataset release and a weight file, if someone wrote the profile. And an individual draft that carries AI usage preferences in SIP and SDP [I13] shows how a consent signal could travel with a call - relevant only if call audio were ever proposed as training data, which nothing in the ULBC record does.

INT Whether training a neural speech codec falls within the vocabulary's training category is open to argument: the decoder synthesises audio, but it reconstructs a transmitted signal rather than generating new content. The distinction matters only for web-crawled material, since the attachment mechanisms address publishers and crawlers; a licensed corpus is governed by its licence, not by a preference signal.

Preferences and receipts attach to content; the codec bodies have their own signal.

7.5 Usage signals and provenance: the SEI message, C2PA, CAWG and TDMRep

SRC Video coding now carries a usage signal of its own. Rec. ITU-T H.274 version 4, approved in January 2026, includes an AI usage restrictions request SEI message. Its six categories distinguish any AI application, training and inference, each for AI in general and for generative AI, and a restriction can be scoped to commercial, non-commercial, official government or research use. Setting the training category signals that the content *"ought not to be used for any AI training"* JVET-AQ0086, AI usage restrictions request SEI, NOTE 3; a note adds that *"no conformance requirements for receiving systems are imposed by this Specification"* JVET-AQ0086, AI usage restrictions request SEI, NOTE 1 [J07] [K11] [K12]. JVET is also studying a message referencing C2PA content credentials [J08].

CTX The content-provenance stack outside the codec bodies points the same way. The C2PA specification reached version 2.4 on 1 April 2026 with an AI disclosure assertion; its earlier training and data-mining assertion now lives with the Creator Assertions Working Group, whose version 1.1 records whether content may be used for training [K17] [K18]. The W3C TDM Reservation Protocol expresses a text-and-data-mining reservation in machine-readable form, and ITU-T F.748.77 asks that evaluation data for foundation models be traceable [I28] [K19] (Table 35).

INT WELL SUPPORTED · All of these are forward-looking signals attached to content: they say

what may not be done with it. None records what a model was trained on. For a ULBC corpus they matter at intake - a curator assembling or auditing a set would exclude content that carries a restriction - which makes them an instrument for the curator, whether that is Mozilla, a proponent or, at the fifth level of Figure 6, 3GPP itself.

INT ESTABLISHED · The SEI message and the IETF vocabulary share a posture: a request, not a requirement. H.274 imposes no conformance requirement on receivers, and AIPREF leaves enforcement out of scope. Standards bodies have so far been willing to standardise how a preference is expressed and unwilling to make obeying it a condition of conformance - which is also the line between what SA4 can specify and what only law can require.

Read one by one, the bodies differ; read against one list, they form a pattern.

7.6 The pattern across bodies, and the layer none has taken on

CNT ESTABLISHED · Figure 8 reads 8 bodies or efforts against eight provisions. A common training set for primary results is used by 3; disclosure of training material is required by 6; separation of training from test material by 6; a licence condition on training data by 2; rights or access for the body itself by 2; normative trained parameters by 1, with a second partly; reproducibility controls by 2. Legal or liability governance of the data: none. A cell is counted only where the provision is stated; a partial reading is printed and never added.

INT WELL SUPPORTED · Two layers are the most recurrent practices in the 8 efforts reviewed: separation and disclosure, each in 6 of 8. A common set is used where controlled comparison is the purpose - in explorations and in the one call whose result was a standard with normative weights - and it travels with licence limits and leakage checks. These similarities concern technical instruments. They do not imply that the bodies decide the same way, standardise the same model artefact, or could adopt a rule on the same timetable (Table 60, Figure 13). What no body has done is take on legal or liability governance of a training corpus. That is the layer PCG has been asked about, and it has no precedent to copy.

INT ESTABLISHED · Reading these bodies against each other needs one more discipline. The same apparent level of support does not mean the same standardisation maturity: a condition of entry to a competition binds submissions to that competition and is not a clause of the later standard; an adopted test model is convergence sufficient for joint experimentation and not an approved text; a committee draft is expert maturity and not national-body approval; and a 3GPP Pdoc can be highly mature technically and still be open on a disputed point.

Table 59 sets out what each observable signal establishes, what it does not, and the words this brief uses for it, so that a milestone in one body is never counted as the equivalent of a milestone in another.

INT WELL SUPPORTED · That discipline cuts both ways for ULBC. It means the training-data instruments in Tables 35 to 40 are available evidence that such rules can be written and applied - JVET's disclosure and re-training conditions are in force for a call now open, and JPEG AI's common set was run through to a published standard. It also means none of those dates forecasts a 3GPP date, and none of those bodies has been asked the question PCG was asked. The public MPEG record read here does not even establish which model object a learned coding standard should make normative (section 7.3), so the nearest comparators are still open on a question ULBC has to close.

INT The absence has more than one explanation, and the record does not choose between them. Only JPEG AI has normative weights, so elsewhere the data never reaches the standard; ISO/IEC and the ITU carry standards through legal structures unlike 3GPP's Organisational Partners; and a speech corpus differs from an image corpus less in how personal it is - faces, places and behaviour travel in pictures too - than in the fact that its speaker, language and acoustic composition is also the technical variable being compared. Each fits the pattern; none is established by it.

INT CONDITIONAL · A reading for 3GPP follows either way. The technical layers - separation, disclosure and, if SA4 wants it, a common set for primary results - have precedent and can be decided in SA4; the governance layer has none and sits with PCG. The two tracks Singapore created therefore match the division other bodies' practice suggests, whatever PCG decides. *Signal that would weaken this reading:* a PCG outcome that assigns a technical rule, such as separation, to the legal track, or a legal one to SA4.

What this part established

JPEG AI mandated a common set and made its parameters normative; JVET uses common sets in exploration and disclosure with access for re-training in its call; MPEG maintains common conditions; the IETF's neural codec work separates test from training.

Separation and disclosure are the most recurrent practices, each in 6 of 8 efforts reviewed; no body has taken on legal or liability governance of a training corpus (Figure 8). That layer is PCG's, without precedent.

These are technical instruments, and the bodies that use them do not decide alike: 3GPP seeks consensus with a 71% vote fallback, JVET passes its output through two parent approvals, MPEG and JPEG through ISO/IEC ballots, the IETF through rough consensus. A mechanism that works in one is not a rule 3GPP can adopt on the same clock (Table 60, Figure 13).

Hands forward: the law that bears on a speech corpus - part 8.

Table 35. Usage signals and provenance instruments

Instrument	Body and version	What it expresses	Force	Source
AI usage restrictions request SEI	ITU-T H.274 version 4, January 2026	six AI-use categories, four contexts	a request; no receiver conformance	[J07] [K11] [K12]
C2PA specification 2.4	C2PA, 1 April 2026	provenance manifests; AI disclosure assertion	voluntary	[K17]
Training and data-mining assertion 1.1	Creator Assertions Working Group	whether content may be used for training	voluntary	[K18]
JPEG Trust, second edition	ISO/IEC 21617-1:2026	provenance, attribution, rights declarations	a standard; use voluntary	[K10]
TDM Reservation Protocol	W3C Community Group report	a machine-readable mining reservation	not a W3C standard	[I28]
Robots Exclusion Protocol	RFC 9309	crawler access rules	recognised in the EU code of practice	[I05] [L04]
Assessment of foundation models	ITU-T F.748.77	traceable evaluation data	a recommendation	[K19]
Authenticity standards map	IEC, ISO and ITU, July 2026	provenance, rights and watermarking standards	a report	[K15]

Signals that travel with content and say what may not be done with it. None records what a model was trained on.

Figure 8. Eight bodies, eight provisions: the training-data and model-artefact practices visible in the public record

	Common training set for primary results	Disclosure of training material	Training kept apart from test material	Licence condition on training data	Rights or access granted to the body	Trained parameters normative in the public...	Legal or liability governance of the data	Reproducibility controls
	3 of 8	6 of 8	6 of 8	2 of 8	2 of 8	1 of 8	0 of 8	2 of 8
JPEG AI (learning-based image coding)	●	●	●	●	●	●	○	●
JVET NNVC exploration	●	●	●	●	●	—	○	●
JVET beyond-VVC joint call (2026)	○	●	●	●	●	—	○	●
MPEG Video Coding for Machines	●	●	●	●	—	—	○	●
MPEG Feature Coding for Machines	●	●	●	●	—	—	○	●
MPEG AI-based point cloud coding	○	●	—	●	—	—	○	●
IETF mlcodec (Opus neural extensions)	○	○	●	○	—	●	○	—
IETF AIPREF (usage preferences)	—	—	—	●	—	—	○	—
3GPP SA4 ULBC, at the SA#113 cutoff	⊙	⊙	⊙	⊙	⊙	●	⊙	⊙

● stated ● partly or as quoted ○ not stated — not applicable ⊙ open at the cutoff

Read from the records in Tables 35 to 40. The last row is 3GPP for comparison and is not counted. Counts in the head are stated cells only; partial cells are drawn and never added. 'Not stated' means not established in the public records read for this edition, not that a body has decided against the provision.

The cross-body pattern in one view, for readers from other standards bodies and for PCG. Separation and disclosure are the most recurrent practices in the eight efforts reviewed; common sets appear where controlled comparison is the purpose; no body governs the legal or liability side of a corpus. The columns compare instruments, not the decision processes behind them (Table 60, Figure 13).

Table 36. JPEG AI: the training-data rules, and the standard they produced

Item	What the record states	Source
Common training set	Proponents must use the JPEG AI training dataset; any subset used is described	[K05]
Additional data	Allowed if described in detail with its licensing conditions, included in the submission, and with rights granted to JPEG for standardisation	[K05]
Licence, as the call states it	CC0, permissive; usable for standardisation	[K05]
The dataset	5,264 training and 350 validation images; test set withheld; released on request	[K06]
Conditions	A common training and test conditions document	[K09]
Parameters in the standard	Trainable model parameters are an integral part of ITU-T T.840.1 and ISO/IEC 6048-1:2025, held outside the text	[K07] [K08]
Provenance	JPEG Trust, second edition 2026: attribution assertions, rights declarations	[K10]
Sharing	JPEG AI training sets made available to JVET members	[J03]
Liability for the set	No statement in the documents read	-

The one learned coding standard with a mandated common set and normative parameters. Context for 3GPP, and a record for JPEG.

Table 37. JVET's 2026 joint call: the training-data provisions

Provision	Wording or record	Where	What it does
Separation	"No part of test sequences shall be used as a training set"	coding conditions	no test sequence in any training set
Disclosure	"it is mandatory to report which material was used to train corresponding parts of the algorithm"	coding conditions	mandatory reporting of training material
Scripts	"may include training scripts or equations used to derive parameters of a tool or algorithm"	source code and IPR	training scripts can be part of the technology provided
Availability	"training materials used in the preparation of a proposal would be made available for further investigation and verification in JVET"	source code and IPR	material open to investigation and verification
Fallback	"re-training would need to be conducted using alternative material that is available to JVET"	source code and IPR	re-training on material JVET holds
Template	"describe training materials and how they were used to obtain the learned parameters"	JVET-AQ2022	a described training process
Test material	"the JVET database of test sequences is only available to accredited members of JVET"	JVET-AQ0004	test sequences held apart
Dates	approved by the ITU-T SG21 plenary on 15 July 2026; main packages due end of October; evaluation January 2027	[K02] [K03]	the same quarter as ULBC's first tests

The nearest existing text to a transparency rule, in the call's own words. The joint call is evaluated in January 2027, the month SP-260907 pencils in for ULBC's first contracts.

Table 38. MPEG-AI and related ISO/IEC work

Document	Subject	Stage as read	Training-data provision	Source
ISO/IEC 23888-1	Vision and scenarios (technical report)	published July 2026	-	[K04]
ISO/IEC 23888-2	Video Coding for Machines	DIS ballot closed July 2026	recommended training sets; public data under commercial-use licences, as TTA and the Apex export quote the call	[K04] [J11] [P02]
ISO/IEC 23888-4	Feature Coding for Machines	committee draft	common test and training conditions, reissued every meeting	[K04]
ISO/IEC 23888-5	AI-based point cloud compression	work item	guidelines on training data, as the Apex export cites them	[J12]
ISO/IEC 23888-6	Media authenticity and provenance	call open; proposals due 13 January 2027	-	[K04] [K03]
ISO/IEC 15938-17	Neural network coding	second edition 2024	compresses trained networks for distribution	[K14]
MPEG and JVET	Call for training materials	January 2024	widens the pool of data the bodies can use	[K13]

Stages as the public pages show them at the cutoff. The Video Coding for Machines condition is carried as TTA and the Apex export quote it. The public MPEG material read for this edition does not establish whether a future learned coding specification will normatively fix model structure alone, trained weights as well, or a separately held model artefact; a blank cell here should not be read as a decision against normative weights.

Table 39. The training sets in JVET's NNVC common test conditions

Set	Items	Where hosted	Note in the record
BVI-DVC	193	data.bris.ac.uk	-
TVD	86	multimedia.tencent.com	-
UGC	130	media.withyoutube.com	-
BVI-AOM	239	opencontent.netflix.com	-
vimeo90K triplet	73,191	toflow.csail.mit.edu	-
vimeo90K septuplet	91,701	toflow.csail.mit.edu	-
RWTH	100	vqa.lfb.rwth-aachen.de	proposed for removal: never used, may overlap CTC or CFP test sequences
DIV2K train and valid	800 and 100	data.vision.ee.ethz.ch	academic research purposes only; kept with a comment
JVET-AH0298	5	vqa.lfb.rwth-aachen.de (gaming)	-
JVET-AH0120	9	vqa.lfb.rwth-aachen.de	-
JPEG AI natural	5,264 / 350 / 50	vqa.lfb.rwth-aachen.de	JPEG AI training set made available to JVET
JPEG AI extra	3,029 / 202 / 36	vqa.lfb.rwth-aachen.de	JPEG AI training set made available to JVET

As the NNVC ad hoc group lists them. Four carry a note in the record: one limited to academic use, one at risk of overlapping test sequences, two provided by JPEG.

Table 40. IETF documents bearing on training data

Document	Group and status	What it does	Bearing on ULBC	Source
draft-ietf-mlcodec-opus-dred-07	mlcodec; working-group last call	normative decoder and weights; training non-normative	the same step as ULBC, without data rules	[115] [116]
draft-ietf-mlcodec-opus-speech-coding-enhancement-04	mlcodec; working-group document	bars training on test clips; discourages their source datasets	a separation rule, and a caution on TTA's corpus	[117]
draft-ietf-aipref-vocab-07	AIPREF; working-group document	two categories: training AI models, search	web-crawled content only	[102]
draft-ietf-aipref-attach-05	AIPREF; working-group document	Content-Usage header and robots.txt rule; would update RFC 9309	as above	[103] [105]
RFC 9943	SCITT; Proposed Standard, June 2026	signed statements in append-only transparency services	a log for dataset and weight statements, if profiled	[106]
draft-ietf-scitt-scrapi	SCITT; approved, awaiting publication	reference API	as above	[107]
RFC 9969	IAB; Informational, May 2026	report of the AI-CONTROL workshop	the origin of AIPREF	[108]
vCon documents	vCon; working-group and individual drafts	a container for recorded conversations; lawful basis; consent revocation	consent records for recorded speech	[110] [111] [112]
draft-altanai-aipref-realtime-protocol-bindings	individual	AI usage preferences in SIP and SDP	a consent signal inside a call	[113]
Provenance and attestation drafts	individual; no consensus	training data to output attestation; dataset integrity records	design references only	[123] [124] [125]
Web Bot Auth	working group	cryptographic identification of crawlers	accountability of a collector	[109]

Status as the datatracker shows it at the cutoff. Individual drafts carry no IETF consensus and are listed as design references only.

8 Law and regulation that bear on a speech corpus

For: OP counsel, regulators, policy staff.

8.1 Eight layers of rights in one recording

SRC Each input names the rights it has in mind. The legal paper lists copyright in the underlying recordings, performer rights, database rights, confidentiality, contractual licensing restrictions and personal-data protection for voice; TTA lists licensing and intellectual property, privacy, provenance and the rights 3GPP needs to use the material; ATIS asks for licensing terms and documented privacy details [P02] [P03] [P04].

CTX A single recording of read speech can carry eight distinct layers: copyright in the text read and in the sound recording; the performer's rights in the performance; database rights in the collection; the contract - licence and platform terms - under which the collection is obtained; confidentiality where data is held as a secret; personal data in an identifiable voice; biometric data where the voice is processed to identify the speaker; and personality or publicity rights in the voice as an attribute of the person [A20] [A19] [L12] [V02] [L10] [U05] [C05] (Figure 9).

CNT ESTABLISHED · Of the eight layers, all three inputs name three - copyright, the licence terms and personal data. The legal paper alone names three more: performer's rights, database rights and confidentiality. None of the inputs names the last two, biometric data and personality rights in the voice - the two layers on which courts and legislators have acted most since 2024.

INT WELL SUPPORTED · A permissive licence disposes of at most the copyright and related-rights layers its licensor holds, and only if the licensor held them. It disposes of nothing in the personal-data, biometric or personality layers. TTA's two conditions - permissive licensing and documented privacy - are therefore two conditions, and the second cannot be met by meeting the first; ATIS keeps them apart as two requirements for the same reason.

CTX Two decisions show why consent to publish and consent to train are different consents. The Beijing Internet Court held in April 2024 that a licence to a person's recordings did not authorise use of the voice for AI synthesis, and a Berlin court held in August 2025 that an AI imitation of a voice actor's recognisable voice infringed his personality right [C13] [C05]. In New York, voice actors' publicity claims against a voice-cloning service survived a motion to dismiss in July 2025 [C12]. None concerns a codec; each concerns what a speaker agreed to.

Of the eight layers, copyright is where jurisdictions differ most.

8.2 Copyright and text-and-data mining across jurisdictions

CTX In the European Union, Article 4 of the Digital Single Market Directive permits text and data mining of lawfully accessible works unless the rightholder has reserved the right, in machine-readable form for content online [L11]. The general-purpose AI Code of Practice commits signatories to respect robots.txt under RFC 9309 and other machine-readable protocols adopted by standardisation bodies [L04] [I05], and the Commission consulted on such protocols over the winter [L05]. Policy is moving: the Parliament's resolution of 10 March 2026 calls for remuneration and a presumption of use, the Commission gathered evidence for a copyright review in May and June, and the Court of Justice heard its first generative-AI copyright case, C-250/25, in March [L18] [L20] [L21] (Table 41).

CTX National courts have begun to draw lines. In Germany the Hamburg Higher Regional Court held in December 2025 that building a training dataset fell within the mining exceptions and that a reservation had to be machine-readable at the time of collection; the Federal Court of Justice heard the appeal on 3 September 2026 and will pronounce on 17 December [C02] [C01]. The Munich Regional Court held in November 2025 that lyrics memorised in model weights are reproduced, and ruled against a music generator in July 2026 [C03] [C04]. In England the High Court found in November 2025 that model weights which did

not store the works were not an infringing copy, a point now under appeal, and the United Kingdom's exception remains limited to non-commercial research [C06] [C07] [A02] [A01].

CTX In the United States fair use is decided case by case [U02]. The Copyright Office's report on training, still in pre-publication form, rejects a categorical answer [U01]. A California court held in June 2025 that training on lawfully acquired books was fair use while keeping pirated copies was not, and the case settled with final approval in July 2026 [C08]; a second court found for the developer on the record before it and left a distribution claim open [C09]. An appeal on a non-generative tool was argued in June 2026, the Justice Department filed a statement on fair use on 1 September, and record labels were allowed in August to add an anti-circumvention claim over audio taken from a video platform [C10] [C11] [C14].

CTX In Asia the positions differ again. Japan's Article 30-4 permits uses that are not for enjoyment of the work, within limits [A05]; Korea's culture ministry issued a case-by-case fair-use guide for training in February 2026, in English in May [A10]; Singapore's computational data analysis exception requires lawful access [A15]; China's measures for public generative-AI services require lawful sources and respect for intellectual property [A11] [A12]; and an Indian government working paper has proposed a mandatory blanket licence for training, for consultation [A17].

INT WELL SUPPORTED · The divergence matters most where there is no licence. For a corpus under an express licence the exceptions matter less, because the licence grants the use; for collected web audio, broadcasts or data of unrecorded origin, the same corpus can be lawful to assemble in Japan or Singapore and need a licence in the United Kingdom. The spread therefore bites hardest under the proponent-data option, where 3GPP cannot see how material was collected, and least under a named, licensed corpus.

INT CONDITIONAL · The memorisation rulings raise a question a codec can answer by test. A decoder reconstructs speech from a transmitted bitstream rather than generating content on request, but whether it can emit training utterances independently of its input is measurable. *Signal*: a selection test or deliverable that measures it - which would give implementers evidence no licence can supply.

Copyright attaches to the recording; data protection attaches to the voice.

8.3 Voice as personal data, as biometric data and as an attribute of the person

CTX Under the GDPR a voice recording is personal data where the speaker is identifiable; it becomes biometric data only through technical processing that allows or confirms unique identification, and special-category data only when processed for that purpose [L10]. The European Data Protection Board's guidance on voice assistants requires a separate legal basis for training on recordings and notes how hard voice is to anonymise [L09]; its opinion on AI models holds that a model trained on personal data is not automatically anonymous, and that a controller deploying a model must assess whether it was lawfully developed [L08]. The United Kingdom applies the same biometric test [A04] (Table 43).

CTX Elsewhere voice is reached through several instruments. Illinois lists the voiceprint as a biometric identifier requiring a written release [U05]; the Federal Trade Commission treats voice recordings as biometric information [U08]; Tennessee protects the voice, real or simulated [U09]; and a federal voice-and-likeness bill cleared the Senate Judiciary Committee on 18 June

2026 without being enacted [U07]. China lists the voiceprint as sensitive personal information [A13]; Japan's 2026 amendment adds an exception for statistical uses including AI training data and new biometric rules [A08]; India's rules bring most obligations into force in May 2027 [A18].

INT WELL SUPPORTED · Codec training does not of itself process voice to identify anyone, and so does not by that fact enter the biometric categories. A corpus with speaker identifiers, a speaker-similarity measure in testing - a usual way to check that a codec keeps who is speaking - or demographic and health annotations can bring processing into them. The metrics SA4 chooses to verify robustness across speaker demographics therefore carry personal-data consequences of their own.

INT CONDITIONAL · The one artefact every option shares is the published weights. Whichever corpus is used, the specification will carry weights trained on recorded voices, and under the Board's opinion such a model is anonymous only if extracting personal data from it is insignificant, which is assessed case by case. The question whether the weights themselves are personal data therefore arises under every option, proponent data included, and none of the three inputs treats the weights apart from the data. *Signal*: whether the new PCG group's terms name the weights.

Rights and data protection bind whoever handles the data; the AI Act binds whoever places the system.

8.4 The AI Act after the Omnibus: what reaches a codec and what does not

CTX The AI Act entered into force in August 2024. Obligations for general-purpose models have applied since 2 August 2025 and the Commission's enforcement powers since 2 August 2026; the Omnibus, Regulation (EU) 2026/1744, in force since 27 July 2026, moved the high-risk dates to 2 December 2027 for Annex III systems and 2 August 2028 for systems in products under Annex I [L01] [L02]. The Article 50 transparency duties apply from 2 August 2026, with Commission guidelines of 20 July and a code of practice on AI-generated content [L06] [L07] (Table 45).

CTX The training-data duties sit in Article 53 and bind providers of general-purpose models: a copyright policy that honours reservations under the mining exception, and a public summary of training content in the Commission's template, which asks for large public datasets, licensed private datasets, collected web data with crawler and domain information, user data and synthetic data, and for how reservations were respected [L03] [L04].

INT WELL SUPPORTED · A task-specific speech codec is unlikely to meet the general-purpose criteria, which turn on generality, and it is not an Annex III use; the Annex I route runs through product legislation and, after the Omnibus, not before August 2028. The Act's most concrete point of contact is Article 50: providers of systems generating synthetic audio must mark it, with exceptions for standard editing and for systems that do not substantially alter the input, and the guidelines give technical compression and noise reduction as examples of standard editing [L06].

INT CONDITIONAL · A codec that reconstructs the speaker's own speech falls on the editing side. A decoder at one kilobit per second that altered what was said, or shifted who seemed to say it, would test the boundary. The evidence that matters under the Act is therefore fidelity - content and speaker preservation measured in selection - rather than training documentation. *Signal*: fidelity measures among the ULBC-3 performance requirements.

CTX Outside the Union, California has required developers of generative AI systems to post summaries of their training data since 1 January 2026 [U03]; Colorado's replacement law takes effect in January 2027 [U04]; Korea's framework act has applied since 22 January 2026 [A09]; China's labelling rules for AI-generated content apply since September 2025 [A14]; Japan's promotion act and its July 2026 basic plan are policy rather than obligation [A06] [A07]; and a United States executive order of December 2025 seeks a national framework without pre-empting state law by itself [U06].

INT ESTABLISHED · TTA's proposal to model a transparency declaration on the EU template is the voluntary use of a format, not compliance with a duty - and a sound choice of format, because regulators already read it. SP-260912's statement that the Act can require documentation relating to training data is accurate as a general statement; the legal paper itself qualifies it by model and system class [P03].

The AI Act regulates the system; liability law decides who answers when it fails.

8.5 Liability, product safety and cyber resilience

CTX The revised Product Liability Directive, (EU) 2024/2853, must be transposed by 9 December 2026, and a corrigendum published in May 2026 fixes the date from which it reaches products: those placed on the market or put into service after 8 December 2026 [L13] [L23]. Software, including AI systems, is a product; manufacturers of components, software components included, can be liable; defectiveness takes account of a product's capacity to keep learning and of cybersecurity; courts may order disclosure of evidence; and Article 10 sets rebuttable presumptions of defect and causation, including where technical complexity makes proof excessively difficult [L13]. The proposed AI Liability Directive was withdrawn in October 2025 [L15] (Table 46).

CTX Two further instruments apply around a codec. The General Product Safety Regulation, applicable since 13 December 2024, asks the safety assessment to consider evolving, learning and predictive functionalities [L14]. The Cyber Resilience Act's reporting obligations for actively exploited vulnerabilities and severe incidents apply from 11 September 2026 through ENISA's single reporting platform, with full application in December 2027 [L16] [L17].

INT WELL SUPPORTED · The exposure sits with manufacturers of devices and network equipment and with suppliers of codec software as a component. The legal paper's view that the OPs do not place products on the market is the conventional one; whether distributing reference weights could make a standards body a component manufacturer is a question the base briefing puts to counsel [B02]. Either way, dataset documentation is what a manufacturer would need to answer a disclosure order or rebut a presumption - which is why the legal paper can call due diligence an industry requirement without choosing a scenario.

INT CONDITIONAL · The Cyber Resilience Act already requires manufacturers to run coordinated vulnerability disclosure, and the legal paper's mechanism for datasets would be its upstream counterpart. Whether a weakness in a dataset or a trained model - poisoned data, audio crafted to mislead the decoder - counts as a vulnerability under the Act is not settled in anything read for this brief. *Signal*: Commission or ENISA guidance on AI components under the Act.

One instrument sits closer to 3GPP than any statute: the IPR Policy it already applies.

8.6 The ETSI IPR Policy, and what it does not reach

CTX 3GPP has no single IPR policy of its own. Its Working Procedures put ownership and rights in two different places: Article 46 gives the Organisational Partners joint ownership, copyright included, of the Technical Specifications and Reports 3GPP produces, and Article 55 binds each Individual Member to the IPR policy of its own Organisational Partner, through which essential-IPR declarations are made [G11] [G25]. ETSI's policy is therefore one of seven, and it is read here because it is the one whose database carries most of the declarations and whose text is public.

CTX That policy, Annex 6 of ETSI's Rules of Procedure, defines an IPR as any intellectual property right conferred by statute law other than trade marks, and excludes rights relating to confidential information and trade secrets; it defines as essential an IPR without which equipment or methods complying with a standard cannot, on technical grounds, be made or used; its licensing undertakings concern essential IPR; and it addresses copyright in contributions and deliverables separately [G10] (Table 47).

INT WELL SUPPORTED · Set against the ULBC artefacts, the policy reaches less than it appears to. Patents on the codec's methods fall under the familiar disclosure and licensing regime. The weights and reference software fall under the copyright arrangements for contributions and deliverables. Rights in the training recordings are not needed to implement the standard - implementers copy the weights, not the data - and so are not essential in the policy's sense. And a dataset protected as a trade secret lies outside the definition of an IPR altogether.

INT ESTABLISHED · The instruments the partnership already applies to third-party rights in its standards therefore do not reach the training data. That is why all three inputs speak of requirements and due diligence rather than of IPR declarations, and why the legal paper restates the contribution principles - good-faith copyright, disclosure of essential IPR, authorisation for third-party material - before explaining why a dataset changes their scale, and leaves data held as secret

know-how expressly unaddressed [P03]. Nor should the ETSI policy be read as the whole answer for the other six partners: what each requires of its own members is that partner's instrument, and the compatibility of the seven is itself a PCG question.

INT CONDITIONAL · A second gap sits beside the first, and no input names it. Article 46 gives the partners joint ownership of the specifications and reports; a weight file too large to print in a specification would be published beside it, incorporated by reference. What instrument makes that binary a stable 3GPP artefact - who owns it, who hosts it, who grants implementers the right to copy and embed it, who keeps the superseded versions - is not answered by a rule written for documents. *Signal:* whether the ULBC deliverables list in ULBC-6 names a file and a licence for it, or only a specification (section 6.1).

INT CONDITIONAL · One link could bring the data back inside. If rights in training data were held to flow into the weights, publication of the weights would engage them directly. The English High Court found that weights which did not store the works were not a copy, a point under appeal; the Munich Regional Court found memorised lyrics reproduced in weights [C06] [C07] [C03]. *Signal:* the English appeal, and the German Federal Court of Justice's decision in December, or a reference to the Court of Justice.

What this part established

A recording carries eight layers of rights; the inputs together name six, and none names biometric data or personality rights in the voice (Figure 9).

Copyright regimes diverge most where there is no licence, so the proponent-data option carries the widest spread. The weights, published under every option, raise the personal-data and rights-flow questions the inputs leave open.

The AI Act touches a codec mainly through fidelity; liability and cyber law reach manufacturers, for whom documentation is evidence; the ETSI IPR Policy does not reach training data (Table 47).

A voice recording is personal data wherever the speaker can be identified and becomes biometric only through processing that identifies, so the metrics SA4 chooses for selection decide part of the legal question (Table 43).

The dates outside 3GPP fall first: the Product Liability Directive's transposition in December 2026 and the German Federal Court's decision the same month, both before selection would begin (Figure 10).

None of these instruments was written for a standards body, and none of them names one. They reach the manufacturer who places a product on a market, the controller who processes a recording and the party that copied a work; a partnership that publishes a specification is not, on the texts read here, any of the three. That is why the exposure a normative corpus would create runs to the implementer rather than to 3GPP, and why the question PCG was asked is about what the partnership takes on rather than what the law already puts on it (Table 46, Table 47).

What this part established — part 8

The divergence between jurisdictions is itself a constraint on the option space. A rule written to satisfy the widest reading has to satisfy the German position on memorisation, the United States fair-use analysis, the Japanese non-enjoyment exception and the Chinese scope limit at once, and the OPs sit in seven of these environments (Table 46, section 4.8).

Hands forward: product, supply chain and assurance - part 9.

Figure 9. Eight layers of rights in one recording, and which input names each

TTA	Legal	ATIS			What carries it
☐	☐	☐	Personality and publicity rights in the voice	the voice as an attribute of the person	court decisions 2024-2026; ELVIS; the NO FAKES bill
☐	☐	☐	Biometric data, where processed to identify	a voiceprint or speaker embedding	GDPR Art. 9; BIPA; GB/T 45574-2025
☒	☒	☒	Personal data in the voice	an identifiable speaker in every recording	GDPR; PIPL; APPI; DPDP; PIPA
☐	☒	☐	Confidentiality and know-how	non-public data held as a secret	trade-secret and contract law
☒	☒	☒	Contract: licence and platform terms	how the collection is obtained and used	CC0, CC BY 4.0; the distributor's terms
☐	☒	☐	Database rights	the collection as an investment	Directive 96/9/EC in the EU
☐	☒	☐	Performer's rights	the speaker's performance as fixed	WPPT; national related rights
☒	☒	☒	Copyright in the recording and the work read	the text read aloud and the sound recording	Berne; WPPT; national copyright law

Filled markers show the layers each PCG input names in its own text. The stack is an ordering of layers, not of weight. Sources for the right-hand column are in Tables 41, 43 and 44.

What a permissive licence does and does not settle, for counsel and curators. All three inputs name copyright, licence terms and personal data; the legal paper alone names three more; none names biometric data or personality rights in the voice.

Table 41. Copyright and text-and-data mining, by jurisdiction

Jurisdiction	Rule for training	Recent decisions or policy	For a speech corpus	Sources
European Union	Mining of lawfully accessible works unless reserved, machine-readably for online content	Code of practice; opt-out consultation; Parliament resolution; copyright review; C-250/25 heard	Licensed corpora unaffected; collected audio must respect reservations	[L11] [L04] [L05] [L18] [L20] [L21]
Germany	The same exceptions, national form	Hamburg appeal Dec 2025; Munich Nov 2025 and Jul 2026; Federal Court decision 17 Dec 2026	Machine-readability judged at collection; memorisation is reproduction	[C02] [C03] [C04] [C01]
United Kingdom	Exception for non-commercial research only	Getty v Stability, Nov 2025, under appeal; government report Mar 2026	Commercial training needs licences	[A02] [C06] [C07] [A01] [A03]
United States	Fair use, case by case	Copyright Office report; Bartz; Kadrey; Ross on appeal; Justice Department statement Sep 2026	Acquisition matters as much as use	[U02] [U01] [C08] [C09] [C10] [C11] [C14]
Japan	Non-enjoyment uses permitted, within limits	Agency guidance	Relatively permissive; limits on exploiting datasets sold for AI	[A05]
Korea	Fair use, case by case	Culture ministry guide, Feb 2026 (English May)	Documented reasoning favoured	[A10]
Singapore	Computational data analysis with lawful access	-	Lawful access is the condition	[A15]
China	Lawful sources and respect for IP, for public generative services	GB/T 45652-2025 on training-data security	Scope limited to public services	[A11] [A12]
India	No dedicated exception	Working paper proposing a blanket licence	Consultation only	[A17] [A16]


Context, as the texts or reliable summaries state it at the cutoff; not legal advice. The fourth column is this brief's reading for a speech corpus.

Table 42. What governments are deciding, by mechanism rather than by country

Jurisdiction	Training permitted by default	Reservation or opt-out	Licensing or remuneration	Public disclosure of training content	Regulator access	Personal and biometric constraints	Product and cyber evidence
European Union	permitted for lawfully accessible works	yes, machine-readable	under debate; a Parliament resolution asks for remuneration	yes, for general-purpose models	yes, through the AI Office	GDPR, and biometric rules where identification is the purpose	product liability, safety and cyber duties
United Kingdom	research only; commercial training needs a licence	not settled	four options consulted on; no reform until objectives are met	no	no	UK GDPR	product law outside the EU regime
United States	fair use, decided case by case	no	no	no federal duty	no	state biometric and voice statutes	state and sectoral
California	as United States	no	no	yes, for generative systems since January 2026	no	state statutes	state law
Japan	non-enjoyment uses permitted, within limits	no	no	no	no	personal information rules, with a 2026 amendment	product law
Korea	fair use, decided case by case	no	no	no	no	personal information rules	framework act on AI since January 2026
China	lawful sources and IP respect, for public generative services	no	no	labelling of generated content	yes, on request	voiceprint as sensitive personal information	sectoral
India	no dedicated exception	no	proposed in a working paper	no	no	digital personal data rules, most duties from May 2027	product law

The same eight jurisdictions read across seven mechanisms. A corpus assembled once has to satisfy the strictest cell in every column it touches.

Figure 10. The regulatory and judicial calendar around the ULBC selection



Filled diamonds are dates passed at the cutoff; open diamonds are dates ahead. The axis is linear; labels alternate above and below the lane and no event is moved in time. Sources in Tables 41, 43, 44, 45 and 46.

Outside dates beside 3GPP's own, for a reader timing a decision. The Product Liability Directive's transposition and the German Federal Court's decision both fall before ULBC's test contracts.

Table 43. Voice as personal data, biometric data and an attribute of the person

Jurisdiction	Voice as personal data	Biometric rule	Voice as a personal attribute	Sources
European Union	Personal data where identifiable	Special category only if processed to identify	Member-state personality law	[L10] [L09] [L08]
United Kingdom	Personal data where identifiable	The same technical test	Digital-replica right under consideration	[A04] [A01]
United States, federal	FTC treats voice recordings as biometric information	-	NO FAKES bill reported, not enacted	[U08] [U07]
Illinois	-	Voiceprint a biometric identifier; written release	-	[U05]
Tennessee	-	-	Voice protected, real or simulated	[U09]
New York	-	-	Publicity claims survived against a voice-cloning service	[C12]
Germany	-	-	AI imitation of a recognisable voice infringed a personality right	[C05]
China	Personal information	Voiceprint listed as sensitive	Recording licence did not cover AI synthesis	[A13] [C13]
Japan	Personal information; 2026 amendment eases statistical uses	New biometric rules in the 2026 amendment	Tort treatment of voice under study	[A08] [A07]
India	Digital personal data; most duties from May 2027	-	-	[A18]
Korea	Personal information	-	-	[A09]

Where a voice enters the special or biometric categories, and where it is protected as part of the person. Context; not legal advice.

Table 44. Court decisions on training data and voice, 2024-2026

Court and date	Case	Holding as reported	Bearing on ULBC	Source
Beijing Internet Court, Apr 2024	AI voice synthesis	A licence to recordings did not authorise AI synthesis of the voice	Consent to record is not consent to train	[C13]
Berlin Regional Court II, Aug 2025	Voice actor	AI imitation of a recognisable voice infringed a personality right	Voice as an attribute of the person	[C05]
S.D.N.Y., Jul 2025	Lehrman and Sage v Lovo	Publicity and consumer claims survived; voice not copyrightable	Consent terms decide	[C12]
N.D. Cal., Jun 2025; settled Jul 2026	Bartz v Anthropic	Training on lawfully bought books fair use; pirated library not	Acquisition matters	[C08]
N.D. Cal., Jun 2025	Kadrey v Meta	Fair use on the record; distribution claim open	Evidence of market harm decides	[C09]
High Court, Nov 2025	Getty v Stability	Weights that did not store the works not an infringing copy; appeal permitted	Whether rights flow into weights	[C06] [C07]
Munich Regional Court I, Nov 2025	GEMA v OpenAI	Memorised lyrics reproduced in weights	A decoder should not reproduce training content	[C03]
Hamburg Higher Regional Court, Dec 2025	Kneschke v LAION	Dataset building within the mining exceptions; reservation judged at collection	Timing of opt-outs	[C02]
Third Circuit, argued Jun 2026	Thomson Reuters v Ross	Pending	First appellate view on training	[C10]
Munich Regional Court I, Jul 2026	GEMA v Suno	Memorised works infringe; mining defence rejected; not final	Training abroad no shield	[C04]
D. Mass., Aug 2026	UMG and Sony v Suno	Anti-circumvention claim over platform audio allowed	Sourcing from video platforms	[C14]
Federal Court of Justice, 17 Dec 2026	Kneschke v LAION	Pronouncement due	Watch item	[C01]

Holdings as reported by the courts or reliable summaries; several are under appeal. None concerns a codec; each bears on a question a codec's corpus raises.

Table 45. The AI Act after the Omnibus, and comparable instruments elsewhere

Provision	Who it binds	Applies from	Reaches a ULBC codec?	Source
European Union				
General-purpose models, Art. 53	providers of general-purpose models	2 Aug 2025; enforcement from 2 Aug 2026	unlikely: the criteria turn on generality	[L01] [L02]
Training-content summary template	general-purpose providers	24 Jul 2025	a format usable voluntarily	[L03]
Code of practice, copyright chapter	signatories	10 Jul 2025	opt-out practice for collected data	[L04]
Transparency, Art. 50	providers and deployers of synthetic audio	2 Aug 2026	through fidelity: standard editing is excepted	[L06] [L07]
High-risk, Annex III	providers of listed uses	2 Dec 2027	not a listed use	[L02]
High-risk, Annex I	products under listed safety legislation	2 Aug 2028	only as a safety component	[L02]
Data governance, Art. 10	high-risk providers	with the high-risk dates	a benchmark for coverage documentation	[L01]
Elsewhere				
California AB 2013	developers of generative AI systems	1 Jan 2026	untested for a codec	[U03]
Colorado SB 26-189	developers and deployers of decision tools	1 Jan 2027	unlikely	[U04]
Korea, Framework Act on AI	generative and high-impact AI	22 Jan 2026	labelling unaddressed for a codec	[A09]
China, labelling measures	generated content	1 Sep 2025	unaddressed for decoded speech	[A14]

What applies, to whom and from when. The fourth column is this brief's reading for a task-specific codec; it is context and not legal advice.

Table 46. Liability, product safety and cyber resilience

Instrument	Applies from	What it does	Who it reaches	Bearing on the dataset	Source
Product Liability Directive (EU) 2024/2853	products placed after 9 Dec 2026	software a product; component liability; disclosure; rebuttable presumptions	manufacturers, including software-component makers	documentation as evidence	[L13]
General Product Safety Regulation (EU) 2023/988	13 Dec 2024	safety assessment covers learning functionality	manufacturers and distributors	behaviour, not provenance	[L14]
AI Liability Directive (proposal)	withdrawn Oct 2025	-	-	no fault-based AI instrument	[L15]
Cyber Resilience Act (EU) 2024/2847	reporting 11 Sep 2026; full 11 Dec 2027	vulnerability and incident reporting; disclosure policy	manufacturers of products with digital elements	an analogue for dataset remediation	[L16] [L17]
Data Act (EU) 2023/2854	12 Sep 2025	user access to connected-product data	device makers	a possible sourcing route	[L22]

The instruments that decide who answers when a product fails, and what documentation is worth to a manufacturer. Context; not legal advice.

9 Product, supply chain and assurance

For: implementers, operators, regulators.

9.1 Who carries which exposure along the chain

SRC Two of the Madrid papers look at the chain from its downstream end. SP-260912 argues that an implementer of ULBC would inherit compliance obligations arising from training data over which it had no control and into which it may have no visibility [M03]; SP-260919 argues that operators rely on 3GPP, not on one company, to maintain what they deploy [M05]. The legal paper states the implementer's position plainly: industry participants are unlikely to rely on a dataset that does not document third-party rights and regulatory compliance or whose maintenance is uncertain [P03].

INT WELL SUPPORTED · Along the chain from a recording to a call, exposure is heaviest at the two ends (Table 50). Upstream are the speakers and rightholders whose consent is at stake; downstream are the manufacturers who ship products and answer to product and data law. Neither is represented at SA4 as such. Operators sit beyond the manufacturers and carry the service for its lifetime, which is why their concern is maintenance and continuity rather than compliance. The Organisational Partners sit in the middle, as publishers of the specification and its weights.

INT WELL SUPPORTED · Documentation flows down the chain and obligations flow up it. A manufacturer needs documentation from 3GPP; 3GPP can pass on only what a curator or a proponent gives it. Under the proponent-data option that chain runs through the disclosure of a competitor; under an external corpus it runs through the curator's records and terms; under a 3GPP corpus it runs through the partnership itself. The legal paper's 'industry requirement' is this chain described from the implementer's end.

INT CONDITIONAL · What a manufacturer will ask for is predictable from the instruments in Part 7: the rights basis for each use, the personal-data basis, and evidence that the decoder does not reproduce training content or alter speech. If 3GPP can supply a declaration that answers those, the option chosen matters less to implementers than the base briefing's scenarios suggest. If it cannot, implementers will seek the same evidence from the proponent directly. *Signal:* implementer contributions to the PCG group or to ULBC-6.

What a manufacturer needs is documentation, and the formats for it exist.

9.2 Documenting a dataset: the instruments that already exist

SRC Both the legal paper and TTA reach for existing forms. The legal paper likens dataset provenance work to establishing bills of material for software, calls for standardised dataset documentation such as datasheets or transparency records, and names regulators' frameworks and templates for training data as reference points; TTA names the EU template for the public summary of training content [P03] [P02].

CTX The forms exist and are current (Table 48). SPDX 3.0.1 has a Dataset profile - collection process, availability, intended use, known bias, sensitive personal data, declared and concluded licence - and relationship types for a model trained on or tested on a dataset [N09] [N10]. CycloneDX, published by Ecma as ECMA-424 with a second edition in December 2025, carries a machine-learning bill of materials [N11].

Croissant 1.1, of January 2026, adds provenance and usage policies to dataset metadata, beside its responsible-AI vocabulary [N12] [N13]. Datasheets for datasets and the Data and Trust Alliance's provenance standards give the human-readable and the metadata forms [N14] [N15].

CTX Quality and security standards frame the same material. ETSI EN 304 223, the European Standard of December 2025 on baseline cyber security for AI models and systems, asks under its eighth principle for documented sources of training data, the time they were obtained and cryptographic hashes of shared model components [N01] [N02]; ETSI's data supply-chain report and its September 2026 data-quality framework, ISO/IEC 5259 on data quality for machine learning, ISO/IEC 8183 on the data life cycle and ISO/IEC 42001 on AI management complete the set [N03] [N04] [N06] [N07] [N08], and the Commission's request to the European standards bodies includes dataset governance and quality [N05].

INT WELL SUPPORTED · The bill-of-materials analogy therefore has a direct implementation. SPDX and CycloneDX already carry dataset and model profiles, and SPDX can state that a model was trained on a named dataset. A ULBC declaration would need no new format; it would need a profile - a choice of fields and a rule for who fills each - and the choice between the formats is secondary to the choice of fields, which section 9.3 sets out.

The formats exist; the fields are the decision.

9.3 A ULBC dataset declaration, field by field

INT WELL SUPPORTED · A declaration for ULBC can be assembled from what the inputs ask and what the instruments already define. Table 49 sets out 14 fields, expanding the base briefing's dataset bill of materials [B02]: identity and version, source and provenance, a rights inventory, permitted uses, personal data, biometric status, geography and transfers, composition and coverage, separation and quality, the link to the training recipe, maintenance, incident handling, assurance evidence and a public summary - each with the instrument that defines it and the party that would fill it.

CNT ESTABLISHED · Against the 14 fields, TTA asks for 8 outright and 1 in part, the legal paper for 8 and 2 in part, ATIS for 4, and SP-260919 for 1 with 3 in part. Three fields are asked for by none of them: biometric status, geography and transfers and training recipe link.

INT WELL SUPPORTED · The three unasked fields are the ones other parts of this brief found load-bearing. Biometric status decides whether the processing enters the special categories (section 8.3); geography decides whether a corpus can be moved or hosted at all under transfer rules; and the recipe link decides whether the data delivers maintenance or only a necessary condition for it (section 6.5).

INT ESTABLISHED · The declaration works under every option. What changes between options is who fills each field - the proponent, the curator or 3GPP - and who checks it. It is therefore the one mechanism PCG could adopt before settling common against proprietary, and the one on which the three inputs already overlap most.

INT CONDITIONAL · If the new PCG group starts from requirements, a declaration of this kind is where its work would naturally land. If it decides the scenario first, the declaration follows the scenario rather than preceding it. *Signal that would weaken this reading:* a first output of the group that names a corpus without naming the information any corpus must carry.

Documentation answers lawyers; assurance answers the people who depend on the call.

9.4 Emergency calls, lawful interception and what a decoder must support

SRC ULBC serves IMS voice calls over NB-IoT non-terrestrial access through geostationary satellites [G01], at the two bit rates SA4 has fixed, 1.0 and 1.4 kbps [M01]. Its performance requirements stand at 15% and its test conditions include input levels, noise types and signal-to-noise ratios [M01]. SA3's lawful-interception specifications, TS 33.127 and TS 33.128, define the architecture and protocols through which intercepted communications are delivered [G17] [G18].

INT ESTABLISHED · Nothing in the records read for this brief establishes a requirement that a codec's training data be disclosed for emergency or lawful-interception purposes, and the base briefing makes the same point [B02]. The credible links between a learned codec and those regimes run through the decoder, not the corpus: whether a captured bistream can later be decoded with an identified, preserved set of weights, and whether reconstruction can be distinguished from loss concealment or implementation behaviour.

INT WELL SUPPORTED · Table 51 sets out the assurance questions a learned codec raises, with the evidence that would answer each, whether training data is involved and who owns it. Of the questions listed, only those on coverage and on availability of the data after a licence change involve the corpus; the rest concern the weights, their versions, their conformance and their behaviour, and SA4 and SA3 can answer them without resolving the dataset question.

INT CONDITIONAL · For emergency use the fidelity question of section 8.4 and the robustness question of section 6.3 converge: a codec that rebuilds speech at one kilobit per second must be shown not to alter content under stress, noise and impaired speech. *Signal:* test conditions for stressed or impaired speech in ULBC-3 or ULBC-8a, and any liaison from SA3 on version identification.

What this part established

Exposure is heaviest at the two ends of the chain; documentation flows down it and obligations flow up (Table 50).

The formats exist. A ULBC declaration of fourteen fields needs a profile, not a new format, and three of its fields are asked for by no input (Table 49).

Emergency and interception questions run through the weights and their versions rather than through the corpus (Table 51).

Hands forward: the option space for 3GPP - part 10.

Table 47. The ETSI IPR Policy against the ULBC artefacts

Artefact	Rights that may attach	Reached by the IPR Policy	Why	Source
Codec methods	patents	yes: disclosure and licensing undertakings for essential IPR	essential if needed to implement the standard	[G10]
Weights and reference software	copyright	through the provisions on contributions and deliverables	contributed and published	[G10]
Training recordings	copyright, performer and database rights	not as essential IPR	implementers copy the weights, not the data	[G10]
A dataset held as a secret	trade secret, confidentiality	no	excluded from the definition of IPR	[G10]
Personal data in the recordings	data-protection rights	no	not intellectual property	[L10]
Rights flowing into the weights	unsettled	only if courts so hold	the English and Munich courts differ on their facts	[C06] [C03]

Why the partnership's existing instrument for third-party rights does not reach training data. Definitions as the Policy states them; the reading is this brief's.

Table 48. Dataset documentation and data-quality instruments

Instrument	Body and version	What it records	Declaration fields it serves	Source
SPDX 3.0.1 Dataset profile	SPDX; ISO/IEC 5962 edition 2 in progress	collection, availability, intended use, bias, sensitive data, licences	identity, rights, personal data	[N09]
SPDX relationships	SPDX 3.0.1	trained on, tested on	recipe link, separation	[N10]
CycloneDX ML-BOM	Ecma, ECMA-424, 2 nd edition Dec 2025	models and datasets as components	identity, recipe link	[N11]
Croissant 1.1 and RAI	MLCommons, Jan 2026	dataset metadata, provenance, usage policies	provenance, uses, composition	[N12] [N13]
Datasheets for datasets	Communications of the ACM, 2021	composition, collection, uses, maintenance	composition, maintenance	[N14]
Data Provenance Standards	Data and Trust Alliance, v1.0.0	22 metadata fields on source, provenance and use	provenance, uses	[N15]
Training-content template	European Commission, Jul 2025	public summary of training content	transparency summary	[L03]
AB 2013 disclosure items	California, in force Jan 2026	sources, types, IP status, personal data, dates	transparency, geography	[U03]
ETSI EN 304 223	ETSI, Dec 2025	sources and time of training data; hashes of model components	provenance, recipe link	[N01] [N02]
ETSI TR 104 180	ETSI, Sep 2026	data-quality metrics including lineage	quality	[N04]
ISO/IEC 5259-1	2024	data quality for analytics and ML	quality	[N06]
ISO/IEC 8183	2023	data life cycle	maintenance	[N07]
ISO/IEC 42001	2023	AI management system	assurance	[N08]
ETSI GR SAI 002	2021	data supply-chain security	provenance, incident	[N03]

The formats a ULBC declaration could use without inventing one. The fourth column maps each to the fields of [Table 49](#).

Table 49. A ULBC dataset declaration, field by field

Field	Minimum content	TTA	Legal paper	ATIS	Verizon paper	Instruments	Filled by
Identity and version	name, maintainer, release, manifest with hashes	stated	stated	not stated	partly	SPDX Dataset; Croissant; the corpus release identifier [N09] [N12]	proponent or curator
Source and provenance	source categories, acquisition, collection dates and places	stated	stated	not stated	partly	D&TA provenance fields; Croissant PROV-O; EN 304 223 Principle 8 [N15] [N12] [N01]	proponent or curator
Rights inventory	copyright, related, database, contract and confidentiality positions	stated	stated	stated	not stated	SPDX declared and concluded licence; ODRL rights expressions [N09] [K10]	proponent or curator; checked by 3GPP
Permitted uses	training, evaluation, standardisation, implementation, deployment, retraining, redistribution	stated	stated	stated	not stated	ODRL policies in Croissant 1.1; CAWG assertion [N12] [K18]	rightholder, through the licence
Personal data	presence, legal basis, notices, retention, deletion handling	stated	stated	stated	not stated	SPDX sensitive-data fields; Croissant RAI [N09] [N13]	curator; controller records
Biometric status	speaker identifiers, voiceprints, embeddings, health or demographic labels	not stated	not stated	not stated	not stated	SPDX sensitive-data fields; GDPR Art. 9 assessment [N09] [L10]	curator; checked by 3GPP
Geography and transfers	where collected, hosted and processed; transfer basis	not stated	not stated	not stated	not stated	none dedicated; California AB 2013 asks when data was collected [U03]	curator
Composition and coverage	hours, speakers, languages, conditions, known gaps	partly	partly	not stated	stated	Datasheets for Datasets; Croissant RAI; EU training-content template [N14] [N13] [L03]	proponent or curator
Separation and quality	checks against test material, filtering, duplicates	stated	not stated	not stated	partly	SPDX testedOn; ISO/IEC 5259; JVET call rule [N10] [N06] [J01]	proponent; checked by test laboratory
Training recipe link	preprocessing, code, seeds, framework, quantisation, weight hash	not stated	not stated	not stated	not stated	SPDX trainedOn; CycloneDX ML-BOM; EN 304 223 hashes [N10] [N11] [N01]	proponent
Maintenance and availability	responsible entity, retention, availability, versions	stated	stated	not stated	not stated	the legal paper's framework elements [P03]	curator or 3GPP
Incident and remediation	reporting channel, correction, notification, revalidation	not stated	stated	not stated	not stated	coordinated vulnerability disclosure practice under the CRA [L16]	curator or 3GPP
Assurance evidence	level of technical and legal review, sign-off, known limits	not stated	stated	stated	not stated	none dedicated; EN 304 223 audit trail [N01]	verifier
Public transparency summary	summary in a format regulators already read	stated	partly	not stated	not stated	EU training-content template; California AB 2013 [L03] [U03]	proponent or curator
Asked for outright	of 14	8	8	4	1		

Fourteen fields, expanded from the base briefing's dataset bill of materials. The four input columns show which document asks for each field: stated, partly, or not. Three fields are asked for by none.

Table 50. Who carries which exposure along the chain

Party	What it handles	Layers that reach it	What it needs documented	Who speaks for it here
Speakers and rightholders	recordings, performances, texts read	copyright, performer, personal data, personality	the scope of consent and licence	no one at SA4; the curator's terms
Corpus curator	the collection and its distribution	database, contract, personal data as controller	contributor terms, provenance	Mozilla's terms, as TTA reads them
Proponent	training, weights, recipe	every layer, for its own data	a declaration and a recipe	the proponents in SA4
SA4 and test laboratories	the selection	separation; confidentiality of test material	test plans and conditions	SA4
Organisational Partners	the specification and the weights	copyright in deliverables; possibly personal data in weights	terms, disclaimers, evidence	the legal paper, ATIS, TTA
Implementers	products carrying the decoder	product liability, safety, cyber, AI Act transparency	a declaration; fidelity evidence	SP-260912 speaks of them
Operators	a service for decades	continuity; emergency obligations	a maintenance path; version identification	SP-260919
Users and authorities	calls, emergency services, interception	fidelity and evidence	decoder versions	none in the record read

The chain in Figure 7, read by party. Exposure is heaviest at the two ends, neither of which is represented at SA4 as such.

Table 51. Assurance questions a learned codec raises

Question	Evidence that would answer it	Involves training data	Owner
Intelligibility across languages, accents, ages and impaired speech	listening tests under ULBC-3 conditions	coverage only	SA4
Behaviour under stress, noise and unfamiliar speech	test conditions and characterised failure	coverage only	SA4
Content and speaker fidelity	fidelity measures in selection	no	SA4; bears on Art. 50
Decoder conformance	bit-exact or tolerance-based test vectors	no	SA4
Which weights a device is using	a version field or profile signalling	no	SA4 with SA2 and SA3
Decoding a historical capture	every published weight set preserved	no	3GPP and the OPs
Reconstruction against loss concealment	documented decoder behaviour	no	SA4 and SA3
Availability of the data after a licence change or deletion	maintenance fields in the declaration	yes	curator; 3GPP

For operators, regulators and public-safety authorities. Most questions concern the weights and their versions, not the corpus.

10 The option space for 3GPP

For: PCG, OP legal staff, strategists.

10.1 Six arrangements, read against seven criteria

SRC Three sets of options are on file. The legal paper's three levels run from requirements with proponent data, through an external dataset referenced for all, to a dataset 3GPP creates and maintains [P03]. The base briefing lists five, from deferring to SA4 entirely to a hybrid in which a selected proponent deposits data or maintenance artefacts under controlled terms [B02]. And the Madrid papers add sequencing, requirement dimensions and transparency as ways to reach any of them [M03] [M05].

INT WELL SUPPORTED · Table 52 reads six arrangements - the status quo and five options - against seven criteria: assurance to implementers, fairness of the comparison, independence of maintenance from the proponent, freedom to use better data, exposure of the Rel-20 schedule, the burden on 3GPP in the legal paper's terms, and precedent elsewhere. No option leads on every criterion. Figure 11 places them on the two that trade most directly: the burden 3GPP takes on, and how far maintenance can proceed without the proponent.

INT WELL SUPPORTED · Fairness follows maintenance. The arrangements under which the winner trained on data 3GPP can name are the same ones that keep the comparison clean, so the fairness case and the maintenance case - argued separately on the reflector and at Madrid - are one case, and the arrangements that answer one answer the other. Legal assurance does not follow either: it depends on the declaration and its verification, which can be added to any option.

INT CONDITIONAL · Escrow is the one arrangement that raises maintenance independence without raising 3GPP's data burden to the legal paper's medium level,

because it secures access by contract rather than by custody. It is also untested in 3GPP and would need confidentiality, access and competition terms that no document read here drafts. *Signal*: any contribution proposing a deposit by the selected proponent.

CTX The base briefing offers draft wording for PCG: technical need and suitability for SA4; proportionate baseline criteria from interested OP legal representatives, with counsel as needed; governance scaled to whether 3GPP specifies requirements, references a dataset or creates one; and SA4 continuing under any interim criteria the OPs agree [B02]. This brief neither adopts nor opposes it; it notes that the wording is compatible with options A, B, C and E, and not with sequencing SA4 behind PCG.

Every option needs someone to check it.

10.2 Who could verify, and what each would need

SRC Verification is proposed in three forms. ATIS proposes that a group of OP legal representatives verify that datasets meet the requirements [P04]. The legal paper holds that contributor self-declarations alone would likely be insufficient for a 3GPP dataset, proposes a review combining technical and legal assessment, warns that 3GPP's legal resources are limited and suggests external counsel [P03]. TTA doubts that a proponent's self-declaration alone mitigates the risk for implementers [P02]; and SA4 put legal and licensing aspects out of its own scope [M02].

INT ESTABLISHED · Verification is three activities, not one (Table 53). Technical verification checks separation from test material, coverage and, where required, reproduction. Documentary verification checks that a declaration is complete and consistent. Legal verification judges whether the rights and the personal-data basis suffice for the uses declared. They need different people, and only the third needs lawyers.

INT WELL SUPPORTED · The technical checks already have natural owners in the process: the independent

laboratories that run the listening tests, and cross-checkers of the kind JVET assigns to re-train candidate models [J05]. The documentary check needs a checklist and a secretariat. The legal check is the one the legal paper warns 3GPP lacks the resources for, and it is the one ATIS's group is aimed at.

INT CONDITIONAL · A split along these lines would let SA4 own the technical checks now and leave the legal check to the OP group once it exists, so that the Rel-20 schedule depends on the legal track only at the point where a candidate's rights must be confirmed. *Signal*: whether the new PCG group's terms include verification or requirements only.

The legal check is the one the partnership lacks; counsel is the one it can obtain.

10.3 What an external counsel could be asked

SRC The legal paper suggests that seeking advice from external counsel on liability may be appropriate [P03], and the base briefing lists fifteen questions such counsel could be asked, from which legal entity would be treated as publisher of a common dataset to how deletion requests should be handled after training [B02]. PCG#56bis opened, as PCG meetings do, with an antitrust and competition law compliance statement [P01].

INT WELL SUPPORTED · Table 54 groups the questions under eight headings and marks where this brief's reading adds to them: the role of the OPs, the rights needed for each use, personal data in the weights, biometric processing through test metrics, cross-border transfers, deletion after training, product liability, and competition law. Two questions are new against the base briefing's list: whether published weights are personal data (section 8.3), and whether rights in training data flow into weights (section 8.6).

INT ESTABLISHED · Both new questions arise under every option, because every option ends in weights the OPs publish. That makes them the questions to put first: their answers bound the OPs' exposure whichever

scenario PCG chooses, where the scenario questions bound it only for the scenario chosen.

CTX Competition law belongs on the list for a practical reason. A common dataset, a shared declaration or an escrow among competing proponents involves exchanging or withholding information between companies that will compete in the market the codec serves, and the terms of access to an escrow would decide who can maintain the standard. Standards bodies handle such questions routinely; the point is that the terms would have to be written, and none has been.

The last question is how far any answer reaches beyond this codec.

10.4 Precedent: a rule for ULBC or for 3GPP's AI work

SRC Every document that discusses scope beyond ULBC takes a position on precedent. SP-260912 warns that *"how this dependency is handled will set the precedent for future 3GPP work involving AI"* SP-260912, section 2 [M03]. TTA calls ULBC *"the first case of specifying an AI/ML model in 3GPP"* and asks PCG to continue studying other AI/ML functionality [P02]; the legal paper treats SA4 as the first use case, to anticipate similar questions elsewhere [P03]. On the reflector a participant asked SA4 to state that *"this decision SHALL NOT set a precedent for other codecs"* SA4 reflector [W04], message 9 [W04].

CTX 3GPP's other AI/ML work uses training data differently. The radio framework describes data collection feeding model training and inference, including network-side collection for offline training [G13]; the core-network protocol group describes vertical federated learning, in which participants train jointly without exposing raw data [G14]; the architecture group is studying enhanced core-network support for AI/ML in Rel-20 [G15]; and the management group is specifying AI/ML management in a third phase [G16] (Table 55).

INT WELL SUPPORTED · The difference matters less than it looks, because the dependency travels even where the artefact does not. Each of those specifications assumes a model whose behaviour someone can rely on, and a model's behaviour is a property of its training material: a radio model trained on one deployment's measurements is not a model for another; a federated model is only as good as the local data of the partners that trained it, which none of them can inspect; an analytics output is a function of what the network functions feeding it collected; and a management interface that requests training, reports performance and asks for retraining is, in each of those operations, making a statement about a dataset (Table 55). None of those groups holds the material and none has a rule about it. What each has instead is a dependency on whoever does - which is the position ULBC puts to the partnership first, and the reason a rule written here is read in RAN and CT as well as in SA4.

INT ESTABLISHED · ULBC is the only item in the 3GPP record read here in which trained parameters are standardised. Elsewhere 3GPP specifies data semantics, collection, signalling and management, and the training data stays with operators and vendors - in federated designs deliberately so. But the transfer question is a ladder rather than a line (Figure 12, Table 56). The full package - a declaration, a verifier, a maintenance package - attaches where the partnership publishes a trained artefact. Single governance items attach lower down, because a standard that fixes what is collected, how a model is transferred or which conditions a test runs under already depends on shared assumptions about provenance, versioning, integrity and the separation of test material from training.

INT WELL SUPPORTED · Both precedent arguments are therefore partly right. SP-260912 is right that the first case sets a pattern; the reflector participant is right that a ULBC decision need not bind other codecs, which could be designed and trained differently. The two meet if PCG states the threshold at which its outcome applies rather than leaving precedent to be read from a single case.

INT CONDITIONAL · *Signal*: whether PCG frames its outcome as specific to ULBC, as a rule for standardised models - the fifth rung of the ladder - or as a rule for all 3GPP AI/ML work. The third would reach work whose training data 3GPP never sees, and would be the outcome least consistent with the way 3GPP's other AI/ML specifications are built. *What would weaken this reading*: a PCG text that names governance items rather than a scope - identifiers, provenance, separation - which would transfer down the ladder without reaching for custody at the top of it.

10.5 What 3GPP has already approved about training data

CNT The question can be put to 3GPP's own approved record. An export of every 3GPP temporary document whose visible text names AI, a model and training data, and whose recorded decision is approved, agreed or endorsed, returns 1,000 records between August 2021 and August 2026 [B06]: 674 approved, 259 agreed and 56 endorsed, across 248 meetings, 43 specifications, 71 work items and 300 recorded sources. 266 of them are change requests to a specification rather than study material (Table 66).

CNT The work is not evenly spread. Management and orchestration carries 362 of the 1,000 records and the single most-changed specification, TS 28.105 with 207 of them [G19] [G20]; the core network and its analytics carry 166 [G21], the radio groups 150 [G22], security 50, application enablement 47, service requirements 45, and the plenary itself 100. SA WG4 carries 80, of which 29 are the media AI study [G23] and 4 are agreed versions of the ULBC study report itself [G24].

INT ESTABLISHED · What that record standardises is the handling of training data, not the data. Its vocabulary is operational: training data in 728 records, model training in 430, data collection in 194, retraining in 108, a dataset in 98, validation in 92 and federated learning in 85. The interfaces request training, describe the data a model was trained on, report how a model performs and ask for it to be retrained; the data itself stays with the operator or vendor that holds it (Table 67).

CNT ESTABLISHED · The words a rights discussion would need are absent. In the text the export carries - each record's title, its keyword chips and the excerpt that matched - a licence, copyright, intellectual property, personal data, data protection, provenance, ownership, attribution and open source appear in none of the 1,000. 'Rights' appears in 2, in both cases in its ordinary sense - the right network function, not available right away. What does appear is narrower and procedural: consent in 25 records, which is authorisation to collect or use network data; privacy in 18, mostly in the security study of AI/ML services; regulation in 6; and anonymisation once, as an enumerated data-processing operation in an interface description.

INT WELL SUPPORTED · That is the measure of what PCG was asked for. Five years of approved 3GPP work on AI and training data has needed no vocabulary of licensing, provenance or personal data, because in all of it the partnership specifies how data is handled by others and never holds or distributes the data itself. ULBC is the first case in which the partnership's own deliverable would carry the residue of someone's

recordings, which is why the question arrives at PCG rather than at a working group, and why nothing in the approved record answers it.

INT WELL SUPPORTED · Read the other way round, the same count says how widely the answer would be read. Every family in Table 57 - radio, core and analytics, management, security, service requirements, application enablement - already specifies a model whose behaviour follows from data the group does not hold, and the 1,000 records are what that dependence looks like written down: requests to train, descriptions of what a model was trained on, reports of how it performs, calls for retraining. None of those groups has needed a rule about the material, because none of them publishes it. A delegate in RAN or CT therefore has no stake in ULBC as a codec and a direct stake in it as the first place the partnership writes down what it requires of training data (Table 55, Figure 12).

INT CONDITIONAL · *Signal*: the vocabulary of any text PCG's new group produces. A rule written in the operational vocabulary above would sit inside the existing record and reach it; a rule written in the vocabulary of licences and rights would be new to 3GPP, and would apply to whatever the group's threshold names - which is the precedent question section 10.4 leaves open. *What would weaken this reading*: a record outside this export - the full text of a specification rather than its temporary documents - that does carry the rights vocabulary.

What this part established

No option leads on every criterion; fairness follows maintenance, and escrow secures maintenance by contract rather than custody (Figure 11).

Verification is technical, documentary and legal, and only the last needs lawyers (Table 53). Two questions for counsel arise under every option and belong first (Table 54).

A ULBC outcome transfers only where a model is standardised, and PCG can say so (Table 55).

3GPP's own approved record on AI training data - a thousand decisions over five years - specifies how data others hold is handled, and carries no vocabulary of licences, provenance or personal data at all (Tables 57 and 58).

Hands forward: how the consensus is forming, and who is standing where - part 11.

Figure 11. Six arrangements placed by burden on 3GPP and by maintenance independence



Both axes are ordinal and carry no ticks or distances; positions are this brief's reading of Table 52 and are interpretation. The corner names describe regions, not recommendations.

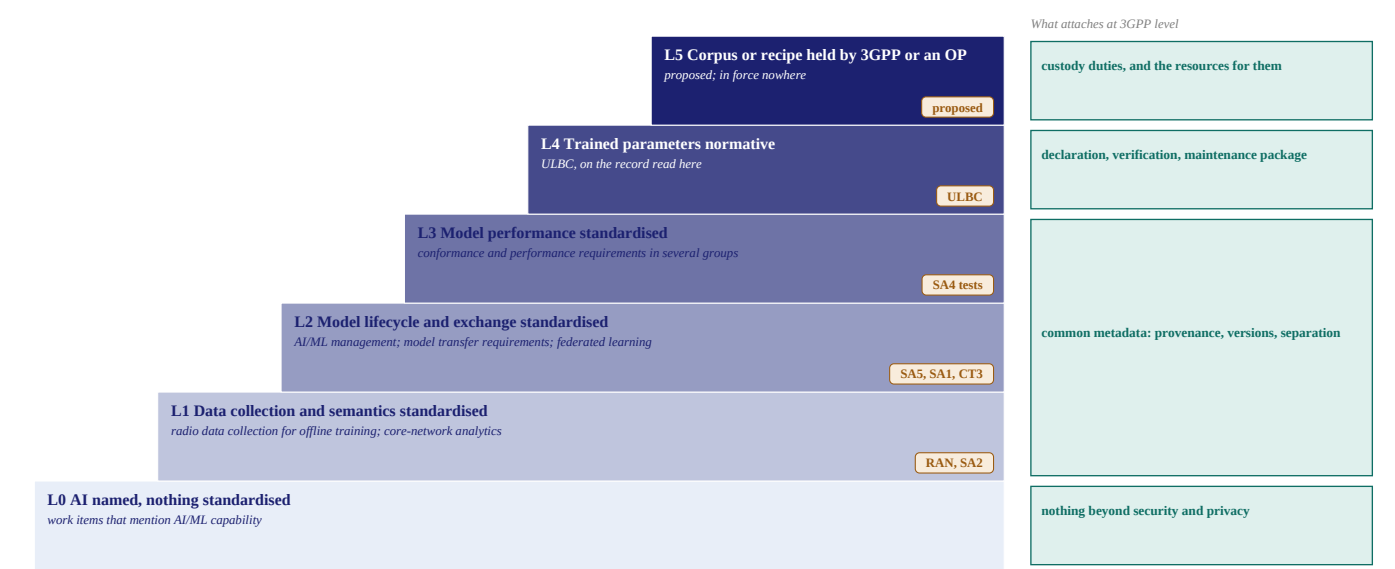
The trade the options make, for PCG and OP counsel. Independence rises with burden, except for escrow, which secures access by contract rather than custody.

Table 52. Six arrangements read against seven criteria

Arrangement	Implementer assurance	Fair comparison	Maintenance independence	Freedom to use better data	Rel-20 exposure	Burden on 3GPP	Precedent
No rule	low	low	low	full	none	none	earlier audio practice, as message 11 describes it
A. Disclosure and separation	low to medium	medium	low	full	low	more limited	JVET joint call
B. Verified legal baseline	medium to high	medium	low	full	medium: verification before selection	more limited to medium	none found
C. External common set	medium to high	high	medium to high, with the recipe	limited for primary results	medium: a corpus to choose	medium	JPEG AI; JVET NNVC
D. 3GPP corpus	high, if resourced	high	high	limited	high	high	none found
E. Escrow for the winner	medium	medium	high	full	low to medium	more limited, plus contract terms	none in the bodies read

Interpretation throughout: each cell is this brief's ordinal reading, with the legal paper's own burden terms in the seventh column. No cell is a score, and no row is totalled.

Figure 12. When training data becomes a 3GPP question: six levels of dependence



The six levels are this brief's ordinal scale for how far a standard depends on a trained model; step height carries order only. The chips name where 3GPP work sits today. Read from Table 56.

The transfer question drawn as a ladder rather than a line, for PCG and for any group that wonders whether a ULBC rule would reach it. The chips say where 3GPP work stands today: RAN and SA2 at collection, SA5, SA1 and CT3 at lifecycle and exchange, SA4 at performance, ULBC alone at normative parameters. The full package attaches only at the level where the partnership publishes a trained artefact; single governance items attach much earlier, which is where most groups would meet a rule written here.

Table 53. Who could verify, and what each could check

Candidate	Technical	Documentary	Legal	Resource basis	Proposed or described by
SA4	separation, coverage	partly	no: out of its scope	the working group	SP-260909
Test laboratories	separation checks, test conditions	no	no	paid by the candidates	this brief's reading
Cross-checker, as in JVET	reproduction by re-training	no	no	a member other than the proponent	JVET practice [J05]
Secretariat with a checklist	no	yes	no	support staff	this brief's reading
Group of OP legal representatives	no	yes	yes	OP legal staff	ATIS
The new PCG ad hoc group	framing	framing	framing	PCG	SP-260912; SP-260919
External counsel	no	no	advice	a budget	the legal paper
Self-declaration	partly	partly	no	the proponent	insufficient alone: TTA; the legal paper

Verification split into its three kinds. Only the legal check needs lawyers, and it is the one the legal paper warns the partnership lacks the resources for.

Table 54. Questions an external counsel could be asked

Heading	Question	Arises under	In this brief	Source of the question
Role of the OPs	Does publishing normative weights make an OP a distributor or component supplier under product and data law?	every option	8.5	new
Personal data in the weights	Are published weights personal data, given the risk of extraction?	every option	8.3	new
Rights into weights	Can rights in the training recordings reach the weights?	every option	8.6	new
Biometric processing	Do speaker-similarity metrics or speaker labels make the processing biometric?	every option	8.3	new
Rights for each use	What licence scope covers standardisation, evaluation, implementation, deployment and retraining?	every option	8.1	base briefing
Performer and phonogram rights	How are they separated from copyright in the text read?	every option	8.1	base briefing
Disclaimers	Which risks cannot be disclaimed between businesses?	every option	4.3	base briefing
Representations	What can OPs state about provenance without creating a warranty?	B, C, D	4.3	base briefing
Transfers	Which transfer rules apply to a corpus hosted or curated across borders?	C, D	9.3	base briefing
Deletion after training	How is a withdrawal handled once weights exist?	C, D, E	8.3	base briefing
A restricted corpus	How does the standard respond if a court or regulator restricts part of a corpus?	C, D	8.2	base briefing
Competition	What access and information terms do a common set, a shared declaration or an escrow among competitors need?	C, D, E	10.3	base briefing

Twelve questions under their headings. Those marked new arise from this brief's reading and apply under every option; the rest are from the base briefing's list.

Table 55. Every TSG already specifies models: what each depends on, and whether a ULBC rule would transfer

Group	Work	What 'training data' means there	What depends on it	Model standardised	Would a ULBC rule transfer	Source
SA4	ULBC-MED	a speech corpus behind a normative codec	the codec's quality, its language and noise coverage, and whether it can be retrained at all	yes: the weights	-	[G01]
RAN	AI/ML for the air interface	measurements collected for training and inference	a model trained on one deployment's measurements is not a model for another	no: collection and signalling	only above a normativity threshold	[G13]
CT3	AI/ML data; vertical federated learning	local datasets, trained without exposing raw data	the joint model is only as good as each partner's local data, and no partner can inspect another's	no	no: the data stays local	[G14]
SA2	core-network AI/ML, phase 2	analytics data in the network	an analytics output is a function of what the network functions feeding it collected	no	no	[G15]
SA5	AI/ML management, phase 3	model lifecycle and performance data	interfaces that request training, report performance and ask for retraining are each making a statement about a dataset	no	no	[G16]
SA1	requirements for transferring and splitting models	the model as a thing moved between parties	a receiver can rely on a transferred model only if it can evaluate what it was trained to do	no	no	[B06]
SA6	application-layer enablers for training and model sharing	data held by applications and their members	an enabler that lets applications train and share models presumes the material behind them is usable	no	no	[B06]
SA3	AI/ML security; lawful interception	not training data: threats to AI/ML services, and intercepted content	a protection is written for data that someone else holds	no	no: the link is decoding	[G17] [G18]

Why this is not only SA4's question. A model's behaviour is a property of the material it learned from, so each of these specifications rests on training data that the group writing it neither holds nor governs - the material stays with operators and vendors. Only ULBC standardises the trained parameters themselves, which is why the rule is written there first and read everywhere else. The fourth column is this brief's reading of what each group's own work rests on; the rest is the record.

Table 56. Six levels of dependence, and what attaches at each

	Level of dependence	What 3GPP standardises	Where this exists today	What attaches at 3GPP level	Whose question it is
L0	AI named, nothing standardised	nothing about data or models	work items that mention AI/ML capability	nothing beyond the usual security and privacy duties	the vendor or operator
L1	Data collection and semantics standardised	what is collected, how it is labelled, and the roles that collect it	radio data collection for offline training; core-network analytics	provenance, quality, retention and separation as common metadata; the data itself stays local	the working group, technically
L2	Model lifecycle and exchange standardised	model identifiers, transfer, retraining triggers and management	AI/ML management; model transfer requirements; federated learning	versioned lineage, integrity and compatibility metadata across groups	cross-group coordination
L3	Model performance standardised	tests a model must pass, and the conditions they run under	conformance and performance requirements in several groups	separation of test from training material; coverage and robustness conditions	the working group, technically
L4	Trained parameters normative	the weights themselves, as part of the specification	ULBC, on the record read here	declaration of the data and the recipe, verification, a maintenance package, immutable identifiers	the partnership, with the working group
L5	Corpus or recipe held by 3GPP or an OP	custody of the material itself	proposed; in force nowhere	controller duties: retention, deletion, transfer, access and the resources to run them	the Organisational Partners

Where a training-data rule would bite, level by level, and the answer for a group that has never held training material but depends on one that has. Most 3GPP AI/ML work sits at levels 1 to 3, where single governance items attach - provenance, versions, the separation of test material from training. Only the top two levels raise questions the working group cannot answer on its own.

Table 57. The approved 3GPP record that names AI model training data, by group family

Group family	Records	Most-changed specification	What its records specify	Where the training data sits
Management and orchestration (SA5)	362	TS 28.105 (206)	interfaces that request training, describe the data a model was trained on, report how it performs and ask for retraining	with the operator running the management system
Core network and analytics (SA2, CT)	166	TR 23.700-80 (19)	analytics functions, provisioning of models between network functions, and federated training that never moves the data	with the network functions that collect it
Radio access (RAN1-5)	150	TR 38.843 (18)	collection of measurements for training and inference, and life-cycle management of models over the air interface	with the device and the network that collect it
Security (SA3)	50	TR 33.738 (8)	threats to AI/ML services and the protections against them, including authorisation to use collected data	unchanged: the studies protect data others hold
Service requirements (SA1)	45	TR 22.876 (8)	requirements for transferring and splitting models, and the 5G system's support for them	with whoever supplies the model
Application enablement (SA6)	47	-	enablers that let applications train, share and transfer models across the application layer	with the application and its members
Media and codecs (SA4)	80	TR 26.927 (20)	studies of AI in media services, and the agreed versions of the ULBC study report itself	the open question this brief reads
Plenary decisions (TSG SA)	100	-	approval of the work items, study reports and revisions above	not addressed in the decision
All records	1,000	TS 28.105 (207)	266 of them change a specification; the rest are study material	outside 3GPP in every family above

What five years of approved 3GPP decisions on AI training data are about, counted from the export. Every family here - radio, core, management, security, service requirements and application enablement - specifies how a model is trained, moved, managed or protected, and in every one of them the data itself stays with someone else. That is the dependency a rule written for ULBC would be read against.

Table 58. One fixed vocabulary against 1,000 approved records

Term	Records	What the hits are	Share of the record
------	---------	-------------------	---------------------

The words the approved record uses, and the words it does not. A term with no hits is kept and reported as none, because the absence is the finding.

		Operational vocabulary	
training data	728	the keyword the export selects on; present in the record's own text	72.8%
model training	430	training as an operation the network requests or performs	43.0%
data collection	194	collecting measurements or network data for training	19.4%
retraining	108	retraining after drift, defect or release change	10.8%
dataset	98	a dataset as an object an interface names or a model is trained on	9.8%
validation	92	validation data and validation of a trained model	9.2%
federated learning	85	training across participants without moving the data	8.5%
model transfer	41	moving a model between a network and a device	4.1%
labelling	37	labels attached to training data	3.7%
quality	34	data quality as a property a management interface reports	3.4%
test data	13	data used to test rather than train a model	1.3%
		Procedural vocabulary	
consent	25	authorisation to collect or use network or user data	2.5%
privacy	18	privacy as a threat class in the security studies	1.8%
regulation	6	regulatory context named in a study's framing	0.6%
third party	5	a third party as a source of data or a model	0.5%
bias	2	biased outcomes as a trustworthiness concern in management studies	0.2%
anonymisation	1	an enumerated data-processing operation in an interface description	0.1%
		Rights and custody vocabulary	
curation	none	no record describes curating a corpus	-
licence or licensing	none	no record states a licence for training material	-
copyright	none	no record raises copyright in training material	-
rights	2	twice, and in its ordinary sense: the right function, right away	0.2%
intellectual property	none	no record raises intellectual property in the data	-
personal data	none	no record calls training data personal data	-
data protection	none	no record names a data-protection duty	-
provenance	none	no record asks where the data came from	-
ownership	none	no record asks who owns it	-
attribution	none	no record requires attribution	-
open source	none	no record names an open-source or public corpus	-

11 How the consensus is forming, and who is standing where

For: everyone interpreting the record.

11.1 How each body decides, and why the same word can mislead

CTX The same technical proposal can meet a different decision threshold in each of these bodies, because they do not reach agreement through one common machine. 3GPP is consensus-first rather than unanimity-based: under Article 25 of the Working Procedures a Technical Specification Group or working group shall endeavour to reach consensus, and if consensus cannot be achieved the Chair may decide to take a vote, which 71% of the votes cast approves. Article 19 gives PCG the same structure at partnership level, with one vote per Organisational Partner and the same threshold [G26]. The Organisational Partners then publish the result as their own deliverable under Article 46 [G25]. Nothing in that machine decides a question of law, which is why the referral went sideways to PCG rather than upwards to a plenary (Table 60).

INT ESTABLISHED · A sustained objection is therefore consequential without being a veto. It can prevent a consensus finding and consume meeting and schedule margin while the group tries to resolve it, and that is a real cost against a calendar the Chair already calls thin. But the formal rules provide a fallback, so this brief writes *consensus remains open* or *unresolved objection* rather than *blocked*, unless a meeting record itself uses that word. Whether a group moves from consensus-seeking to the vote fallback is itself an observable procedural event, and a different one from reaching agreement (Table 59).

CNT ESTABLISHED · The rule and the practice are not the same thing, and this brief can say how far apart they sit in the record it reads. Across the 1,000 approved 3GPP records in the keyword export [B06], every recorded disposition is one of four - 674 approved, 259 agreed, 56 endorsed and 11 partially approved - and none is recorded as carried by a vote. Searching all 1,257 records this brief holds, including the SA WG4 reports, the reflector thread, the PCG inputs and the Madrid papers, returns no mentions of a vote, a ballot or an abstention. The Article 25 fallback is a rule this corpus does not show in use.

INT WELL SUPPORTED · Two readings of that are open, and the record does not choose between them. On the first, the fallback is available and has simply not been needed, so an unresolved objection delays a decision and nothing more. On the second, a body that in practice settles by agreement is one in which a sustained objection on a contribution operates closer to a stop than to a delay, because the route around it is not one this record shows being taken; on that reading the difference between 3GPP and a body that ballots is larger than the written procedures suggest, since MPEG, JPEG and IEEE reach a point where numbers decide and this corpus shows 3GPP reaching no such point. What the record establishes is narrower than either reading: the fallback exists, and nothing here shows it used. *Signal*: a meeting report that records a vote being called on a ULBC Pdoc. *What would weaken this reading*: a vote recorded in any 3GPP report outside this corpus, which would show the count above is a property of what was read rather than of the practice.

CTX The ISO/IEC coding bodies separate expert convergence from formal approval. MPEG and JPEG working groups operate by expert consensus; the parent subcommittee is composed of national-body delegations, and at the enquiry and final stages

numerical criteria take over - approval where two-thirds of the P-members voting are in favour and no more than a quarter of all votes cast are negative [K28]. ISO/IEC also states plainly that sustained opposition must be addressed but is not a right to veto [K29]. JVET adds a second layer: its Terms of Reference make all group decisions by consensus as determined by the Chairs, and all of its output then undergoes the approval process of each parent body [K27] - a study group's, normally by the alternative approval process [K32], and a subcommittee's ballot.

CTX The internet bodies use another architecture again. An IETF working group runs on rough consensus judged by its Chairs: an objection that has been considered and answered does not prevent progression, and a headcount does not settle the question either way [I30]. Its neural speech-codec work carries a normative prohibition on training with the test clips in a draft that is, by its own status, work in progress [I15] [I16]. W3C also centres consensus and sustained objection but keeps a formal-objection record and a vote available once consensus methods fail [I31]; its reservation work is a community report rather than a Recommendation [I28]. IEEE and ETSI expose their thresholds most directly, at 75% of an IEEE SA standards ballot [K30] and 71% of the votes cast in an ETSI Technical Group [K31].

INT ESTABLISHED · Three consequences follow. First, a call condition, a working-group requirement and a partnership-level rule are different procedural objects and should not be compared by elapsed time alone: a call of proposals can carry entry conditions long before the resulting technology completes any approval path, which makes a call condition a different object from a 3GPP working-group requirement, not a faster version of one. Second, every learned-coding body in this record has placed its training-data controls at the competition or experiment stage, where they can be checked against submissions [J01]. Third, a disputed 3GPP rule can consume schedule margin while consensus is sought; escalation to the 71% fallback exists, but using it is a different procedural event from reaching consensus (section 4.8).

CTX The record behind ULBC contains two lineages, and they are not one cycle (Figure 14, Table 61). In the first, 3GPP writes requirements and runs a contest to select a codec of its own: AMR was launched at ETSI SMG#23 in October 1997 and selected in September 1998 [G33]; AMR-WB ran from a work item at TSG SA#2 in March 1999 through qualification at TSG SA#8 to selection at TSG SA#10 in December 2000 and characterisation through 2001 [G34]; EVS opened on 25 March 2010, qualified candidates at SA4#72bis in March 2013, selected at SA4#80bis in August 2014 and had its specifications approved at TSG SA#65 that September [G27] [G29] [G32]; IVAS opened on 8 September 2017, moved from Rel-17 to Rel-18 in December 2021, and was selected at SA4#125 in August 2023 and approved at SA#101 [G28] [G30] [G31]. ULBC belongs to this lineage.

CTX In the second, a joint team of ITU-T and ISO/IEC develops a codec 3GPP did not invent, and 3GPP evaluates and profiles it afterwards. JVT developed H.264 from 2001 and completed its first version in May 2003 [K33]; JCT-VC first met in Dresden in April 2010 and H.265 was approved in April 2013 [K34]; an informal exploration team ran from October 2015 to October 2017 [K35] before the formal JVET was created on 27 October 2017 and carried H.266 to approval on 29 August 2020 [K36]. 3GPP's own act came later each time: H.263 mandatory in Release 1999 [G37], H.264 profiled in Rel-6, HEVC evaluated in TR 26.906 and taken into MTSI in 2015 [G38], and VVC characterised in TR 26.955 - a study record rather than a normative service profile [G39]. The names are often

confused: H.264 came from JVT and H.265 from JCT-VC, neither from JVET.

INT ESTABLISHED · One episode in that second lineage bears directly on the process question. In September 2004 the paper taking the H.264 adoption changes to TSG SA#25 records that consensus had not been reached at the SA WG4 meeting in Prague on two issues, that the changes were therefore not presented to the plenary by SA WG4, and that supporting companies brought them to TSG SA directly - covering messaging, streaming, file storage, conversational services and 3G-324M [G36]. An unresolved objection can therefore change the route a document takes and the time it takes, without ending the work. It is the closest thing in this record to a worked example of what section 1.4 describes as a process dependency.

INT ESTABLISHED · Put in years, the record says something the generation gaps hide. On the native side the run has lengthened: AMR took a year from launch to selection, AMR-WB 1.8 years to selection and 2.8 to characterisation, EVS 4.5 years to approved specifications, IVAS 6.0 years to selection. On the video side each joint team was comparatively quick - JVT reached H.264 in 1.5 years, JCT-VC reached H.265 in 3.0, and H.266 took 4.9 counting the two-year exploration before the formal team - but 3GPP's own act came afterwards and took longer each time: 1.3 years after H.264 was approved, 2.2 after H.265, 3.7 after H.266. End to end, from a joint team's first work to the 3GPP deliverable, the three video generations ran 2.8, 5.2 and 8.5 years, and the last of those ends in a characterisation rather than a normative service profile (Figure 14, Table 62). ULBC's recorded plan is one year.

INT CONDITIONAL · Three clocks have to be kept apart, and Table 62 keeps them apart: the interval between codec generations, the time a team spent developing one, and the time 3GPP took to select or adopt one. The gap between H.264 and H.265 is about ten years, but JCT-VC's core work ran about three. None of these programmes was a learned codec whose weights would be normative, and each differed from ULBC in scope, starting maturity, candidate structure and requirements. Their durations are context, not a forecast, and nothing here establishes that the recorded ULBC plan can or cannot be met. What they do illustrate is the shape of the cycle: requirements, candidate development, qualification and selection, independent testing, verification and specification completion (Table 64). Two of them also show that a selection need not end in one winner - AMR-WB+ and Enhanced AACPlus were both recommended, because no consistent overall ranking was possible [G35]. *Signal*: whether the qualification and selection Pdocs reach agreed versions at SA4#138 and SA4#139. *What would weaken this reading*: a compressed selection that closes those documents in one meeting, which would show the ULBC cycle is not comparable to any of these precedents.

11.2 What crosses between bodies: the liaison record

CNT ESTABLISHED · What actually crosses between these bodies is narrower than the overlap of their subjects, and it is traceable. Table 63 lists what this brief could follow: a joint call issued by two bodies at once; an SEI message proposed in one body to carry another body's content credentials; a usage-restriction message that began as a signalling proposal two years earlier; datasets named in an internet draft and reused as a criterion in a codec competition; and, inside 3GPP, a referral from a plenary to the partnership and a draft liaison back to a working group.

CTX The traffic runs in one direction more often than the other. Technique and test conditions move freely: JVET and MPEG share a call, MPEG's common conditions are reused, the IETF's separation rule is the same rule SA4's survey supported. Rights vocabulary does not move at all - no instrument in the record carries a corpus-clearance obligation from one body to another - and the one place it appears, in content credentials and usage signals, it arrives from outside the coding bodies altogether, from C2PA and W3C [J08] [I28].

INT WELL SUPPORTED · That asymmetry is the mechanism behind the pattern this brief keeps finding. Bodies copy what they can verify. A test-separation rule can be checked by running the test; a disclosure requirement can be checked by reading the submission; a claim about who holds rights in a recording cannot be checked by any body that does not hold the recording. The absence of a legal layer across eight bodies is not an oversight but a property of what a standards body can enforce (section 7.6).

INT ESTABLISHED · For 3GPP the practical reading is that borrowing is available on the technical side and unavailable on the legal side. The separation rule, the disclosure condition, the re-training cross-check and the artefact manifest can be lifted almost as they stand from JVET, JPEG AI and the IETF. The declaration, the verifier and the custody question have no donor, and are the part PCG would have to write.

INT CONDITIONAL · *Signal*: a liaison statement in either direction between 3GPP and a coding body on training data specifically. None exists in the record read here. *What would weaken this reading*: a rights-side instrument crossing between bodies - a common licence template for training corpora, say - which would show the legal layer can travel after all.

11.3 The company lens: seven roles, and what each asks for

CNT The company lens is not one lens. The record shows the same organisations in different roles: 7 organisations argue in the reflector thread, 24 answer the survey anonymously, 15 appear in the codec record as source, co-source or editor, and three Organisational Partners speak for members rather than for themselves (Table 29, Table 24). A company can hold two positions at once without inconsistency, because its exposure differs by role.

INT WELL SUPPORTED · Seven roles run through the record, and each asks for something different. A codec proponent wants a comparison it can win and a corpus it need not publish. A competing proponent wants the data advantage removed. An operator wants the codec to outlive the company that trained it. A device or chipset implementer wants a licence to embed the weights, a version it can pin and evidence it can show a regulator. A test laboratory wants material that has not leaked into training. A dataset curator wants its licence terms respected by a standard that names it. And a rightholder or speaker wants the recording's terms to survive the training (Table 65).

INT ESTABLISHED · Read that way, the argument in the reflector thread is not a disagreement about facts. Every position in it is consistent with the same evidence and with a different seat: the objection that timing is not a technical reason, the reply that a common corpus does not guarantee performance across languages, the point that access to data is not sufficient to replicate training, and the request that a decision set no precedent for other codecs are each true of the seat that made them (Table 16).

INT WELL SUPPORTED · The economics differ by role too, and they are what the positions ultimately track. A common corpus moves cost from every proponent to

whoever assembles and licenses it. A declaration moves cost to the proponent and to whoever verifies it. Escrow moves cost to the winner and to a custodian, and only after selection. Custody by 3GPP moves cost to the partnership permanently. No option is free, and the option space in part 9 is in part a question about who pays (Table 52).

INT CONDITIONAL · *Signal*: whether implementers - not proponents - contribute to PCG's new group or to ULBC-6. Implementer contributions would show the evidence question has reached the seat that has to ship the product; their absence would leave the rule written by the parties competing to supply the model. *What would weaken this reading*: a proponent asking for a stricter declaration than its competitors, which would show the roles do not predict the positions as neatly as the record suggests.

11.4 The government lens: five questions, eight answers

CTX Governments are answering five questions about training data, and answering them differently. May a model be trained on lawfully accessible material by default, or only with a licence? Can a rightholder reserve, and how must the reservation be expressed? Must anything be disclosed publicly about what a model was trained on? What must be available to a regulator even if it is not public? And what continues to bind once the trained model is inside a product (Table 42)?

CTX The answers do not cluster by region as neatly as the debate suggests. The European Union permits mining of lawfully accessible works with a machine-readable reservation and adds a public training-content summary for general-purpose models [L11] [L03]; the United Kingdom's exception reaches non-commercial research only and its government has consulted on four options without choosing [A01] [A02]; the United States decides fair use case by case while California requires disclosure from generative systems [U02] [U03]; Japan permits non-enjoyment uses within limits [A05]; Korea published a fair-use guide [A10]; China requires lawful sources for public generative services and treats a voiceprint as sensitive personal information [A11] [A13]; and India has a proposal, not a law [A17].

INT ESTABLISHED · For a corpus that must serve one global standard, the operative rule is the strictest cell in every column it touches. A dataset assembled under Japan's exception and used to train a codec deployed in the Union has to answer the Union's reservation and personal-data questions; a corpus of United States origin used in a product placed on the European market carries the product-evidence duties of that market. That is why the legal paper's instinct - settle the requirements before the scenario - is sound whatever one thinks of its three levels (section 4.3).

INT WELL SUPPORTED · What governments are optimising for is visible behind the mechanisms. Where innovation and access dominate, the default is permission with a reservation. Where rightholder remuneration dominates, licensing proposals appear. Where accountability dominates, disclosure and regulator access appear. Where sovereignty dominates, lawful source and local processing appear. A standards body cannot satisfy all four at once with a single corpus, and it does not have to: a declaration can carry what each regime needs to see, which is the argument for documentation over custody (Part 8).

INT CONDITIONAL · *Signal*: the German Federal Court of Justice's decision in December and the English appeal in *Getty v Stability*, which will move the reservation and memorisation questions in two of the three regimes any global corpus must satisfy [C01]

[C06]. *What would weaken this reading*: convergence - a mutual-recognition arrangement or an international instrument on training-data transparency - which would make one corpus lawful everywhere and remove the reason for the strictest-cell rule.

11.5 Why positions sit where they do: six standing conditions

INT WELL SUPPORTED · Positions in this record are better predicted by six standing conditions than by argument. The political condition is where a body's members sit and what their governments are optimising for. The legal condition is the instrument the actor already has - an IPR policy, a licence, a regulator's file. The economic condition is who pays for the option. The technical condition is what the actor can verify. The supply-chain condition is who has to show evidence to whom. And the vertical condition is what the codec is for: emergency-capable voice over satellite reaches public-safety authorities that no other codec argument reaches (Table 65).

INT ESTABLISHED · Read against the three Organisational Partner inputs, the six conditions account for the differences the papers themselves do not explain. The partner that publishes the specification writes about the publisher's exposure; the partner whose members compete writes about the fairness of the competition; the partner whose legal culture runs on documented compliance writes a requirement that can be verified. None of the three argues from its own jurisdiction's law, and none needs to: the instrument each holds is enough to explain the paper it wrote (section 4.8).

CTX The supply-chain condition is the one that will grow. A trained decoder ships inside devices and networks placed on markets that increasingly ask for traceable evidence: product liability from December 2026, cyber-resilience reporting since September 2026, and market-surveillance interest in what a model does and how it was produced [L13] [L16]. The implementer that has to answer those questions is not represented in any of the three PCG inputs, and is the seat with the least discretion (Table 50, Table 65).

INT WELL SUPPORTED · The vertical condition is specific to this codec and is under-argued. ULBC-MED is specified for IMS voice over non-terrestrial narrowband access; the traffic it carries includes emergency calls, and its intelligibility under stress is a public-safety property rather than a quality-of-experience one (section 9.4). A dataset rule that improves coverage of accents, languages and noisy channels is therefore not only a fairness measure; it is the mechanism by which the codec works for the speakers who most need it to work.

INT CONDITIONAL · *Signal*: whether any Madrid or PCG#57 text names an actor outside the proponents - an implementer, a regulator, a public-safety authority - as someone the evidence is for. *What would weaken this reading*: a position in the record that cuts against its holder's six conditions, which would show that argument is doing more work here than standing.

What this part established

The bodies decide by different machines and on different clocks: a condition of entry can be written in weeks, a requirement in a specification takes a release, and a rule about rights takes the seven partners (Table 60).

What crosses between bodies is technique and test conditions; rights vocabulary does not travel, and where it appears it arrives from outside the coding bodies (Table 63).

Seven company roles ask for six different things, and each position in the reflector thread is consistent with the seat that made it (Table 65).

Governments are answering five questions differently, so a single corpus is governed by the strictest cell in every column it touches (Table 42).

Two reflector threads are now printed in one format: the working group arguing over the wording of its own report in August, and the plenary arguing over whether to instruct it in September. Each row carries what the message does, this brief's reading and the sender's own words (Table 16, Table 8).

Madrid is the first occasion in this record on which the consensus-first machine ran to its end inside one meeting: a text was drafted four times, the objections were recorded at each step, and no vote was called (Table 60, Table 8).

Hands forward: what the record establishes, and what could follow - part 12.

Figure 13. Five decision architectures: where each family of bodies puts its decisive gate



Lanes are families of bodies that reach agreement the same way; boxes are the gates a proposal passes, in order. Arrows are sequence, not duration, and box width carries no quantity. The lanes are not ranked for speed or for rigour: the figure exists so that a milestone in one lane is not read as the equivalent of a milestone in another. Rules as the cited procedures state them [G26] [K27] [K28] [K30] [K31] [K32] [I30] [I31].

Why the same word carries different weight in each body. A proposal that has cleared the decisive gate in one lane may be years from the artefact at the end of another, and the gate a body treats as decisive is what section 11.1 compares. 3GPP and ETSI seek consensus and keep a numerical fallback; MPEG and JPEG place expert agreement before member-body ballots; JVET separates the joint team's technical consensus from the approval machinery of its two parents; the IETF asks whether objections have been answered rather than how many voted; IEEE makes the ballot itself the record of consensus.

Table 59. What a signal of progress establishes, and what it does not

Observable signal	What it establishes	What it does not establish	How this brief writes it
Many contributions, many companies	interest, contest and the salience of the issue	consensus, direction, or the likelihood of adoption	<i>activity increased</i>
A 3GPP TDoc is presented or noted	the issue reached the agenda and entered the record	agreement with what it proposes	<i>presented, or noted</i>
A 3GPP Pdoc, CR or text is agreed	agreement at that procedural level, by consensus or, if used, by the vote fallback	resolution of the PCG or legal questions outside that group	<i>name the group and the disposition</i>
A 3GPP objection remains unresolved	consensus has not been reached on that point	a permanent veto: Article 25 provides the vote fallback [G26]	<i>consensus remains open</i>
PCG establishes an ad hoc group	the partnership recognises work that has to be done	a policy decision, or a date	<i>the issue remains under PCG work</i>
JVET adopts a test model or an experiment condition	technical convergence sufficient for joint experimentation	a Recommendation or an International Standard	<i>adopted for exploration</i>
JVET issues a call condition	a binding condition of entry to that competition	a normative clause of the later standard	<i>condition of entry</i>
An MPEG or JPEG working group reaches committee draft	expert work has matured into committee-stage text	national-body approval	<i>committee-stage maturity</i>
An ISO/IEC enquiry draft passes	the national-body thresholds were met: two-thirds in favour, no more than a quarter negative [K28]	unanimity, or the absence of technical objection	<i>state the stage and the threshold</i>
ITU-T Consent is given	the study group regards the draft mature enough to enter the approval process [K32]	final approval	<i>consented to start the approval process</i>
An IETF working-group last call passes	the Chairs judge rough consensus sufficient to progress [I30]	unanimity, or a majority of any poll	<i>rough consensus to progress</i>
An IEEE SA ballot reaches 75%	the defined numerical threshold was met [K30]	that every negative was resolved	<i>give the approval and return figures</i>
An ETSI Technical Group vote passes	the formal threshold was reached [K31]	unanimity, or that consensus was reached	<i>approved by vote, where a vote was used</i>

The vocabulary this brief holds itself to when it reports another body's progress. The last column is the phrasing used in the text; where a stronger word appears it is because the record itself uses it. Contribution counts, survey tallies and ballot percentages are not interchangeable, and none of them is a decision.

Table 60. How each body decides, whether it votes, and what that means for a rule about training data

Body	How a decision is made	Voting: whether, what kind, and on what condition	What it can bind	What can move, and how fast	What it cannot do
3GPP working group	contribution and discussion; the group shall endeavour to reach consensus on all issues [G26]	Yes, as a fallback. If consensus cannot be achieved the Chair may decide to take a vote; 71% of the votes cast approves, abstentions excluded, with a quorum of Voting Members at the meeting. There is no tie to break: short of 71% the proposal is not approved and the item stays open (Art. 25, 26) [G26]	the competition and the deliverables it produces	a Pdoc can change at the next meeting if consensus is reached or the fallback is used; an unresolved objection can leave the rule open while consensus is sought	decide what the partnership warrants
3GPP TSG plenary	approval of the group's output, or endorsement of a way forward, on the same consensus-first basis [G26]	Yes, as a fallback, on the same terms and at the same 71% of votes cast [G26]	the work item, its schedule and its scope	one meeting, three times a year	settle a question of law
3GPP PCG and the OPs	the partners meet as partners and shall endeavour to reach consensus; the OPs publish and own the specifications jointly [G25] [G26]	Yes, as a fallback. One vote per Organisational Partner, 71% of votes cast, quorum half the partners (Art. 19). With seven partners voting, 71% needs five in favour, so no single partner holds a veto and no simple majority carries [G26]	the partnership's own conduct and exposure	unbounded: PCG#56 bis_Ad Hoc set no date and PCG#57 has none	write technical requirements
JVET (ITU-T and ISO/IEC)	exploration experiments, ad hoc groups and calls; all group decisions are made by consensus as determined by the Chairs [K27]	No, not in the joint team. The Chairs' finding of consensus is what settles a disputed point. Voting comes later and elsewhere: all output undergoes the approval process of each parent body, and where the two parents differ the ToR leaves them to resolve it between themselves [K27]	the conditions of entry to its own calls, and the text of its standards	a call or experiment condition can be set before the resulting standard completes either parent's approval path	allocate liability for the material a proponent used
MPEG (ISO/IEC JTC 1/SC 29)	working documents and common test and training conditions, agreed by expert consensus in the working group [K28]	Not in the working group; yes in the parent. A draft is approved where two-thirds of the P-members voting are in favour and no more than a quarter of all votes cast are negative. A negative vote is answered in a comment-resolution round rather than by a re-run, and sustained opposition must be addressed but is not a right of veto [K28] [K29]	common test and training conditions, and licence conditions on a call's data	a condition can be set well before the standard it feeds completes its ballot stages	clear rights in a corpus
JPEG (ISO/IEC and ITU-T)	as MPEG, in the same expert-consensus framework, with a mandated common training set for its learned codec [K28]	As MPEG, and joint texts pass an ITU-T path as well, where approval at a study group meeting is to be unopposed and, failing that, carries if no more than one Member State present opposes [K28] [K32]	the training set and the trained parameters of its standard	already done: T.840.1 was approved in March 2025	answer for the images it trained on
IETF	rough consensus in a working group; Internet-Drafts expire and are work in progress until published [I30]	No. The Chairs judge whether the objections raised have been answered. A hum or a poll may inform that judgement but does not decide it, and a headcount settles nothing either way: an objection considered and answered does not prevent progression, and a large majority does not close an unanswered one [I30]	what a compliant implementation must and must not do	a draft revision: the test-clip prohibition is already normative text	enforce a preference signal
W3C community work	a community group report, not a Recommendation; the Recommendation track itself turns on consensus [I31]	Not in a community group. On the Recommendation track a vote is available only after consensus methods have failed, one vote per represented organisation, and a Formal Objection is recorded and may be reviewed above the group rather than outvoted [I31]	nothing: it publishes a vocabulary others may use	a revision	make a reservation legally effective
IEEE SA	working-group ballot and recirculation, then an IEEE SA standards ballot [K30]	Yes, and the ballot is the record of consensus. Approval needs 75% of those voting to approve or disapprove, subject to a ballot-return requirement; negatives are answered and the draft recirculated rather than re-run, and an excess of abstentions invalidates the ballot [K30]	the text of its own standards	a ballot cycle, with comment resolution and recirculation before each recount	reach a corpus held by another body's members
ETSI Technical Group	the group shall endeavour to reach Consensus on all issues, during a meeting or between meetings [K31]	Yes, as a fallback. The Chair may call a weighted vote where Consensus cannot be achieved; an ordinary technical proposal passes at 71% of the votes cast, and where 71% is not reached the vote is recalculated on the votes of Full members only - the one explicit re-computation rule among these bodies [K31]	its own deliverables, and its IPR policy	a meeting cycle, with additional national procedures for some deliverables	bind 3GPP, of which it is one Organisational Partner among seven

Table 61. The codec programmes behind ULBC, in two lineages

Programme	Media	Where it was developed	Window, as the record dates it	What 3GPP did with it	Release
<i>Every programme that materially shaped how SA WG4 selects or adopts a codec, and what 3GPP did with each. The first band is the lineage ULBC belongs to: the partnership sets requirements, runs a qualification and selection contest and publishes its own specifications. The second is the lineage 3GPP joins late: a joint ITU-T and ISO/IEC team develops the codec and 3GPP evaluates and profiles it for its own services. Two of the entries in the first band show that a selection need not end in one winner, and one in the second shows what happens when a working group does not reach consensus on an adoption (section 11.1).</i>					
Native to 3GPP: the partnership selects its own codec					
GSM FR, HR and EFR	speech	ETSI GSM and SMG11	1990s	inherited into the 3GPP GSM specifications; the practice of reference code, test sequences and bit-exact conformance starts here	GSM to R99
AMR	speech	ETSI SMG, then 3GPP	feasibility 1996-97; launched at SMG#23 in October 1997; qualification and selection 1998; SMG11 consensus September 1998 [G33]	adopted as the mandatory speech codec for early 3G multimedia telephony	R98, R99
AMR-WB	speech	ETSI SMG11 with 3GPP SA WG4	work item at TSG SA#2 in March 1999; qualification at TSG SA#8 in June 2000; selection testing July to October 2000; selection approved at TSG SA#10 in December 2000; characterisation through 2001 [G34]	native wideband speech codec; the stage sequence ULBC now runs	Rel-5
AMR-WB+ and Enhanced aacPlus	audio	3GPP SA WG4, on MPEG-derived technology	competition 2003; evaluation and selection 2003-04 [G35]	both were recommended: no consistent overall ranking was possible, so selection produced two codecs with usage guidance rather than one winner	Rel-6
DSR extended advanced front end	speech front end	ETSI Aurora, evaluated by 3GPP SA WG4	2002-04	external technology evaluated by SA WG4 and taken into the speech-enabled services specifications	Rel-6
EVS	speech and audio	3GPP SA WG4	work item 25 March 2010; the SA WG4 building block ran to 12 December 2014 and the parent item to 15 March 2016; qualification SA4#72bis March 2013; selection SA4#80bis August 2014; specifications approved TSG SA#65 September 2014 [G27] [G29] [G32]	native codec programme: about four and a half years from opening to codec approval	Rel-12
IVAS	immersive speech and audio	3GPP SA WG4	work item 8 September 2017 to 6 June 2024, moved from Rel-17 to Rel-18 in December 2021; public collaboration May 2022; selected at SA4#125 August 2023; approved at SA#101 [G28] [G30] [G31]	native codec extension of the EVS family: close to six years from opening to selection	Rel-18
ULBC-MED	neural speech	3GPP SA WG4	work item recorded from 4 June 2026 with a target end of 6 June 2027; the plan places selection testing on 19 March 2027 and approval in June [G01] [M01]	the current case, and a recorded plan rather than an achieved duration	Rel-20

Programme	Media	Where it was developed	Window, as the record dates it	What 3GPP did with it	Release
Every programme that materially shaped how SA WG4 selects or adopts a codec, and what 3GPP did with each. The first band is the lineage ULBC belongs to: the partnership sets requirements, runs a qualification and selection contest and publishes its own specifications. The second is the lineage 3GPP joins late: a joint ITU-T and ISO/IEC team develops the codec and 3GPP evaluates and profiles it for its own services. Two of the entries in the first band show that a selection need not end in one winner, and one in the second shows what happens when a working group does not reach consensus on an adoption (section 11.1).					
Developed in a joint ITU-T and ISO/IEC team, then evaluated by 3GPP					
H.263 and MPEG-4 Visual	video	ITU-T VCEG; ISO/IEC MPEG	H.263 first edition March 1996, second edition February 1998; MPEG-4 Visual published December 1999	H.263 mandatory and MPEG-4 Visual optional in the Release 1999 circuit-switched multimedia path, and again in the Rel-4 streaming profile [G37]	R99, Rel-4
H.264, or MPEG-4 AVC	video	JVT: ITU-T VCEG with ISO/IEC MPEG [K33]	team formed 2001; first version completed May 2003; the ISO/IEC twin registered March 2002, approved for the final stage April 2003 and published November 2003 [K37]	evaluated and profiled by SA WG4 for messaging, streaming, file storage and conversational services; the adoption changes reached the plenary by an unusual route in 2004 [G36]	Rel-6
H.265, or HEVC	video	JCT-VC: ITU-T VCEG with ISO/IEC MPEG [K34]	first met Dresden, 15-23 April 2010; H.265 approved 13 April 2013	evaluated for 3GPP services in TR 26.906 while the standard was still being finalised, and integrated into MTSI at SA WG4#83 and TSG SA#68 in 2015 [G38]	Rel-12
H.266, or VVC	video	an informal exploration team from October 2015 [K35], then the formal JVET from 27 October 2017 [K36]	exploration October 2015 to October 2017; call for proposals October 2017; proposals evaluated April 2018; first version July 2020; H.266 approved 29 August 2020; the ISO/IEC twin published February 2021 [K38]	characterised alongside EVC and AV1 in TR 26.955, which is study and evaluation rather than a normative service profile [G39]	Rel-18 study

Table 62. Three clocks, kept apart: what each span measures and why it does not transfer

Programme	What is being measured	Span as recorded	Why the span does not transfer
AMR-WB	work item March 1999 to selection December 2000, characterisation through 2001	about two to three years, depending on the endpoint	the endpoints differ: selection and characterisation are different milestones, and the codec was not learned
EVS	work item March 2010 to codec approval September 2014	about four and a half years	one joint candidate in the end, a mature signal-processing baseline, and no model to train
IVAS	work item September 2017 to selection August 2023	close to six years, with a work item spanning almost seven	release planning changed mid-programme, and a public collaboration replaced a contest
JVT, for H.264	team formed 2001 to first version May 2003	about two years	a joint team's technical convergence, not a 3GPP selection; the ISO/IEC ballot ran separately into November 2003
JCT-VC, for H.265	first meeting April 2010 to approval April 2013	about three years	extensions continued for years afterwards, so the first edition is not the whole programme
Exploration and JVET, for H.266	exploration October 2015, formal team October 2017, approval August 2020	about five years, of which two were exploration	the exploration phase has no 3GPP equivalent, and the generation gap since HEVC is not the development time
H.264 into 3GPP	H.264 approved May 2003 to the change requests reaching TSG SA#25 in September 2004; 2.8 years from the joint team's formation to that plenary	1.3 years after approval, 2.8 years end to end	the 3GPP act is a service profile, not a codec selection, and it is a separate clock from the joint team's
H.265 into 3GPP	H.265 approved April 2013 to integration into MTSI at SA WG4#83 and TSG SA#68 in 2015; 5.2 years from the first JCT-VC meeting	2.2 years after approval, 5.2 years end to end	3GPP began evaluating in TR 26.906 before the standard was finished, so the two clocks overlap and cannot simply be added
H.266 into 3GPP	H.266 approved August 2020 to characterisation in TR 26.955 in 2024; 8.5 years from the exploration team's first meeting	3.7 years after approval, 8.5 years end to end	the 3GPP deliverable is a characterisation and evaluation record, not a normative service profile, so the end point is not the same kind of event as for H.264 or H.265
ULBC-MED	work item June 2026 to a target end of June 2027, with selection testing planned for 19 March 2027	about one year, as planned	a recorded plan and not an achieved duration; the comparison establishes neither that it can be met nor that it cannot


The interval between codec generations, the time a team spent developing one, and the time 3GPP took to select or adopt one are three different measurements, and this brief does not substitute any of them for another. The three video rows carry both clocks: the lag after the codec was approved, and the whole span from the joint team's first work to the 3GPP deliverable - 2.8, 5.2 and 8.5 years for H.264, H.265 and H.266. The native rows run 1.8 years for AMR-WB to selection, 4.5 for EVS to approved specifications and 6.0 for IVAS to selection. Every span here is context for the density of the ULBC calendar. None is a forecast, and the last row is a recorded plan rather than an achieved duration.

Table 63. The liaison record: what crossed between bodies, and when

When	From	To	What crossed	What it shows
Jan 2023	MPEG's call for video coding for machines	JVET	a question about whether the datasets behind the call could be used at all [J14]	dataset licensing enters a coding body as a question, not as a rule
Apr 2024	a JVET contribution	the VSEI specification	a proposal to signal that video was used for AI-driven processes [J18]	the first signalling of AI use of content inside a codec standard
Jan 2025	a JVET contribution	the VSEI specification	an AI usage restrictions message, carried into version 4 [J17] [K06]	a restriction expressed in the bitstream rather than in a licence
Jul 2026	JVET and MPEG together	proponents of the beyond-VVC call	disclosure of training material, availability for verification, scripts, and a re-training fallback [J01]	the conditions of entry become the training-data rule
Jul 2026	C2PA content credentials	JVET AHG9	a proposal to carry the credentials in an SEI message [J08]	provenance vocabulary arrives from outside the coding bodies
Jul-Aug 2026	the IETF's neural speech work	its own test conditions	a prohibition on training with the test clips, and named public corpora [I15] [I16]	the separation rule is normative text in one body before it is a rule in another
Jun 2026	TSG SA#112	SA WG4 and PCG	the codec's approval, an edited guidance line and a referral of the legal question [S06] [S07]	the split of one question into a technical and a partnership half
Sep 2026	three Organisational Partners	PCG#56 bis_Ad Hoc	three inputs on datasets, requirements and OP exposure [P02] [P03] [P04]	the partnership's first reading of the question
Sep 2026	SA WG4	TSG SA#113	the report, the status report and a draft liaison [M01] [M02]	the technical half returns to the plenary while the legal half stays open

Only movements this brief could follow in a document. Technique and test conditions travel; rights vocabulary does not, and where it appears it arrives from outside the coding bodies.

Figure 14. Thirty years of codec standardisation in two lineages, how long each took, and how long 3GPP took after it



One row per programme, so no label can collide with another. A bar is the window in which the programme was materially active as the record dates it, not the life of the codec after publication; an outlined bar is a window still open at the cutoff. Bar length is elapsed time and nothing else, and the figure beside each bar states that time in years so it need not be measured off the axis. Years are from the dates in Table 61, rounded to one decimal. The fourth band re-plots intervals shown above from a different start, so its bars are not additional programmes. A hollow diamond and a dashed stem mark a handover: the diamond sits on the date the codec was approved in the body that developed it, and the stem drops to the row where 3GPP acted on it. It records that a handover happened and carries no duration. The third band is drawn as an outline on a pale ground: acting on a codec developed elsewhere is different work from selecting one's own.

Two lineages, and neither is quick. In the first band 3GPP writes requirements and runs a contest to select a codec of its own, and the run has lengthened: AMR took a year to selection, AMR-WB 1.8, EVS 4.5 to approved specifications, IVAS 6.0 to selection. ULBC plans 1.0. In the second band a joint ITU-T and ISO/IEC team develops a video codec 3GPP did not invent - JVT reached H.264 in 1.5 years, JCT-VC reached H.265 in 3.0, and H.266 took 4.9 counting the two-year exploration that preceded the formal team. The fourth band is the part usually left out: after the codec exists, 3GPP's own act took 1.3 years for H.264, 2.2 for H.265 and 3.7 for H.266, so end to end the three video generations ran 2.8, 5.2 and 8.5 years from the joint team's first work to the 3GPP deliverable - and the last of those is a characterisation in TR 26.955 rather than a normative service profile. The diamonds mark the three handovers themselves - the instant a codec approved elsewhere becomes something 3GPP has to act on - and they are the only points in this figure where work crosses from one lineage to the other. Nothing crosses the other way: no codec 3GPP selected for itself has been taken up by a joint video team. ULBC did not arrive through a handover either. It was opened as 3GPP's own work item for IMS voice over NB-IoT non-terrestrial access via GEO satellite - a bitrate and delay budget no existing codec in either lineage was built for - and it is the first item in either lineage whose trained parameters would be normative, which is why its predecessors constrain its calendar less than the bars alone suggest. None of these numbers forecasts ULBC: they are different programmes with different scope, starting maturity and candidate structure, and none standardised a trained model (Table 62).

Table 64. EVS and IVAS milestone by milestone, as the record dates them

Milestone	EVS	IVAS	What the comparison carries
Work item opened	25 March 2010 [G27]	8 September 2017 [G28]	the portal's own start date for the work item, not for the study that preceded it
Release as recorded	Rel-12 [G27]	Rel-18, moved from Rel-17 in December 2021 [G28]	release planning changed during IVAS; a recorded move is a schedule event, not a delay measure
Collaboration or candidate structure	competing candidates, then one joint candidate [G29]	a public collaboration established in May 2022 developed a joint candidate [G30]	candidate structure differs between the two, and ULBC has twenty declared candidates
Qualification	SA4#72bis, March 2013 [G29]	not separately recorded in the sources read here	ULBC plans qualification testing from January 2027 [M01]
Selection	SA4#80bis, August 2014 [G29]	SA4#125, August 2023 [G30] [G31]	ULBC plans selection tests from 19 March 2027 and a winner in May [M01]
Specifications approved	TSG SA#65, September 2014 [G29]	SA#101, September 2023 [G30] [G31]	ULBC targets TSG SA#116 in June 2027 [M01]
Later work in the same item	floating-point codec completed at TSG SA#66, December 2014; work item closed 15 March 2016 [G27] [G29]	work item closed 6 June 2024 [G28]	a work item stays open past codec approval, so its span is not the development time
Opening to codec approval or selection	about four and a half years, opening to approval	close to six years, opening to selection; the work item spans almost seven	context for the density of the ULBC calendar; not a forecast, and not a claim that the recorded plan can or cannot be met

The two nearest native comparators in detail. EVS and IVAS differ from ULBC in scope, starting maturity, candidate structure and requirements, and neither standardised a trained model, so their durations do not transfer. What they show is the shape of the cycle this partnership has run before: requirements, candidate development, qualification, selection, independent testing and specification completion. Nothing here establishes that the recorded ULBC plan can or cannot be met.

12 What the record establishes, and what could follow

For: everyone acting on the brief.

12.1 What the record already establishes

CNT ESTABLISHED · Table 66 is the first of the four layers: sixteen findings, each with the record that carries it, its quantitative support, its tier and its class. Twelve are established - direct statements of a record or counts from one - and four are well supported, because they

rest on a reported outcome, a reading of a legal text or this brief's own analysis. None is conditional; the conditional readings are Table 67's.

INT ESTABLISHED · The established core is narrow and firm. SA4 found no consensus; its members nearly all support keeping test material out of training; the three OP inputs agree that datasets must meet legal requirements; the rules that govern qualification and selection are unwritten; no Pdoc concerns training data; and the question SA4 was asked had been edited before it was asked. Every one of these can be checked against a document in Table A1 without reading anything else in this brief.

INT WELL SUPPORTED · What the record does not establish is as specific. It does not show what PCG the new PCG ad hoc group will conclude, or when; it does not show why each objecting company holds the position it holds, only what it said; and it does not show whether the guidance question will return at TSG SA#114. Each is listed with its disposition in Table 77, and each is a place where a later edition of this brief would change.

What is established bounds what can be read forward.

12.2 Conditional readings, each with what would weaken it

INT CONDITIONAL · Table 67 carries the second and third layers: nine conditional readings, each written in the same five parts - the dimension, the assumption it rests on, what follows if the assumption holds, what follows if it fails, and the observable signal - with a sixth column naming the observation that would weaken it. None carries a probability, and none is ranked.

INT CONDITIONAL · Three of the readings turn on documents that will exist within six weeks. The Rel-20 schedule turns on the disposition of SP-260912 at Madrid; PCG's next step on the terms of its new group, which the PCG#57 record should show; and technical requirements in SA4 on the sampling-rate decision in ULBC-4. Each can be read off a record without interpretation, which is why they are listed first.

INT CONDITIONAL · The others turn on wording rather than decisions - what 'transparent' is taken to mean, whether the weights are named, at what threshold a PCG outcome applies - and they will be settled, if at all, in the text of liaisons, terms of reference and requirements. A reader tracking the question after Madrid can use the sixth column as a checklist: each entry is a sentence that will or will not appear in a record.

The conditional readings combine into three scenarios.

12.3 Three scenarios for Madrid and PCG#57

INT CONDITIONAL · Table 68 bundles the conditional readings into three scenarios. They are descriptions of internally consistent combinations, not forecasts, and each names the signal that would indicate it and the observation that would invalidate it. *Parallel tracks*: Madrid notes the SA4 report and the sequencing paper, SA4 proceeds, and PCG's group drafts requirements alongside. *Sequenced*: Madrid endorses SP-260912 and PCG#57 becomes the gate for SA4's rules. *Requirements first*: Madrid sends SP-260919's two points to SA4.

INT CONDITIONAL · The three differ most in where the Rel-20 margin is spent. Under parallel tracks it stays with SA4 and the legal risk travels towards selection;

under sequencing it passes to PCG, whose next meeting is online and whose group has no published terms; under requirements first it is spent on new test conditions, whose cost the funding estimate would have to absorb. The scenarios are not exclusive at every point - Madrid could send a liaison and endorse sequencing - and the table marks where they combine.

INT CONDITIONAL · Whichever scenario the record shows, two observations cut across all three: whether SA4 writes the separation rule, and whether PCG's group names the weights. Both can happen under any scenario, and neither depends on how Madrid disposes of the four documents in front of it.

Scenarios are tested by what appears in records, which the next section lists.

12.4 Observable signals to watch after Madrid

INT The fourth layer is a list of things that will or will not appear in a record. Table 69 gives twelve dated observables from the Madrid plenary to the launch of the selection tests, each with the forum, what to watch, the record in which it will appear and the section of this brief whose reading it tests.

INT CONDITIONAL · The densest stretch is October 2026. Within ten days PCG#57 meets online, JVET and MPEG meet in Hangzhou, and the joint call's main submissions fall due, so the governance track, the technical track and the nearest outside precedent all produce records in the same fortnight. A reader who can follow only one should follow PCG#57, because it is the first record that can show what the new group was asked to do.

CTX Three of the dates are outside 3GPP and fixed by law or by a court: the Product Liability Directive's transposition deadline on 9 December, the German Federal Court of Justice's pronouncement on 17 December, and the AI Act's Annex III date in December 2027 [L13] [C01] [L02]. None changes what 3GPP must do; each changes what implementers will ask it for.

The same signals read differently from different seats.

12.5 Eleven seats, eleven readings of the same record

INT Table 74 reads the same record from eleven seats, each with what the record means for that seat, where in this brief to start and what to watch. The seats are readerships, not parties: a company may occupy several, and no seat is attributed to a named organisation.

INT WELL SUPPORTED · The seats divide on the question they bring. Operators, curators and PCG delegates bring the maintenance and governance question; proponents, test laboratories and SA4 leadership bring the calendar; implementers, chipset vendors and regulators bring evidence - what must be shown about a decoder before it ships. The first group reads Parts 3 and 9 first, the second Part 4, the third Parts 7 and 8.

INT WELL SUPPORTED · Only one observable appears on every seat's list in substance, if not in wording: the terms of the new PCG group. Whatever else changes after Madrid, those terms decide whether the dataset question is answered as a requirement that any corpus must meet or as a choice of corpus.

What this part established

Sixteen findings, nine conditional readings, three scenarios and twelve dated signals, kept in four layers (Tables 66 to 69).

The terms of the new PCG group are the one observable that matters to every seat (Table 74).

What is established is largely about process - who asked what, what was answered, and what was never put. Madrid has now answered one of the two things the readings turned on: the plenary endorsed no sequence, and sent nothing. What remains open is the new PCG group's terms (Table 67).

No seat needs the corpus question settled before it can act, and each has one document to watch (Table 74).

Three of the established findings are about the process rather than the dataset: that 3GPP is consensus-first with a 71% vote available where consensus cannot be reached; that cross-SDO precedent supports the feasibility of a technical mechanism without supporting an inference that 3GPP can adopt it on the same calendar; and that at Madrid a liaison was drafted four times and noted, which is the fallback not being used in a case where consensus was not reached (Table 66).

The schedule conditional turns on unresolved agreement rather than on how many documents remain: four selection Pdocs have no agreed version and the Chair calls the margin thin, so the observable is now the agreed versions at SA4#139 in February rather than any plenary disposition - not a completion percentage (Table 67).

What this part established — part 12

Every seat's position was stated before Madrid and none moved at it: the five companies in both meeting records hold the same ground in September that they held in August (Table 9).

Hands forward: what a normative dataset would move, and for whom - part 13.

Table 65. Six standing conditions behind each actor's position

Actor	Political standing	Legal instrument it holds	Economic exposure	What it can verify	Supply-chain duty	Vertical interest
Codec proponent	competes in a global selection	holds its corpus under contract or trade secret	bears the cost of any disclosure	can verify its own results only	must satisfy the buyer that ships its model	wants the codec adopted at all
Competing proponent	as above, without the data	may have no comparable corpus	gains from a common set	cannot see the incumbent's advantage	the same buyers	the same
Organisational Partner	answers to its members and its region	holds an IPR policy written for patents	would carry the cost of custody	cannot verify a licence	publishes what implementers rely on	neutral between codecs
Operator	buys from several vendors	faces service obligations	pays for the whole service life	can test but not retrain	answers to regulators and to subscribers	emergency calling and coverage
Device or chipset vendor	ships worldwide	needs a licence for the weights it embeds	carries product liability and update duties	needs a version it can pin	shows evidence to its own customers	battery, memory and conformance
Test laboratory	paid by the process	handles material under contract	priced per condition tested	can detect leakage if the rule exists	reports to SA4	validity of the comparison
Dataset curator	publishes under a public licence	grants what it can grant and no more	funded by projects or platforms	documents provenance	asks that its terms travel with the data	wants the corpus used well
Rightsholder or speaker	represented by law rather than in the room	holds copyright, performer, personality and data-protection rights	unremunerated in every option read here	cannot inspect a model	reaches the chain only through the law	consent and identity
Regulator or authority	national or regional mandate	applies product, safety, cyber and data-protection law	does not pay	reads evidence rather than data	asks the manufacturer, not the standards body	public safety and market surveillance

Read from the record; no motive is attributed to any organisation, and no cell is a statement about a named company.

Table 66. What the record establishes: the first layer

Finding	Where recorded	Quantitative support	Tier	Class
SA4 does not identify a consensus need for technical training-dataset requirements and proposes no work beyond the approved item	SP-260909 [M02]	13 for, 9 against, 2 without a position, of 24	Established	SRC
Keeping qualification, selection and characterisation material out of training has near-unanimous support	SP-260919, quoting S4-261674 [M05]	20 of 24 (83.3%)	Established	SRC
The survey's additional figures reach TSG SA only through a company paper	SP-260909 annex; SP-260919 [M02] [M05]	two annex figures	Established	SRC
The SA#112 guidance to SA4 was edited from 'common' to 'technical' training dataset	TSG SA#112 draft report [S06]	one edited word	Established	SRC
All three OP inputs require that datasets meet legal requirements, naming licensing and privacy	PCG#56bis(26)02 to 04 [P02] [P03] [P04]	3 of 16 elements stated by all three	Established	CNT
PCG#56bis reached no consensus between common and proprietary data and created an ad hoc group on the legal aspects	PCG#56bis draft report [P06]	the meeting's own minutes, decision D-PCG56-Bis/01	Established	SRC
TSG SA#113 sent no guidance to SA4: one liaison statement, four drafts, noted	TSG SA#113 register and report [M07] [M08]	7 ULBC documents at item 3.4, 4 noted and 3 revised	Established	CNT
The qualification and selection rules and both test plans have no agreed version	SP-260907 [M01]	4 of 11 Pdocs	Established	CNT
No Pdoc concerns training data	SP-260907 [M01]	0 of 11	Established	CNT
None of the 8 bodies read here governs the legal or liability side of a training corpus	Figure 8; Table 36 to Table 35	0 of 8	Well supported	CTX
The ETSI IPR Policy does not reach rights in training data that implementers never copy	ETSI IPR Policy [G10]	definitions in clauses 15.6 and 15.7	Well supported	CTX
Two of the eight rights layers in a recording are named by none of the inputs	Figure 9	2 of 8	Well supported	INT
The published weights are the one artefact common to every option	sections 8.3 and 10.1	every option ends in weights the OPs publish	Well supported	INT
Process comparability across these bodies is limited: the public record supports comparing technical instruments, but the decision paths differ, so cross-SDO precedent supports the feasibility of a mechanism and not an inference that 3GPP can adopt it on the same calendar	Working Procedures Art. 19, 25 [G26]; JVET ToR [K27]; ISO/IEC [K28]; RFC 7282 [I30]	8 bodies, five decision architectures (Figure 13)	Established	CTX
No 3GPP decision in this corpus is recorded as carried by a vote: every disposition in the approved-record export is agreed, approved, partially approved or endorsed, and searching every record read for this brief returns no mention of a vote, a ballot or an abstention	the approved-record export [B06]; the ledger, Pdoc register, thread and Madrid papers	0 of 1,257 records	Established	CNT
3GPP is consensus-first rather than unanimity-based: a group shall endeavour to reach consensus, and where consensus cannot be achieved the Chair may take a vote carried by 71% of the votes cast	3GPP Working Procedures, Art. 19 and Art. 25 [G26]	71% of votes cast; one vote per Organisational Partner at PCG	Established	SRC

Every finding with the record that carries it, its support, its tier and its class. No finding here is conditional.

Table 67. Conditional readings, with what would weaken each

Dimension	Assumption	If it holds	If it fails	Signal	What would weaken it
The Rel-20 schedule	TSG SA does not endorse sequencing SA4 behind PCG	SA4 finalises the Pdocs at SA4#138 and the January contracts proceed	finalisation of the qualification and selection rules waits for a PCG conclusion	the disposition of SP-260912	an endorsed sequencing proposal
Where the schedule risk sits	one or more material rules remains under unresolved objection through SA4#138 and SA4#139, and the Article 25 vote fallback continues not to be used, as no record in this corpus shows it being used	the 19 March selection-test launch is exposed even where most drafting is complete, because the rules that bind the competition would not be agreed when it starts; the May and June dates move with it	the recorded work plan can proceed on its dates	agreed Pdoc versions and dispositions after Shanghai, Calgary and SA4#139, and whether any report records a vote being called	the four selection Pdocs reaching agreed versions at SA4#138, or a recorded vote closing a disputed rule
PCG's next step	the new ad hoc group is asked to draft requirements	a first text is possible at or after PCG#57	the division between the three inputs returns unchanged	the group's terms of reference	terms that ask the group to choose a scenario
Technical requirements in SA4	the design constraints settle, starting with sampling rate	SA4 can ask the question again against fixed parameters	any requirement has to be written as conditional	the sampling-rate decision in ULBC-4	SA4 declining requirements after the constraints settle
The separation rule	SA4 writes the near-unanimous separation into its rules	ULBC acquires its first training-data rule without any legal decision	the most widely supported rule stays unwritten	separation text in ULBC-5a, ULBC-5b or the test plans	rules adopted without it
The meaning of 'transparent'	it is defined as disclosure plus access for verification	the maintenance case is partly met without a common set	publication or custody requirements raise 3GPP's burden	the wording of any liaison from Madrid	a definition requiring publication of the data
Weights as the common object	PCG's group treats weights and data as separate objects	personal-data and rights-flow questions are addressed under every option	the OPs' exposure through the weights stays unaddressed	whether the group's terms name the weights	terms limited to datasets
The precedent threshold	PCG states the threshold at which its outcome applies	the rule reaches standardised models only	a blanket rule reaches AI/ML work whose data 3GPP never sees	the framing of the PCG#57 record	a rule declared for all 3GPP AI/ML work
Evidence for implementers	a dataset declaration is adopted in some form	the option chosen matters less to implementers than the scenarios suggest	implementers seek the evidence from proponents directly	implementer contributions to PCG or to ULBC-6	a corpus named without information requirements

The second and third layers. Each reading is written in the same five parts, with a sixth naming the observation that would weaken it. None carries a probability or a ranking.

Table 68. Three scenarios, as bundles of the conditional readings

Scenario	Madrid	SA4 and PCG	Consequence	Signals	What would invalidate it
Parallel tracks	notes SP-260909 and SP-260912; may or may not send SP-260919's two points to SA4	SA4 proceeds to Shanghai and Calgary; the PCG group drafts requirements alongside	the Rel-20 plan holds; legal risk is carried towards selection, with disclosure possibly added	no endorsement of sequencing; Pdoc progress at Shanghai	an endorsed sequencing proposal
Sequenced	endorses SP-260912's first proposal	SA4 progresses the Pdocs but cannot finalise the qualification and selection rules until PCG concludes; PCG#57 becomes the gate	the plan holds if PCG concludes in October; otherwise the January contracts and March selection tests move	an endorsement recorded; a liaison to SA4 citing PCG	SP-260912 noted or revised without endorsement
Requirements first	endorses SP-260919 and sends a liaison to SA4	SA4 considers requirement dimensions and transparency at Shanghai and Calgary; PCG continues on the legal side	new content in ULBC-3 and the selection rules, with test-cost consequences	a liaison with SP-260919's wording; contributions on dimensions	no liaison from Madrid

Descriptions of consistent combinations, not forecasts. They are not exclusive at every point: Madrid could, for instance, send a liaison and endorse sequencing.

Table 69. Observable signals to watch after Madrid

Date	Forum	What to watch	Where it will be recorded	Tests section
15-18 Sep 2026	TSG SA#113, in the TSG#113 round at Madrid	dispositions of SP-260909, SP-260912, SP-260913 and SP-260919; any liaison to SA4	TSG SA#113 report	1.3
28-30 Sep 2026	SA4 Audio SWG ad hoc, Shanghai	progress on ULBC-3, ULBC-4 and ULBC-5a; sampling rate; any contribution on training material	ad hoc report and Pdoc versions	5.2
14 and 28 Oct 2026	Audio SWG telcos	Pdoc progress towards SA4#138	telco notes	5.1
17-23 Oct 2026	JVET#44 and MPEG 156, Hangzhou	NNVC training-set list; registrations for the joint call	meeting reports	7.2
21 Oct 2026	PCG#57, online	report of the new ad hoc group; its terms of reference; any request for counsel	PCG#57 report	4.6
22 Oct 2026	OP#56, online, as the base briefing lists it	any OP-level statement on liability or resources	OP meeting record	10.2
26 Oct 2026	JVET joint call	main submission packages, with training disclosures	JVET documents	7.2
Nov 2026	SA4#138, Calgary	Pdoc finalisation, funding estimate, qualification decision, deliverables in ULBC-6	SA4#138 report	6.1
9 Dec 2026	EU member states	transposition of the Product Liability Directive	national law	8.5
17 Dec 2026	German Federal Court of Justice	pronouncement in I ZR 281/25	court decision	8.2
Jan 2027	SA4; JVET in Brisbane	test-laboratory contracts; joint call evaluation	SA4 and JVET records	6.2
19 Mar 2027	SA4	selection tests launched under whatever rules then apply	SA4 record	6.2

The fourth layer, in date order. Each is something that will or will not appear in a record; the last column names the section whose reading it tests.

13 If the dataset becomes normative: what it would move

For: IPR and licensing readers, research and policy staff.

13.1 What would change the moment the weights are normative

INT WELL SUPPORTED · Every option in part 9 is described by what it costs the partnership. This part reads them the other way: if the weights are normative and the material behind them is governed by a 3GPP rule, what moves outside 3GPP. Three things change at once. The specification acquires an input whose lawfulness someone has to answer for. The competition acquires an entry condition that is not technical. And the licensing question acquires a second instrument, because a patent licence has never had to cover the copying of a corpus (Table 70).

CTX The starting point is that a trained model is not one legal object but several. The decoder architecture and the inference method are the kind of technical teaching a patent claim reads on. The weight file is a set of numbers derived from recordings, and what governs its copying is copyright, database right, contract and the terms the recordings came with, none of which is patent law (Figure 9, Table 47, Table 51). The two instruments have different holders, different terms and different remedies, and nothing in the record read here joins them.

INT ESTABLISHED · That separation is why the question could not be settled in SA WG4 and is now at PCG. A working group can write a requirement; only the partners can decide what the partnership warrants about an input to a specification they publish jointly under Article 46 [G25]. Everything in this part follows from that single fact and is written as context and reading rather than as advice: whether any instrument applies to a given actor, product or corpus is a question for counsel, as section 14.4 states.

13.2 Standard-essential patents, and an artefact patents do not reach

CTX A standard-essential patent is one whose claims cannot be avoided by an implementation of the standard. Each Organisational Partner's own IPR policy asks its members to declare essential IPR in good time and to undertake to license it on fair, reasonable and non-discriminatory terms; the ETSI policy is one of

seven and is read here as the example the record gives [G25] [P05]. That machinery has one subject: claims. It says nothing about whether an implementer may copy the material a specification carries.

INT WELL SUPPORTED · A normative weight file therefore adds a second gate rather than a new class of essential claim. Claims reading on the architecture, the quantisation, the entropy coding or the inference path are essential exactly as they are for a classical codec. What is new is that an implementer must also be able to copy, embed and ship the weights, and the right to do that comes from copyright and contract rather than from a FRAND undertaking. An implementer that has cleared every declared patent has not thereby cleared the artefact (Table 70).

INT CONDITIONAL · *Signal*: whether any PCG text distinguishes the licence on the specification from the licence on the material it carries. *Holds*: the two gates are named separately and the second is answered once, for every implementer. *Fails*: the second gate is left to each implementer's own counsel, and the evidence question in part 8 returns as a licensing question. *What would weaken this reading*: an OP IPR policy that already treats a published artefact as covered by the undertaking on the specification; none read here does.

INT WELL SUPPORTED · Three consequences follow for holders and for licensees. A declaration is made against a specification, so a specification carrying a trained artefact would be declared against as before, while the dataset behind it is declared nowhere in the partnership's machinery. A maintenance obligation - regenerate the codec if the corpus becomes unusable - is not an IPR undertaking and would need some other instrument to carry it (section 6.4). And a claim that reads only on a training method is an open question here: the standard publishes the trained result, not the training, so whether such a claim is essential to it is not something the record read for this brief answers.

13.3 Innovation, research budgets and product strategy

INT WELL SUPPORTED · The choice between a common corpus and proprietary data is an innovation choice before it is a legal one, and it cuts in opposite directions for different firms. A common set removes data as a source of advantage: every candidate trains on the same material, and the comparison measures architecture and training craft. Proprietary data keeps data as a source of advantage, which rewards whoever has spent years collecting it and raises the cost of entry for everyone

else. Neither is neutral, and the survey split of 13 to 9 sits on exactly this line (Table 20).

INT WELL SUPPORTED · For research and development the two regimes ask for different budgets. Under disclosure and verification the spending moves to reproducible pipelines - documented provenance, versioned corpora, scripts a third party can run, the artefacts part 5 describes as the recipe (Table 38). Under a proprietary regime it moves to acquisition: recording, licensing and curating speech in more languages and conditions. A firm that has built for one cannot switch to the other inside a release cycle, which is why the sequencing question in section 5.3 is read as a competitive question as well as a calendar one.

CTX For products the exposure is downstream and does not wait for the standard. A device or network vendor that embeds the weights is placing a product on markets that increasingly ask for traceable evidence about it: the revised Product Liability Directive treats software as a product from 9 December 2026, cyber-resilience reporting began on 11 September 2026, and market surveillance asks the manufacturer rather than the standards body [L13] [L16] (Table 46, Table 53). Whatever PCG decides, the implementer is the one asked.

INT CONDITIONAL · *Signal*: whether any Madrid or PCG#57 text names an obligation that falls on an implementer rather than on a proponent. *What would weaken this reading*: a decision that keeps every obligation inside the competition, which would leave the product chain exactly where it is today.

13.4 Regulation, policy and the geopolitics of one global corpus

CTX A corpus assembled once for a global standard has to be lawful in every jurisdiction the standard reaches, and the seven Organisational Partners sit in seven of them. The European Union permits mining of lawfully accessible works with a machine-readable reservation and adds a public training-content summary for general-purpose models; the United Kingdom's exception reaches non-commercial research only; the United States decides fair use case by case, with state biometric statutes beneath it; Japan permits non-enjoyment uses within limits; Korea published a fair-use guide in February 2026; China requires lawful sources for public generative services and treats a voiceprint as sensitive personal information; and India has a proposal rather than a law (Table 47, Table 50).

INT WELL SUPPORTED · The operative rule for one global corpus is therefore the strictest cell in every column it touches, and that makes the choice of corpus a choice about jurisdiction as well as about quality. It also makes it a choice about representation: the languages, accents and recording conditions in the material decide which speakers the codec works for, and a corpus assembled under the law of one region is not automatically a corpus that covers the world's speech (section 6.3).

INT WELL SUPPORTED · For policy makers the interesting feature of this case is that a standards partnership is being asked to answer a question legislatures have answered differently. Whatever PCG writes will be the first statement by a global telecommunications partnership about what it requires of training material, and it will be read outside 3GPP - by the coding bodies that already impose conditions of entry, by the assurance committees drafting dataset documentation, and by regulators reading the AI Act's training-content duties (Table 54, Table 49).

INT CONDITIONAL · *Signal*: whether the new PCG group's terms name a jurisdiction, a standard of care, or neither. *Holds*: a requirement written to the strictest cell travels, and the partnership has made an implicit legal choice explicit. *Fails*: a requirement written to no jurisdiction leaves the strictest-cell problem with each proponent. *What would weaken this reading*: a mutual-recognition arrangement for training-data transparency, which no instrument read here provides.

13.5 Fourteen seats: interests, hard lines and tradeable ground

INT WELL SUPPORTED · Table 71 reads the same decision from fourteen seats outside the room as well as inside it, and gives each what it wants, what it fears, what it can trade and what it cannot. The seats are kinds of actor, not companies: a firm may occupy several at once, and no cell attributes a position to any named organisation. Figure 15 places the same seats by how much they can move the outcome and how much they carry if it moves.

INT ESTABLISHED · The pattern the table shows is an asymmetry, not a disagreement. The seats that can move the decision - the partners, the plenary, the proponents - are largely not the seats that carry it. The implementer, the operator, the test laboratory and the speaker carry the consequences of a normative artefact and have no vote on whether it becomes one; the rightsholder and the speaker are represented by law rather than in the room, as Table 19 already records.

INT WELL SUPPORTED · The hard lines are fewer than the disagreement suggests. Read across the table, only four are stated in absolute terms: no proponent will disclose a corpus it holds under contract or trade secret without a term that protects it; no Organisational Partner will take custody of material it cannot verify; no implementer will ship weights it cannot licence; and no regulator will accept a fidelity claim it cannot see evidence for. Everything else in the table is negotiable, which is what makes a requirement before a corpus - the first conditional outlook on page one - the arrangement most seats can live with.

INT CONDITIONAL · *Signal*: whether any Madrid or PCG#57 text names an actor from the lower half of Figure 15 - an implementer, a regulator, a speaker. *Holds*: the decision is being written for the people who carry it. *Fails*: it is being written for the people in the room, and the carrying seats meet it as a fait accompli at qualification. *What would weaken this reading*: a contribution from an implementer or an operator at PCG#57, which would show the lower half is represented after all.

13.6 Industrial verticals: the first-order effect, and the one after it

INT WELL SUPPORTED · A seat is not a market. Table 72 breaks the same decision down by industrial vertical and separates two things usually run together: what a normative dataset does to a vertical at once, and what follows from that a step later. The second-order column is where most of the cost sits, and it is the column no document read here discusses.

INT WELL SUPPORTED · Three examples carry the pattern. For **chipset vendors** the first-order effect is that the weights become a fixed target: their size, precision and operator set decide the silicon. The second-order effect is that a maintenance rule - the option most seats can live with - turns that fixed target into a moving one, because a part must be sized for a model nobody has trained yet. A rule written at PCG about a corpus becomes a die-area decision two years before anyone trains anything. For **UE vendors** the first-order effect is a second licence to obtain; the second-order effect is that product-liability and cyber duties make the weight file a maintained artefact for the product's life [L13] [L16], so a codec stops being an integration and becomes a supported dependency. For **operators** the first-order effect is that the service outlives the company that trained the model; the second-order effect arrives before the standard does, in procurement language that begins to name regeneration rights, escrow and version support.

CTX Four verticals the work item exists for, or would depend on, are absent from the record. ULBC-MED is specified for IMS voice over NB-IoT non-terrestrial access through geostationary satellites [G01]: the codec will meet a channel with impairments a terrestrial corpus does not contain, and devices with a decade of battery and no update path. An IMS codec also follows the session onto **Wi-Fi**, where voice certification has no place to record which model was tested. And a common corpus has to be hosted, which makes a **cloud or hyperscale provider** the ecosystem's custodian by default - a dependency, with residency and transfer terms attached, that the standard never chose. No ULBC document read for this brief is sourced by a satellite operator, an IoT module maker, a Wi-Fi vendor or a cloud provider.

INT WELL SUPPORTED · Four more verticals turn on the same rule and appear in no document read here. **Streaming and content platforms** hold the cleared recorded speech a corpus would want, and the clearance machinery a standards body does not have; a partnership that states what a corpus must satisfy writes, in effect, a licence form the catalogue business does not yet issue. **Voice biometrics and financial services** carry the sharpest second-order effect in the table: a learned codec reconstructs speech rather than carrying a waveform, so the features a verifier relies on are a property of the model, and a maintenance rule that changes the weights becomes a re-enrolment programme in a bank. **Automotive** meets a type-approval regime that records a software build rather than a model version, in products with fifteen-year lives and no weekly patch. And **defence and mission-critical** procurement asks where the material came from before it accepts an artefact at all, which turns the jurisdiction question of section 13.4 into an accreditation one. Two more sit beside them: **enterprise voice and contact centres**, which hold recorded speech at scale under consent written for service delivery rather than for training, and **accessibility and assistive technology**, where a corpus that under-represents atypical or impaired speech makes the codec worst exactly where intelligibility is the product (Table 72, Figure 16).

INT CONDITIONAL · *Signal*: whether a satellite, IoT, automotive, accessibility or public-safety input appears

at PCG#57 or in the SA4 requirement work. *Holds*: the corpus requirement is written for the channel the codec is for. *Fails*: it is written for the speech the proponents already hold. *What would weaken this reading*: a coverage requirement in ULBC-3 that already names impaired and satellite-carried speech, which the versions read here do not.

INT WELL SUPPORTED · Two verticals are structurally different from the rest. A **research and licensing innovator** contributes invention and licenses it rather than shipping product: it holds patents and not corpora, and a corpus is not patentable subject matter (Table 70). Nothing in its own instrument changes. But if a learned codec's advantage moves from architecture to data, the leverage a standard returns moves with it - from those who invent the method to those who hold the material - and the balance between them is settled by a decision about datasets rather than by anything in an IPR policy. **AI model developers and data curators** sit at the other end of the same shift: their corpora are the advantage a common set removes, and a telecommunications partnership stating what it requires of training material hands regulators and buyers a comparator for disclosure elsewhere. Neither effect is raised in any document read for this brief.

13.7 Four lenses on the same effect: liability, IP, dependency and position

INT WELL SUPPORTED · Table 73 reads the same twenty verticals again through four lenses, because a first-order effect and a second-order effect are not the same kind of exposure and should not be weighed as if they were. Legal and compliance exposure asks who answers for the material to a regulator. Research and intellectual property asks whose advantage the rule creates or removes. Supply-chain dependency asks who has to wait for someone else. Strategic position asks what happens to speed to market and to whatever moat a firm has.

INT ESTABLISHED · The four lenses do not rank the verticals the same way, and that is the finding. Legal exposure concentrates at the two ends of the chain - the curator who holds the recordings and the vendor who ships the product - exactly as Table 53 already shows for the chain as a whole. Intellectual-property exposure concentrates in the middle, among the firms whose corpora or methods are the competitive asset. Supply-chain dependency runs one way only: everything downstream of the trained artefact waits on whoever can regenerate it. And strategic position moves in opposite directions for the two kinds of innovator, which is why one decision looks like an opening to one firm and a closing to another.

CTX Read against the law, the lenses also explain why the deadlines outside 3GPP matter more to some verticals than to others. The revised Product Liability Directive and the Cyber Resilience Act reach whoever places the product on the market, not whoever trained the model [L13] [L16]. The AI Act's training-content duties reach the model developer [L11] [L03]. Data protection reaches whoever holds the recordings (Table 49, Table 51). No instrument read here reaches the standards body that chose the corpus, which is the asymmetry every row of the table turns on.

INT CONDITIONAL · *Signal*: whether any Madrid or PCG#57 text allocates a duty by lens - to the holder of the material, to the publisher, or to the implementer - rather than leaving all three to the general law. *Holds*: the allocation is deliberate and can be planned against. *Fails*: it is decided later by whichever regulator asks first. *What would weaken this reading*: a legal opinion tabled at PCG that already allocates the duties; the paper read here sets out levels of involvement rather than allocations [P03].

What this part established

A patent licence and a licence to copy a trained artefact are different instruments; the partnership's declaration machinery has only ever addressed the first (Table 70).

A common corpus removes data as a source of advantage and a proprietary one keeps it, so the choice is an innovation choice before it is a legal one (Table 20).

One global corpus is governed by the strictest cell in every column it touches, which makes the choice of corpus a choice about jurisdiction (Table 50).

Fourteen seats, four hard lines, and an asymmetry: the seats that can move the decision are largely not the seats that carry it (Table 71, Figure 15).


Twenty verticals, each with a first-order effect and a second-order one; the second is where the cost sits, and no document read here discusses it (Table 72).

Under four lenses the same verticals rank differently: legal exposure at the ends of the chain, IP exposure in the middle, dependency all one way (Table 73).

What this part established — part 13

Hands forward: the method - part 14.

Figure 15. Who can move the outcome, and who carries it: fourteen seats by influence and exposure



Two ordinal scales, read from the record: distance carries order only, and no disc is a named organisation. Read from Table 71.

The asymmetry in one picture, for anyone deciding whom a rule has to be written for. The seats that can move the decision are largely not the seats that carry it: the implementer that ships the weights, the operator that pays for the service life, the laboratory that must detect leakage and the speaker whose recording is in the corpus all sit high on exposure and low on influence, and none of them has tabled a document at either meeting.

Table 70. Standard-essential patents against a normative dataset: what each instrument reaches

Question	What answers it today	What a normative dataset would change	What stays open
What makes a patent essential	a claim that an implementation of the specification cannot avoid; each OP's own IPR policy asks for timely declaration and an undertaking to license on fair, reasonable and non-discriminatory terms [G25]	nothing in the test: claims on the architecture, quantisation and inference path read on a learned decoder as they do on a classical one	whether a claim reading only on a training method is essential to a standard that publishes the trained result
What protects the corpus	copyright in the recordings, performers' rights, the database right, contract and platform terms, and data-protection law - none of them patent law	the same instruments, now attached to an input of a published specification	who holds them once the artefact is the partnership's
What an implementer needs	a licence on FRAND terms from the declared essential-patent holders	that, and a separate right to copy, embed and ship the weights themselves	whether OP publication carries that right, since Article 46 is written for documents [G25]
Where a declaration is made	against a specification, in each Organisational Partner's own database	against a specification that now carries a trained artefact	the dataset behind it is declared nowhere in the partnership's machinery
What the FRAND undertaking covers	the licensing of essential claims	unchanged: no IPR policy read here undertakes to make any input available	how a maintenance obligation would be carried, since it is not an IPR undertaking
Who can be asked	the declared holders	and whoever holds rights in the recordings, who has given no undertaking	consent and remuneration, unaddressed in every option read here
Whether the material can be licensed at scale	catalogue and platform licences exist for viewing, performance and distribution, not for training a standardised model	a standards body would need a licence form no catalogue licence currently issues	who negotiates it, on whose behalf, and at what price
What a dispute would be about	essentiality, infringement and the terms of a licence	and, separately, whether the material behind the artefact was lawfully used	which forum hears the second question, and against whom

For licensing, IPR and research readers. A patent licence and a licence to copy an artefact are different instruments with different holders, and the partnership's declaration machinery has only ever addressed the first. Context and this brief's reading, not advice: whether any instrument applies to a given actor, product or corpus is a question for counsel.

Table 71. Fourteen seats: what each wants, what each fears, and what each will not trade

Seat	What it wants	What it fears	Negotiable	Hard line	What to watch
Proponent holding a private corpus	to win with a codec it has already trained	disclosing material it holds under contract or trade secret	a described corpus, an escrowed copy, a third-party verifier	publishing the material itself	the selection rules, ULBC-5b
Proponent without one	a comparison its architecture can win	losing to data rather than to design	a described corpus, if the description is checkable	a comparison in which one candidate's data cannot be examined at all	the qualification rules, ULBC-5a
Organisational Partner, as publisher	a specification it can stand behind	custody of material it cannot verify, and the cost of keeping it	a declaration its own lawyers verify; a maintenance package held elsewhere	custody without a legal basis and the resources to run it	the new PCG group's terms
TSG SA and SA WG4, as roles	a selection that starts on time	a dependency on a body with no date	the order of work	selecting on rules that are not agreed	the TSG SA#113 report
Mobile operator	a codec that outlives the company that trained it	a fixed weight file nobody can regenerate late in the service life	who holds the recipe, and where	no path to regeneration at all	the wording on 'transparent'
Device or chipset vendor	weights it can pin, licence and ship	a second licensing gate discovered after the design is frozen	the form of the licence and who grants it	shipping weights it cannot licence	whether PCG separates the two gates
Essential-patent holder	declaration and terms that work as they do today	an undertaking read as reaching material it does not hold	declaring against a specification that carries an artefact	an IPR undertaking extended to a corpus	any PCG text that names licensing
Implementer as licensee	one answer to the rights question, not one per vendor	clearing every patent and still not clearing the artefact	where the answer comes from	carrying the residual risk alone	a rights field in any declaration
Test laboratory	conditions it can test against	leakage it has no rule to detect	how many languages, and which	certifying a comparison it cannot verify	ULBC-8a and ULBC-8b
Dataset curator or platform	its terms respected by whoever names its corpus	redistribution under terms it never granted	attribution, versioning and access conditions	redistribution its licence does not permit	whether a corpus is named at all
Rightsholder, performer or speaker	consent, attribution and a say in reuse	a voice in a global specification with neither	nothing here is theirs to trade: they are represented by law, not in the room	use the law does not permit	any text that names consent
Content platform as rightsholder at scale	a licence form for training that its catalogue can actually issue	material used for training under terms written for viewing	price, scope, attribution and term	training use its own licences never granted	whether any text names a licence for training rather than a requirement
Regulator or public-safety authority	evidence that emergency speech stays intelligible under stress	a fidelity claim it cannot inspect	the form the evidence takes	accepting a claim with no evidence behind it	fidelity measures in ULBC-3
National or regional policy maker	a rule its own law can live with	a global artefact lawful in some regions only	the standard of care	a rule that forces a choice between markets	whether the group's terms name a jurisdiction

The decision read from every seat it reaches, inside the room and outside it. The seats are kinds of actor rather than companies - a firm may occupy several at once - and every cell is this brief's reading of the record, with no motive attributed to any named organisation. Only four hard lines are stated in absolute terms; the rest of the table is negotiable ground.

Table 72. Twenty verticals: the first-order effect, the second-order effect, and whom the record names

Vertical	What it builds or runs	First-order effect	Second-order effect	Whom the ULBC record names here
----------	------------------------	--------------------	---------------------	---------------------------------

For anyone reading this from a market rather than from a meeting. The first-order column is what a normative dataset does at once; the second-order column is what follows from that, and it is where most of the cost sits and where no document read here looks. The last column is a fact about the record - whom it names in that vertical - and never a statement about any company's position; where a vertical is illustrated by firms the record does not name, the cell says so.

Chipset and silicon vendors	the decoder that runs inside a power and memory budget	the weights become a fixed target: their size, precision and operator set decide the silicon	if the weights can be retrained, every retrain is a new conformance target, so the part must be sized for a model that does not exist yet - a maintenance rule at PCG becomes a die-area decision two years earlier	Qualcomm, Samsung, Huawei, in the record
Infrastructure and base-station vendors	IMS, the media path and the interworking round it	little directly: the codec runs at the endpoints	a decoder that can change version needs the version signalled, negotiated and interworked, so a training-data maintenance rule surfaces as an IMS and transcoding feature	Ericsson, Nokia, Huawei, CMCC
Wi-Fi and short-range vendors	the access an IMS call also runs over	an IMS codec reaches every access the session can use, so a rule about its weights follows the call onto Wi-Fi	voice certification for that access meets a decoder whose version can change, and a certification regime built for fixed codecs has no place to record which model was tested	none in the record read here
UE and handset vendors	the product that ships the weights to market	a second licence to obtain: the right to copy, embed and distribute the artefact	product-liability and cyber duties make the weight file a maintained artefact for the product's life, so a codec becomes a supported dependency rather than an integration	Samsung, vivo, Huawei
Operators and carriers	a voice service with a life measured in decades	the service will outlive the company that trained the model	procurement language changes before the standard does: contracts begin to name regeneration rights, escrow and version support, moving the cost into the supply agreement	Orange, CMCC, and the survey's operator answers
IoT module and device makers	the constrained devices ULBC-MED is written for	codec footprint decides device cost in a market measured in cents	a maintenance rule that assumes retraining reaches devices with a decade of battery and no update path, so the rule's cost lands where the update never arrives	none in the record read here
Satellite and non-terrestrial vendors	the geostationary link the work item exists for [G01]	the codec is selected for a channel whose latency and impairments a terrestrial corpus does not contain	if the corpus omits satellite-carried speech the winner is tuned for conditions it will not meet, and coverage becomes a deployment question rather than a fairness one	none in the record read here
Cloud and hyperscale providers	where training runs and where a corpus would physically sit	a common corpus has to live somewhere, with an availability commitment behind it	custody the partnership will not take is taken by whoever hosts it: residency, transfer and egress terms become governance terms, and the standard acquires a dependency it never chose	none in the record read here; the type is illustrated by the large cloud platforms
AI model developers and data curators	the corpora and the models trained on them	a partnership rule about training material becomes a published reference point	a telecommunications partnership stating what it requires of a corpus gives regulators and buyers a comparator for training-content disclosure elsewhere, which is leverage nobody asked for and nobody controls	the four corpora TTA names [P02]; no curator tables a document
Gaming and interactive platforms	speech at scale, and the services that generate it	their material is the kind a common corpus would be assembled from	a rule about corpora becomes a reference for their own dataset governance, and for what they will and will not contribute to a standards body	Tencent, in the record as the Chair's organisation
Media and audio technology companies	candidate codecs and the training behind them	the corpus rule decides whether a data advantage counts at all	research budgets move - under disclosure toward reproducible pipelines, under a proprietary regime toward acquisition - and what is worth patenting moves with them	Dolby, Fraunhofer, VoiceAge, Ericsson, Nokia
Streaming and content platforms	catalogues of recorded speech and dialogue, licensed and rights-cleared	they hold the cleared material a corpus would want, and the clearance machinery a standards body does not have	a partnership that states what a corpus must satisfy writes, in effect, a licence form the catalogue business does not yet issue - and whoever issues it first sets the price of recorded speech for training	none in the record read here; the type is illustrated by the large streaming services
Automotive and in-vehicle systems	hands-free voice and regulated emergency calling, in products with fifteen-year lives	a type-approval regime that records a software build meets a decoder whose model version can change	a retrain reaches a fleet through a recall or an over-the-air campaign, so a maintenance rule written for handsets lands on the homologation calendar of an industry that cannot ship a patch in a week	none in the record read here
Enterprise voice and contact centres	business calling, conferencing and recorded customer conversations that interwork with IMS	they hold recorded speech at scale, under consent written for service delivery rather than for training	a partnership rule about training material becomes the question their own customers ask about call recordings, and consent captured for quality assurance does not carry across to model training	none in the record read here
Voice biometrics and financial services	speaker verification over the telephone channel	a learned codec reconstructs speech rather than carrying a waveform, so the features a verifier relies on become a property of the model and its training material	change the weights and enrolled voiceprints may no longer match, which turns a maintenance rule at PCG into a re-enrolment programme in a bank - the sharpest second-order effect in the table, and the vertical least visible in the record	none in the record read here

Vertical	What it builds or runs	First-order effect	Second-order effect	Whom the ULBC record names here
----------	------------------------	--------------------	---------------------	---------------------------------

For anyone reading this from a market rather than from a meeting. The first-order column is what a normative dataset does at once; the second-order column is what follows from that, and it is where most of the cost sits and where no document read here looks. The last column is a fact about the record - whom it names in that vertical - and never a statement about any company's position; where a vertical is illustrated by firms the record does not name, the cell says so.

Accessibility and assistive technology	hearing aids, captioning and speech aids, where intelligibility is the product	a corpus that under-represents atypical, impaired or accented speech makes the codec worst exactly where it matters most	accessibility duties attach to the service rather than to the standard, so a coverage decision taken inside a corpus arrives as a compliance question for whoever sells the service	none in the record read here
Defence and mission-critical networks	mission-critical and government voice carried on 3GPP networks	an artefact whose provenance is undisclosed is hard to accept in an accredited or sovereign deployment	a corpus assembled under one region's law becomes a procurement obstacle in another, so the jurisdiction question of section 13.4 returns as an accreditation and export question	none in the record read here
Research and licensing innovators	invention and contributions rather than shipped goods	nothing in their own instrument: a corpus is not patentable subject matter	if a learned codec's advantage moves from architecture to data, the leverage a standard returns moves from those who invent the method to those who hold the material - the second-order effect with the longest reach and no document behind it	InterDigital, in the reflector thread; firms such as Ofinno and Wilus are named here as examples of the type and appear nowhere in the record read for this brief
Test laboratories and certification bodies	the qualification and selection tests	separation of test from training material becomes something to check, not assume	a checkable rule creates an accreditation market: someone must be trusted to say the separation held, and no body read here appoints them	named by role in SP-260907; none by name
Public-safety and emergency services	the calls that must stay intelligible under stress	intelligibility becomes a property of the training material, not only of the algorithm	evidence duties follow: a claim about emergency fidelity has to be demonstrable to a regulator that reads evidence rather than data	none in the record read here

Figure 16. The impact hierarchy: one decision, five industry groups, twenty verticals, and the effect that follows the first one



Depth carries structure and nothing else: a leaf lower on the page is not larger, later or more important than one above it, and the order inside a group is the order of Table 72. The groups are this brief's grouping of verticals that share a kind of exposure; the squares are its reading of Table 73 on a three-step scale.

The whole analysis in one structure, for a reader looking for their own industry before reading about it. The root is the decision; the second column groups the verticals that share a kind of exposure; each leaf carries the second-order effect - the one that follows the obvious one - and the four squares say how heavily that vertical is exposed on each lens. Read the squares across a row for one industry's profile, and down a column for where a single kind of exposure concentrates.

Table 73. The same verticals under four lenses: liability, intellectual property, dependency and position

Vertical	Legal and compliance exposure	R&D and intellectual property	Supply-chain dependency	Strategic position
Chipset and silicon vendors	low: it neither collects nor holds the material	a fixed artefact it cannot optimise around; licence terms it does not negotiate	locked to whoever can regenerate the weights	wins if the weights are frozen early; loses if they move after tape-out
Infrastructure and base-station vendors	low, until version identity becomes evidence	no data advantage either way; interworking effort is new engineering	carries version negotiation for every endpoint it serves	neutral: sells into whichever regime wins
Wi-Fi and short-range vendors	low: it carries the session, not the model	none: no proprietary corpus is at stake	inherits a decoder version it does not control	affected only through certification, and late
UE and handset vendors	high: product liability from 9 December 2026, cyber reporting since 11 September 2026, and market surveillance asks the manufacturer [L.13] [L.16]	pays a licence it cannot cross-license with patents it holds	depends on the proponent for every retrain over the product's life	speed to market turns on whether the rights question is answered once or per vendor
Operators and carriers	moderate: service obligations, emergency calling, and subscriber data	no data advantage; regeneration rights are a contract term, not an asset	the deepest lock-in in the chain: a decades-long service on one trained artefact	gains from any rule that makes the codec outlive its author
IoT module and device makers	moderate, and unmanaged: liability attaches to devices with no update path	cost, not IP: the footprint is the product decision	cannot absorb a retrain; a maintenance rule reaches them as a scrap decision	loses under any rule that assumes the field can be updated
Satellite and non-terrestrial vendors	low directly; high indirectly, if intelligibility fails on their link	no corpus of their own; the impaired-speech data that would help may not exist	depends on a corpus assembled by others for a channel that is not theirs	the case the codec is for, and the case least represented in the room
Cloud and hyperscale providers	high if they host: residency, transfer and processor duties attach to whoever holds it	hosting is not an IP position; the advantage is the relationship, not the rights	becomes the ecosystem's single point of custody by default	gains a standards-scale dependency without having asked for it
AI model developers and data curators	the AI Act's training-content duties already reach them; a 3GPP rule adds a comparator [L.11] [L.03]	their proprietary corpora are the advantage a common set removes	supplies the material the rest of the chain depends on	a disclosure standard set elsewhere becomes the benchmark they are measured against
Gaming and interactive platforms	privacy exposure on speech they already hold; consent is the live question	large corpora, no standards obligation to share them	a potential contributor rather than a dependant	chooses whether to be a supplier of data to standards work at all
Media and audio technology companies	moderate: copyright and performer rights in what they recorded	the clearest R&D swing in the table: data advantage against method advantage	supplies the artefact the whole chain then depends on	the moat is the corpus under one regime and the architecture under the other
Streaming and content platforms	high: performer, dubbing and neighbouring rights in catalogue speech, and consent for any synthetic voice built from it	the catalogue is the asset; licensing it for training is a new line of revenue or a new liability, and the same contract decides which	supplies rather than depends	a rule written elsewhere becomes the template for their own licensing
Automotive and in-vehicle systems	high: type approval, product liability and emergency-calling duties at once	none of its own: it buys the codec	the least updatable fleet in the table	the longest product life here, and the least tolerance for a moving artefact
Enterprise voice and contact centres	high: recordings, consent and cross-border transfer, on material already held	moderate: the recordings are an asset only if consent permits the use	depends on the codec and on whoever can retrain it	gains a template, and risks a retro-fit consent problem on what it already holds
Voice biometrics and financial services	high: a voiceprint is biometric data before it is a security control	low: the verifier is theirs, the codec is not	the most brittle dependency here: a version change can invalidate enrolments	exposed to a decision it has no seat in and no document read here mentions
Accessibility and assistive technology	moderate: accessibility duties run to the service, not to the standard	none; coverage is bought, not owned	depends entirely on what the corpus contains	coverage is the whole product, and coverage is a corpus decision
Defence and mission-critical networks	high, and sovereign: accreditation asks where the material came from	low; procurement rather than invention	may refuse an artefact it cannot inspect, a dependency exercised as a veto	the one buyer that can decline the standard outright
Research and licensing innovators	low: no product, no data custody	patents unaffected; a corpus is not patentable, so no new declarations arise	outside the maintenance chain entirely	leverage falls as data displaces method, and no instrument they hold offsets it
Test laboratories and certification bodies	low, unless they certify what they cannot verify	method, not IP: the detection technique is the asset	the ecosystem depends on them to say the separation held	a new accreditation market, if the rule is written to be checkable
Public-safety and emergency services	the duty runs to them, not from them	none; they buy rather than build	depends on every layer above for evidence they can read	the one seat that can make fidelity non-negotiable, and it has tabled nothing

14 Method, counting rules and limits

For: anyone checking a number.

14.1 Scope and documentary base

CNT **ESTABLISHED** · The documentary base is 68 source records in 7 bands, set out in [Table A1](#) and summarised in [Table 75](#). As one partition: 29 were read in full and 4 in part; 6 are read through the companion brief's verified datasets; 12 through another document that quotes or reports them; 9 as the Apex export quotes them; and 8 are known by title, status or listing only. The six classes sum to the ledger. The documentary cutoff is 18 September 2026; this edition, dated 19 September 2026, adds one record dated before that cutoff and read after it - the meeting register of PCG#56 bis_Ad Hoc [P05] - and says so wherever a count moved because of it.

SRC The bands follow the chain. The closed record of TSG SA#111 and #112 is read through the exports and verified dispositions of C20004166 [B04]. The SA4 record is read through the Chair's report, the status report and the public reflector thread, supplied as a message [M01] [M02] [W04]. The three PCG#56bis inputs and the agenda were read in full, as were the four

Madrid documents that could be opened [P01] [M03] [M05]. The other standards bodies are read through ten JVET documents and the public pages of MPEG, JPEG, ITU and the IETF [J01] [K04] [I01].

CTX External context - law, regulation, other standards, corpora - is carried in 213 references in 10 bands, each with its address printed in full. Official texts were used where they could be reached; where only a court's press note, a law firm's summary or a secondary report was available, the reference says so, and every paragraph resting on such material carries the context mark.

What the evidence is decides how it can be counted.

14.2 How a number is counted, and what the count excludes

CNT **ESTABLISHED** · Every number printed in this brief is computed from the packaged files by one module, and every rule is stated once in [Table 76](#). Survey shares are printed with their base, and two bases exist: all 24 respondents, and those who expressed a position. A figure one document reports about another is attributed to the document that reports it, never to the one it describes.

CNT **ESTABLISHED** · Matrix cells are stated, partly stated or not stated; only stated cells are counted, and

partial readings are printed but never added to them. Dates carry their precision: a date the record gives only to the month is drawn across the month and never given a day.

CNT **ESTABLISHED** · Roles are not organisations. The SA WG4 Chair, the Audio SWG Chair, the rapporteur and the Pdoc editors are cited by role, and no role-authored document or message is counted as an organisation's activity; the thread's organisations are identified on the basis the thread itself carries - a signature, an address domain, a display name - or, for 2 messages, from the sender the thread names, with the organisation supplied for this brief; none is inferred. No participant in the standards process is named: people are cited by the role the record gives them, and court cases by their published names.

INT Two limits follow and are disclosed. The survey counts one answer per company, whatever that company's role, and its anonymity means no answer can be attributed; counts from it are therefore statements about the tally, not about who holds which view. And a matrix of what inputs state is a reading of their text: a different reader could mark a partial cell as stated or not stated, which is why the basis of every cell is printed beside it in [Table 23](#).

Counting is one half of the discipline; citing is the other.

14.3 The record and the reading: how this brief cites

SRC Four marks head every paragraph that carries a claim - SRC for a statement a source record makes, CNT for a count by a stated rule, INT for this brief's reading, CTX for context the 3GPP record does not establish - and a paragraph takes the weakest class it depends on. Confidence is stated where a finding is first made: established, well supported or conditional. Every conditional reading names its signal and what would weaken it.

CNT **ESTABLISHED** · Wording taken from a record is a "verbatim quotation" - the one term this brief uses for it - and is set in green italic with its source in the same short form everywhere, so source text is distinguishable at sight from this brief's own. The 136 quotations come from 22 sources and each was verified as an exact substring of the converted source text, after collapsing whitespace only; they are listed in [Table A2](#). A bracketed key - [P02], [M05], [L13] - leads to the ledger row or the reference that carries it, and every such key in the text is a live link.

INT Quotation is kept short and always separated from reading: a sentence of this brief never puts its own interpretation inside quotation marks, and a quotation is never extended beyond what the source says. Where a record's own number appears inside a quotation it is reproduced exactly, and the numeric conventions of this brief apply only outside it. Every interpretation names the documents it reads and carries the keys, anchors and references that reach them, each a live link.

A brief that states how it cites must also state what it does not claim.

14.4 Boundary conditions and what this brief does not claim

INT This brief is standards and policy analysis, not legal advice. Its legal and regulatory material describes instruments and decisions as their texts or reliable summaries state them, as of 18 September 2026; whether any of them applies to a given actor, product or corpus depends on facts this brief does not have, and is a question for counsel.

INT **ESTABLISHED** · It does not claim to know what PCG#56bis concluded beyond what two company papers report, what SP-260913 asks, what the new PCG

ad hoc group's terms are, or how any company answered the SA4 survey. It does not rank the inputs, attribute a motive to any organisation, or read an organisation's absence from a document as a view. [Table 77](#) lists every record this brief could not read and what was done instead.

INT **ESTABLISHED** · **The cross-SDO process boundary.** Other standards bodies are compared here for the technical instruments visible in their public records. The comparison does not assume that they share consensus rules, approval paths, model artefacts or schedules, and elapsed time in one body is not used as a forecast for another ([Table 60](#), [Table 59](#), [Figure 13](#)). Working discussion that is not public is not treated as a source record. Where such discussion points to a gap, it is used only to test whether the public record supports a stronger claim than this brief already makes - which is why [section 7.3](#) states what the MPEG material read here does not establish, rather than filling the gap from another source.

INT **ESTABLISHED** · **What the historical codec comparison carries.** The EVS and IVAS programmes in [Table 61](#) are recorded context for the density of the ULBC calendar. They differ from ULBC in scope, starting maturity, candidate structure and requirements, and neither standardised a trained model, so their durations are not a forecast and nothing in this brief infers from them that the recorded ULBC plan can or cannot be met.

INT It does not forecast what Madrid, PCG or SA4 will decide. The conditional readings and scenarios describe what would follow from stated assumptions, and each names the record that would confirm or weaken it; none is a prediction, and none carries a probability.

What this part established

68 source records, 136 verified quotations and 213 references; every count by a rule stated once; no participant named; no forecast ([Tables 75](#) to [77](#)).

Six read classes partition the ledger, so every count says how the records behind it were read; 29 of the 68 were read in full ([Table 75](#)).

What could not be read is printed rather than smoothed over, with what was done instead and what it changes ([Table 77](#)).

What this part established — part 14

The brief is standards and policy analysis. Legal material is described as the texts or reliable summaries state it at the cutoff; whether an instrument applies to a given actor, product or corpus is a question for counsel.

Table 74. Eleven seats, eleven readings of the same record

Seat	What the record means for it	Start at	Watch
Operator	maintenance independence over a decades-long service	5.4, 6.4	the liaison wording on 'transparent'
Declared codec proponent	the rules and calendar of the competition	5.1, 5.3, 6.2	ULBC-5a, ULBC-5b and sequencing
Implementer, device or network	the evidence needed to ship and defend a product	9.1, 9.3	a declaration; ULBC-6 deliverables
Chipset vendor	the size of the weights and the conformance tolerance	5.2, 6.1	complexity metrics in ULBC-4
OP legal counsel	exposure, verification and resources	4.3, 10.2, 10.3	the PCG group's terms; counsel
PCG delegate	the option space and what each input adds	4.5, 10.1	PCG#57
SA4 leadership, as roles	a dependency on a body with no date	1.4, 5.2	TSG SA#113 guidance
Test laboratory	separation checks and the number of test languages	6.2, 6.3	ULBC-8a and ULBC-8b
Regulator or public-safety authority	fidelity and behaviour under stress	8.4, 9.4	fidelity measures in ULBC-3
Dataset curator or researcher	the terms on which a corpus could be named	4.7, 9.2	a corpus named, and its access terms
RAN or CT delegate, and the standards researcher	a precedent for every model their own groups already specify	1.5, 10.4, 10.5	the threshold PCG names, and whether it is written for ULBC or for models

Readerships, not parties; no seat is attributed to a named organisation.

Table 75. The documentary base, by band and by how each record was read

Band	Records	In full	In part	Companion datasets	Apex export	Through another	Title only
Closed record: TSG SA #111 Fukuoka and #112 Singapore	12	-	-	5	-	-	7
SA WG4 #137-e and its reflector, August 2026	13	1	-	-	-	12	-
PCG#56 bis_Ad Hoc, 3 September 2026	6	6	-	-	-	-	-
TSG SA #113 Madrid, held 15-18 September 2026	10	8	1	-	-	-	1
TSG SA reflector, 16-18 September 2026	2	2	-	-	-	-	-
Other standards bodies: JVET, MPEG and JPEG records	19	8	2	-	9	-	-
Base briefings and companion reports	6	4	1	1	-	-	-
All records	68	29	4	6	9	12	8

One read class per record; the six classes sum to the ledger. Table A1 gives every record with its address.

Table 76. Every counting rule, stated once

Rule	Where applied	What it keeps or excludes	Count under it
Survey shares carry their base	Tables 13 and 66; Figure 5	both bases printed	24 respondents; 22 with a position on need
Reported figures stay with the reporting document	Section 3.5; Table 13	annex figures attributed to SP-260919	2 figures
Matrix cells: stated, partly, not stated	Tables 23 and 49; Figure 8	partial readings printed, never added	48 input cells; 64 body cells
Dates carry their precision	Figure 1; Table 3	a month-only date is drawn across its month	8 month-only dates
Roles are not organisations	Tables 16, 27 and 30	role messages and documents not counted for any organisation	2 Chair messages; 11 editor assignments
Organisation named by the thread, or supplied and marked	Table 16	no organisation inferred from a position or a domain guess	9 from the thread; 2 supplied
A timestamp is converted only where a zone is recorded	Table 16	a quoted header is rendered in the quoting client's own locale	4 of 11 carry a zone
One read class per record	Table 75	six classes that sum to the ledger	68 records
Verbatim: exact substring after whitespace collapse	Table A2	no other normalisation	136 verbatim quotations
Codec rows from the companion exports	Table 14	rows naming the codec only	22 rows

The rules are applied in the data module and nowhere else, so no measure downstream can reclassify a row.

Table 77. What this brief could not read, and what was done instead

Record	What could not be done	Why	What was done instead	Effect
SP-260988 and SP-261002	text not read	the two late drafts were circulated in the meeting drafts folder and superseded within the week	read through the meeting report, which records what each changed and what was said	the drafts are described, not quoted
PCG#56 bis_Ad Hoc register	read late	the spreadsheet could not be parsed where it was first listed	read in full for this edition; it sources every document and records no disposition	the attribution of the legal paper is now recorded, not inferred
TSG SA#113 report	a draft	the meeting report is approved at the following plenary	read as a draft and cited as one; every disposition cross-checked against the meeting's own register and the Chair's notes	wording may change before TSG SA#114 approves it
S4-261382, S4-261674, S4-261712	not read directly	not among the documents available	read through SP-260909, SP-260912, SP-260919 and the thread	annex figures via SP-260919
SA4#137-e contributions	not read	not among the documents available	read through the base briefing	positions via a secondary source
SA4 public list archive	the archive page was not paged through	the thread was already complete in the message as circulated on the public list	the thread read message by message from that circulated copy	none: thread read in full
MPEG Video Coding for Machines call	not opened	distributed as an archive	the condition as TTA and the Apex export quote it	carried as quoted
Some legal and standards texts	read through public summaries	the full texts of some standards are published for sale, and some judgments are published only in summary	official notices, ministry statements and public legal commentary, each marked as context	may lag the primary text
SA4#138 and later dates	month only	the record gives months	drawn across the month	none
The new PCG ad hoc group	no terms available	none published	treated as unknown	conditional readings only

The exceptions register. Each row says what could not be done, why, what was done instead, and what it changes.

Appendix A. Source ledger and verbatim ledger

Part 1 gives every record this brief reads, sectioned by band, with how it was read and the address at which it can be found; every anchor in the text leads to its row. Part 2 reproduces every quotation with its source and place, each verified as an exact substring of the source text.

Table A1. Source ledger: every record this brief reads

Anchor	Kind and date	Identity as cited, and where	What it carries	How read
<i>Part 1 of the appendix. Every anchor in the text leads to its row here; the address is printed in full under the identity.</i>				
		Closed record: TSG SA #111 Fukuoka and #112 Singapore		
[S01]	LS out; Mar 2026	SP-260324 LS on Support of ULBC in Release 20 (revision of SP-260283, SP-260257) www.3gpp.org/ftp/tsg_sa/TSR/SA/TSRGS_111_Fukuoka_2026-03/Docs/SP-260324.zip	Approved at SA#111: the plenary asks SA WG4 about Rel-20 support for the ultra-low bitrate codec	disposition and export row, via C20004166
[S02]	discussion; Mar 2026	SP-260258 Views and Suggestions on progressing NB IoT GEO voice www.3gpp.org/ftp/tsg_sa/TSR/SA/TSRGS_111_Fukuoka_2026-03/Docs/SP-260258.zip	SA#111 discussion: whether Rel-20 carries the full codec or a phased approach; study recorded as 43% complete	disposition, via C20004166
[S03]	draft TR; Jun 2026	SP-260421 TR 26.940 v1.0.0 www.3gpp.org/ftp/tsg_sa/TSR/SA/TSRGS_112_Singapore_2026-06/Docs/SP-260421.zip	The FS_ULBC study report, approved at SA#112	disposition and export row, via C20004166
[S04]	discussion; Jun 2026	SP-260609 Principles for standardizing an AI based codec for ULBC www.3gpp.org/ftp/tsg_sa/TSR/SA/TSRGS_112_Singapore_2026-06/Docs/SP-260609.zip	Where the training-dataset concern was raised at SA#112, per PCG#56bis(26)02; noted	title and status only; text not read
[S05]	discussion; Jun 2026	SP-260623 Thoughts on proposed principles for standardizing an AI based codec for ULBC www.3gpp.org/ftp/tsg_sa/TSR/SA/TSRGS_112_Singapore_2026-06/Docs/SP-260623.zip	The counter-paper tabled beside SP-260609; noted	title and status only; text not read
[S06]	other; Jun 2026	SP-260641 Guidance for SA WG4 ULBC Work (revised to SP-260671) www.3gpp.org/ftp/tsg_sa/TSR/SA/TSRGS_112_Singapore_2026-06/Docs/SP-260641.zip	The guidance whose wording SA#112 edited from common to technical training dataset	disposition and export row, via C20004166
[S07]	WID new; Jun 2026	SP-260696 New WID on Ultra Low Bitrate Speech Codec for Voice over GEO (revision of SP-260640, SP-260417) www.3gpp.org/ftp/tsg_sa/TSR/SA/TSRGS_112_Singapore_2026-06/Docs/SP-260696.zip	The approved ULBC-MED work item, Rel-20	disposition and export row, via C20004166
[S08]	discussion; Jun 2026	SP-260611 New WID on Ultra Low Bitrate Speech Codec for IMS Voice Call over NB-IoT NTN www.3gpp.org/ftp/tsg_sa/TSR/SA/TSRGS_112_Singapore_2026-06/Docs/SP-260611.zip	An alternative WID text at SA#112; noted	title and status only
[S09]	discussion; Jun 2026	SP-260619 Clarification on ULBC codec bitrates support in the Rel-20 ULBC-MED Work Item www.3gpp.org/ftp/tsg_sa/TSR/SA/TSRGS_112_Singapore_2026-06/Docs/SP-260619.zip	Bitrate scope, co-signed by satellite operators and ESA; noted	title and status only
[S10]	discussion; Jun 2026	SP-260602 Considerations and proposals for ULBC www.3gpp.org/ftp/tsg_sa/TSR/SA/TSRGS_112_Singapore_2026-06/Docs/SP-260602.zip	Noted at SA#112	title and status only
[S11]	discussion; Jun 2026	SP-260627 Views and Suggestions on Rel-20 ULBC www.3gpp.org/ftp/tsg_sa/TSR/SA/TSRGS_112_Singapore_2026-06/Docs/SP-260627.zip	The study rapporteur's view; noted	title and status only
[S12]	report; Jun 2026	SP-260614 SA WG4 chair status report on ULBC www.3gpp.org/ftp/tsg_sa/TSR/SA/TSRGS_112_Singapore_2026-06/Docs/SP-260614.zip	The Chair's status report at SA#112; noted	title and status only
		SA WG4 #137-e and its reflector, August 2026		
[W01]	survey; Aug 2026	S4-261382 Survey on technical considerations related to training and validation datasets for ULBC candidate codecs www.3gpp.org/ftp/tsg_sa/WG4_CODECTSGS4_137-e/Docs/S4-261382.zip	The written survey; legal and licensing placed outside SA4	quoted through SP-260912 and the base briefing
[W02]	survey results; Aug 2026	S4-261674 ULBC Training and Validation Datasets - Survey results www.3gpp.org/ftp/tsg_sa/WG4_CODECTSGS4_137-e/Docs/S4-261674.zip	The fuller tally, excluded from the report to SA; its annex figures are quoted by SP-260919	through SP-260909 and SP-260919
[W03]	report; Aug 2026	S4-261712 SA4 Chair proposed report to SA on training dataset requirements (revision of S4-261675) www.3gpp.org/ftp/tsg_sa/WG4_CODECTSGS4_137-e/Docs/S4-261712.zip	The draft that became SP-260909	through the reflector thread and SP-260909
[W04]	reflector thread; Aug 2026	Re: Revised proposed response to SA on the ULBC training data requirements https://lists.etsi.org/scripts/wa.exe?A0=3GPP_TSG_SA_WG4	The public debate over how the report to SA should state the reasons on each side	read in full (message as supplied)
[W05]	contribution; Aug 2026	S4-261644 Dolby reply to survey on technical considerations related to training and validation datasets www.3gpp.org/ftp/tsg_sa/WG4_CODECTSGS4_137-e/Docs/S4-261644.zip	A published survey reply against SA4 developing requirements or defining common data	excerpt as quoted in the reflector thread
[W06]	contribution; Aug 2026	S4-261572 Response to survey on need for technical training dataset requirements for ULBC www.3gpp.org/ftp/tsg_sa/WG4_CODECTSGS4_137-e/Docs/S4-261572.zip	A published survey reply	cited by the reflector thread and the base briefing
[W07]	contribution; Aug 2026	S4-261500 Process for establishing common ULBC training data www.3gpp.org/ftp/tsg_sa/WG4_CODECTSGS4_137-e/Docs/S4-261500.zip	Access, identification, versioning and freezing of a common pool	through the base briefing
[W08]	contribution; Aug 2026	S4-261501 Draft call for candidate datasets for a common ULBC training and validation data pool www.3gpp.org/ftp/tsg_sa/WG4_CODECTSGS4_137-e/Docs/S4-261501.zip	Candidate-dataset metadata: identity, rights pointer, use, overlap with test material	through the base briefing
[W09]	contribution; Aug 2026	S4-261509 Training data requirements for ULBC candidate submission www.3gpp.org/ftp/tsg_sa/WG4_CODECTSGS4_137-e/Docs/S4-261509.zip	Compliance with PCG/SA legal requirements as a qualification prerequisite	through the base briefing
[W10]	contribution; Aug 2026	S4-261510 Common training data for ULBC candidate codecs www.3gpp.org/ftp/tsg_sa/WG4_CODECTSGS4_137-e/Docs/S4-261510.zip	Common data to separate model merit from data advantage	through the base briefing
[W11]	contribution; Aug 2026	S4-261511 Invitation and process for defining common ULBC training data www.3gpp.org/ftp/tsg_sa/WG4_CODECTSGS4_137-e/Docs/S4-261511.zip	A timetable toward a stable common pool by SA4#138	through the base briefing
[W12]	contribution; Aug 2026	S4-261643 Considerations on ULBC-MED standardization www.3gpp.org/ftp/tsg_sa/WG4_CODECTSGS4_137-e/Docs/S4-261643.zip	Legal arrangements for selection testing, including multi-party NDAs	through the base briefing
[W13]	contribution; Aug 2026	S4-261703 ULBC collaboration www.3gpp.org/ftp/tsg_sa/WG4_CODECTSGS4_137-e/Docs/S4-261703.zip	A shared baseline: architecture, training code, common data, evaluation	through the base briefing
		PCG#56 bis Ad Hoc, 3 September 2026		
[P01]	agenda; 3 Sep 2026	PCG#56bis(26)01 Agenda www.3gpp.org/ftp/PCG/PCG_56_Bis/Docs/PCG56Bis_01.zip	Agenda item 3: 3GPP involvement and legal requirements for AI/ML training datasets	read in full
[P02]	contribution; 3 Sep 2026	PCG#56bis(26)02 Issues with AI model training datasets within 3GPP (ref: ULBC codec) and Way Forward www.3gpp.org/ftp/PCG/PCG_56_Bis/Docs/PCG56Bis_02.zip	The case for a common, public, permissively licensed ULBC corpus selected by SA4	read in full
[P03]	contribution; 3 Sep 2026	PCG#56bis(26)03 Legal Considerations Relating to Training Datasets for AI models www.3gpp.org/ftp/PCG/PCG_56_Bis/Docs/PCG56Bis_03.zip	Two scenarios, six rights categories, three governance levels, liability and external counsel	read in full
[P04]	slides; 3 Sep 2026	PCG#56bis(26)04 ATIS recommendations on legal requirements for AI/ML training datasets www.3gpp.org/ftp/PCG/PCG_56_Bis/Docs/PCG56Bis_04.zip	Two legal requirements, an OP legal group to verify them, and a common dataset as one means	read in full
[P05]	TDoc list; 7 Sep 2026	TDoc_List_Meeting_PCG#56 bis Ad Hoc.xlsx, the meeting's own document register https://www.3gpp.org/ftp/PCG/PCG_56_Bis/Docs/	The four documents with source, type, agenda item, recorded status and the reservation and upload timestamps	read in full (all four rows parsed into the packaged dataset)
[P06]	meeting report; 18 Sep 2026	PCG57_xx_Draft Summary minutes, decisions and actions from 3GPP PCG Meeting#56-Bis www.3gpp.org/ftp/PCG/PCG_56_Bis/Report/PCG57_xx%20-%20PCG56-Bis%20Draft%20Meeting%20Report%20v.1.3.docx	The meeting's own record: the three inputs as presented, the exchange between the Organisational Partners, the Chair's closing summary, decision D-PCG56-Bis/01 and action A-PCG56-Bis/01	read in full; the outcome of PCG#56bis is a recorded outcome from this edition
		TSG SA #113 Madrid, held 15-18 September 2026		
[M01]	report; 8 Sep 2026	SP-260907 SA WG4 chair status report on ULBC-MED www.3gpp.org/ftp/tsg_sa/TSR/SA/TSRGS_113_Madrid_2026-09/Docs/SP-260907.zip	Work plan to TSG SA#116, twenty declared proponents, agreements, open issues, Pdoc maturity	read in full
[M02]	report; 7 Sep 2026	SP-260909 Training Data Requirements for ULBC Candidate Codecs: SA WG4 Chair report to TSG SA#113 www.3gpp.org/ftp/tsg_sa/TSR/SA/TSRGS_113_Madrid_2026-09/Docs/SP-260909.zip	The answer to SA#112: no consensus need for technical training-dataset requirements, no work beyond the WID	read in full

Anchor	Kind and date	Identity as cited, and where	What it carries	How read
--------	---------------	------------------------------	-----------------	----------

Part 1 of the appendix. Every anchor in the text leads to its row here; the address is printed in full under the identity.

[M03]	discussion; 8 Sep 2026	SP-260912 Next steps on ULBC www.3gpp.org/ftp/tsg_sa/TSG_SA/TSGS_113_Madrid_2026-09/Docs/SP-260912.zip	Finalise the Pdocs governing qualification and selection only after PCG concludes; SA4 to re-evaluate a common dataset	read in full
[M04]	LS out; 8 Sep 2026	SP-260913 [DRAFT] Guidance on ULBC, the first of four drafts www.3gpp.org/ftp/tsg_sa/TSG_SA/TSGS_113_Madrid_2026-09/Docs/SP-260913.zip	Three bullets of guidance to SA WG4; technical requirements for datasets, maintainability of the specifications, and a work plan that takes account of PCG	read in full through the revision SP-260947 and the meeting report
[M05]	discussion; 8 Sep 2026	SP-260919 Ensuring a deployable and maintainable ULBC codec www.3gpp.org/ftp/tsg_sa/TSG_SA/TSGS_113_Madrid_2026-09/Docs/SP-260919.zip	An operator's case for transparent training data and verifiable requirement dimensions, with an LS to SA4	read in full
[M06]	CR pack; 7 Sep 2026	SP-260790 CR Pack on WI FS ULBC www.3gpp.org/ftp/tsg_sa/TSG_SA/TSGS_113_Madrid_2026-09/Docs/SP-260790.zip	The SA WG4 change requests on the ULBC study, approved in a block at agenda item 20.4	title only
[M07]	meeting report; 18 Sep 2026	SP-113_Draft_Report_v003, the draft report of TSG SA#113 www.3gpp.org/ftp/tsg_sa/TSG_SA/TSGS_113_Madrid_2026-09/Report/SP-113_Draft_Report_v003.zip	The plenary's own account of agenda item 3.4; the abstract, the discussion and the recorded status of each of the seven ULBC documents, with the companies named as they spoke	read in full; every ULBC entry parsed into the packaged datasets
[M08]	document register; 18 Sep 2026	TdocsByAgenda.htm, the meeting's own document register www.3gpp.org/ftp/tsg_sa/TSG_SA/TSGS_113_Madrid_2026-09/TdocsByAgenda.htm	All 299 documents tabled at TSG SA#113 with agenda item, type, purpose, source, comment and recorded result; the source of every disposition counted in this part	read in full; all 444 rows parsed into the packaged dataset
[M09]	chair notes; 18 Sep 2026	ChairNotes_Puneet_09-18-1059.doc, the TSG SA Chair's running notes www.3gpp.org/ftp/tsg_sa/TSG_SA/TSGS_113_Madrid_2026-09/INBOX/Chair_Notes/ChairNotes_Puneet_09-18-1059.doc	The same document-by-document dispositions as they stood when the meeting closed	read in full; used to confirm the register rather than as an independent source
[M10]	index; 17 Sep 2026	SA_Index_2026.xlsx, the TSG SA document index for 2026 www.3gpp.org/ftp/tsg_sa/TSG_SA/TSGS_113_Madrid_2026-09/SA_Index_2026.zip	The year's TSG SA documents with their meeting, type and status	read in part: the TSG SA#113 rows
		TSG SA reflector, 16-18 September 2026		
[M11]	reflector thread; 16-18 Sep 2026	DRAFT LS out on ULBC guidance: the TSG SA list thread of 16-18 September 2026 https://list.etsi.org/scripts/wa.exe?A0=3GPP_TSG_SA	27 messages from 9 organisations across the three days the guidance was drafted, each carrying the sender's organisation	read in full, message by message; all 27 parsed into the packaged dataset
[M12]	draft LS; 16 Sep 2026	SP-260947 [DRAFT] Guidance on ULBC, the second of the four drafts, as circulated on the list www.3gpp.org/ftp/tsg_sa/TSG_SA/TSGS_113_Madrid_2026-09/INBOX/DRAFTS	The three bullets of guidance, the action on SA4 and the dates of TSG SA#114 and #115; the two later drafts of the same text are read through the meeting report	read in full
		Other standards bodies: JVET, MPEG and JPEG records		
[J01]	call for proposals; Jul 2026	JVET-AQ2021 Joint Call for Proposals on Video Compression with Capability beyond VVC jvet-experts.org/doc_end_user/documents/43_Geneva/wg11/JVET-AQ2021-v1.zip	No training on test sequences; mandatory report of training material; training material available to JVET or re-training	read in full
[J02]	template; Jul 2026	JVET-AQ2022 Draft Template for Proposal Description Documents jvet-experts.org/doc_end_user/documents/43_Geneva/wg11/JVET-AQ2022-v1.zip	A section for learning-based elements and the training materials behind them	read in full
[J03]	AHG report; Jul 2026	JVET-AQ0011 JVET AHG report: Neural network-based video coding (AHG11) jvet-experts.org/doc_end_user/documents/43_Geneva/wg11/JVET-AQ0011-v2.zip	Mandates on training materials, content owners and bit-exact reproducibility	read in full
[J04]	teleconference report; Jun 2026	JVET-AQ0042 [AHG11] [AHG14] Teleconference on NNVC jvet-experts.org/doc_end_user/documents/43_Geneva/wg11/JVET-AQ0042-v2.zip	The NNVC training-set list and two licence and overlap clarifications	read in full
[J05]	EE summary; Jul 2026	JVET-AQ0023 EE1: Summary report of exploration experiment on neural network-based video coding jvet-experts.org/doc_end_user/documents/43_Geneva/wg11/JVET-AQ0023-v3.zip	All EE1 proposals trained on the CTC training sets; bit-exact reproducibility tests	read in full
[J06]	AHG report; Jul 2026	JVET-AQ0004 JVET AHG report: Test material and visual assessment (AHG4) jvet-experts.org/doc_end_user/documents/43_Geneva/wg11/JVET-AQ0004-v1.zip	Copyright-restricted test-sequence database, licence information per sequence	read in full
[J07]	draft text; Jul 2026	JVET-AQ0086 Initial proposed text for a potential WD of VSEI v5 jvet-experts.org/doc_end_user/documents/43_Geneva/wg11/JVET-AQ0086-v2.zip	The AI usage restrictions request SEI message and the text-description purposes for copyright and AI marking	read in part (SEI semantics)
[J08]	contribution; Jul 2026	JVET-AQ0166 AHG9: Support of C2PA Content Credentials Reference SEI Message jvet-experts.org/doc_end_user/documents/43_Geneva/wg11/JVET-AQ0166-v1.zip	A proposal to carry or reference C2PA manifests at the codec layer	read in full
[J09]	contribution; Jul 2026	JVET-AQ0200 [AHG11] Training-Time Optimization for LOP7 jvet-experts.org/doc_end_user/documents/43_Geneva/wg11/JVET-AQ0200-v2.zip	A worked instance of the mandatory training-information template	read in full
[J10]	TuC; Aug 2026	JVET-AQ2032 Technologies under consideration for future extensions of VSEI (version 13) jvet-experts.org/doc_end_user/documents/43_Geneva/wg11/JVET-AQ2032-v4.zip	Records that the AI usage restrictions SEI moved from TuC into the draft VSEI v4 text	read in part
[J11]	CFP; MPEG 139	w21823 Call for Proposals on Video Coding for Machines (as cited in the Apex export) www.mpeg.org/wp-content/uploads/mpeg_meetings/139_OnLine/w21823.zip	Training dataset publicly available with commercial-use licence; no training on validation data	as quoted in the Apex export
[J12]	guidelines; MPEG 138	w21568 AI-based point cloud compression guidelines (as cited in the Apex export) www.mpeg.org/wp-content/uploads/mpeg_meetings/138_OnLine/w21568.zip	Proprietary training allowed with extra results; public data availability and redistribution to be clarified; training details reported	as quoted in the Apex export
[J13]	CTTC; MPEG 153	w25972 Common test and training conditions for FCM (as cited in the Apex export) www.mpeg.org/wp-content/uploads/mpeg_meetings/153_OnLine/w25972.zip	A common set of datasets for training, with licensing arrangements for access	as quoted in the Apex export
[J14]	dataset paper; Jan 2023	JVET-AC0050 (VCM dataset usability and licences, as cited in the Apex export) jvet-experts.org/doc_end_user/documents/29_Teleconference/wg11/JVET-AC0050-v1.zip	Datasets reviewed for usability and licence; authors grant no licences	as quoted in the Apex export
[J15]	AHG report; Jan 2026	JVET-AO0011 NNVC AHG report (as cited in the Apex export) jvet-experts.org/doc_end_user/documents/41_Teleconference/wg11/JVET-AO0011-v4.zip	Training of NNVC algorithms cannot be reproduced exactly; same data and procedure typically within 0.1%	as quoted in the Apex export
[J16]	contribution; Jan 2026	JVET-AO0127 (NNVC training reproducibility, as cited in the Apex export) jvet-experts.org/doc_end_user/documents/41_Teleconference/wg11/JVET-AO0127-v1.zip	Non-deterministic trainer behaviour and environment mismatch prevent identical training	as quoted in the Apex export
[J17]	contribution; Jan 2025	JVET-AK0114 (AI usage restrictions SEI, as cited in the Apex export) jvet-experts.org/doc_end_user/documents/37_Geneva/wg11/JVET-AK0114-v2.zip	Restriction categories including no AI training; concerns on category clarity	as quoted in the Apex export
[J18]	contribution; Apr 2024	JVET-AH0112 AHG9: Signalling video usage for AI-driven processes in SEI message jvet-experts.org/doc_end_user/documents/34_Rennes/wg11/JVET-AH0112-v2.zip	Whether coded video may be used for AI training was unclear; signalling proposed	as quoted in the Apex export
[J19]	workshop; Aug 2026	w26730 AI traffic characterisation workshop material (as cited in the Apex export) www.mpeg.org/wp-content/uploads/mpeg_meetings/155_Geneva/w26730.zip	Sharing of data, logs and software needs more discussion	as quoted in the Apex export
		Base briefings and companion reports		
[B01]	base briefing; 10 Sep 2026	3GPP AI Training Dataset Governance, Version 3 (10 September 2026 edition) -	The briefing this report expands: PCG history, cross-TSG analysis, five-level model, legal survey	read in full
[B02]	base briefing; 11 Sep 2026	3GPP AI Training Dataset Governance, Version 3 (11 September 2026 edition) -	The same briefing with the SA4#137-e record, JVET and MPEG status and the IETF drafts	read in full
[B03]	record export; 11 Sep 2026	Apex Standards AI export, four questions on AI training-data governance -	MPEG, JVET and IETF passages on training-data rules, reproducibility and AI-use signalling	read in full
[B04]	companion brief; 9 Sep 2026	C20004166 3GPP Plenary Consensus Trajectory: Fukuoka to Singapore, and the Madrid Decision Set -	The plenary datasets from which the closed ULBC record is extracted	datasets re-read

Anchor	Kind and date	Identity as cited, and where	What it carries	How read
--------	---------------	------------------------------	-----------------	----------

Part 1 of the appendix. Every anchor in the text leads to its row here; the address is printed in full under the identity.

[B05]	companion brief; Sep 2026	C20004152 JVET Codec Consensus Formation: Santa Eulària to Geneva -	The JVET record around the beyond-VVC call and VSEI version 4	read in part
[B06]	record export; 11 Sep 2026	Apex Standards keyword export: 1,000 3GPP records naming AI, model and training data, decided approved, agreed or endorsed -	The approved 3GPP record on AI model training data, August 2021 to August 2026	read in full (all 1,000 rows parsed into the packaged dataset)

Table A2. Verbatim ledger: every quotation as recorded

Key	Anchor	Where in the source	Verbatim quotation as recorded
-----	--------	---------------------	--------------------------------

Part 2 of the appendix. Every quotation is an exact substring of its source after collapsing whitespace, verified at build; the reading applied to it is in the section that cites it, never inside the quotation.

			TSG SA draft reports, via C20004166, 3 verbatim quotations
V76	[S06]	TSG SA#112 draft report, disposition of SP-260641	"It was decided to change 'common training dataset' to 'technical training dataset'."
V77	[S06]	TSG SA#112 draft report, disposition of SP-260641	"It was decided to leave this to the SA WG4 Chair"
V78	[S02]	TSG SA#111 draft report, disposition of SP-260258	"whether this will be the full ULBC Codec is still to be determined"
			SA4 reflector, 27-28 Aug 2026, 11 verbatim quotations
V51	[W04]	message 9	"this decision SHALL NOT set a precedent for other codecs"
V52	[W04]	message 3 (Ericsson)	"A common training data does not guarantee a certain performance, e.g. across different languages."
V53	[W04]	message 3 (Ericsson)	"Access to common training data is not sufficient to replicate training."
V54	[W04]	message 3 (Ericsson)	"if re-training is foreseen, it raises questions on maintaining the interoperability of the standard"
V55	[W04]	message 3 (Ericsson)	"Why should non-public data be forbidden to be used, if it can improve the codec performance?"
V56	[W04]	message 11 (Nokia)	"ML codecs do not present a specific new technical dataset problem"
V57	[W04]	message 9	"There are means to consider private datasets in conjunction with a common public dataset"
V58	[W04]	message 7 (Dolby)	"the timeline aspect was definitely not the reason"
V59	[W04]	message 8 (Samsung)	"the survey responses were anonymous and only the Chair has visibility of all 24 replies"
V60	[W04]	message 6 (SA WG4 Chair)	"3GPP (including SA4 and SA) is contribution driven"
V61	[W04]	message 2 (Samsung)	"is not a technical reason but a general concern about the tight Rel-20 timeline"
			SP-260907, 5 verbatim quotations
V33	[M01]	slide 4	"Completing ULBC WI for Rel-20 is still possible, but the safety margin is thin."
V34	[M01]	slide 5	"Bitrates: 1.0 kbps and 1.4 kbps, and only these two"
V35	[M01]	slide 2	"A total of 20 companies declared their interest in submitting a ULBC candidate"
V36	[M01]	slide 7	"Overall completion rate for the WI was determined by the Rapporteur to be at 10%"
V37	[M01]	slide 6	"the metrics themselves still need to be defined"
			SP-260909, 8 verbatim quotations
V26	[M02]	slide 3	"Anonymized responses were requested and received in private"
V27	[M02]	slide 4	"legal and licensing aspects were out of scope"
V28	[M02]	slide 6, conclusion	"does not identify a consensus need for technical training dataset requirements"
V29	[M02]	slide 6	"the two main reasons include timing/schedule impacts and no technical need"
V30	[M02]	slide 6	"they split between informal guidance and a mandatory requirement, mainly for fair comparison and robustness"
V31	[M02]	slide 8, annex	"It does not form part of the feedback provided to TSG SA in this report."
V32	[M02]	slide 4	"The survey was designed to be comprehensive and deliberately went beyond the SA request"
V75	[M02]	slide 2, the SA#112 guidance as quoted	"SA WG4 should provide feedback to TSG#113 if any additional work is needed"
			PCG#56bis(26)01, 1 verbatim quotation
V01	[P01]	agenda item 3	"3GPP involvement and legal requirements for AI/ML training datasets"
			PCG#56bis(26)02, 8 verbatim quotations
V02	[P02]	section 2, problem statement	"embed the outcome of the training process into the codec definition"
V03	[P02]	section 2	"3GPP currently has no established governance framework specific to AI training data in its specifications"
V04	[P02]	section 3.2	"Defining a dataset does not mean that 3GPP must own or host the dataset"
V05	[P02]	section 3.3	"creating a long-term dependency for the maintenance of a 3GPP specification"
V06	[P02]	section 3.3	"selection of the dataset based on technical suitability remains within the responsibility of SA4"
V07	[P02]	section 5, proposal	"the first case of specifying an AI/ML model in 3GPP"
V08	[P02]	section 2	"a best-effort mechanism to reduce such risks"
V09	[P02]	section 4.1	"None of the other public voice dataset has an explicit contributor consent framework and documentation"
			PCG#56bis(26)03, 11 verbatim quotations
V10	[P03]	section 1	"This contribution does not address which of those options 3GPP should choose."
V11	[P03]	section 2	"legal due diligence cannot be treated as a secondary issue"
V12	[P03]	section 2	"Robust dataset due diligence should therefore be considered an industry requirement."
V13	[P03]	section 2	"similar to that of the establishment of appropriate Bills of Material for software"
V14	[P03]	section 3	"3GPP cannot formally provide guaranties"
V15	[P03]	section 3	"contributor self-declarations alone would likely be insufficient"
V16	[P03]	section 3	"a mechanism comparable to Coordinated Vulnerability Disclosure (CVD) could be considered"
V17	[P03]	section 3	"3GPP legal resources are limited and such activities would constitute a new area of responsibility"
V18	[P03]	section 4	"seeking advice from external counsel may be appropriate"
V19	[P03]	section 4	"a de facto determinant element on the final product"

Key	Anchor	Where in the source	Verbatim quotation as recorded
-----	--------	---------------------	--------------------------------

Part 2 of the appendix. Every quotation is an exact substring of its source after collapsing whitespace, verified at build; the reading applied to it is in the section that cites it, never inside the quotation.

V20	[P03]	section 2	"while reversing the burden of proof on manufacturers"
			PCG#56bis(26)04, 5 verbatim quotations
V21	[P04]	slide 2, first requirement	"Available to be used for standardisation, implementation, and deployment purposes with appropriate licensing terms"
V22	[P04]	slide 2, second requirement	"Dataset obtained with transparent and publicly documented privacy details"
V23	[P04]	slide 2, verification	"a group of legal representatives from all OPs"
V24	[P04]	slide 2	"All ULBC candidate codecs must be trained with a dataset that meets the legal requirements."
V25	[P04]	slide 2	"A common dataset could be a means to ensure"
			SP-260912, 7 verbatim quotations
V38	[M03]	section 1	"no consensus was reached on whether a common or a proprietary training dataset approach"
V39	[M03]	section 1	"a new PCG ad hoc group has been established"
V40	[M03]	section 3, proposal	"should proceed only after PCG has reached a conclusion on the legal issues"
V41	[M03]	section 3, proposal	"SA4 is encouraged to re-evaluate the need for a common dataset"
V42	[M03]	section 2	"how this dependency is handled will set the precedent for future 3GPP work involving AI"
V43	[W01]	section 2, quoting S4-261382	"are expected to be handled at the appropriate level outside SA4"
V44	[M03]	section 2	"the open question in PCG concerns only how the legal requirements should be satisfied"
			SP-260919, 6 verbatim quotations
V45	[M05]	section 2	"non-disclosed proprietary training data carries a structural maintenance risk for the specification itself"
V46	[M05]	section 2, quoting S4-261674	"near-unanimous support (20 of 24 companies)"
V47	[M05]	section 2, quoting S4-261674	"approximately 60% were in favor"
V48	[M05]	section 3, proposal	"The codec training dataset should be transparent such that 3GPP can maintain its specifications"
V49	[M05]	section 2	"the principal reservation among those not in favor was not a technical objection"
V50	[M05]	section 3, proposal	"the number of languages / language families, speaker demographics, acoustic environments"
			JVET-AQ0004, 1 verbatim quotation
V71	[J06]	test material	"the JVET database of test sequences is only available to accredited members of JVET"
			JVET-AQ0011, 2 verbatim quotations
V66	[J03]	AHG11 mandates	"Promote the call for training materials, distribute it, and actively communicate with content owners."
V67	[J03]	AHG11 mandates	"Investigate bit-exact reproducibility of NN-based methods on various platforms"
			JVET-AQ0023, 5 verbatim quotations
V70	[J05]	introduction	"For all proposals, training sets specified in [1] were used."
V79	[J05]	EE1 test conditions	"Candidates for adoption to NNVC are required to undergo training cross-check"
V80	[J05]	EE1 test conditions	"training scripts (including model quantizing algorithm) must be made available"
V81	[J05]	EE1-1.1 results	"After re-training of EE1-1.1.1 the BD-rate difference (Y component) of 0.1%"
V82	[J05]	introduction	"Comparison is done between tests which use the same sub-set of training data."
			JVET-AQ0042, 2 verbatim quotations
V68	[J04]	clarification on NNVC CTC, question 1	"This training set is for academic research purposes only."
V69	[J04]	clarification on NNVC CTC, question 2	"has a chance to overlap with test sequences"
			JVET-AQ0086, 2 verbatim quotations
V72	[J07]	AI usage restrictions request SEI, NOTE 1	"no conformance requirements for receiving systems are imposed by this Specification"
V73	[J07]	AI usage restrictions request SEI, NOTE 3	"ought not to be used for any AI training"
			JVET-AQ0200, 1 verbatim quotation
V83	[J09]	training configuration note	"the reported training time and memory usage are representative only on Ampere-class or newer GPUs"
			JVET-AQ2021, 5 verbatim quotations
V62	[J01]	coding conditions for submissions	"No part of test sequences shall be used as a training set"
V63	[J01]	coding conditions for submissions	"it is mandatory to report which material was used to train corresponding parts of the algorithm"
V64	[J01]	source code and IPR considerations	"re-training would need to be conducted using alternative material that is available to JVET"
V65	[J01]	source code and IPR considerations	"may include training scripts or equations used to derive parameters of a tool or algorithm"
V84	[J01]	source code and IPR considerations	"training materials used in the preparation of a proposal would be made available for further investigation and verification in JVET"
			JVET-AQ2022, 1 verbatim quotation
V74	[J02]	learning-based elements	"describe training materials and how they were used to obtain the learned parameters"
			PCG56bis-REPORT, 9 verbatim quotations
V93	[P06]	Chair's closing summary	"No clear conclusion on the legal analysis provided"
V94	[P06]	Chair's closing summary	"An ad-hoc group in PCG will be created for the clarification of the legal aspects of the training of AI data models"
V95	[P06]	ETSI in discussion	"there is no consensus on having a common dataset in SA4 and ETSI"
V96	[P06]	SA Chair's summary	"no consensus on the way forward for a common training dataset and left to SA4 to decide"
V97	[P06]	ETSI in discussion	"the legal risks associated with non common datasets are higher"
V98	[P06]	ETSI in discussion	"Legal and technical decisions should be separate"
V99	[P06]	CCSA in discussion	"Rel-20 should not be delayed"
V100	[P06]	TTA in discussion	"no consensus in SA4 technically speaking, we cannot leave the decision to SA4 since they cannot have consensus"
V101	[P06]	decision D-PCG56-Bis/01	"Creation of an ad-hoc group on legal aspects of AI for training datasets"
			SA-REFLECTOR-20260918, 28 verbatim quotations
V102	[M11]	message 4, Nokia	"the text goes beyond what can reasonably represent a consensus view"

Key	Anchor	Where in the source	Verbatim quotation as recorded
Part 2 of the appendix. Every quotation is an exact substring of its source after collapsing whitespace, verified at build; the reading applied to it is in the section that cites it, never inside the quotation.			
V103	[M11]	message 21, Nokia	"Before indicating that SA is hopeful of a consensus-based decision in SA4, we should first have consensus within SA on that position"
V104	[M11]	message 13, Dolby	"there is no consensus on what problem technical training dataset requirements will solve"
V105	[M11]	message 5, Ericsson	"SA4 do not need any guidance from SA at this point in time"
V106	[M11]	message 26, Samsung	"SA4 is not making sufficient progressing and it is the responsibility of SA to guide"
V107	[M11]	message 17, Dolby	"Guidance received by SA may also be understood as corrective directions to SA4"
V108	[M11]	message 10, Orange	"the LS simply means that SA overrules some technical decisions and even decisions that have been or to be made by SA4"
V109	[M11]	message 25, Ericsson	"Even if you soften the wording a lot it is not is not helping SA4 to progress and will add confusion what SA plenary consensus is"
V113	[M11]	message 1, Samsung	"I tried to reflect the discussion on Monday with some meaningful guidance that will allow SA4 to progress the technical work"
V114	[M11]	message 2, Dolby	"on 10.10.10.10 it is not accessible for remote participants"
V115	[M11]	message 3, BT	"Can I suggest you look here?"
V116	[M11]	message 6, Novamint	"the draft LS goes beyond what was discussed and that, at this stage, there is no need for an LS from the SA plenary"
V117	[M11]	message 7, Nokia	"There is no consensus in SA on providing guidance to spend precious agenda time allocated to this work item"
V118	[M11]	message 8, Ericsson	"It may in reality rather consume SA4 meeting time to try to interpret what SA plenary wanted to give guidance on"
V119	[M11]	message 9, Samsung	"Samsung views technical dataset requirements and maintainability as critical issues to the development of ULBC"
V120	[M11]	message 11, Samsung	"I modified the first two bullets based on the feedback I received"
V121	[M11]	message 12, Verizon	"I assume you agree we should develop maintainable specifications?"
V122	[M11]	message 14, Samsung	"by communicating the same points will at least keep everyone on the same page"
V123	[M11]	message 15, Nokia	"updating a codec does not necessarily preserve full backward compatibility"
V124	[M11]	message 16, Samsung	"we are only advising SA4 to"
V125	[M11]	message 18, Verizon	"SA4 should allocate time to discuss a set of"
V126	[M11]	message 19, Samsung	"I have applied the change requested by VZ to the first bullet"
V127	[M11]	message 20, Samsung	"It is significantly reworded to be of"
V128	[M11]	message 21, Nokia	"this represents a non-consensus-based indication to SA4, which is not usually appropriate to report in a liaison statement"
V129	[M11]	message 22, Samsung	"we will remain open to any text proposals that can find a way forward"
V130	[M11]	message 23, vivo	"I further update the LS (highlighted in yellow) considering the comments by Alessio"
V131	[M11]	message 24, Samsung	"Many thanks for your constructive suggestions, these are good and we can take them on-board"
V132	[M11]	message 27, Ericsson	"I do not see that the SA discussion have been conclusive and that the information provided help SA4"
			SP-113-REPORT, 12 verbatim quotations
V85	[M07]	Qualcomm on SP-260909	"the use of 'consensus' in this is not in line with the use of consensus in 3GPP"
V86	[M07]	SA WG4 Chair on SP-260909	"there is no basis for obtaining a consensus or way forward at this point"
V87	[M07]	plenary on SP-260912	"the work in SA WG4 should not be stopped while discussing this issue and should aim to meet the schedule for the work"
V88	[M07]	abstract of SP-260912	"no consensus was reached on whether a common or a proprietary training dataset approach is appropriate"
V89	[M07]	Nokia on SP-261002	"each time this LS is presented there are the same comments received"
V90	[M07]	Qualcomm on SP-261002	"the text of the third bullet is misleading as to the scope of the PCG discussions and the responsibilities of SA WG4 regarding AI work"
V91	[M07]	close of agenda item 3.4	"PCG has decided to further review these aspects in October 2026 and updated information is expected to be reported at the December TSG SA#114 meeting"
V92	[M07]	Ericsson on SP-260913	"an SA WG4 colleague had considered this draft and did not think this would be useful to SA WG4 at this time"
V133	[M07]	MediaTek on SP-260913	"the first 2 bullets are debatable as they are related to SA WG4 discussions"
V134	[M07]	Apple on SP-261002	"it would be good to make an official statement to SA WG4 that the PCG is working on the AI training sets legal implications"
V135	[M07]	Verizon on SP-260913	"whether there were any companies who would not want to have maintainable specifications"
V136	[M07]	Samsung at the close of item 3.4	"the divergence of views on technical requirements for training datasets and the codec maintenance aspects for AI-based codec solutions"
			SP-260947, 3 verbatim quotations
V110	[M12]	section 1, first bullet	"SA4 should strive to develop a set of technical requirements for datasets that any ULBC candidate codecs shall be trained/tested against"
V111	[M12]	section 1, second bullet	"SA4 should consider/evaluate the maintainability of the specs when using AI training datasets"
V112	[M12]	section 1, third bullet	"ULBC work plan (e.g., governing candidate qualification and selection phase) should take into account any potential guidance on the training dataset and legal requirements expected from PCG"

Methodology and interpretation criteria

What each number is read from, the one criterion applied to every source, what it does not capture, and the direction of the resulting error.

What is read from where. Survey counts are read from SP-260909 as the SA WG4 Chair reports them, and the annex figures from SP-260919; the Pdoc register, the work plan and the offline calls from SP-260907; the thread's shape from the supplied message; the closed plenaries' codec rows and dispositions from the verified datasets of C20004166; PCG positions from the three inputs themselves. Where two sources disagree, the one closer to the record prevails: a document's own text over another document's description of it, and a working group's report over a company's report of that group.

The criterion. One rule for every source: a statement is attributed to the document that makes it. A company's account of a meeting is reported as that company's account; a survey figure carried by a company paper is attributed to the paper, not to the survey it describes. The rule was fixed before any document was read for this edition.

What it buys. A reader can see which document carries each fact and weigh it; no fact is promoted by being repeated in several documents; and a reported outcome is never presented as a recorded one.

What it does not capture, and how much. The TSG SA#113 report is a draft and is read as one; the two late drafts of the guidance are read through it rather than directly; of 68 source records 29 were read in full and 8 are known by title, status or listing only; and none of the 24 survey answers can be attributed to a company.

The direction of the error. Each Madrid paper reports PCG#56bis in the course of an argument and so reports what its argument uses; a point that neither argument uses may be missing from both, so the reported outcome is a floor on what was discussed rather than a full account. Legal context resting on secondary summaries may lag the primary texts; every such paragraph carries the context mark.

Source hygiene: which documents are attributed to an organisation

Roles are not organisations. The SA WG4 Chair, the Audio SWG Chair, the ULBC-MED rapporteur and the Pdoc editors act for the working group; their documents and messages are cited by role and never counted as an organisation's activity. Where a role holder's company appears on a slide or in a signature, this brief does not carry it into any organisation measure.

An organisation is named only where the record names it as a source, co-source, editor or sender. In the reflector thread the organisation is taken from the thread itself - a signature, an address domain or a display name - or, for 2 messages, from the sender the thread names, supplied for this brief; none is inferred.

The legal paper carries no Source line; it is attributed to ETSI only because two Madrid papers do so, and its positions are read as the paper's rather than as any Organisational Partner's policy. No participant in the standards process is named anywhere in this brief; court cases are cited by their published names.

How a reported outcome is read: the criterion, what it supports, and what it does not

A recorded outcome is one a body's own record states: SA4's conclusion in SP-260909, a disposition in a plenary report. A reported outcome is one a participant describes: PCG#56bis as SP-260912 and SP-260919 describe it. The two are never merged, and a reported outcome never reaches the established tier on its own.

The survey is read as a tally, not as a map of positions. Its answers are anonymous and one per company whatever the company's role, so no count in this brief

says which company holds which view, and the coincidence of two identical splits is described as consistent with, not evidence of, the same companies on each side.

Update triggers: what would prompt a new edition

- The TSG SA#113 report, with its dispositions of SP-260909, SP-260912, SP-260913 and SP-260919 and any liaison to SA4.
- The text of SP-260913, which this edition did not read.
- Minutes of PCG#56bis, and the PCG#57 record, including the new ad hoc group's terms of reference.
- Pdoc versions from the Shanghai ad hoc and SA4#138, especially ULBC-3, ULBC-4, ULBC-5a, ULBC-5b and ULBC-6.
- The German Federal Court of Justice's decision of 17 December 2026, the English appeal in Getty v Stability, and the Court of Justice in C-250/25.
- Commission or ENISA guidance on AI components under the Cyber Resilience Act, and the Commission's list of generally agreed opt-out protocols.

What triggers a re-derivation

The approved TSG SA#113 report, the PCG#57 record, the new PCG ad hoc group's output, and new Pdoc versions. Every quantity in this brief is computed at build time, so a corrected or added input re-derives the prose rather than leaving a stale figure in it.

What this brief does not read

The SA4#137-e contributions themselves, the survey synthesis S4-261674, and the working-group discussion behind the Pdocs; each reaches this brief only through documents that report it, as Table 77 records.

What is outside the record entirely

The terms of any corpus licence as they will stand when SA4 selects, the outcome of pending court cases, national transposition of EU directives, and the commercial terms on which any candidate would be licensed. All bear on the outlook and none appears in the 3GPP record.

Citation and forecast practice

Every source identity in this brief - an anchor key, a meeting document number, a JVET document, a specification, a companion brief - is a live link to the ledger row or reference that carries it, and every reference prints its address in full. A citation form is generated from the ledger rather than typed, so two places citing the same source cite it identically. A source no entry covers is left plain and counted, so the gap between what the brief cites and what its back matter carries is reported at every build. No conditional reading carries a probability, a ranking or a committed date.

References

The list is sectioned by kind of source, and every entry prints in full the address at which it was read. Keys are stable: a bracketed key in the text leads to its entry here, and the records this brief read directly are in the source ledger, Table A1.

3GPP, PCG and ETSI

- [G01] 3GPP Portal, work item 1120096 ULBC-MED, Ultra Low Bitrate Speech Codec for IMS Voice Call over NB-IoT NTN via GEO Satellite (Rel-20, SA WG4). portal.3gpp.org/desktopmodules/WorkItem/WorkItemDetails.aspx?workItemId=1120096
- [G02] 3GPP file server, PCG#56 bis_Ad Hoc document folder: the four documents PCG56Bis_01 to PCG56Bis_04 and the meeting's TDoc list. https://www.3gpp.org/ftp/PCG/PCG_56_Bis/Docs/
- [G03] 3GPP file server, PCG meeting folders (no PCG_57 folder at the cutoff). <https://www.3gpp.org/ftp/PCG/>
- [G04] 3GPP file server, TSG SA#113 Madrid document folder (SP-260790, SP-260907, SP-260909, SP-260912, SP-260913, SP-260919). www.3gpp.org/ftp/tsg_sa/TSG_SA/TS GS_113_Madrid_2026-09/Docs/
- [G05] 3GPP file server, TSG SA#112 Singapore document folder. www.3gpp.org/ftp/tsg_sa/TSG_SA/TS GS_112_Singapore_2026-06/Docs/

- [G06] 3GPP file server, TSG SA#111 Fukuoka document folder. www.3gpp.org/ftp/tsg_sa/TSG_SA/TS GS_111_Fukuoka_2026-03/Docs/
- [G07] 3GPP file server, SA WG4 #137-e document folder (S4-261382, S4-261674, S4-261712 and the company contributions). www.3gpp.org/ftp/tsg_sa/WG4_CODEC/TS GS4_137-e/Docs/
- [G08] SA WG4 Chair, S4-261712 v2 draft of the report to TSG SA#113, as linked from the reflector. www.3gpp.org/ftp/tsg_sa/WG4_CODEC/TS GS4_137-e/Inbox/Drafts/Plenary/ULBC/S4-261712%20ULBC%20Training%20Data%20Requirements%20-%20SA4%20Chair%20report%20to%20SA113-v2.pdf
- [G09] 3GPP_TSG_SA_WG4 mailing-list archive (public; the thread of 27-28 August 2026 is cited from the message as supplied for this brief). https://list.etsi.org/scripts/wa.exe?A0=3GPP_TSG_SA_WG4
- [G10] ETSI, Rules of Procedure, Annex 6: ETSI Intellectual Property Rights Policy (definitions in clauses 15.6 and 15.7; copyright in clause 9). <https://www.etsi.org/images/files/IPR/etsi-ipr-policy.pdf>
- [G11] 3GPP Working Procedures (archived HTML), responsibilities of PCG, OPs and TSGs. www.3gpp.org/ftp/Information/Working_Procedures/archive/2019-08-23/3GPP_WP.htm
- [G12] ETSI portal, 3GPP Project Coordination Group. <https://portal.etsi.org/new3g/management/pcg.htm>
- [G13] 3GPP Highlights, Issue 10: RAN intelligence framework (data collection, training, inference). www.3gpp.org/images/newsletters/3GPP_Highlights_Issue_10_WEB.pdf
- [G14] 3GPP Highlights, Issue 11: CT3 AI/ML data and vertical federated learning. www.3gpp.org/images/newsletters/3GPP_Highlights_Issue_11_FLIPBOOK.pdf
- [G15] 3GPP TR 23.700-04, study on core network enhanced support for AI/ML, phase 2 (Rel-20). portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=4400
- [G16] 3GPP TR 28.882, AI/ML management, phase 3 (Rel-20). portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=4429
- [G17] 3GPP TS 33.127, lawful interception architecture and functions. portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3182
- [G18] 3GPP TS 33.128, lawful interception protocol and procedures. portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3183
- [G19] 3GPP TS 28.105, management and orchestration; AI/ML management - the specification most often changed in the approved training-data record. https://www.3gpp.org/ftp/Specs/archive/28_series/28.105/
- [G20] 3GPP TR 28.908, study on AI/ML management. https://www.3gpp.org/ftp/Specs/archive/28_series/28.908/
- [G21] 3GPP TS 23.288, architecture enhancements for 5G System to support network data analytics services. https://www.3gpp.org/ftp/Specs/archive/23_series/23.288/
- [G22] 3GPP TR 38.843, study on AI/ML for the NR air interface. https://www.3gpp.org/ftp/Specs/archive/38_series/38.843/
- [G23] 3GPP TR 26.927, study on AI/ML in 5G media services (SA WG4, FS_A14Media). https://www.3gpp.org/ftp/Specs/archive/26_series/26.927/
- [G24] 3GPP TR 26.940, the ULBC study report, whose agreed versions appear in the approved training-data record. https://www.3gpp.org/ftp/Specs/archive/26_series/26.940/
- [G25] 3GPP Working Procedures (2025), Article 46 on joint ownership and copyright of Technical Specifications and Reports, and Article 55 binding Individual Members to their Organisational Partner's IPR policy. www.3gpp.org/ftp/tsg_ran/WG4_Radio/TS GR4_115/Templates/3GPP%20Working%20Procedures_2025.pdf
- [G26] 3GPP Working Procedures (2025), Article 19 on PCG decision making and Article 25 on TSG and WG decision making: the group shall endeavour to reach consensus, and if consensus cannot be achieved the Chair may decide to take a vote, approved by 71% of the votes cast; Article 26 sets the meeting-vote mechanics and quorum. www.3gpp.org/ftp/Information/Working_Procedures/3GPP_WP.htm
- [G27] 3GPP Portal, work item 470030 EVS_codec, Codec for Enhanced Voice Services (Rel-12): start 25 March 2010, end 15 March 2016. portal.3gpp.org/desktopmodules/WorkItem/WorkItemDetails.aspx?workItemId=470030
- [G28] 3GPP Portal, work item 770024 IVAS_Codec, EVS Codec Extension for Immersive Voice and Audio Services (Rel-18): start 8 September 2017, end 6 June 2024, recorded release change Rel-17 to Rel-18 on 21 December 2021. portal.3gpp.org/desktopmodules/WorkItem/WorkItemDetails.aspx?workItemId=770024
- [G29] 3GPP TR 26.952, EVS codec performance characterization: qualification at SA4#72bis in March 2013, selection at SA4#80bis in August 2014, specifications approved at TSG SA#65 in September 2014, floating-point work completed at

TSG SA#66 in December 2014. www.etsi.org/deliver/etsi_tr/126900_126999/126952/13.00.00_60/tr_126952v130000p.pdf

[G30] 3GPP, 'IVAS Codec for the NG 3GPP Voice and Audio Services': work item launched September 2017, public collaboration established May 2022, codec selected at SA4#125 in August 2023, TSG SA approval September 2023. <https://www.3gpp.org/technologies/ivas-2023>

[G31] 3GPP TR 26.997, IVAS codec development: SA4#125 selected the IVAS codec candidate and the SA#101 plenary approved the selection and the set of IVAS specifications. www.etsi.org/deliver/etsi_tr/126900_126999/126997/18.00.00_60/tr_126997v180000p.pdf

[G32] 3GPP Portal, work item 470032 EVS_codec-SA4, the SA WG4 Stage 2/3 building block of the EVS codec work item: start 25 March 2010, end 12 December 2014, Rel-12. portal.3gpp.org/desktopmodules/WorkItem/WorkItemDetails.aspx?workItemId=470032

[G33] 3GPP TR 26.975 and the AMR project history, SP-000021: feasibility 1996-97, standardisation launched at ETSI SMG#23 in October 1997, candidate qualification and selection through 1998, SMG11 consensus on the selected solution in September 1998. www.3gpp.org/ftp/tsg_sa/tsg_sa/TSGS_07/Docs/PDF/SP-000021.pdf

[G34] 3GPP TR 26.976 presentation, SP-010692, the AMR-WB project history: work item proposed at TSG SA#2 in March 1999, qualification results at TSG SA#8 in June 2000, selection testing July to October 2000, selection approved at TSG SA#10 in December 2000, characterisation through 2001. www.3gpp.org/ftp/TSG_SA/TSG_SA/TSGS_14/Docs/PDF/SP-010692.pdf

[G35] 3GPP SP-040061, the SA WG4 PSS and MMS audio codec selection: AMR-WB+ and Enhanced aacPlus both met the requirements, no consistent overall ranking was possible, and both were recommended with usage guidance. www.3gpp.org/ftp/tsg_sa/tsg_sa/TSGS_23/Docs/PDF/SP-040061.pdf

[G36] 3GPP SP-040661 to TSG SA#25, Palm Springs, 13-16 September 2004: consensus was not reached at the SA WG4 meeting in Prague on two issues concerning H.264, the change requests were therefore not presented to the plenary by SA WG4, and supporting companies brought them to TSG SA directly. www.3gpp.org/ftp/tsg_sa/tsg_sa/TSGS_25/Docs/PDF/SP-040661.pdf

[G37] 3GPP SP-000394, the SA WG4 Release 1999 status: H.263 mandatory and MPEG-4 Visual optional in the circuit-switched multimedia path. www.3gpp.org/ftp/tsg_sa/tsg_sa/TSGS_09/docs/PDF/SP-000394.pdf

[G38] 3GPP TR 26.906, evaluation of HEVC for 3GPP services (Rel-12), and the change request integrating H.265 into MTSI agreed at SA WG4#83 and approved at TSG SA#68. portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=1483

[G39] 3GPP TR 26.955, video codec characteristics for 5G services (Rel-18): VVC, EVC and AV1 are characterised, which is a study and evaluation record rather than a normative service-profile adoption. www.etsi.org/deliver/etsi_tr/126900_126999/126955/18.00.00_60/tr_126955v180000p.pdf

[G40] 3GPP file server, TSG SA#113 Madrid meeting report folder: SP-113_Draft_Report_v003, the draft report of the meeting. www.3gpp.org/ftp/tsg_sa/TSG_SA/TSGS_113_Madrid_2026-09/Report/SP-113_Draft_Report_v003.zip

[G41] 3GPP file server, TSG SA#113 Madrid document register by agenda item, with the recorded result of every document tabled. www.3gpp.org/ftp/tsg_sa/TSG_SA/TSGS_113_Madrid_2026-09/TdocsByAgenda.htm

[G42] 3GPP file server, TSG SA#113 Madrid Chair's notes as the meeting closed. www.3gpp.org/ftp/tsg_sa/TSG_SA/TSGS_113_Madrid_2026-09/INBOX/Chair_Notes/ChairNotes_Puneet_09-18-1059.doc

[G43] 3GPP file server, TSG SA document index for 2026. www.3gpp.org/ftp/tsg_sa/TSG_SA/TSGS_113_Madrid_2026-09/SA_Index_2026.zip

[G44] 3GPP file server, PCG#56 bis_Ad Hoc draft meeting report v1.3: summary minutes, decisions and actions. www.3gpp.org/ftp/PCG/PCG_56_Bis/Report/PCG57_xx%20-%20PCG56-Bis%20Draft%20Meeting%20Report%20v1.3.docx

[G45] 3GPP_TSG_SA mailing-list archive (public); the thread of 16-18 September 2026 on the draft guidance is read from it). https://list.etsi.org/scripts/wa.exe?A0=3GPP_TSG_SA

[G46] 3GPP file server, TSG SA#113 Madrid drafts folder, in which each revision of the guidance was circulated. www.3gpp.org/ftp/tsg_sa/TSG_SA/TSGS_113_Madrid_2026-09/INBOX/DRAFTS

[G47] 3GPP, Release 20 page: stage 1 frozen June 2025, stage 2 final freeze September 2026, stage 3 targeted March 2027, ASN.1/OpenAPI freeze June 2027. www.3gpp.org/specifications-technologies/releases/release-20

[G48] 3GPP, 'Timeline for Release 21', approved at the June 2026 plenaries: stage 1 freeze March 2027, stage 2 freeze June 2028, stage 3 freeze December 2028, ASN.1/OpenAPI freeze March

2029. <https://www.3gpp.org/news-events/3gpp-news/rel21-timeline>

[G49] 3GPP, Release 21 page. www.3gpp.org/specifications-technologies/releases/release-21

JVET, MPEG, JPEG, ITU and content provenance

[K01] ITU-T, Joint Video Experts Team page (JVET#43 Geneva, JVET#44 Hangzhou, joint call beyond VVC). www.itu.int/en/ITU-T/studygroups/2025-2028/21/video/Pages/jvet.aspx

[K02] ITU-T SG21 TD PLEN-0348, approval of the Joint Call for Proposals on Video Compression with Capability beyond VVC, 15 July 2026. <https://www.itu.int/md/T25-SG21-260706-TD-PLN-0348/en>

[K03] MPEG 155 press release (Geneva, 13-17 July 2026): joint call, provenance call, next meetings. www.mpeg.org/wp-content/uploads/2026/08/agw3254SN26474Press-Release-MPEG155.pdf

[K04] MPEG-AI (ISO/IEC 23888) project pages: Part 2 Video Coding for Machines, Part 4 Feature Coding for Machines, Part 6 media authenticity and provenance. <https://www.mpeg.org/standards/MPEG-AI/>

[K05] JPEG, Final Call for Proposals for JPEG AI (WG1 N100095), training-data requirements. [ds.jpeg.org/documents/jpegai/wg1n100095-094-REQ-Final_Call_for_Proposals_for_JPEG_AI.pdf](https://www.mpeg.org/documents/jpegai/wg1n100095-094-REQ-Final_Call_for_Proposals_for_JPEG_AI.pdf)

[K06] JPEG AI dataset page (training, validation and withheld test sets). <https://jpeg.org/jpegai/dataset.html>

[K07] ITU-T T.840.1 summary: trainable model parameters are an integral part of the Recommendation | ISO/IEC 6048-1. www.itu.int/dms_pubrec/itu-t/rec/t/T-REC-T.840-1-202503-!#!SUM-HTML-E.htm

[K08] ISO/IEC 6048-1:2025, JPEG AI core coding system. <https://www.iso.org/standard/88911.html>

[K09] JPEG, JPEG AI Common Training and Test Conditions (WG1 N100106). ds.jpeg.org/documents/jpegai/wg1n100106-094-ICQ-JPEG_AI_Common_Training_and_Test_Conditions.pdf

[K10] ISO/IEC 21617-1:2026, JPEG Trust edition 2 (provenance, attribution and rights declarations). <https://www.iso.org/standard/91405.html>

[K11] ITU-T H.274 (VSEI), version approved January 2026. www.itu.int/ITU-T/recommendations/rec.aspx?rec=H.274

[K12] MPEG 152 press release: VSEI edition with the AI usage restrictions request SEI message to FDIS. <https://www.mpeg.org/152nd-meeting-of-mpeg/>

[K13] MPEG emerging technologies: call for training materials for neural-network video coding (w23582). <https://www.mpeg.org/emerging-technologies/>

[K14] ISO/IEC 15938-17:2024, neural network coding (NNC). <https://www.iso.org/standard/85545.html>

[K15] IEC, ISO and ITU (AMAS), Technical Report on AI and Multimedia Authenticity Standards, edition 2, July 2026. www.worldstandardscooperation.org/wp-content/uploads/2026/07/IEC-ISO-ITU-Technical_Report_on_AI_and_Multimedia_Authenticity_Standards_Edition2.pdf

[K16] IEC, ISO and ITU (AMAS), policy paper on the emerging legal landscape for generative AI across regions, July 2026. www.worldstandardscooperation.org/wp-content/uploads/2026/07/IEC-ISO-ITU-Policy_paper_on_Emerging_legal_landscape_for_generative_AI_across_different_regions.pdf

[K17] C2PA Technical Specification 2.4 (1 April 2026). [spec.c2pa.org/specifications/specifications/2.4/specs/C2PA_Specification.html](https://specifications.c2pa.org/specifications/2.4/specs/C2PA_Specification.html)

[K18] Creator Assertions Working Group, training and data mining assertion, version 1.1. <https://cawg.io/training-and-data-mining/1.1/>

[K19] ITU-T F.748.77 (12/2025), assessment criteria for foundation models (evaluation-data traceability). <https://www.itu.int/rec/T-REC-F.748.77>

[K25] MPEG 155 (Geneva, 13-17 July 2026): SC 29/WG 2 call for proposals on media authenticity and provenance indication; proposals due 13 January 2027, evaluated at MPEG 157. <https://www.mpeg.org/155th-meeting-of-mpeg/>

[K24] ITU-T Y.3181 (09/2022), architectural framework for a machine learning sandbox in future networks including IMT-2020. <https://www.itu.int/rec/T-REC-Y.3181>

[K23] ITU-T Y.3176 (10/2020), machine learning marketplace integration in future networks including IMT-2020. <https://www.itu.int/rec/T-REC-Y.3176>

[K22] ITU-T Y.3172 (06/2019), architectural framework for machine learning in future networks including IMT-2020: a pipeline whose nodes include a data collector and a preprocessor. <https://www.itu.int/rec/T-REC-Y.3172>

[K21] ITU, 'Beyond VVC: call for proposals on future video coding', 28 July 2026: ITU, ISO and IEC; registration 1 August to 1 September 2026; evaluation at the 45th JVET meeting, January 2027; target standard 2029. www.itu.int/hub/2026/07/beyond-vvc-call-for-proposals-on-future-video-coding/

[K20] ITU-T Study Group 21, JVET page: the Joint Video Experts Team of ITU-T Q6/21 (VCEG) and ISO/IEC JTC 1/SC 29; JVET#43 Geneva, 7-15 July 2026; JVET#44 Hangzhou, 17-23 October 2026. www.itu.int/en/ITU-T/studygroups/2025-2028/21/video/Pages/jvet.aspx

[K27] JVET Terms of Reference: all group decisions are made by the consensus of the JVET as determined by the Chairs, and all output of the JVET undergoes the approval process of each parent body. www.itu.int/en/ITU-T/studygroups/2017-2020/16/Documents/video/JVET-ToRv2-TD-PLN-0476.pdf

[K28] ISO, stages and resources for standards development: a draft International Standard is approved where two-thirds of the P-members of the committee are in favour and not more than one quarter of the total votes cast are negative (ISO/IEC Directives Part 1, 2.6). www.iso.org/stages-and-resources-for-standards-development.html

[K29] ISO/IEC Guide 59:2019, clause 4.4.4: consensus requires that substantial issues be addressed, and sustained opposition cannot be interpreted as a right to veto. <https://www.iso.org/standard/71917.html>

[K30] IEEE SA Standards Board Operations Manual, clause 5.4.3: an IEEE SA standards ballot requires a minimum of 75% of those voting to approve or disapprove to approve the draft, subject to ballot-return requirements. <https://standards.ieee.org/about/policies/opman/sect5/>

[K31] ETSI Technical Working Procedures, clause 1.7: a Technical Group shall endeavour to reach Consensus, and if Consensus cannot be achieved the Chair may decide to take a vote, an ordinary technical proposal passing at 71% of the votes cast under the weighted-vote rules. <https://portal.etsi.org/Resources/ETSI-Directives>

[K32] ITU-T Recommendation A.8 (01/2024), the alternative approval process: Consent, a four-week Last Call, comment resolution and Additional Review, and approval at a study group meeting where no more than one Member State present opposes. <https://www.itu.int/rec/T-REC-A.8-202401-1>

[K33] ITU-T, Joint Video Team page: JVT was formed in 2001 by ITU-T VCEG and ISO/IEC MPEG, and the first version of H.264 was completed in May 2003. www.itu.int/en/ITU-T/studygroups/com16/video/Pages/jvt.aspx

[K34] ITU-T, Joint Collaborative Team on Video Coding: JCT-VC was created in 2010 by ITU-T VCEG and MPEG, first met in Dresden 15-23 April 2010, and H.265 was approved on 13 April 2013. www.itu.int/en/ITU-T/studygroups/2013-2016/16/pages/video/jctvc.aspx

[K35] ITU-T, Joint Video Exploration Team: the informal exploration team ran eight meetings from 19-21 October 2015 in Geneva to 18-24 October 2017 in Macao, before the formal team was created. www.itu.int/en/ITU-T/studygroups/2017-2020/16/Pages/video/prejvet.aspx

[K36] ITU-T, Joint Video Experts Team: the formal JVET was created on 27 October 2017, the final call for proposals issued in October 2017, proposals evaluated in April 2018, the first version of VVC completed in July 2020 and H.266 approved on 29 August 2020. www.itu.int/en/ITU-T/studygroups/2017-2020/16/Pages/video/jvet.aspx

[K37] ISO/IEC 14496-10, Advanced Video Coding: registered as a new project on 13 March 2002, the enquiry draft approved for the final stage on 9 April 2003, published 27 November 2003 - a ballot sequence separate from the joint team's technical convergence. <https://www.iso.org/standard/37729.html>

[K38] ISO/IEC 23090-3, Versatile Video Coding, first edition published February 2021, the ISO/IEC twin of ITU-T H.266. <https://www.iso.org/standard/73022.html>

IETF and W3C

[I01] IETF AI Preferences (AIPREF) working group: charter, scope and out-of-scope items. <https://datatracker.ietf.org/wg/aipref/about/>

[I02] draft-ietf-aipref-vocab-07, A Vocabulary for Expressing AI Usage Preferences (August 2026). <https://www.ietf.org/archive/id/draft-ietf-aipref-vocab-07.txt>

[I03] draft-ietf-aipref-attach-05, Associating AI Usage Preferences with Content in HTTP (August 2026). <https://datatracker.ietf.org/doc/draft-ietf-aipref-attach/>

[I04] AIPREF interim meeting minutes, 24-26 August 2026. ietf-wg-aipref.github.io/wg-materials/interim-26-08/minutes.html

[I05] RFC 9309, Robots Exclusion Protocol (September 2022). <https://www.rfc-editor.org/info/rfc9309>

[I06] RFC 9943, An Architecture for Trustworthy and Transparent Digital Supply Chains (June 2026). <https://www.rfc-editor.org/info/rfc9943>

[I07] draft-ietf-scitt-scrapi, SCITT Reference APIs (approved, in the RFC Editor queue). <https://datatracker.ietf.org/doc/draft-ietf-scitt-scrapi/history/>

[I08] RFC 9969, Report from the IAB Workshop on AI-CONTROL (May 2026). <https://www.rfc-editor.org/info/rfc9969>

[I09] IETF Web Bot Auth working group. <https://datatracker.ietf.org/wg/webbotauth/about/>

[I10] IETF vCon working group documents (core container, privacy primer). <https://datatracker.ietf.org/wg/vcon/documents/>

[I11] draft-howe-vcon-lawful-basis, lawful basis and purpose grants for vCon. <https://datatracker.ietf.org/doc/draft-howe-vcon-lawful-basis/>

[I12] draft-howe-vcon-lifecycle, consent revocation and deletion tracking for vCon. <https://datatracker.ietf.org/doc/draft-howe-vcon-lifecycle/>

[I13] draft-altanai-ai-pref-realtime-protocol-bindings, AIPREF preferences in SIP and SDP. datatracker.ietf.org/doc/draft-altan-ai-ai-pref-realtime-protocol-bindings/

[I14] IETF mlcodec working group documents. <https://datatracker.ietf.org/wg/mlcodec/documents/>

[I15] draft-ietf-mlcodec-opus-dred-07, Deep Audio Redundancy for Opus (4 August 2026, in working-group last call). <https://datatracker.ietf.org/doc/draft-ietf-mlcodec-opus-dred/>

[I16] draft-ietf-mlcodec-opus-dred-07, full text (normative decoder; weights distributed outside the document; Appendix C). [www.ietf.org/archive/id/draft-ietf-mlcodec-opus-dred-07.txt](https://datatracker.ietf.org/doc/draft-ietf-mlcodec-opus-dred-07.txt)

[I17] draft-ietf-mlcodec-opus-speech-coding-enhancement-04, Integration of Speech Codec Enhancement Algorithms into Opus (July 2026). [www.ietf.org/archive/id/draft-ietf-mlcodec-opus-speech-coding-enhancement-04.txt](https://datatracker.ietf.org/doc/draft-ietf-mlcodec-opus-speech-coding-enhancement-04.txt)

[I18] draft-reilly-ai-pref-compliance, verifiable compliance records for AI usage preferences. <https://datatracker.ietf.org/doc/draft-reilly-ai-pref-compliance/>

[I19] draft-wallace-ai-pref-grant-binding, verifiable-credential grants lifting AI usage preferences. datatracker.ietf.org/doc/draft-wallace-ai-pref-grant-binding/

[I20] draft-hood-ai-pref-earmark, signed usage preferences embedded in content. <https://datatracker.ietf.org/doc/draft-hood-ai-pref-earmark/>

[I21] draft-howe-vcon-provenance, provenance of generated vCon content. <https://datatracker.ietf.org/doc/draft-howe-vcon-provenance/>

[I22] draft-dawkins-scitt-ai-article50, SCITT receipts for EU AI Act Article 50. <https://datatracker.ietf.org/doc/draft-dawkins-scitt-ai-article50/>

[I23] draft-sharif-ai-model-lifecycle-attestation, attestation from training data to inference output. datatracker.ietf.org/doc/draft-sharif-ai-model-lifecycle-attestation/

[I24] draft-messous-eat-ai, EAT profile with training-data provenance claims (expired). <https://datatracker.ietf.org/doc/draft-messous-eat-ai/>

[I25] draft-reilly-sentinel-protocol, integrity records for datasets, training and inference. <https://datatracker.ietf.org/doc/draft-reilly-sentinel-protocol/>

[I26] draft-williams-ai-content-tagging, human- or AI-generated tagging for clean datasets (expired). datatracker.ietf.org/doc/draft-williams-ai-content-tagging/

[I27] draft-akhavain-moussa-ai-network, dataset descriptors and discovery for AI training. datatracker.ietf.org/doc/draft-akhavain-moussa-ai-network/

[I28] W3C TDM Reservation Protocol, Final Community Group Report (May 2024). www.w3.org/community/reports/tdmrep/CG-FINAL-tdmrep-20240510/

[I29] IETF 127, San Francisco, 14-20 November 2026: the meeting at which the neural speech-codec drafts next progress. <https://www.ietf.org/meeting/127/>

[I30] RFC 7282, On Consensus and Humming in the IETF: rough consensus is not unanimity, an objection that has been considered and answered does not prevent progression, and a headcount does not determine consensus. <https://www.rfc-editor.org/rfc/rfc7282.txt>

[I31] W3C Process Document: consensus and sustained objection, the recording of Formal Objections, and a vote only after consensus methods have failed. <https://www.w3.org/policies/process/>

Security, data-quality and dataset-documentation standards

[N01] ETSI EN 304 223 V2.1.1 (2025-12), baseline cyber security requirements for AI models and systems (Principle 8: document data, models and prompts). www.etsi.org/deliver/etsi_en/304200_304299/304223/02.01.01_60/en_304223v020101p.pdf

[N02] ETSI TS 104 223 V1.1.1 (2025-04), securing AI: baseline requirements. www.etsi.org/deliver/etsi_ts/104200_104299/104223/01.01.01_60/ts_104223v010101p.pdf

[N03] ETSI GR SAI 002 V1.1.1 (2021-08), data supply chain security. www.etsi.org/deliver/etsi_gr/SAI/001_099/002/01.01.01_60/gr_SAI002v010101p.pdf

[N04] ETSI press release, data-quality measurement framework (TR 104 180), September 2026. www.etsi.org/newsroom/press-releases/etsi-launches-new-framework-to-measure-data-quality-a-key-enabler-for-trustworthy-ai/

[N05] European Commission, AI Act standardisation (CEN-CENELEC JTC 21 request, including dataset governance and quality). digital-strategy.ec.europa.eu/en/policies/ai-act-standardisation

[N06] ISO/IEC 5259-1:2024, data quality for analytics and machine learning, part 1. <https://www.iso.org/standard/81088.html>

[N07] ISO/IEC 8183:2023, data life cycle framework. <https://www.iso.org/standard/83002.html>

[N08] ISO/IEC 42001:2023, AI management system. <https://www.iso.org/standard/42001>

[N09] SPDX 3.0.1 Dataset profile (collection process, availability, intended use, known bias, sensitive data). <https://spdx.github.io/spdx-spec/v3.0.1/model/Dataset/Dataset/>

[N10] SPDX 3.0.1 relationship types (trainedOn, testedOn). spdx.github.io/spdx-spec/v3.0.1/model/Core/Vocabularies/RelationshipType/

[N11] ECMA-424, CycloneDX bill of materials, 2nd edition December 2025 (machine-learning BOM). ecma-international.org/publications-and-standards/standards/ecma-424/

[N12] MLCommons Croissant 1.1 (29 January 2026), with PROV-O provenance and usage policies. docs.mlcommons.org/croissant/docs/croissant-spec-1.1.html

[N13] MLCommons Croissant RAI vocabulary 1.0. docs.mlcommons.org/croissant/docs/croissant-rai-spec.html

[N14] Datasheets for Datasets, Communications of the ACM 64(12), 2021. <https://dl.acm.org/doi/10.1145/3458723>

[N15] Data and Trust Alliance, Data Provenance Standards v1.0.0 (July 2024). <https://www.dtaalliance.org/work/data-provenance-standards>

[N16] IEEE 2801-2022, quality management of datasets for medical AI. <https://standards.ieee.org/ieee/2801/7459/>

European Union law and policy

[L01] Regulation (EU) 2024/1689, the AI Act. eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32024R1689

[L02] Regulation (EU) 2026/1744, the AI Omnibus amending the AI Act timelines (in force 27 July 2026). <https://eur-lex.europa.eu/eli/reg/2026/1744/oj>

[L03] European Commission, template for the public summary of training content of general-purpose AI models (24 July 2025). digital-strategy.ec.europa.eu/en/faqs/template-general-purpose-ai-model-providers-summarise-their-training-content

[L04] European Commission, General-Purpose AI Code of Practice (10 July 2025), copyright chapter. digital-strategy.ec.europa.eu/en/policies/contents-code-gpai

[L05] European Commission, consultation on protocols reserving rights from text and data mining (December 2025 - January 2026). digital-strategy.ec.europa.eu/en/consultations/commission-launches-consultation-protocols-reserving-rights-text-and-data-mining-under-ai-act-and

[L06] European Commission, guidelines on the Article 50 transparency obligations, C(2026) 5054 (20 July 2026). ai-act-service-desk.ec.europa.eu/sites/default/files/2026-07/guidelines_on_the_implementation_of_the_transparency_obligations_for_certain_ai_systems_under_article_50_of_the_ai_act_bzptwqh0k0i91dlddap41psfy_131215.pdf

[L07] European Commission, code of practice on transparency of AI-generated content (final, 10 June 2026). digital-strategy.ec.europa.eu/en/policies/code-practice-ai-generated-content

[L08] EDPB Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models (17 December 2024). www.edpb.europa.eu/system/files/2024-12/edpb_opinion_202428_ai-models_en.pdf

[L09] EDPB Guidelines 02/2021 on virtual voice assistants. www.edpb.europa.eu/our-work-tools/our-documents/guidelines/guidelines-022021-virtual-voice-assistants_en

[L10] Regulation (EU) 2016/679, the General Data Protection Regulation (Art. 4(14), Art. 9). <https://eur-lex.europa.eu/eli/reg/2016/679/oj>

[L11] Directive (EU) 2019/790 on copyright in the Digital Single Market (Art. 4 text and data mining). <https://eur-lex.europa.eu/eli/dir/2019/790/oj>

[L12] Directive 96/9/EC on the legal protection of databases. eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:3196L0009

[L13] Directive (EU) 2024/2853 on liability for defective products (transposition by 9 December 2026). <https://eur-lex.europa.eu/eli/dir/2024/2853/oj>

[L14] Regulation (EU) 2023/988, General Product Safety Regulation (applies from 13 December 2024). <https://eur-lex.europa.eu/eli/reg/2023/988/oj>

[L15] Withdrawal of the AI Liability Directive proposal, OJ C/2025/5423 (6 October 2025). eur-lex.europa.eu/legal-content/EN/TXT/?uri=OJ:C_202505423

[L16] European Commission, Cyber Resilience Act reporting obligations (from 11 September 2026). digital-strategy.ec.europa.eu/en/policies/cra-reporting

[L17] ENISA, CRA single reporting platform launched. www.enisa.europa.eu/news/the-cra-single-reporting-platform-is-launched

[L18] European Parliament resolution of 10 March 2026 on copyright and generative AI. www.europarl.europa.eu/doceo/document/TA-10-2026-0066_EN.html

[L19] EUIPO study on generative AI from a copyright perspective (May 2025). www.euipo.europa.eu/en/news/euipo-releases-study-on-generative-artificial-intelligence-and-copyright

[L20] European Commission, call for evidence on the review of EU copyright rules (May-June 2026). digital-strategy.ec.europa.eu/en/news/commission-seeks-views-review-eu-copyright-rules

[L21] European Commission IP Helpdesk, first CJEU hearing on generative AI and copyright, C-250/25 Like Company v Google Ireland. intellectual-property-helpdesk.ec.europa.eu/news-events/news/cjeu-grand-chamber-rules-music-sampling-and-pastiche-first-cjeu-hearing-generative-ai-and-copyright-2026-04-24_en

[L22] European Commission, Data Act explained (applies from 12 September 2025). digital-strategy.ec.europa.eu/en/factpages/data-act-explained

[L23] Corrigendum to Directive (EU) 2024/2853, published 7 May 2026: the regime applies to products placed on the market or put into service after 8 December 2026. [eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32024L2853R\(01\)](https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32024L2853R(01))

Court decisions and proceedings, 2024-2026

[C01] Bundesgerichtshof, I ZR 281/25 (Kneschke v LAION): hearing 3 September 2026, pronouncement 17 December 2026. www.bundesgerichtshof.de/SharedDocs/Termine/DE/Termine/IZR281-25.html

[C02] OLG Hamburg, 5 U 104/24 (10 December 2025), Kneschke v LAION, commentary. www.harte-bavendamm.de/en/ip-blog/hamburg-higher-court-specifies-text-and-data-mining-restrictions-for-ai-training

[C03] LG München I, 42 O 14139/24 (11 November 2025), GEMA v OpenAI, commentary. www.osborneclerke.com/insights/gema-a-vs-openai-ai-memorisation-reproduction-relevant-copyright-law-and-tdm-exception-does

[C04] LG München I, 42 O 763/25 (31 July 2026), GEMA v Suno. <https://www.gema.de/de/w/suno-entscheidung-2026>

[C05] LG Berlin II, 2 O 202/24 (20 August 2025), AI imitation of a voice actor's voice. <https://www.urheberrecht.org/news/7818/>

[C06] Getty Images v Stability AI [2025] EWHC 2863 (Ch), judgment of 4 November 2025. <https://www.judiciary.uk/judgments/getty-images-v-stability-ai/>

[C07] Getty Images v Stability AI: permission to appeal on the infringing-copy point (December 2025). ipkitten.blogspot.com/2026/01/permission-to-appeal-granted-in-getty.html

[C08] Bartz v Anthropic (N.D. Cal.): final approval of the class settlement, July 2026. authorsguild.org/news/court-grants-final-approval-anthropic-copyright-settlement/

[C09] Kadrey v Meta (N.D. Cal.): fair-use ruling of 25 June 2025 and the remaining distribution claim. chatgptiseatingtheworld.com/2026/03/27/kadrey-v-meta-gets-scheduling-order-summary-judgment-motion-hearing-for-distribution-claim-wont-be-heard-until-feb-25-2027/

[C10] Thomson Reuters v Ross Intelligence: Third Circuit oral argument, 11 June 2026. www.bakerbots.com/thought-leaders/hip/publications/2026/july/third-circuit-hears-oral-argument

[C11] In re OpenAI copyright litigation (S.D.N.Y.): Department of Justice statement on fair use, 1 September 2026. ipwatchdog.com/2026/09/03/doj-sides-with-openai-warns-obstacles-to-ai-development-threaten-national-security/

[C12] Lehrman and Sage v Lovo (S.D.N.Y., 10 July 2025): voice-cloning claims under New York publicity law. www.loeb.com/en/insights/publications/2025/07/lehrman-v-lovo-inc

[C13] Beijing Internet Court (23 April 2024): licence to recordings did not authorise AI voice synthesis. www.kingandwood.com/cn/en/insights/latest-thinking/china-s-first-case-regarding-ai-generated-voice-infringement.html

[C14] UMG and Sony v Suno (D. Mass.): anti-circumvention claim over audio taken from YouTube, August 2026. www.musicbusinessworldwide.com/umg-and-sony-accuse-suno-of-pirating-youtube-via-stream-ripping-to-train-its-ai/

United States law and policy

[U01] U.S. Copyright Office, Copyright and Artificial Intelligence, Part 3: Generative AI Training (pre-publication, May 2025). www.copyright.gov/ai/Copyright-and-Artificial-Intelligence-Part-3-Generative-AI-Training-Report-Pre-Publication-Version.pdf

[U02] 17 U.S.C. 107, fair use. <https://www.govinfo.gov/link/uscode/17/107>

[U03] California AB 2013, Generative AI Training Data Transparency Act (in force 1 January 2026), commentary. www.goodwinlaw.com/en/insights/publications/2026/01/alerts-other-industries-californias-ab-2013-takes-effect

[U04] Colorado SB 26-189 replacing the Colorado AI Act (effective 1 January 2027), commentary. www.seyfarth.com/news-insights/colorado-enacts-artificial-intelligence-replacement-law.html

[U05] Illinois Biometric Information Privacy Act, 740 ILCS 14/10 (voiceprint). www.ilga.gov/documents/legislation/ilcs/docume nts/074000140k10.htm

[U06] Executive Order 14365, Ensuring a National Policy Framework for Artificial Intelligence (11 December 2025). www.federalregister.gov/documents/2025/12/16/2025-23092/ensuring-a-national-policy-framework-for-artificial-intelligence

[U07] NO FAKES Act of 2026, S. 4591 (reported by the Senate Judiciary Committee, 18 June 2026). <https://www.congress.gov/bill/119th-congress/senate-bill/4591>

[U08] Federal Trade Commission, policy statement on biometric information (18 May 2023). www.ftc.gov/legal-library/browse/policy-statement-federal-trade-commission-biometric-information-section-5-federal-trade-commission

[U09] Tennessee ELVIS Act (in force 1 July 2024), commentary. www.dwt.com/blogs/artificial-intelligence-law-advisor/2024/04/tennessee-elvis-act-ai-voice-replica

United Kingdom, Asia-Pacific and international instruments

[A01] UK Government, report and impact assessment on copyright and artificial intelligence (18 March 2026). www.gov.uk/government/publications/report-and-impact-assessment-on-copyright-and-artificial-intelligence

[A02] UK Copyright, Designs and Patents Act 1988, section 29A (text and data analysis for non-commercial research). <https://www.legislation.gov.uk/ukpga/1988/48/section/29A>

[A03] UK Data (Use and Access) Act 2025, section 136. <https://www.legislation.gov.uk/ukpga/2025/18/section/136>

[A04] ICO, biometric data guidance: key data protection concepts. ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/awful-basis/biometric-data-guidance-biometric-recognition/key-data-protection-concepts/

[A05] Japan Agency for Cultural Affairs, General Understanding on AI and Copyright (Article 30-4). www.bunka.go.jp/english/policy/copyright/pdf/94055801_01.pdf

[A06] Japan, Act on the Promotion of Research, Development and Utilisation of AI (Act No. 53 of 2025). <https://www.japanesela wtranslation.go.jp/en/laws/view/5066/en>

[A07] Japan Cabinet Office, AI Basic Plan (14 July 2026). www8.cao.go.jp/cstp/ai/ai_plan/aipplan_eng_20260714.pdf

[A08] Japan APPI amendment (Act No. 56 of 2026), commentary. www.aoshearman.com/en/insights/ao-shearman-on-data/amen dments-to-the-act-on-the-protection-of-personal-information-p romulgated

[A09] Republic of Korea, Framework Act on AI (in force 22 January 2026), commentary. www.cooley.com/news/insight/2026/2026-01-27-south-koreas-ai-basic-act-overview-and-key-ta keaways

[A10] Korea MCST, English edition of the guide on fair use in generative AI training (18 May 2026; Korean original February 2026). <https://www.mcst.go.kr/english/policy/pressView.jsp?p Seq=635>

[A11] China CAC, Interim Measures for the Management of Generative AI Services (2023). https://www.cac.gov.cn/2023-07/13/c_1690898326795531.htm

[A12] China GB/T 45652-2025, security specification for generative-AI pre-training and fine-tuning data. std.samr.gov.cn/gb/search/gbDetailed?id=33D40F1160BE5D92E06397BE0A0A5B93

[A13] China GB/T 45574-2025, sensitive personal information (voiceprint listed), commentary. digitalpolicyalert.org/event/31186-national-standard-gbt-45574-2025-information-security-te chnologysecurity-requirements-for-processing-sensitive-person al-information-enters-into-force

[A14] China, labelling measures and GB 45438-2025 for AI-generated content (from 1 September 2025), commentary. www.ioeb.com/en/insights/publications/2025/03/chinas-ai-labelling-measures-and-mandatory-national-standards-take-effect-september-1

[A15] Singapore Copyright Act 2021 (sections 243-244, computational data analysis). <https://sso.agc.gov.sg/Act/CA2021>

[A16] India, AI Governance Guidelines (MeitY, November 2025). static.pib.gov.in/WriteReadData/specificdocs/documents/2025/nov/doc2025115685601.pdf

[A17] India, DPIIT working paper on generative AI and copyright (December 2025), commentary. www.ikigailaw.com/article/655/exploring-the-dpiits-working-paper-on-generative-ai-and-copy right

[A18] India, Digital Personal Data Protection Act 2023 and Rules 2025, summary. <https://www.dlapiipdataprotection.com/?t=law&c=IN>

[A19] WIPO Performances and Phonograms Treaty. <https://www.wipo.int/en/web/treaties/ip/wppt/index>

[A20] Berne Convention for the Protection of Literary and Artistic Works. <https://www.wipo.int/en/web/treaties/ip/berne/index>

Speech corpora and licences

[V01] Mozilla Data Collective, Common Voice Scripted Speech 26.0 dataset page (licence and re-sharing terms). mozilla.org/datacollective.com/datasets/cm4im4lxy00tun07cjkupeg

[V02] Mozilla Data Collective, terms for consumers (updated 6 May 2026). <https://mozilla.org/datacollective.com/terms/consumers>

[V03] Common Voice dataset changelog (release sizes and languages). github.com/common-voice/cv-dataset/blob/main/datasets/scripted-speech/CHANGELOG.md

[V04] OpenSLR SLR12, LibriSpeech ASR corpus (CC BY 4.0). <https://www.openslr.org/12/>

[V05] University of Edinburgh DataShare, CSTR VCTK corpus v0.92. <https://datashare.ed.ac.uk/handle/10283/3443>

[V06] CSTR VCTK dataset card (licence statement). huggingface.co/datasets/CSTR-Edinburgh/vctk

[V07] Microsoft DNS Challenge repository README (component datasets and licences). github.com/microsoft/DNS-Challenge/blob/dfa9c9c80e5d67f5b1477ec53f5a2b5a9d77938/README.md

[V08] Creative Commons CC0 1.0 Universal. <https://creativecommons.org/publicdomain/zero/1.0/>

[V09] Creative Commons Attribution 4.0 International. <https://creativecommons.org/licenses/by/4.0/>

[V10] Common Voice privacy notice. <https://commonvoice.mozilla.org/privacy/en.html>

Apex Standards reports and base briefings

[R01] Apex Standards C20004166, 3GPP Plenary Consensus Trajectory: Fukuoka to Singapore, and the Madrid Decision Set (9 September 2026).

[R02] Apex Standards C20004152, JVET Codec Consensus Formation: Santa Eulària to Geneva (3 September 2026).

[R03] Apex Standards C20004064, JVET-VVC pre-meeting review and outlook, Geneva (1 July 2026).

[R04] Apex Standards C20004128, 3GPP SA2 6G Architecture Consensus Evolution: Malta to Prague in 2026 (1 September 2026).

Glossary, in two layers

The first layer is written for a reader outside 3GPP and says what each term means for a decision, a product, a market or a timeline. The second is written for a reader inside the process and gives the term as the record uses it. Where a term has a wider industry meaning, the definition here is the narrower one this brief relies on.

For the reader outside the process

3GPP. The partnership of seven standards organisations that writes mobile standards; its specifications are what phones and networks are built to. **Section 4.1** explains who inside it decides what.

Organisational Partner (OP). One of the seven bodies - ARIB, ATIS, CCSA, ETSI, TSDSI, TTA, TTC - that hold 3GPP's legal personality and publish its specifications; the ones who would carry any liability (**section 4.3**).

PCG. The Project Coordination Group, where the partners settle matters of the partnership itself; it received the dataset question in June 2026 (Part 3).

ULBC. Ultra-low-bitrate speech codec: a neural codec for voice calls over satellite narrowband IoT links, at 1.0 and 1.4 kbps, planned for approval in June 2027 (Part 4).

Normative weights. Trained numbers that are part of the standard itself, so every compliant device uses the same ones; what makes training data matter here (**section 6.1**).

Training data. The recordings a model learns from. For ULBC they would shape the standardised weights, and carry rights and personal data with them (Part 7).

Common dataset. One corpus every candidate trains on for the results that are compared; TTA's proposal (**section 4.2**).

Selection. The competition in which candidate codecs are tested by independent laboratories and one is chosen; its rules are still being written (**Table 27**).

Escrow. A deposit of data and training recipe with a trusted party, released only for maintenance; one of the options in **Table 52**.

Declaration. A structured statement of what a dataset is, where it came from and what may be done with it; **Table 49** sets out one for ULBC.

Fidelity. Whether a decoder reproduces what was said, and who said it, without altering either; the point at which the AI Act touches a codec (**section 8.4**).

Text and data mining exception. In European Union law, permission to mine lawfully accessible works unless the rightholder has reserved the right in machine-readable form;

where a corpus carries an express licence the exception matters less, because the licence grants the use (**section 8.2**).

Training-content summary. The public list of what a general-purpose AI model was trained on, in the European Commission's template. TTA proposes it as the shape of a 3GPP declaration, which is the use of a format rather than compliance with a duty (**section 8.4**).

Conditional reading. A statement about what would follow if something outside the record held; never a forecast (**Table 67**).

For the reader inside the process

TSG SA. Technical Specification Group Service and System Aspects, the plenary that owns SA4's work. Its #113 meeting sits on 15-18 September 2026 inside the TSG#113 round at Madrid, 14-18 September, where RAN#113 and CT#113 meet in the same week [B04].

SA WG4 (SA4). The codec working group. Its conclusion on training data is SP-260909; its Chair's status report is SP-260907.

ULBC-MED. Work item 1120096, Rel-20, approved at TSG SA#112 on SP-260696.

Pdoc. A permanent document of the working group recording agreed rules, plans and requirements for the selection; eleven exist for ULBC (**Table 27**).

Qualification and selection. The two test phases of the competition; qualification may be dropped in favour of a direct selection test.

PCG#56 bis, Ad Hoc. The PCG ad hoc electronic meeting of 3 September 2026 on AI/ML training datasets, under the name its own register gives it; written PCG#56bis after first use, as the Madrid papers write it.

Reflector. A 3GPP mailing list; the SA4 thread of 27-28 August 2026 is cited as [W04].

Separation. Keeping qualification, selection and characterisation material out of training; supported by 20 of 24 survey respondents as SP-260919 reports.

NNVC. JVET's neural-network-based video coding exploration, which compares tools trained on common sets.

Joint call beyond VVC. The JVET and MPEG call for proposals of July 2026, which requires disclosure of training material.

AIPREF. The IETF working group on AI usage preferences; it standardises signals, not enforcement.

IPR, essential. In the ETSI IPR Policy, an IPR without which a compliant product cannot be made; training data rights are not essential in that sense (**Table 47**).

Re-training cross-check. JVET's requirement that a party other than the proponent retrains a proposed tool and reproduces the reported gain within a stated tolerance (**section 6.5**).

Biometric data. Under the GDPR, data resulting from technical processing that allows or confirms unique identification. A voice recording is personal data; it becomes biometric through that processing rather than by being a voice (**section 8.3**).

Scope of this report

This brief reads the 3GPP, PCG, SA WG4, JVT and other standards records it names, as they stood on 18 September 2026, together with the public legal and standards sources listed in the references. It reproduces no document: short passages are quoted with their source attached, and every count is derived from data the package carries. The legal and regulatory material is described, not applied, and nothing in this brief is legal advice. TSG SA#113 had not been held at the cutoff; every statement about it concerns documents tabled for it. The records are listed in [Table A1](#).

Legal disclaimer

The information contained in this report is provided strictly on an "AS IS" basis, without any representations or warranties, express or implied, including but not limited to warranties of accuracy, completeness, fitness for a particular purpose, or non-infringement. References to third-party individuals, organisations, products, standards, or services are included solely for informational purposes and do not imply any affiliation with, or endorsement by, Apex Standards. Apex Standards assumes no responsibility or liability for errors, omissions, or any consequences arising from the use of this report, and undertakes no obligation to update, amend, or revise this report following its publication.

About APEX STANDARDS

Apex Standards provides high-efficiency analytical tools and expert advisory services for global standardisation and patenting. We offer advanced software platforms that analyse complex data across 3GPP, IEEE, IETF, ITU, Open RAN, video codec, JEDEC, AI infrastructure, CXL, HBM and related technology ecosystems. We equip governments, research institutes, and corporate innovators with the rigorous standards and intellectual property intelligence needed to drive R&D excellence.

Website: www.apexstandards.com

Email: support@apexstandards.com