

Epoch rollover and the ETSI corpus

A first corpus-wide look at bounded time representations and their horizons,
in support of ETSI PWI WKI_ID 79366

Christian (fukami) Horchert

Review and input: Trey Darley

66302 of 67122 deliverables screened · data snapshot 2026-07-23 15:37

Licensed [CC BY 4.0](#)

Published at <https://lab.fukami.eu/epoch-scan/>

First published 2026-09-20 · revised 2026-09-21; revisions recorded in Appendix B

Scope of this document. All 66302 text-bearing deliverables of the 67122-document ETSI corpus have been scanned, which is the Published subset: no Historical or On Approval deliverable carried extractable text into the scan, so the screen says nothing about those 697. Every place quoted below was checked against the specification text before being listed. Everything named here is a **candidate for expert review** by the owning technical body, not a finding against a standard, and no statement is made about documents that were not flagged or not read. Figures are generated from the analysis pipeline and are regenerable; scoring, validation, and pattern-registry detail are kept out of this document and live in the project records and in a methods companion published alongside it.

AI-assistance note. AI tooling was used throughout the production of this document, for support in corpus processing and screening, and for drafting support around a deterministic analysis pipeline. Every statement about a specification quoted here passed a human read at the flagged clause; screen outputs were candidates, never conclusions. Responsibility for the content rests with the author, and the method can be re-run from the pipeline code and checked against the per-candidate evidence records; reproduction is on-premises (see “What this look does and does not establish”).

Summary

[ETSI PWI WKI_ID 79366](#) proposes a fitness check of the ETSI corpus for epoch-rollover exposure: places where a deliverable fixes a bounded time representation, such as a 32-bit second count or a 16-bit day count, and its text says nothing about the date at which that representation runs out. This document supports the proposal. It reports a first corpus-wide look, produced by a weighted text search over the 66302 text-bearing deliverables of the 67122-document corpus, followed by careful reading at the most exposed places, and it offers first numbers on what a proper check would involve. Everything named below is a candidate for review by the owning group, not a finding against a standard.

ETSI’s own work programme confirms the class of problem. [EN 300 468](#) (DVB Service Information) stores dates in a 16-bit day count whose last representable day is 2038-04-22; the published text is silent about this. Two sibling specifications re-specify the same field, and a ring of further documents imports it by reference, including the DVB software-update scheduling fields, so the remediation channel itself rides the wrapping field. The DVB maintainers have opened a work item ([WKI 78272](#), on approval) to clarify exactly this behaviour, independently of this scan.

Further places flagged for expert review (the full list is under “Places worth expert review”):

- [TS 103 221-2](#) (lawful interception X2/X3, mandated for 5G LI): the defining clause’s own NOTE acknowledges a horizon and defers it, but the horizon it names is the POSIX type’s rather than the one the clause mandates, which is two unsigned 32-bit fields running to 2106. One for ETSI TC LI maintenance.
- [TS 129 061](#) (charging): a 32-bit 1900-era time carried in RADIUS accounting records with no era field; the 2036 era boundary then depends on context the record does not carry. One for 3GPP, via liaison.
- Dated normative references to **NTPv3 (RFC 1305)** and the **RFC 868 Time Protocol**, both with documented 2036 horizons, in several DVB and cable families including both DOCSIS generations. Clause-level dependency review.
- A **long-tail horizon catalogue** ($\approx$ 2081 DOCSIS 3.1, 2132 DAB, 2217 DARC and DRM): worth recording, no urgent action.

First numbers for planning (more under “What a fitness check would involve”): the screen flags 1129 specification families whose text combines time encodings with a rollover horizon left silent or ambiguous; 302 of them define a time encoding themselves, and 139 do so strongly enough to form the head a check would read first. So far, 142 of the 223 families queued for reading have been read and recorded. On these numbers, a full check of the families that define a time encoding themselves is on the order of a few hundred family reads.

This document is meant to support the decision on the work item. It shows the kind of constructs involved, gives first numbers on the volume a check would face, and lists places that may merit expert review; the 3GPP-maintained items would go upstream through the liaison channel, and the approach can be reused by other technical bodies on their own corpora.

Background

ITU-T Study Group 17 (SG17) agreed the text of [XSTP.epoch](#)¹, the *Technical Paper on Global Coordination Requirements for 2038-class rollover events (including but not limited to 2036, 2038, 2106)*, previously circulated as X.EPOCH, at its June 2026 plenary (Geneva, 1-10 June 2026; SG17 management approval by correspondence, 3 July 2026, TD182/WP3). As an ITU-T Technical Paper it is an informative deliverable rather than a Recommendation. SG17 circulated the agreed text by liaison statement (SG17-LS161), for action to the [FIRST Time Security SIG](#) and for information to ETSI TC CYBER among other bodies, inviting suggestions for next steps in awareness-raising, projects, or specific standards, with a reply deadline of 7 September 2026. The step this document supports is [ETSI PWI WKI ID 79366](#), a proposed work item for an epoch-rollover fitness check of the ETSI corpus, an ETSI-side counterpart to the taxonomy’s sector view. XSTP.epoch groups rollover events into failure-mode classes; in this document those classes appear only as tags on the search patterns (the fm column in Appendix A), and the observations are presented by construct.

The ETSI corpus encodes, references, and constrains time in thousands of deliverables. Fixed-width time representations have known exhaustion dates, among them the 2038 signed-32-bit boundary, GPS 1024-week rollovers, MJD range limits, two-digit years, and certificate validity semantics. This scan asks one question across the whole corpus: *which ETSI-published deliverables normatively fix a bounded, epoch-relative time representation yet leave its rollover horizon silent or ambiguous?* Where a specification is actively maintained or a work item is open, a clarification can travel with ordinary maintenance. The intervention path is not the same for every document: it differs for historical, defunct-deployment, and upstream-maintained cases, and the candidates below state which one applies.

A specification is not itself “exposed”. Editing its text is neither necessary nor sufficient for safety in the field. It is not necessary, because a decoder can often rebuild the era from context, so a silent specification need not lead to a real failure. It is not sufficient either, because a specification that leaves a field’s width unspecified


¹<https://propertools.be/commons/xstr-epoch/>

can still be implemented unsafely, so a real failure need not show up as silence. [TS 103 770](#) is an example of that second case, since its Unix time has no stated width. What the scan looks for is therefore narrower: it finds the **normative origin** of a bounded time constraint, and whether that origin acknowledges its own horizon. That is the place where a coordinated fix would begin, and it is not a claim about how much risk is deployed.

The screen documented here is not the fitness check. It is a prioritisation aid that generates candidates; the check itself is clause-level review with standard-specific expertise, performed by or with the owning technical bodies. What the work item would involve is bounded: a few hundred families to read rather than tens of thousands of documents, outputs in the form of a reviewed candidate list with clause references, a horizon catalogue and a reusable search vocabulary, and, for anything an expert confirms, possible ways forward from a corrigendum to a liaison item (see “What a fitness check would involve”).

The problem space: bounded time in standards text

Standards fix time in text: a field width here, an epoch there, a validity window, a week counter. Each such choice quietly creates a horizon, the date at which the representation runs out or wraps. Two boundary behaviours must not be conflated. In **range exhaustion**, a field fails at a fixed calendar date with no way back: a 16-bit day count decodes day 65 536 as 1858. In **modulo-era ambiguity**, the value wraps by design and stays decodable as long as era context is available, as with NTP’s 1900-era seconds or the 10-bit GPS week; such a construct fails only when its era context is wrong or lost, which can be years after the calendar anchor. The cases in this section were each checked against the specification text at the flagged clause, and they are presented as review candidates for the owning groups. The bounded constructs read in this section, placed on one axis:



One construct, three widths, three horizons

The clearest teaching example in the corpus is the **broadcast Modified Julian Date**, the count of days since 1858-11-17 carried in the date field of broadcast signalling. The construct is broadcast-wide, and the **horizon is the width the defining standard picked**.

2038 is EN 300 468’s width choice, not a property of broadcast MJD. DVB-SI stores MJD in 16 bits, so its last representable day is MJD 65 535 = 2038-04-22, and the next day, stored as its low 16 bits, decodes as 1858, a silent ~180-year backward jump. DAB stores the same MJD in 17 bits with a stated decimal range to 99 999, while DARC and DRM use the full 17-bit binary range, which runs to 2217. Every sibling that widened the field bought a century-plus of headroom. This also becomes a parser-differential surface wherever shared

tooling, gateways, or validators process more than one MJD-bearing format: “MJD” then names incompatible fields of different widths, and a recogniser that assumes the wrong one mis-decodes the other system’s dates.

The family around the definer. TS 103 197 (SimulCrypt) and TS 102 474 (IPDC service purchase and protection, whose rights-validity windows are a defunct DVB-H deployment) both **re-specify** the 16-bit construct natively. An inheritor ring (TS 102 323, TS 102 006, TR/TS 101 211, TS 102 727, and others) imports it by reference, with TS 102 006 importing it into the **software-update scheduling** fields. The exposure is family-shaped: one definer, and a blast radius that includes the remediation channel itself. That shape also bounds any remediation: the construct has one normative origin, two native re-specifications that would need edits of their own, and a ring of importers enumerable from the references.

A fix at the definer is already under way. The DVB maintainers have opened a work item whose scope names exactly this behaviour; it is on EN Approval at the time of writing:

EN 300 468 (DVB Service Information), on EN Approval since 2026-06-16.

Work-item scope, verbatim: *“This WI will add a clarification on MJD date and time behaviour in 2038 wrap-around, as well as the support for T/UWA-005 signalling.”*

This is the MJD/2038 case that the screen ranked at the head of its list. The work item corroborates that the maintainers consider a clarification needed, independent of this scan; the published edition is silent on the wrap, and the work item’s scope makes the clarification explicit.

Disclosure can be present, and wrong

A silence check alone would miss a second class: documents that **do** mention a horizon and name the wrong one. TS 103 221-2, the lawful-interception internal interface mandated for 5G LI by TS 133 128, defines its PDU timestamp as two 32-bit **unsigned** integers (seconds and nanoseconds), which puts the wire horizon at 2106. The same clause’s NOTE acknowledges a horizon and defers its resolution to a future version, but scopes it to “the timespec format as defined in POSIX.1-2017” and gives the year as 2038. That sentence is about the POSIX type rather than about the wire, and the conflation is the defect: 2038 is a signed `time_t` property the mandated encoding does not have, the cited POSIX edition fixes no width for that field at all, and the horizon the interface does have is never stated anywhere in the clause. An implementer is warned about a year the interface will not reach and is left unwarned about the one it will. The disclosure is present and it points at the wrong year. This case surfaced through a by-hand pass over documents the automated screen had set aside precisely because they mention rollover, and it is arguably the most actionable class of all, since maintenance intent already exists. The locus is an in-house conversation at ETSI TC LI, not a liaison item.

Records outlive the wire

The 3GPP-transposed part of the corpus (published by ETSI under the 120-149 ranges, authored and maintained at 3GPP) was read at its most exposed head to completion, 122 specification families. The reassuring part first: the feared signed-32-bit `time_t` class is essentially absent there; 3GPP time is NTP-1900 counts, BCD calendar strings, GPS weeks, and text dates, with deliberately widened 40- and 64-bit epochs for newer 5G features.

One case is worth an expert look, and it is a **record**, not a wire message. TS 129 061 carries a 32-bit 1900-era time (3GPP-User-Location-Info-Time) in RADIUS accounting records with no era field, so a record that outlives its context loses the unambiguous reading of a value whose era rolls over in 2036. GTPv2-C, next to it, carries the same 32-bit value on the wire, yet its

charging record re-encodes the value as a calendar date in the CDR (TS 132 298), so the wire wraps in 2036 while the retained record does not. Whether a stored value crosses an era boundary therefore depends on the **storage** specification, not only the protocol one. A sibling observation: TS 129 274 names the 2036 era

boundary beside every affected field, and ships no era field; awareness in text never became disambiguation on the wire. All of these are candidate review locations for 3GPP, carried through the liaison channel; they are not ETSI-owned work.

Dependencies with documented horizons

Some exposure enters by reference rather than by definition. Two dependency candidates stand out, both of the same shape: a normative reference to a time protocol whose own documentation states a rollover horizon and defers the handling.

- **RFC 1305 (NTPv3), dated normative pins.** Four ETSI families of the DVB, cable-LI, and IPCablecom class reference NTPv3 normatively, three of them pinned to the 1992 text. RFC 1305 documents its 2036 rollover and delegates the handling, noting that “some external means will be necessary”, and it was later obsoleted by RFC 5905.
- **RFC 868 (Time Protocol) in DOCSIS, both generations.** EN 302 878-4 (DOCSIS 3.0) and TS 103 311-3 (DOCSIS 3.1, 2017) require the cable modem to obtain time of day via the Time Protocol, which is 32-bit unsigned seconds since 1900 and whose serviceability the RFC itself bounds at 2036. The dependency survived the generation change, so it is not a legacy leftover. These are CableLabs-authored texts with an ETSI badge, so the contribution tone applies as it does for the 3GPP set.

The mirror image of this class is just as useful: the trust-services and e-signature families (EN 319 4xx, CAdES/XAdES) define **no** local time encoding at all; their time semantics live in RFC 3161/5652/5280 profiles. For them, naming the definer is the finding: the review question moves to the referenced format, and nothing about the ETSI text is thereby cleared. The NFV Security series is a similar concentration of time **dependencies** (PKI, TSA, NTP, PTP, GNSS across 6 external time-format families): these documents consume time services rather than define encodings, so they read as a dependency-review cluster, not as candidates in their own right.

Deployment evidence: horizons arm, then fire

The Telstra outage (2026-07-08). The 10-bit GPS week construct that this scan flags (in TETRA, SES A-GNSS, and the transposed GPS-week cluster) produced a live national incident during this report’s drafting. The external investigation by Technology Audit Partners, dated 24 August 2026 and [published by Telstra](#)² with its findings accepted, gives the account: a planned change to a Melbourne timing chassis brought the unit back online, its GPS card reset, and because the card was missing a firmware update that the supplier had published bulletins about, the receiver resolved the broadcast week into the previous epoch and computed a date 1 024 weeks in the past. The resulting outage reached voice, messaging, data, and emergency calling. The arithmetic can be checked independently of that report: 2026-07-08 minus 1 024 weeks is 2006-11-22, both are Wednesdays, and both truncate to 10-bit week 378 (full weeks 2426 versus 1402), which agrees with the November 2006 date the investigation records. A separate ACMA investigation into Telstra’s regulatory compliance remains open.

Two features of that account bear on the reading offered here. The defect was created by a rollover in April 2019 and lay dormant until a restart in July 2026, so the boundary date is when a defect of this class **arms**, while it **fires** only at the next loss of the volatile state masking it, years after the calendar anchor. The wrong date then did not stay on one device: the investigation describes a timing loop in which sources that appear independent are in fact validating one another, so the available sources converged and clients accepted the majority reading. Agreement among sources that share a root cause is not evidence of correctness. Finite-time horizons in long-lifetime products are also being treated by courts as remediable product defects. XSTP.epoch’s own Table 4 example is the RATP/Alstom case, a range-exhaustion sibling of this modulo class, in which the Tribunal Administratif de Paris ordered remediation

²<https://www.telstra.com.au/exchange/what-we-ve-learned-from-the-external-investigation-of-our-july-o>

of a post-2037 date-representation defect in rolling stock (2025-11-13). For long-lifetime products, the document stage is the earliest and least contentious place to address a horizon.

Places worth expert review

The flagged places, in one list. Each row was checked against the specification text at the flagged clause, by careful readers who are not specialists in these standards; each is a pointer for the group named in the last column, and none is a final verdict.

Specification	Observation	Who would act
EN 300 468 (DVB-SI)	16-bit day count (MJD); last representable day 2038-04-22; text silent, with a misleading adjacent range note (Annex C formula note, 1900–2100)	DVB; live WI 78272 already targets the clarification
TS 103 197 (SimulCrypt)	re-specifies the same 16-bit field in its own head-end protocol; silent	DVB maintenance; same fix class as EN 300 468
TS 102 474 (IPDC SPP)	same 16-bit date in rights-validity windows; silent; deployment defunct	specification-level record only; no field action expected
Inheritor ring (TS 101 211, TS 102 006, TS 102 323, ...)	import the EN 300 468 construct by reference, including the software-update scheduling fields	follows the definer's fix; notify the ring
TS 103 221-2 (LI X2/X3)	unsigned 2×32-bit timestamp, wire horizon 2106; the clause's NOTE defers the POSIX timespec 2038 limit instead, and the mandated encoding's own horizon is never stated	ETSI TC LI, in-house maintenance
TS 129 061 (3GPP charging)	32-bit 1900-era time in RADIUS accounting records; 2036 era boundary; no era field on the record	3GPP, via liaison
TS 129 274 (GTPv2-C)	names the 2036 era boundary beside every affected field; ships no era field	3GPP, via liaison
NTPv3 / RFC 868 references	dated normative pins to time protocols with documented 2036 horizons (four DVB/cable-LI families; both DOCSIS generations)	dependency review by the owning TBs
TS 103 770 (DVB-I)	service-discovery arithmetic on Unix time with no stated width; behaviour is implementation-defined	DVB; pin the width
Long-tail horizons	TS 103 311-3 ≈ 2081; EN 300 401 (DAB) 2132-08-31; ES 201 980	horizon-catalogue entries; no urgent action

Specification	Observation	Who would act
	(DRM) 2217-09-27; EN 300 751 (DARC) 2217-09-27; each verified at the field, far from the near term	
TETRA / SES GNSS assistance	10-bit and modulo-1024 GPS week constructs; decodable while era context is intact (see the deployment box above)	awareness pointer to the owning groups
ESI / e-signature families	no local time encoding defined; semantics live in RFC 3161/5652/5280 profiles	dependency edge, not a clearance

Horizon dates give the last representable day for day-count constructs, with the first wrapped day the next one; for second-resolution boundaries the date names the day the boundary instant falls on. Two further families flagged by the search turned out to be artefacts of the search itself and are recorded as such; reading exists precisely to absorb them.

What a fitness check would involve

How the candidates were produced, stated plainly: every text-bearing deliverable was matched against a registry of time-related phrases (Appendix A); the matches were weighted; documents whose text carries time encodings but no registered horizon mention were flagged; the flagged set was ranked by how much of that time encoding the document **defines itself** rather than borrows; and the top of the ranking was read at the flagged clauses. The search itself analyses nothing; it produces a reading order, and the substance comes from the reading.

There is precedent for exactly this shape of work. The IETF's Year 2000 Working Group screened the RFC corpus in 1999 with a keyword filter and put every hit through manual review, stating outright that the hits "simply identify elements that need to be examined" (RFC 2626). The screen here differs mainly in ranking the candidates and in attributing an encoding to the document that defines it.

The volume picture (bar lengths on a log scale; the first two bars count the whole corpus, the lower three count ETSI-owned families only):

deliverables in corpus		67122
with machine-readable text		66302
ETSI-owned families: encodings present, rollover silent		1129
... which define a time encoding themselves		302
... strongly enough to read first (the head)		139

Reading so far: of the 223 specification families queued across the reading passes, 142 have been read and recorded; the rest remain unread candidates. For scale, a plain keyword search over the same corpus returns 771 families, without any ordering; the weighting exists to give a reading order.

What the reading has bought, and what finishing would cost. The most prominent candidate here, the DVB MJD field, already has a live work item at its definer, so it is fair to ask what the screen added beyond a

finding the maintainers had reached on their own. Three things. The head has been sampled rather than read out, so most of the candidate list is still candidates, and the companion's gate table gives the coverage stratum by stratum. The dependency findings are separate from that work item and sit with different owners, among them the dated NTPv3 and RFC 868 references and the era-less 1900-epoch timestamp in 3GPP charging records. And two of the outputs are corpus-level rather than per-finding, the horizon catalogue in Appendix A.1 and the search vocabulary in A.3, which outlast any single candidate and transfer to another body's corpus.

The cost of finishing is the number a work-item decision actually turns on. At the fifteen-minute clause-read budget used here, reading out the whole head, 139 families, is on the order of a working week, and extending to every family that defines a time encoding of its own, 302 families, is on the order of two. That is the scale being proposed. It is also why the ranking is the contribution rather than the filtering: the same budget spent on all 1129 flagged families is nearer seven weeks, and on the 771 an unordered keyword search returns nearer five, and neither ordering nor attribution comes with them.

Outputs. A check run at this depth produces, per family, a short record with the clause reference, what was seen there, and, where applicable, a pointer to where it could be handled; across the corpus, a horizon catalogue (which formats run out when, at what width) and a reusable search vocabulary that other technical bodies can apply to their own corpora.

Possible ways forward. Whether and how a confirmed candidate is addressed is for the owning group to decide. The mechanisms that exist range from a corrigendum or revision at the defining document, through implementation guidance and notification of dependent documents, to a liaison item where the text is maintained upstream (3GPP, CableLabs, IETF), or a documented decision to accept a distant horizon. The definer-and-ring structure seen in the DVB family lends itself to coordinated handling: a change at the defining document reaches the importer ring through the existing references, while native re-specifications would need edits of their own.

Neighbouring corpora

ETSI's temporal constructs do not stop at ETSI's borders. Deliverables import time types from ITU-T ASN.1, from IETF protocols, and from GPS and broadcast conventions; 3GPP and CableLabs texts are published under ETSI numbers; and the same constructs recur, at different widths, across bodies. For calibration, the same search was therefore also run over neighbouring corpora: the ITU-T corpus (5 734 documents), the RFC series, and a set of web-platform specifications. Three observations from that work belong in this document; the rest is method material.

On shared constructs: the reference structure concentrates particular time vocabularies. The ITU-T documents that ETSI deliverables cite lean strongly toward the ASN.1 time vocabulary, `GeneralizedTime` and `UTCTime`, relative to the ITU-T corpus as a whole, which is consistent with inherited constructs arriving along reference paths. The rates and their ratios are in the data package rather than here: separating citation selection from inheritance would take an ablation over the reference edges that was designed and not run, and without it the figures describe a citation-selected snapshot and settle nothing.

On vocabulary: communities write time differently. The RFC corpus surfaced IETF-native constructs, such as `TimeTicks` and the NTP timestamp idiom, that ETSI-tuned search patterns initially missed, and web-platform time types have no ETSI footprint at all. No single vocabulary list covers the temporal encodings ETSI defines and inherits; a fitness check's search vocabulary has to be built from the corpus's dependency paths, not from any one community's terms.

On awareness: among the ITU-T documents raised as candidates, one instance of epoch-rollover awareness discourse was detected and verified on reading (ITU-T Q.814). That says at least as much about the search

vocabulary’s reach into another community’s language as it says about that community, and it supports no ranking of standards bodies.

What this look does and does not establish

- Everything named is a **candidate for expert review**. The reading behind it is careful but non-specialist; readers without standard-specific experience can miss context an expert would catch, such as companion specifications or deployment practice. Expert review by the owning group is the necessary next step.
- All counts are **lower bounds on candidates**, not counts of confirmed cases. “Silent” means no phrase in the registry matched; “ambiguous” means the encoding is present but leaves its width or era context unstated, as in the TS 103 770 case. The registry is finite and demonstrably incomplete, and verified silence exists only where a clause was actually read.
- The registry was **revised during this work**, using the same documents that were read by hand. Its coverage is therefore calibrated against those reads rather than tested against an independent set, and no out-of-sample check was run.
- **A missing flag is not a clearance**. No statement is made about documents that were not flagged or not read. Coverage is depth-stratified and disclosed; formal artefacts such as ASN.1 modules, XML schemas, and bit-field tables are invisible to a prose scan.
- The screen is deliberately **tolerant of false alarms**: lexical matching fires on unrelated constants and vocabulary, and reading absorbs those cases (two such families are recorded).
- **A wire encoding is not a stored encoding**: the same 32-bit time that wraps on the wire may be re-encoded on retention, and the reverse. Stored-record questions need the storage specification.
- Editing a specification is **neither necessary nor sufficient** for field safety; what the screen locates is where a coordinated fix would begin.
- Reproduction is on-premises: the corpus is not redistributed with this work, so verification means re-running the pipeline against a local corpus checked by content hashes. The pattern registry is reproduced verbatim in Appendix A; the scored outputs, the per-file manifests and the catalogue the corpus figures above reconcile from exist as a separate data package, licensed CC BY 4.0, which carries no document text. It is published at <https://lab.fukami.eu/epoch-scan/> as epoch-scan-data-2026-07.zip, sha256 f30f8ef73b5e3b25f6ed44e85ed336ecbb782e3f1d0a267433ad3c5e45df4acb.

Appendix A: Search vocabulary and pattern registry

A.1 Time constructs by lineage

The constructs that the scan flagged and that were read at the clause, tabulated by lineage. Defined-by and used-by are different relations: the definer fixes the semantics and the width, and carriers import both. The rows record normative semantics; the search surface that finds them is the registry in A.3, and the two must not be conflated. As in the body text, horizon dates give the last representable day for day-count constructs; for second-resolution boundaries the date names the day the boundary instant falls on.

Construct	Defined by	Epoch and unit	Width	Horizon and boundary behaviour	Carried by (examples)	Registry patterns
Broadcast MJD with BCD time (“UTC_time”)	EN 300 468, Annex C	days since 1858-11-17 (MJD), plus 24-bit BCD time of day	16-bit unsigned	last day 2038-04-22 (MJD 65 535); the next day decodes as 1858-11-17	re-specified by TS 103 197 and TS 102 474; imported by TS 102 006 (software-update scheduling), TS 102 323, TR/TS 101 211 and further ring members	mjd
DAB date field	EN 300 401 (FIG 0/10)	days since 1858-11-17 (MJD)	17 bits, stated range 0 to 99 999	2132-08-31 (MJD 99 999); carries leap-second and UTC-form flags	DAB family	mjd
DARC MJD	EN 300 751; format detail via EN 62106, outside the ETSI corpus	days since 1858-11-17 (MJD)	17-bit binary	2217-09-27 (MJD 131 071)	DARC data broadcasting	mjd
DRM MJD	ES 201 980 (integer form); TS 102 349 (fractional form)	days since 1858-11-17; fractional part in TS 102 349	17-bit binary range (the fractional form is carried in a 32-bit integer)	2217-09-27 (MJD 131 071)	DRM broadcast and receiver profiles	mjd
NTP era-0 seconds	IETF: RFC 1305, later RFC 5905	seconds since 1900-01-01	32-bit unsigned	era 0 ends 2036-02-07; the value wraps by design and stays decodable only with era context	TS 129 061 (RADIUS accounting records, no era field); TS 129 274 (wire); dated NTPv3 references in DVB and cable families	ntp_era, epoch_1900, ntp_ts_format
Time Protocol value	IETF: RFC 868	seconds since 1900-01-01	32-bit unsigned	serviceability bounded at 2036 by the RFC itself	EN 302 878-4 and TS 103 311-3 (cable-modem time of day)	ntp_rfcs
LI X2/X3 PDU timestamp	TS 103 221-2, clause 5.3.10	seconds and nanoseconds since 1970-01-01	two unsigned 32-bit fields	wire horizon 2106; the clause’s NOTE defers the POSIX timespec 2038 limit instead	X2/X3 internal interfaces, mandated for 5G LI by TS 133 128	set aside as aware by the screen; found by hand; the four missed surface forms are not in the registry used for this scan

Construct	Defined by	Epoch and unit	Width	Horizon and boundary behaviour	Carried by (examples)	Registry patterns
Width-unspecified Unix time	TS 103 770 (start_unixtime, end_unixtime)	seconds since 1970-01-01	width not stated	implementation-defined; a horizon exists only per implementation	DVB-I service discovery	unix_time
GPS week counter	GPS interface specification, outside the ETSI corpus	weeks since 1980-01-06	10 bits, modulo 1024; almanac reference week 8 bits, modulo 256	wraps every 1 024 weeks (about 19.6 years); decodable while era context is intact (see the deployment box)	TETRA GNSS assistance (TS 100 392-18-4), SES A-GNSS, transposed positioning-test families	gps_week
DOCSIS 3.1 extended timestamp	TS 103 311-3, clause 7.1.5 (CableLabs-authored)	counts at 10.24 MHz since 1970-01-01 TAI (the PTP epoch)	55 effective bits in an eight-byte structure	≈ 2081-06; epoch and method disclosed, horizon unstated	DOCSIS 3.1 timing	bit32_time, unix_time
ASN.1 UTCTime	ITU-T X.680/X.690, outside the ETSI corpus	calendar text, two-digit year	text form	century resolved by profile convention; RFC 5280 encodes dates through 2049 as UTCTime and later dates as GeneralizedTime	certificate and signature families (ESI) via RFC 3161/5652/5280 profiles	utctime
ASN.1 GeneralizedTime	ITU-T X.680/X.690, outside the ETSI corpus	calendar text, four-digit year	text form	no near-term boundary	same families; concentrated along ETSI citation paths (see “Neighbouring corpora”)	generalizedtime
JWT NumericDate	IETF RFC 7519	seconds since 1970-01-01	JSON number, no fixed width	no fixed-width horizon by design	TS 103 479 (PASSporT profile)	epoch_1970

A.2 What the registry is known to miss

The registry is finite and demonstrably incomplete; “silent” in the body text means that no registered phrase matched, nothing more. The misses found so far are recorded with the registry, and the recurring classes are listed here because they show what a fitness check’s search vocabulary has to absorb. Per-pattern firing counts across the corpora are part of the accompanying data set.

- **Prose definitions.** A 1970-epoch definition written out in words, inside an ASN.1 module comment in an ITU-T security document (H.235.0), was missed by earlier registry versions; it was found on reading, and the wording is now registered. US-order and slash date forms such as “January 1, 1970” were likewise unmatched at first; registering them brought in 398 further ETSI-published documents, most of them 3GPP-transposed.
- **Surface forms.** TS 129 274 writes the era boundary as “7th of February 2036”; the ordinal form needed its own pattern (date_ordinal_epoch). The 3GPP era idiom (“era 0”, “era 1”, “prime epoch”) was invisible to early registry versions and was mined from the corpus and registered.
- **Documents set aside as aware.** TS 103 221-2 was missed on four surface forms at once, while its own NOTE kept it out of the silent set; it was found by the by-hand pass described in the body text; the four missed surface forms are not in the registry reproduced in Appendix A, so the scan reported here does not catch them.

- **Formal artefacts.** ASN.1 modules, XML schemas, and bit-field tables are invisible to a prose scan, comment text aside, as the H.235.0 case above shows (see “What this look does and does not establish”).
- **False alarms, absorbed by reading.** One “era” hit was an extraction artefact: the word “camera” split across a line break. Bare maximum-value constants such as 0xFFFFFFFF fire on CRC seeds and identifier sentinels, and a gate reader treats such a hit as corroborating a candidate rather than establishing one. The registry does not enforce that reading rule: int32_max and uint32_max are tier 1 on the native axis, so they do anchor a ranking. 17 of the 131 latest-edition ETSI-owned head families reach the head on those two patterns and nothing else. Two of them were read and cleared, EN 301 140-1 and EN 302 092-1, both at E_{nat} 21.0 and both in the stratified manual sample rather than a gate queue; the rest are unread, so the group is not shown to be free of false alarms. What is shown is that the reading rule and the scoring rule differ.
- **Year-2000 vocabulary scored as awareness.** The y2k pattern sits on the awareness axis, so a document matching it scores $A > 0$ and leaves the exposed-but-silent quadrant. That is the masking mode described for counter-wrap discourse, with one difference: there the matched text was at least about wrapping. 106 documents in the review quadrant, 15 families, are there on y2k alone, and 3 of those families carry $E_{\text{nat}} \geq 20$. All three are lawful-interception handover specifications, TS 101 671, TR 101 876 and the transposed TS 133 108, and in each the match is a filename-encoding legend of the form “00 = year 2000; 04 = month April”, not rollover discourse. TR 101 876’s legend is itself a two-digit year field, so the pattern scored a candidate exposure as evidence of awareness. The pattern is unchanged in the registry reported here.

A.3 Pattern registry

The pattern rows of the machine registry matched against the corpus (patterns.tsv, v3.2), reproduced without the registry’s internal change-log comments. The columns are id → tier → fm → axis → regex, and the regex is matched case-insensitively per zone. tier is a precision prior (1 strict, 2 medium, 3 generic). axis assigns the scoring semantics: aware maps to A, native/inherited/usage map to $E_{\text{nat}}/E_{\text{inh}}/E_{\text{use}}$, and cyclic maps to CY, a reader column scored into neither A nor E. fm tags the XSTP.epoch failure-mode class a pattern relates to; the assignment is a search-design aid, and the per-pattern rationale is kept with the registry’s change records.

```
epoch_rollover 1 FM1 aware epoch[- ]?roll[- ]?over
rollover 1 FM1 cyclic roll[- ]?over
wraparound 1 FM1 cyclic wrap[- ]?around
counter_wrap 1 FM1 cyclic counter (wrap|overflow)
integer_overflow 1 FM1 aware integer overflow(!s? attack).\{0,60\}\b(time|timestamp|clock|date|epoch|seconds|era|year)\b\{0,60\}\b(time|timestamp|clock|date|epoch|seconds|era|year)\b.\{0,60\}integer overflow(!s? attack)
y2038 1 FM1 aware year[- ]?2038|Y2K738|2038 problem
y2036 1 FM1 aware year[- ]?2036
y2106 1 FM1 aware year[- ]?2106
y2k 1 FM1 aware year[- ]?2000|\bY2K\b|millennium bug
gps_wnro 1 FM1 aware week number roll|\bWNR0\b|1024[- ]weeks?
ntp_era 1 FM2 aware NTP era|era number
int32_max 1 FM1 native 2147483647|0x7FFF[- ]?FFFF
uint32_max 1 FM1 native 4294967295|0xFFFF[- ]?FFFF
date_2038 1 FM1 aware 2038-01-19|19 January 2038
date_2036 1 FM1 aware 2036-02-07|7 February 2036
epoch_1970 1 FM2 native 1970-01-01|1 January 1970|Jan(uary)? 1,? 1970|\b0?1/0?1/1970\b
epoch_1900 1 FM2 native 1 January 1900
epoch_1980 1 FM2 native 1980-01-06|6 January 1980
```

```

ctime_types 1 FM1 native \btime_t\b\bclock_t\b|struct timespec|struct timeval
bit32_time 1 FM1 native 32[- ]?bit.{0,20}(time|counter|seconds)
seconds_since 1 FM1 native seconds since (1970|1900|midnight|the epoch|epoch)
unix_time 1 FM1 native (UNIX|Unix|POSIX|epoch) time
epoch_seconds 1 FM1 native epoch seconds
utctime 2 FM2 inherited UTCTime
generalizedtime 2 FM2 inherited GeneralizedTime
not_before_after 2 FM2 inherited notBefore|notAfter
this_next_update 2 FM2 inherited thisUpdate|nextUpdate
signingtime 2 FM2 inherited [Ss]igningTime
validity_period 2 FM7 usage validity period
two_digit_year 2 FM1 native two[- ]?digit year|century (field|digit)
mjd 3 FM1 native Modified Julian|MJD
mjd_truncated 2 FM1 native 1[0-6][[- ]?(bit|bits|lsb).{0,30}(mjd|modified julian)|(mjd|modified julian).{0,30}1[0-6][[- ]?(bit|bits|lsb)|1[0-6] lsb.{0,15}(mjd|julian|date)
dvb_tables 2 FM1 native time and date table|time offset table
dvb_tdt_tot 2 FM1 native \b(TDT|TOT)\b.{0,25}(table|section|dvb|descriptor|\bSI\b|pid)|(time and date|time offset|dvb|\bSI\b).{0,25}\b(TDT|TOT)\b
leap_second 2 FM2 usage leap[- ]?second|leap smear
tai 2 FM2 usage international atomic time|\bTAI[- ]?UTC\b|UTC[- ]?TAI|\bTAI\b.{0,30}(atomic|utc|leap|offset|scale|second)|(atomic|utc|leap|offset|scale).{0,30}\bTAI\b
iso8601 2 FM2 inherited ISO 8601|RFC ?3339
ntp_rfc3 2 FM3 inherited RFC ?5905|RFC ?4330|RFC ?1305|RFC ?868|RFC ?867
ntp 2 FM3 usage Network Time Protocol|\bS?NTP\b
nts 2 FM3 usage Network Time Security|RFC ?8915
ptp 2 FM3 usage IEEE ?1588|802\|.1AS|precision time protocol
rough_time 2 FM3 usage RoughTime
tsa_rfc3161 2 FM7 inherited RFC ?3161|time[- ]?stamp(ing)? (authority|token)|\bTSA\b
rfc5280 2 FM2 inherited RFC ?5280
nitz 2 FM3 usage \bNITZ\b
gnss_timing 2 FM3 usage (GPS|GNSS|Galileo|GLONASS|BeiDou).{0,30}(time|timing|week)
radio_clock 2 FM3 usage DCF77|\bMSF\b|WWVB
rtc 2 FM3 usage real[- ]?time clock|\bRTC\b.{0,25}(clock|time|counter|register|battery|oscillator)|(clock|time|battery).{0,20}\bRTC\b
holdover 2 FM3 usage holdover
clock_sync 3 FM3 usage clock synchroni[sz]ation
time_sync 3 FM3 usage time synchroni[sz]ation
time_distribution 3 FM3 usage time (distribution|transfer)
clock_drift 3 FM3 usage clock (drift|skew|offset)
system_clock 3 FM3 usage (system|wall|monotonic) (time|clock)
timezone 3 FM2 usage time[- ]?zone|daylight saving
timestamp 3 FM2 usage time[- ]?stamp
replay 3 FM7 usage (anti[- ]?)?replay (protection|window)
freshness 3 FM7 usage freshness
cred_lifetime 3 FM7 usage (certificate|key|token) (expiry|lifetime|validity)
toctou 3 FM7 usage TOCTOU|time[- ]of[- ]check
trusted_time 3 FM7 usage (secure|authenticated|trusted) time
signedness 2 FM1 native \b(un)?signed[- ]?((8|16|32|64)[ -]?bit|integer|int|long|binary|number|value)\b.{0,40}\b(time_t|time(stamp)?s?|seconds|epoch)\b\b(time_t|time(stamp)?s?|seconds|epoch)\b.{0,40}\b((un)?
signed[- ]?((8|16|32|64)[ -]?bit|integer|int|long|binary|number|value)|two's[- ]?complement|sign[- ]?bit)\b\b(two's[- ]?complement|sign[- ]?bit)\b.{0,40}\b(time_t|time(stamp)?s?|seconds|epoch)\b\b
gps_week 2 FM1 native (GPS|GNSS|GANS|Galileo|BeiDou|GLONASS)[- ]?week|week number.{0,30}(modulo|mod(ulo)? ?1024|10[- ]?bit|13[- ]?bit)|10[- ]?bit.{0,45}week
longevity 2 FM4 usage \b(design|designed|operational|service|expected) life(time)?\b|in service for|\b[0-9]+[- ]?year (life(time)?|operation|deployment)\b
cadence 3 FM7 usage \bcron\b|annual(ly)? renew|renewal (date|period|cycle)|next renewal|(daily|weekly|monthly|annual) (batch|job|run)|batch (window|processing|schedule)
monotonic 2 FM3 usage monotonic(ally)?[- ]?(increasing|non[- ]?decreasing)|never (go(es)?|run(s)?|move(s)? backward|(time|clock|timestamp)s? (never )?(go(es)?|run(s)?|move(s)?|jump(s)? backward
version_ambiguity 3 FM2 usage \bNTPv[0-9]\b.{0,40}(compat|interop|coexist|migrat)|\b(version|era)\b.{0,20}(ambiguity|disambiguation|backward[- ]compat\b*w*.{0,40}\b(time(stamp)?s?|era|epoch|version field)\b
wrap_inflected 2 FM1 cyclic wrap(s|ped|ping) (around|(back )?to (zero|0))
seq_wrap 2 FM1 cyclic wrapped sequence|sequence number wrap
modulo_2_32 2 FM1 cyclic modulo 2\^(16|32|64)|modulo 65[ ]?536
serial_arith 2 FM1 cyclic serial number arithmetic
paws 2 FM1 cyclic protection against wrapped

```

```

timeticks 2 FM1 inherited TimeTicks|sysUpTime
ntp_ts_format 2 FM2 inherited NTP timestamp format
epoch_field 2 FM2 native epoch (field|number)
hundredths_sec 3 FM1 native hundredths of (a )?second
date_ordinal_epoch 2 FM1 aware [0-9]{1,2}(st|nd|rd|th) of (January|February) (2036|2038)
domhighrestimestamp 2 FM3 native DOMHighResTimeStamp
epoch_timestamp_ms 2 FM2 native EpochTimeStamp
widl_int64_time 2 FM2 native (unsigned )?long long \w*(time|stamp|duration|instant)
epoch_1582 1 FM2 native 1582-10-15|15 October 1582|Oct(ober)? 15,? 1582
ptp_token 3 FM3 usage \bg?PTP\b
idl_int64_time 2 FM2 native typedef (unsigned )?long long \w*(time|duration|stamp|instant)|TimeBase::TimeT

```

Appendix B: Verification record

What was checked, when, and with what outcome. The records named in the last column are project records; they are kept with the pipeline and the per-candidate evidence and are available on request rather than reproduced here. No entry implies expert adjudication; the ceiling on every check is the one stated in “What this look does and does not establish”.

Date	Check	Outcome	Record
2026-07-11	Stratified manual sample: 75 families read as a precision check on the screen	reads recorded; in the 15-family head stratum: 1 confirmed native limit, 1 width ambiguity, 6 safe by reference, 5 registry false alarms, 2 out of scope or safe horizon; too few confirmed cases for a formal precision estimate, false-alarm classes documented	manual-sample record
2026-07-11 to 15	Clause-level reads across five queues (native head, boundary entrants, NTP-boundary addendum, classic DVB/MJD family, 3GPP-transposed head) plus the awareness-set-aside stratum	223 distinct families queued, 142 read and recorded and 8 queued with no extractable text; all verdicts carried as candidates	gate-queue records
2026-07-14	Fresh-context re-read of the highest-consequence candidates	corrections applied to the records (see Appendix C)	review records
2026-07-15	Independence check on the WKI 78272 corroboration: could the cited work item have been induced or tuned for by this scan?	flag shown independent (the surfacing patterns carry no rollover terms; absence of awareness matches is the signal); corroboration role confirmed	validation-integrity record
2026-07-16 and 22	Temporal backtest of the method against a historical rollover-handling case (PKIX/X.509)	design recorded; the pre-check falsified the frozen registry, which matches neither SDO’s fix vocabulary, so no backtest was run and no in-corpus subject exists	backtest records
2026-07-24 to 25	Data-set integrity: per-corpus manifests hashed and sample-verified; the ITU-T citation-neighbourhood ratios reproduced from the packaged data alone	verified; both ratios reproduce	package build records, checksum list

Date	Check	Outcome	Record
2026-07-23 to 27	Adversarial review rounds on this text: two committee-register rounds, two external model reviews, one fresh-context ground-truth review of the annexes; all horizon dates independently recomputed	findings recorded and addressed or adjudicated; arithmetic matched	review records, 2026-07 set
2026-09-08	Re-sourcing of the deployment-evidence block against the commissioned external investigation, dated 24 August 2026 and published by the operator on 2 September with its findings accepted	one detail this document had been printing was refuted and removed; the 1 024-week arithmetic was corroborated by the investigation's own November 2006 date; the separate regulatory proceeding was reframed as an open compliance investigation rather than a root-cause report	review records
2026-09-08 to 09	Cross-cutting checklist over the whole text, thirteen items, covering claim scope, checkable numbers, disclosure of what was designed and not run, and the caveat set	ten held; two were fixed (the UTCTime border rates were printed so the ratio could be checked, and now travel in the data package rather than the text, and the reference-edge ablation is disclosed as designed and not run); one genuine gap was admitted rather than closed, the circular calibration of the registry, which is now a stated limitation	checklist pass record
2026-09-09	Data-package integrity re-check against the copy already circulated, and a rebuild	every data file hashes identically to the circulated copy; two empty tooling directories that the archive had carried were removed, and the build now asserts their absence	package build records, checksum list
2026-09-20	Package rebuilt through the build script, not patched, to drop the working-draft label from README.md	exactly 1 of the 24 entries in SHA256SUMS changed, README.md, so every data file still hashes identically to the circulated copy; zip sha256 751b9f20..., superseded by the 2026-09-21 build below	package build records, checksum list
2026-09-21	Post-publication review of the package found families carrying gate_status = unadjudicated for which	six families move from unadjudicated to read-and-cleared: EN 302 878-4, TS 100 392-18-4, TS 101 733, ES 201 733, TS 101 903, TS 102 470-1. The	package build records, checksum list, gate queues

Date	Check	Outcome	Record
	a gate verdict with an evidence locator existed. The classic gate queue keys some rows in the source-filename form (30287804), the score tables use the catalogue form (302 878-4), and the package join reconciled whitespace only, so those verdicts were dropped. One such key had been patched individually in July; the normalisation itself was fixed and the package rebuilt	confirmed set is unchanged at six and no screening figure in either document moves. 2 of the 24 entries in SHA256SUMS changed, both ETSI score tables; zip sha256 d62ca953..., superseded by the build below	
2026-09-21 (second)	Review of the republished package found its family table selecting a different “latest edition” from the report: it ordered by the source path string while report/queries.sql and the ladder order by publication date. In the GSM-era release tracks a higher version number is often published earlier, so the path order selected a stale edition	468 of the 7 467 ETSI-owned families were taking an older edition than the report’s own figures use, which is why the package returned 1 057 gap families at the latest-edition unit against the report’s 1 067 and moved the head ranking. The selector now matches the report, the package reproduces 1 067, 284 and 131 and the head ranks 1, 3 and 13, and 1 of the 24 SHA256SUMS entries changed, the ETSI family table. Zip sha256 87500b3b..., superseded by the build below	package build records, checksum list
2026-09-21 (third)	Post-publication review re-derived this revision’s figures from the analysis store and from the published package, and tested two statements about reading coverage: that none of the head families resting on the two maximum-value constants had been gate-read, that 150 of the 223 queued families had been read, and how the	the first two were wrong. Two of those families reach the package as read-and-cleared, EN 301 140-1 and EN 302 092-1, read in the stratified manual sample; and the deduped gate counter treated “not screened” as a verdict, so 8 families with no extractable text counted as read and the figure should have been 142 from the start. The counter now splits as the per-stratum columns do and A.2 states the two clearances. No score or ranking changes. The package README.md still described the superseded path-	funnel producer, gate queues, build records

Date	Check	Outcome	Record
	package describes its own edition selector	string selector as a minor imperfection; it now states the publication-date rule and the 468 families the old one mistook. The package was rebuilt for that one line: 1 of the 24 SHA256SUMS entries changed, README.md, so every data file still hashes identically to the circulated copy, and the rebuild reproduces 1 067, 284 and 131 and the head ranks 1, 3 and 13. Zip sha256 ac1a28a5..., superseded by the build below	
2026-09-21 (fourth)	Whether a reader can re-derive from the package the corpus figures this document states: the total, the Published, Historical and On Approval split, and the step from catalogued to text-bearing	not from the package as it stood. The extraction manifest holds only the 66 424 deliverables an extraction record exists for and carries no status column, so the 67 123 total and the 697 set aside by status could not be checked against anything. The package now ships manifests/etsi-catalogue.tsv, one row per catalogued deliverable with its status, and every figure reconciles from it: 67 123 rows, of which 66 426 Published, 673 Historical and 24 On Approval; 66 424 with an extraction record; 66 303 carrying text and scored. It also shows that two catalogue rows are one deliverable under two work-item identifiers, so the corpus holds 67 122 distinct files under 67 123 rows. The package now carries 25 SHA256SUMS entries, the added file and README.md; zip sha256 770c2d2c..., superseded by the build below	package build records, catalogue
2026-09-21 (fifth)	What the catalogue showed once it shipped: whether its rows and the corpus figures in this document are counts of the same thing	they were not. The catalogue is keyed by work-item identifier and the ETSI work programme carries one file under two of them, TS 136 579-1 V15.5.0 under WKI 66069 and 66094, which the analysis store inherits as two rows. Every corpus figure here was a count of rows, so it overstated the corpus by one	package build records, catalogue

Date	Check	Outcome	Record
		document, and the package shipped that file twice in its per-document score table as a literal duplicate row. The figures now count distinct documents and five of them move by one: the corpus total to 67122, Published to 66425, the transposed cohort to 49010 catalogued and 49009 scanned, and the scanned total to 66302. Nothing about scoring, ranking or any family-level figure changes, since the duplicate sits inside one family. The score table is deduplicated and the catalogue carries its work-item identifier, so the two entries are visible rather than collapsed. Zip sha256 f30f8ef7..., the build published at the address on the title page	
standing	Quotation verification: every specification statement quoted here was read at the flagged clause before listing	per-candidate rationale recorded	evidence packets and gate queues

Appendix C: AI assistance and provenance

This annex records how AI tooling was used in the production of this document; the note in the front matter summarises it.

Tools and period. The work ran in July 2026. The scan pipeline itself, covering text extraction, pattern matching, scoring, and figure generation, is deterministic code with no AI component. AI assistance came from Anthropic Claude models used as the working assistant throughout; independent review passes on drafts of this document were additionally run against OpenAI models. The per-pass records, including model identifiers where recorded, are kept with the project records.

Automated stages. Corpus text extraction, pattern matching, scoring and ranking, and figure generation are deterministic pipeline code: the same inputs reproduce the same outputs, and the figures in this document are generated, not transcribed. AI assistance was used in curating the search vocabulary, in ordering and triaging the reading queue, and as drafting support for this document's text.

The reading step. Every statement about a specification quoted in this document passed a human read at the flagged clause before being listed: the clause was read in the deliverable text, the dependency direction was checked, quotations were verified against the source, and a rationale was recorded per candidate. The readers are careful readers, not specialists in the individual standards, which is why everything named in this document is a candidate for expert review rather than a finding. A second, independent read over the highest-consequence candidates produced corrections to the records before this document was drafted.

Limits of reader attribution. Reading is attributed per stratum, not per verdict, and the record cannot be reconstructed to a finer grain. Execution self-documents and supervision does not, so which reads were carried out by the assistant under review and which by the author is recoverable for some queues and not for others. One stratum is recorded differently and says so where it is counted: the awareness-blocked stratum was assistant-read with its verdicts held as candidates pending the gate, of which one was promoted after a gate read. The methods companion prints the per-stratum attribution alongside the queue counts.

Data handling. All scanned documents are publicly published deliverables, processed on the author’s own systems. Excerpts of specification text were handled by the named AI services in the course of reading support and draft review; the scanned material is public, and no claim is made about the services’ internal handling of submitted text. Beyond the names carried in this document’s own title block, no personal data was sent to external services.

Attribution. The working records do not support a meaningful reconstruction of which individual sentences originated with the author and which with a tool, and no such ratio is claimed. Responsibility for the content rests with the author.