

Contents

1. Summary (five lines)
2. Sources
3. Completeness
4. Clocks and display
5. Software and who ran it
6. Coding
7. Ledgers and denominators
8. Exclusions
9. Quotation check
10. What this statement does not show
11. Data-handling annex
12. Reproduction
Signature block
Annex A: ledger tables
Annex B: codebook
Annex C: run manifest
Annex D: second-check results

DRAFT: for a solicitor's review before first use; Factum is a trading name of [Company]

Method statement

Statement of method (witness of fact): Example v Sample (custody of the dog Pixel) [synthetic]

Purpose. This statement says what was done to the material, how, by whom, with what, and what it does not show. It is written to be read in five minutes. The tool and method state the record; the solicitor cites the line and draws the conclusion. Nothing in this statement is an opinion on any person's honesty, motive or character, or on admissibility.

Matter: EX-0001 **Client reference:** [to be completed] **Instructing solicitor/party:** [to be completed]

Statement date: [to be completed] **Package version:** factum 0.1.0a0, run 64eb7d4964a76303 **Prepared by:** Factum, a trading name of [Company] **Handling lane:** private

1. Summary (five lines)

- Supplied: 3 export file(s) holding 194 message rows, from 14 September 2025 to 14 December 2025 on the court clock (Europe/London).
- Produced: an evidence table of 194 rows (minus any withheld), 13 ledger and register tables, a schedule and a report, each row tied to a source line and hash.
- Coding: 67 code rows accepted under codebook chattel-custody v0; coder file.
- Limits: source tier(s) T2, T3; 1 row(s) withheld; known gaps are listed in sections 2 and 3.
- The processing steps were run by an AI system under the signatory's supervision (section 5); the figures are reproducible under section 12.

2. Sources

Source id	Description (app, platform, who exported, when)	File name	SHA-256	Size	Format tier	Declared complete?	Known gaps/edits/deletions
SG-1	sigtop; exported by P-ADA; clock Europe/Copenhagen	inputs/signal-sigtop.txt	96db7ed81d819c44f1463e7afadbea4a5cc5139b4db9ab23da7c42aa1bc664c2	4740 bytes, 242 lines, 35 rows	T2	no	history before 2025-11-01 not in export
WA-1	whatsapp_ios_12h; exported by P-ADA; clock Europe/Copenhagen	inputs/whatsapp-ios-12h.txt	6ef41f15f670fb9e8fa7ebdb1cce6cfd91ffc2a54e9faf97aa7c33ccb9c97da	7775 bytes, 112 lines, 110 rows	T3	yes	no delivery or read times / original text of edited messages not in export

Source id	Description (app, platform, who exported, when)	File name	SHA-256	Size	Format tier	Declared complete?	Known gaps/edits/deletions
WA-2	whatsapp_ios_24h; exported by P-ADA; clock Europe/Copenhagen	inputs/whatsapp-ios-24h-group.txt	c6236ba901c8a3d8c3a68a7caa49dac6e03a691a069ff00a1a678a6d8a105079	2867 bytes, 49 lines, 49 rows	T3	yes	group export from Ada's phone; Ben's and Cal's devices not examined

Source	Format reference / parser status
SG-1	sigtop cmd/sigtop/fmt_text.go sha256 56b9a542c86cae1

Format tiers (a description of the export format, not of reliability): T1 forensic database extraction; T2 platform or vendor export; T3 user flat-text export. This run rests on tier(s) T2, T3. A T3 source cannot show delivery or read times, the original text of edited messages, deleted content or internal message IDs; findings that would need them are not made. Every row and every ledger line carries its tier. Hashes were calculated on receipt and again before processing; originals are held unchanged.

Each file is read as bytes and split into lines. Every line is assigned exactly one class: parsed (part of a message row), metadata (structure only, such as a blank separator), excluded (removed by a range declared in matter.yaml) or unresolved (the parser could not place it). The run checks, for every file, that each line has exactly one class and that the bytes of the lines sum to the file size; a failed check stops the run. The counts below are those totals.

Source	parsed lines / bytes	metadata	excluded	unresolved
SG-1	172 / 4631	70 / 109	0 / 0	0 / 0
WA-1	112 / 7775	0 / 0	0 / 0	0 / 0
WA-2	49 / 2867	0 / 0	0 / 0	0 / 0

Message text is kept in four representations: `text_raw` (the decoded body of the row, with carriage returns and any byte-order mark removed; the byte span of the row in the file and a SHA-256 of exactly those bytes are recorded), `text`, a form normalised for matching, and a display-safe form. Hostile or unusual characters (zero-width, bidirectional controls, replacement characters, mixed-script words) are shown, not removed, in the store; CSV and Markdown outputs write them as visible escapes and neutralise spreadsheet-formula cells. A line that looks like a message header but sits inside a message body is accepted as a new message only if it passes the genuine-header test for that format (the file's own LRM convention, time order, look-ahead against the following headers, and a sender already seen in the file or listed in matter.yaml); 1 header-like line(s) were folded into the preceding message in this run (WA-1: {`'header_like_no_lrm': 1`}). Message text is treated as untrusted data; instruction-like text is flagged by pattern; no model was called in this run.

3. Completeness

Client's completeness declaration (verbatim): "[to be completed from the client's instruction]". Declared per source in the table above (the exporter's declaration, as recorded in matter.yaml). Our checks: date range; message-number or timestamp discontinuities; gaps longer than 7 days; deleted-message and edited-message markers; media referred to but not supplied. Result: 2 deletion placeholder(s), 2 edited message(s), 2 media reference(s) in the rows; the runs of days with no message are listed in the report. These checks detect some gaps; they cannot show that nothing is missing.

4. Clocks and display

Source clock/time zone as declared by the Client, or as shown by the Source: export wall-clock times are read in the zone declared for each source (SG-1: Europe/Copenhagen, WA-1: Europe/Copenhagen, WA-2: Europe/Copenhagen; the zone of the exporting device, not verified against the device). Factum makes no inference. Display convention: the court clock (Europe/London) with the export clock beside it; 24-hour; ISO date. Time database: version 2026e. A time is stored as a set of candidate instants and a precision, never as a single chosen instant where the source does not fix one; the day and month used to group rows in the ledgers are taken from the earlier reading of the interval. Rows by ambiguity: `dst_fold` 5, none 189. Where a wall time occurs twice (the DST fall-back hour) or not at all (the spring-forward gap), the row keeps every reading, displays each, and latencies built on it are ranges. In the decision-window extracts, ordering across sources is marked unresolved where intervals overlap; elsewhere the ledgers order by the lower end of the interval. File

order was not used to narrow intervals. Precision (seconds) by row count: 1s 194. Where a Source has no seconds or time zone, intervals are stated to the resolution available.

5. Software and who ran it

Parser factum-parsers (factum-parsers 0.1); pipeline factum 0.1.0a0; run manifest id 64eb7d4964a76303 (digest 64eb7d4964a76303c8a1003a35cb5212fbc52fef6aeade3c06024d107388cf34); run date: see run_utc in run-manifest.json. Tests passed: [n/n: to be completed from make test].

The processing steps in sections 2 to 9 were executed by an AI system, Claude Code (Sonnet 5.5), operating Factum's pipeline under the supervision of the signatory. The signatory did not personally run those steps. The signatory personally checked: the run manifest; the quotation check in section 9; and a sample of [n] lines ([%] of the table), each compared by hand with its Source line: [n] matched, [n] differed ([notes]). The signatory did not check the remaining lines individually. A second coder was drawn: 25 of 66 code rows (37.9 %), under the rule minimum max(10 %, 25 rows) or all rows if fewer (required 25); the second coder's results have not been recorded. On re-run on [date] the outputs matched byte for byte [to be completed after the re-run].

Pinned invocation (exact): PYTHONPATH=/tmp/cleanrepo python3 -m factum run-all --matter /tmp/EX-0001-sample/M/matter.yaml --out /tmp/EX-0001-sample/03-Output --store /tmp/EX-0001-sample/02-Working/store --operator "Claude Code (Sonnet 5.5)"; code commit 0b99b50212af799e1cdec5a7fb79714ac7e79edb.

Re-running the pipeline on the same inputs, codebook and archived coding decisions reproduces every CSV and Markdown output byte for byte, on the same Python (3.13.17) and time zone database (version 2026e).

Regenerating identical output from a language model is not guaranteed; the archived decisions are what is replayed. The run identifier is a digest, not a signature; signing is a later step.

6. Coding

Codebook pack and version: Codebook: pack chattel-custody version v0, SHA-256 5658ed038d85cbf741332f058bd44c5b438c71f09df5249fa165214388e285df (published schema in Annex B). Unit of coding: one message row. The decisions were loaded from files; Factum does not know how they were made, and the coder and model named below are as declared in those files, not verified. Candidate messages were selected in mode both (lexicon union semantic retrieval) with embedder local-lexical-v1-d4096-n345-w0.5-docs182-109e85115b; candidate threshold 0.3, top-k 25 per code per source. Retrieval is not the method: only coded and confirmed lines become exhibits. Silence ledger: a term is matched lexically; rows without a lexical hit whose similarity to the nearest seed of the term reaches 0.35 (embedder local-lexical-v1-d4096-n345-w0.5-docs182-109e85115b) are listed as possible references to review and never break a silence (semantic silence rule). Coder: file; model or source: synthetic-fixture (hand-authored by synthetic/generate.py; not a model); parameters: {"declared": "none"}. Decisions received 67; accepted 67; rejected 0 ({}); abstained 0; rows set by default rule 21 (of which 21 left uncoded, open and queued for review; a candidate with no decision or only rejected decisions is never severed by default).

Of the 67 accepted code rows, 67 carry character offsets that the system checked by slicing the raw message text (the sliced words must equal the quoted words, for every row kind; a decision whose span fails is rejected). 0 come from batch lists, which declare the whole body and are checked only for row kind. The 21 default rows carry a whole-body span that was not independently verified. Codes express what a message says, never a view on any person's state of mind. Pragmatic ambiguity (for example irony) is flagged on the row and lowers its weight in the register.

AI-assisted content. No: the coding decisions were loaded from files and no model classified any message in this run.

Second check: A second coder was drawn: 25 of 66 code rows (37.9 %), under the rule minimum max(10 %, 25 rows) or all rows if fewer (required 25); the second coder's results have not been recorded.

7. Ledgers and denominators

Every ledger metric has an executable specification, written to the manifest: the eligible population, exclusions, pairing rule, observation window and the treatment of unanswered items. State for each figure the lines it cites. A percentage is never given without its denominator. Intervals are between messages of different senders, or without recorded messages.

Ledger	What it counts	Eligible n (numerator)	Denominator and why	Excluded from the count	Pairing / rule and window	Unanswered
channels	rows meeting the eligibility rule	2	2: entity links between sources and channel-deflection markers	none	identical normalised text (3+ chars) in two sources; duplicate_observation = same conversation exported twice, not independent corroboration; forward_candidate = 8+ words in another channel; whole matter	not applicable
commitments	rows meeting the eligibility rule	7	7: messages coded commitment.*, exportable, canonical	none	state machine per commitment.made in its source: proposed -> accepted/refused/deferred (answer by another party within 72 h) -> withdrawn (commitment.reneged by the speaker) / asserted_performed (commitment.performed); evidence line per transition; rest of the source after the commitment	state stays 'proposed'
conditions	rows meeting the eligibility rule	2	2: messages coded condition.biting or subject.condition, exportable, canonical	none	next condition-coded message in any source; progress = next request.info or answer.granted after the row; whole export	'absent from this export', never 'never'
deletions	rows meeting the eligibility rule	4	4: delete events and edit events (current version or edit marker)	none	subject_context = a subject-coded message within 3 rows; whole export	a gap in IDs is not called a deletion without a placeholder (T3 exports carry no IDs)
gaps	rows meeting the eligibility rule	194	194: all rows of every source	none	calendar days on the export clock with no row; runs of 7 or more days listed with the prior 90-day baseline; first to last row of each source	a quiet run is reported as quiet in this export, not as inactivity
initiative	latency is a range from the two intervals; medians are given for both ends	54	54: subject-coded messages in pair (non-group) sources, exportable, canonical, not forwarded	none	next message (text, media or deleted placeholder) by the other member of the pair; calls do not count as replies; to the end of the source	censored: reported separately, never zero or infinite
latency	latency is a range from the two intervals; medians are given for both ends	54	54: subject-coded messages in pair (non-group) sources, exportable, canonical, not forwarded	none	next message (text, media or deleted placeholder) by the other member of the pair; calls do not count as replies; to the end of the source	censored: reported separately, never zero or infinite
money	rows meeting the eligibility rule	2	2: messages coded money.*	none	none; whole export	not applicable
ownership	rows meeting the eligibility rule	2	2: messages coded ownership.*	none	none; whole export	not applicable

Ledger	What it counts	Eligible n (numerator)	Denominator and why	Excluded from the count	Pairing / rule and window	Unanswered
requests	delay is a range [lo, hi] from the request interval and the answer interval	6 (answered_within 3, none 3)	6: messages coded request.* that are not severed, not forwarded/pasted and are the canonical copy of their entity	none	first later message by an addressee carrying an answer.* code; later repeats of the request do not reset it; group requests start the clock of every member present at the time; 72 h after the request interval	right-censored when the window ends after the last message of the source; otherwise 'none' (not zero, not infinite)
requests_robustness	rows meeting the eligibility rule	6	6: same population as the request ledger	none	answered share re-run under windows 24/72/168 h and under 'first answer-coded reply' v 'any reply by an addressee'; robust if the share varies by 0.10 or less;	
silence	rows meeting the eligibility rule	181	181: text-kind rows of every source	none	regular expression per term over normalised text (lexical rule). Lexical rule: the term's regular expression over the normalised text of every non-severed text row. Semantic rule (local-lexical-v1-d4096-n345-w0.5-docs182-109e85115b): a row with no lexical hit whose cosine similarity to the nearest seed of the term is at least 0.35 is listed as a possible reference to review; it does not break the silence, and the absence statement stands unless a reviewer decides the row refers to the term.; whole export	a term with no lexical hit is reported as absent from this export with the known gaps, never as 'never'; semantic hits are possible references to review, not a break in the silence
windows	rows meeting the eligibility rule	49	49: rows of any source whose interval overlaps a decision window extended by the margin on both sides	withheld_severed 0	all sources on one clock; ordering across sources flagged where intervals overlap; pleaded court-clock dates +/- 48 h	not applicable

The request ledger was re-run under alternative windows (24, 72 and 168 h) and pairing rules; the answered share is rule-dependent (spread 0.5; n = 6 request(s); with so few requests the verdict is indicative only).

Assertion register. 11 pleaded assertion(s) were tested, those of every party by the same rules; by status: confirmed 2, contradicted 4, qualified 4, silent 1. confirmed: the export supports the assertion (for a 'never' assertion: no row is coded as the asserted code and a lexicon or regex search over the normalised text of every non-severed row in scope found nothing; the register states how many rows were searched, how many of those were coded and how many uncoded; with no lexicon the assertion is at most qualified); qualified: weakened by low confidence, a pragmatic flag, a reported rather than endorsed statement or a date boundary; contradicted: a message that cannot be reconciled with the assertion; silent: the export does not speak to it.

8. Exclusions

Severed topics (client instruction): business. 1 row(s) severed by an explicit code (and their duplicate or forwarded copies) are withheld from every export, ledger, register, schedule and report and are not searched by register tests; 0 unread row(s) matching the severed-topic lexicon are held (their words are withheld, they are not severed). Withheld rows remain in the store, which is kept beside, not inside, the output folder. No line ranges were excluded by declaration. Rows are counted once as a second export of the same conversation only when the text is identical, the sender is the same and the times are within 120 seconds on the clock declared for each source (they are not independent corroboration). Privileged or excluded channels: none declared in this run unless listed above. Messages not processed (media, unreadable, non-text): media and call rows are kept as placeholders and are not opened.

9. Quotation check

Every quotation in the Deliverables was compared mechanically with its Source line. Every quotation of six or more words in the report and the schedule was checked against the message text: 17 checked, 17 exact, 0 difference(s) (listed in quotecheck.txt). Corrections made: [to be completed if any].

10. What this statement does not show

It does not show that the Sources are authentic or complete beyond the checks in sections 2 and 3; who wrote or read any message beyond what the Source records; why anything was said or not said; or whether any item is admissible or what weight it should carry. Intervals are counts of recorded time; they have no meaning in themselves. "Evidence table" describes the table's purpose, not its status before the court. This is not expert evidence and the signatory gives no opinion. If the court requires expert evidence it should be obtained under CPR Part 35.

Coder quality. Gold set supplied: 184 adjudicated item(s) from gold.yaml (SHA-256 eefba12d714ef4efbbf32a2d8044035f8435af0e17d8ac3866eb375541e547c9), scored against the code table of this run. Micro precision 1.000, recall 1.000, F1 1.000; macro F1 over 17 code(s) with support 1.000; abstention or uncoded rate 0.250 of 84 item(s) asked; retrieval recall of the pre-filter 0.952 (60 of 63 gold-coded items). The gold set is synthetic and was derived from the codes planted by the generator that wrote the corpus: these figures test the pipeline, they do not measure any model. Per-code figures are in qa/gold-score.csv, by-stratum figures in qa/gold-strata.csv and each false positive and false negative in qa/gold-errors.csv. The figures describe the items in that file and no others; a different matter, language or platform needs its own gold set.

11. Data-handling annex

Where stored: [bucket region, e.g. EU jurisdiction: to be completed]. Access: [operators, dates; access log extract]. Engines: deterministic pipeline on [host/region]; no AI model received any message text in this run; OCR [provider/none]. Masking applied: [yes/no; what]. Files transmitted to any model: none. Retention: matter close + [30] days unless held. Deletion date: [to be completed]; certificate ref: [to be completed from factum certify]. DPA ref: [to be completed]. Controller / processor and lawful basis: [Company] as [controller or processor: to be completed].

12. Reproduction

The Client's expert, or the court's, may reproduce the figures from the Sources using the pipeline version in section 5, on request, on terms of confidentiality. The exact command is `PYTHONPATH=/tmp/cleanrepo python3 -m factum run-all --matter /tmp/EX-0001-sample/M/matter.yaml --out /tmp/EX-0001-sample/03-Output --store /tmp/EX-0001-sample/02-Working/store --operator "Claude Code (Sonnet 5.5)"` at commit 0b99b50212af799e1cdec5a7fb79714ac7e79edb. Run identifier 64eb7d4964a76303 is the first 16 hex digits of the manifest digest 64eb7d4964a76303c8a1003a35cb5212fbc52fef6aeade3c06024d107388cf34 (SHA-256 of the run manifest without its timestamp, invocation and operator).

Signature block

I am responsible for this statement. The facts in it are true to the best of my knowledge and belief. Section 5 states which steps were run by an AI system and which I checked myself; I have stated where I have relied on others.

[Optional where used in proceedings, in the prescribed form (PD 22): "I believe that the facts stated in this statement are true. I understand that proceedings for contempt of court may be brought against anyone who makes, or causes to be made, a false statement in a document verified by a statement of truth without an honest belief in its truth."]

Name: _____ Role: _____ for Factum, a trading name of [Company]

Signature: _____ Date: _____

Annex A: ledger tables

The ledger tables are the CSV files listed in Annex C (ledger_*.csv, register.csv, evidence_table.csv); the denominators are in section 7.

Annex B: codebook

Pack chattel-custody, version v0, SHA-256 5658ed038d85cbf741332f058bd44c5b438c71f09df5249fa165214388e285df. The published schema is the codebook file named by that hash.

Annex C: run manifest

File	SHA-256	Bytes
accounting.csv	0324cdb02e7e0c36dc05880b6aa17b6244203e391908f0bd96643f7f4dcf1be3	736
audience.csv	69b2b3a94a498229fc589319f8bd7785528e561c675473b8a492d94f4a7e66ca	538
denominators.json	3539ad506b6b8848fc97e6fcaed1907e5e98288cf7e32f4fdd9f0fdadb98d9500	6862
evidence_table.csv	8c8f9e050e5bebaea5624d5292b83562ace15fb5c7313d8d4b1511e4fe198442	22247
ledger_channels.csv	cbbdb49a7a4584818ccb6dcf62fab58b78c7188558080557306df5a4f640a4ce	601
ledger_commitments.csv	a633d9c26e1046bf04885f1946472337a88ac92a8285253eb1b0ad748b862149	1846
ledger_conditions.csv	22462bd1b0b5b0e0b5156c4e7fa5bcf076cd8e1b1f0ec0773ce76c32e3a1e151	680
ledger_deletions.csv	19af25e733533968fee62c7ca083b4777530b8a29ab677ce97688e98d9f13a91	910
ledger_gaps_months.csv	dbba3a74f3ab6b830d1ae3fb3379e3b2c955128ae811d9335f5624acb07a4860	237
ledger_gaps_runs.csv	3b7981ec98fab5dae702f3885fa0bbf83608fe6ac731746020ba7fc1cce94692	86
ledger_gaps_sources.csv	839086b5e78014d8ebd8c7a77d41efc1ac85989b4eaaad714bd3eee20a069663	363
ledger_initiative.csv	88c3296141a3e6305e34efe67176c644907682a31005e623e38ca8bb4c3dbffc	799
ledger_latency.csv	386bcd85572c12ec3367f5a17c24e8703a7a7d8909d70a6a97b40d1257d13d26	8932
ledger_latency_summary.csv	2597383ec13d5db76a5399c5be74a03d5b04c8ae2e993fdbcb39da21087200244	139
ledger_money.csv	b06f17eb90638af35c9c8fa6ad437d16be82c66fdc539ff12d62a9a68bc257aa	449
ledger_ownership.csv	65e275b8170e684f79c041db250e39e5fe4719247ea0d12a7bcab54a5a151dfc	461
ledger_requests.csv	6633132aec949dd9f9d50eac8c74c4e331086727c2cd450193130093d6e67118	2376
ledger_requests_robustness.csv	7894f8b85d226f8805ac06386b12837c0a07805aee0ca5936259e272a7ead8fc	384
ledger_silence.csv	62ee52a4335af2834c56fb015adfd8d381c894dc33c6d91df2c1e7d6dc105573	1442
ledger_windows.csv	393c907864a0bef03e581495a1f739119294ef1d51bf7342ec3e0680187af574	10033
ledger_windows_summary.csv	09fd72fae33f4f7f21c366f991a3274e4506a012a2a9ca6bfb10c20637c38090	733
paragraph_issue_map.csv	8d84579e8cb014640317fc3227e0a68188bd5a9626da330773716b4967c887c4	498
qa/gold-errors.csv	d1755dfa80d1711edeb0c4fb2f4614d5c72639b399d67e7a5713d1c667541bcb	19
qa/gold-score.csv	9346a9ef1505c0c38f260f323b3447edac56f575516efb257a25fa473d2d3573	729
qa/gold-strata.csv	30d534571baabf7297539a9c0ed3c32a227405ff4b29dc3691b29a5c5161f738	455
quotecheck.txt	01c8a7591483fff1b97c7f659015c667b7a037ad071df62a6e996f7df2feffdd	87
register.csv	f1ce22ab6a30c3842b464668012753b620dd6e8f4b81caccff60a7ec52f8d1f7	3102
register.md	3d1a76fd9f88329cbc6dde71da64703eee38b3220ebe03ee65527f8e2c0c70d5	10411
register_evidence.csv	1b82fd58782b0e8d60dd5ebb35361bc2719e1454ac5eaa5ae14ab3b7799539cd	3869
report.md	704d0e1f9c122a4546414dcfffc58651e5f69786a799db0647ee23716a99c6740	4893
schedule.md	ae07b09bd473bd04a3d635006176a07ef16a3db1df2238d2a623a26ac80c220a	9059
sources.csv	1d14289696b8ae6afb063cd59178b078f67a50cc320c9be5e81990057bcdde22	1180

Quality statistics: gold set supplied: 184 adjudicated item(s) from gold.yaml (SHA-256 eefba12d714ef4efbbf32a2d8044035f8435af0e17d8ac3866eb375541e547c9), scored against the code table of this run. Micro precision 1.000, recall 1.000, F1 1.000; macro F1 over 17 code(s) with support 1.000; abstention or uncoded rate 0.250 of 84 item(s) asked; retrieval recall of the pre-filter 0.952 (60 of 63 gold-coded items). The gold set is synthetic and was derived from the codes planted by the generator that wrote the corpus: these figures test the pipeline, they do not measure any model. A second model is a check, not the validation: it shares failure modes with the first.

Annex D: second-check results

A second coder was drawn: 25 of 66 code rows (37.9 %), under the rule minimum max(10 %, 25 rows) or all rows if fewer (required 25); the second coder's results have not been recorded.